Gabor Fichtinger
Anne Martel
Terry Peters (Eds.)

LNCS 6892

Medical Image Computing and Computer-Assisted Intervention – MICCAI 2011

14th International Conference
Toronto, Canada, September 2011
Proceedings, Part II

Springer

Lecture Notes in Computer Science 6892

Commenced Publication in 1973
Founding and Former Series Editors:
Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Gabor Fichtinger Anne Martel
Terry Peters (Eds.)

Medical Image Computing and Computer-Assisted Intervention – MICCAI 2011

14th International Conference
Toronto, Canada, September 18-22, 2011
Proceedings, Part II

Springer

Volume Editors

Gabor Fichtinger
Queen's University
Kingston, ON K7L 3N6, Canada
E-mail: gabor@cs.queensu.ca

Anne Martel
Sunnybrook Research Institute
Toronto, ON M4N 3M5, Canada
E-mail: anne.martel@sri.utoronto.ca

Terry Peters
Robarts Research Institute
London, ON N6A 5K8, Canada
E-mail: tpeters@robarts.ca

ISSN 0302-9743 e-ISSN 1611-3349
ISBN 978-3-642-23628-0 e-ISBN 978-3-642-23629-7
DOI 10.1007/978-3-642-23629-7
Springer Heidelberg Dordrecht London New York

Library of Congress Control Number: 2011935219

CR Subject Classification (1998): I.4, I.5, I.3.5-8, I.2.9-10, J.3, I.6

LNCS Sublibrary: SL 6 – Image Processing, Computer Vision, Pattern Recognition, and Graphics

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Preface

The 14th International Conference on Medical Image Computing and Computer Assisted Intervention, MICCAI 2011, was held in Toronto, Canada during September, 18–22, 2011. The venue was the Westin Harbour Castle Hotel and Conference Centre on the waterfront of Lake Ontario in Downtown Toronto, the world's most ethnically diverse city.

MICCAI is the foremost international scientific event in the field of medical image computing and computer-assisted intervention. The annual conference has a high scientific standard by virtue of the threshold for acceptance, and accordingly MICCAI has built up a track record of attracting leading scientists, engineers and clinicians from a wide range of technical and biomedical disciplines. The year 2011 saw a record 819 paper submissions.

The Program Committee (PC) of MICCAI 2011 comprised 53 members. Each of the 819 papers was assigned to two PC members (a primary and a secondary) according to their expertise and the subject matter of the paper. The primary member knew the identity of the authors, but the secondary one did not. Each PC member had about 17 papers as primary and a further 17 as secondary member. The primary PC member assigned at least three external reviewers to each paper, according to their expertise and the subject matter of the paper. The external reviewers provided double-blind reviews of the papers, and authors were given the opportunity to rebut the anonymous reviews. In cases where reviewer opinions differed significantly and/or the rebuttal made it necessary, the primary member initiated a discussion among the reviewers. The primary member summarized the outcome of the discussion in a short report for the secondary. Finally, the secondary member considered all input (the reviews, rebuttal, discussion, primary's report, and, almost importantly, the paper itself) and made a recommendation for acceptance or rejection. The secondary PC member did not know the identity of the authors.

A two-day PC meeting was held with 33 of the PC members present. Each paper received fair consideration in a three-phase decision process.

- First stage: Initial acceptance of papers ranked very high by both the reviewers and the secondary PC member. Initial rejection of papers ranked very low by both the reviewers and the secondary PC member.
- Second stage: groups of five to seven PC members ranked the remaining papers and again selected the best papers and rejected the lowest ranking papers.
- Third stage: a different set of groups selected the best papers from the remaining undecided papers and rejected the rest.

The PC finally accepted 251 papers, giving a 30% acceptance rate.

We are greatly indebted to the reviewers and to the members of the PC for their extraordinary efforts assessing and evaluating the submissions within a very short time frame.

In 2011, attendees saw two changes in the way the program was organized. All accepted papers were presented as posters, and a subset of these were also invited for oral presentation, which were organized in clinical themes rather than by methodology as in earlier years. Poster sessions were organized in their traditional technical themes as in the past.

In addition to the main 3-day conference, the annual MICCAI event hosted an increased number of satellite tutorials and workshops, organized on the day before and the day after the main conference. This year's call for submission for tutorials and workshops led to a record 21 workshops and 8 tutorials accepted by a committee headed by Randy Ellis (Queen's University) and Purang Abolmaesumi (University of British Columbia). The tutorials provided a comprehensive overview of many areas in both the MIC and CAI domains, offering a unique educational forum for graduate students and postdoctoral fellows. The workshops presented an opportunity to present research, often in an early stage, to peer groups in a relaxed environment that allowed valuable discussion and feedback. The workshop subjects highlighted topics that were not all fully covered in the main conference, and thus added to the diversity of the MICCAI program.

In reviewing the proposals for these events, emphasis was given to workshop submissions that provided a comprehensive and interactive forum to address an open problem in MICCAI. We also promoted tutorials that related to an existing sub-discipline of MICCAI with known materials, approaches and open problems to help train new professionals in the field. Among the accepted workshops, several focused on emerging trends in the field of multi-modal statistical atlases, advanced computational and biomechanical models, and high-performance computing. MICCAI 2011 also hosted eight tutorials that spanned a wide spectrum of topics in basic and advanced software development for medical image analysis, algorithms for image segmentation, registration and visualization, as well as those highlighting new techniques in image-guided interventions. We would like to thank the Workshop and Tutorial Committee for their hard work in putting together such a comprehensive and unique program.

Two of the highlights of the conference were the keynote lectures by two Canadian scientists. Dafydd (Dave) Williams, physician, astronaut, medical robotics researcher, and recently, Hospital CEO, opened the conference with a presentation that looked at lessons that the health care system and medical researchers could learn from the challenges of space travel. The second keynote was given by Mark Henkleman, Director of the Mouse Imaging Centre, Toronto Centre for Phenogenomics, who spoke about high-throughput small-animal imaging techniques and quantitative statistical analysis methods for mapping phenotypic changes associated with genetic disease models in mice.

MICCAI 2011 would not have been feasible without the efforts of many people behind the scenes. We are particularly indebted to the local Organizing Committee in London and Toronto consisting of Janette Wallace, Johanne Guillemette,

Jackie Williams, Jade Orkin-Fenster, Debbie Lilley, Shuo Li, Perry Radau, and Raphael Ronen. In addition, we are deeply grateful to the Robarts Research Institute, the University of Western Ontario, Sunnybrook Research Institute, and Queen's University for their support in ensuring the success of this meeting, and to the staff at Springer for their continued high standards aimed at maintaining the MICCAI proceedings as the flagship of the LNCS series.

We thank the MICCAI Society Board for trusting us with the mandate to organize this conference, and to the Board and staff members for valuable and continuous advice and support through all phases of the project.

A special word of thanks goes to our sponsors, who generously provided financial support for the conference as a whole as well as for specific activities. This greatly assisted with the overall organization of the meeting, enabled us to continue offering best paper awards in various categories, and provided travel stipends to a significant number of student participants.

It was our great pleasure to welcome the attendees to Toronto for this year's MICCAI conference along with its satellite tutorials and workshops. Next year, the 15^{th} International Conference on Medical Image Computing and Computer-Assisted Intervention will be held in Nice, France, October 1–5, 2012. We look forward to seeing you all there.

September 2011

Gabor Fichtinger
Anne Martel
Terry Peters

Organization

Co-chairs

Gabor Fichtinger	Queen's University, Kingston, Canada
Anne Martel	Sunnybrook Research Institute, Toronto, Canada
Terry Peters	Robarts Research Institute, London, Canada

Workshop/Tutorial Co-chairs

Purang Abolmaesumi	University of British Columbia, Vancouver, Canada
Randy Ellis	Queen's University, Kingston, Canada

MICCAI Society, Board of Directors

James Duncan (President)	Yale University, Connecticut, USA
Gabor Fichtinger (Treasurer)	Queen's University, Canada
Alison Noble (Executive Director)	Oxford University, UK
Sebastien Ourselin (Secretary)	University College London, UK
Kensaku Mori	Nagoya University, Japan
Daniel Rueckert	Imperial College London, UK
Nicholas Ayache	INRIA, Sophia-Antipolis, France
Kevin Cleary	National Children's Hospital, Washington DC, USA
Polina Golland	MIT, Cambridge, USA
David Hawkes	University College London, UK
Tianzi Jiang	Chinese Academy of Sciences, Beijing, China
Wiro Niessen	Erasmus MC, Rotterdam, The Netherlands
Ichiro Sakuma	University of Tokyo, Japan
William (Sandy) Wells	Harvard Medical School, Boston, USA

Consultants to Board

Alan Colchester	University of Kent, UK
Terry Peters	Robarts Research Institute, Canada
Richard Robb	Mayo Clinic College of Medicine, USA

Program Committee

Purang Abolmaesumi	University of British Columbia, Canada
Elsa Angelini	Telecom Paris, France
Stephen Aylward	Kitware, USA
Christian Barillot	INRIA, France
Wolfgang Birkfellner	Medical University of Vienna, Austria
Albert C.S. Chung	Hong Kong University of Science and Technology, China
Ela Claridge	University of Birmingham, UK
D. Louis Collins	McGill University, Canada
Dorin Comaniciu	Siemens Medical, USA
Christos Davatzikos	University of Pennsylvania, USA
Marleen de Bruijne	Erasmus Medical College, The Netherlands
Hervé Delingette	INRIA, France
James S. Duncan	Yale University, USA
Philip (Eddie) Edwards	Imperial College London, UK
Bernd Fischer	University of Lübeck, Germany
P. Thomas Fletcher	University of Utah, USA
Alejandro Frangi	Pompeu Fabra University, Spain
Polina Golland	Massachusetts Institute of Technology, USA
Nobuhiko Hata	Harvard Medical School, USA
David Hawkes	University College London, UK
David Holmes	Mayo Clinic, USA
Ameet Jain	Philips Medical Systems, USA
Pierre Jannin	INRIA, France
Leo Joskowicz	Hebrew University of Jerusalem, Israel
Ioannis Kakadiaris	University of Houston, USA
Rasmus Larsen	Technical University of Denmark
Anant Madabhushi	Rutgers University, USA
Frederik Maes	University of Leuven, Belgium
Kensaku Mori	Nagoya University, Japan
Parvin Mousavi	Queen's University, Canada
Nassir Navab	Technical University of Munich, Germany
Poul Nielsen	University of Auckland, New Zealand
Wiro Niessen	Erasmus Medical College, The Netherlands
Sebastien Ourselin	University College London, UK
Nikos Paragios	Ecole Central Paris, France
Xavier Pennec	INRIA, France
Josien Pluim	Utrecht University, The Netherlands
Daniel Rueckert	Imperial College London, UK
Ichiro Sakuma	University of Tokyo, Japan
Tim Salcudean	University of British Colombia, Canada
Julia A. Schnabel	University of Oxford, UK
Dinggang Shen	University of North Carolina, USA
John G. Sled	University of Toronto, Canada

Martin Styner	University of North Carolina, USA
Russell Taylor	Johns Hopkins University, USA
Jocelyne Troccaz	CNRS, France
Aaron Ward	University of Western Ontario, Canada
Simon Warfield	Harvard University, USA
Wolfgang Wein	Technical University of Munich, Germany
William Wells	Harvard Medical School, USA
Carl-Fredrik Westin	Harvard Medical School, USA
Guang-Zhong Yang	Imperial College, London, UK
Ziv Yaniv	National Children's Hospital, Washington DC, USA

Local Organizing Committee

Janette Wallace	Robarts Research Institute, London, Canada
Johanne Guilemette	Robarts Research Institute, London, Canada
Debbie Lilley	Robarts Research Institute, London, Canada
Jade Orkin-Fenster	Robarts Research Institute, London, Canada
Jackie Williams	Robarts Research Institute, London, Canada
Chris Wedlake	Robarts Research Institute, London, Canada
Shuo Li	GE Medical Systems/University of Western Ontario, London, Canada
Perry Radau	Sunnybrook Research Institute, Toronto, Canada
Raphael Ronen	Sunnybrook Research Institute, Toronto, Canada

Reviewers

Acar, Burak	Audette, Michel
Achterberg, Hakim	Avants, Brian
Acosta-Tamayo, Oscar	Avila, Lisa
Adluru, Nagesh	Awate, Suyash
Aja-Fernández, Santiago	Axel, Leon
Alexander, Daniel	Babalola, Kolawole
Aljabar, Paul	Bach Cuadra, Meritxell
Allain, Baptiste	Baghani, Ali
Amini, Amir	Baka, Nora
An, Jungha	Balicki, Marcin
Arbel, Tal	Baloch, Sajjad
Aribisala, Benjamin	Barbu, Adrian
Ashburner, John	Barmpoutis, Angelos
Astley, Sue	Barratt, Dean
Atasoy, Selen	Batchelor, Philip

Batmanghelich, Nematollah
Baumann, Michael
Baust, Maximilian
Bazin, Pierre-Louis
Beg, Mirza Faisal
Beichel, Reinhard
Bellec, Pierre
Bello, Fernando
Berger, Marie-Odile
Bergtholdt, Martin
Beuthien, Björn
Bhotika, Rahul
Biesdorf, Andreas
Bloch, Isabelle
Blum, Tobias
Boctor, Emad
Boisvert, Jonathan
Bosch, Johan
Bossa, Matias Nicolas
Boucher, Maxime
Bouix, Sylvain
Boukerroui, Djamal
Bourgeat, Pierrick
Brady, Mike
Brock, Kristy
Brost, Alexander
Brun, Caroline
Bullitt, Elizabeth
Buonaccorsi, Giovanni
Burgert, Oliver
Burlina, Philippe
Burschka, Darius
Caan, Matthan
Cahill, Nathan
Camara, Oscar
Camp, Jon
Cardenes, Ruben
Cardoso, Manuel Jorge
Castro-González, Carlos
Cattin, Philippe C.
Cetingul, Hasan Ertan
Chakravarty, M. Mallar
Chen, Elvis C. S.
Chen, Terrence
Chen, Thomas Kuiran
Chen, Ting
Chen, Xinjian
Cheng, Jian
Cheriet, Farida
Chien, Aichi
Chintalapani, Gouthami
Chinzei, Kiyoyuki
Chitphakdithai, Nicha
Chittajallu, Deepak Roy
Chou, Yiyu
Chowdhury, Ananda
Christensen, Gary
Cinquin, Philippe
Ciuciu, Philippe
Clarkson, Matthew
Clouchoux, Cédric
Colliot, Olivier
Combès, Benoît
Commowick, Olivier
Cootes, Tim
Corso, Jason
Cotin, Stephane
Coulon, Olivier
Coupe, Pierrick
Criminisi, Antonio
Crum, William
Cuingnet, Remi
Daga, Pankaj
Dahl, Anders L.
Dalal, Sandeep
Daoud, Mohammad
Darkner, Sune
Darvann, Tron
Dauguet, Julien
Dawant, Benoit
de Boer, Renske
De Craene, Mathieu
Dehghan, Ehsan
Deligianni, Fani
Demiralp, Cagatay
Demirci, Stefanie
Dequidt, Jeremie
Deriche, Rachid
Descoteaux, Maxime
Desvignes, Michel

Dijkstra, Jouke
DiMaio, Simon
Doignon, Christophe
Dong, Bin
Doorly, Denis
Douiri, Abdel
Dowling, Jason
Drangova, Maria
Drechsler, Klaus
Duchateau, Nicolas
Duchesne, Simon
Dufour, Alexandre
Duriez, Christian
Durrleman, Stanley
Ebrahimi, Mehran
Ehman, Richard
Ehrhardt, Jan
El-Baz, Ayman
Elhawary, Haytham
Ellis, Randy
Engelke, Klaus
Enquobahrie, Andinet
Erdt, Marius
Eskandari, Hani
Eskildsen, Simon
Essert, Caroline
Fadili, Jalal
Fallavollita, Pascal
Fan, Yong
Farag, Aly
Farag, Amal
Faure, François
Fedorov, Andriy
Fei, Baowei
Fenster, Aaron
Feulner, Johannes
Figl, Michael
Fillard, Pierre
Fischer, Gregory
Fitzpatrick, J. Michael
Fleig, Oliver
Florack, Luc
Fonov, Vladimir
Foroughi, Pezhman
Fouard, Celine
Freiman, Moti
Freysinger, Wolfgang
Fridman, Yoni
Funka-Lea, Gareth
Gangloff, Jacques
Gao, Wei
Gao, Yi
Garcia-Lorenzo, Daniel
Garvin, Mona
Gee, Andrew
Gee, James
Geng, Xiujuan
Gerig, Guido
Gholipour, Ali
Giannarou, Stamatia
Gibaud, Bernard
Gibson, Eli
Gilles, Benjamin
Gilson, Wesley
Giusti, Alessandro
Glocker, Ben
Gobbi, David
Goh, Alvina
Goksel, Orcun
Gonzalez Ballester, Miguel Angel
Gooding, Mark
Goodlett, Casey
Graham, Jim
Grau, Vicente
Greenspan, Hayit
Grova, Christophe
Guetter, Christoph
Guevara, Pamela
Guizard, Nicolas
Habas, Piotr
Haegelen, Claire
Hager, Gregory D.
Hahn, Horst
Haider, Clifton
Hamm, Jihun
Han, Xiao
Haneishi, Hideaki
Hans, Arne
Hansen, Michael Sass
Hanson, Dennis

Harders, Matthias
Hastreiter, Peter
Hatt, Chuck
Haynor, David
He, Huiguang
He, Yong
Hedjazi Moghari, Mehdi
Hefny, Mohamed
Heimann, Tobias
Heinrich, Mattias Paul
Heldmann, Stefan
Hellier, Pierre
Heng, Pheng Ann
Hernandez, Monica
Higgins, William
Ho, Harvey
Hoffmann, Kenneth
Hogeweg, Laurens
Hollis, Ralph
Hornegger, Joachim
Hu, Mingxing
Hu, Zhihong
Huang, Junzhou
Huang, Wei
Huang, Xiaolei
Huisman, Henkjan
Hyde, Damon
Idier, Jerome
Iglesias, Juan Eugenio
Imani, Farhad
Ionasec, Razvan
Jagadeesan, Jayender
Jahanshad, Neda
Janoos, Firdaus
Jerebko, Anna
Jia, Hongjun
Jiang, Tianzi
Jiang, Yifeng
Jolly, Marie-Pierre
Jomier, Julien
Joshi, Anand
Joshi, Sarang
Jurrus, Elizabeth
Kårsnäs, Andreas
Kabus, Sven
Kadoury, Samuel
Kainmueller, Dagmar
Kamen, Ali
Kapur, Tina
Karamalis, Athanasios
Karemore, Gopal
Karssemeijer, Nico
Kaynig, Verena
Kazanzides, Peter
Keeve, Erwin
Keil, Andreas
Kerrien, Erwan
Khan, Ali R.
Khurd, Parmeshwar
Kikinis, Ron
Kim, Kio
Kim, Minjeong
King, Andrew
Kiraly, Atilla
Kitasaka, Takayuki
Klein, Stefan
Klinder, Tobias
Koikkalainen, Juha
Konukoglu, Ender
Kowal, Jens
Kruggel, Frithjof
Krupa, Alexandre
Kumar, Rajesh
Kurkure, Uday
Kuroda, Yoshihiro
Kutter, Oliver
Kwok, Ka-Wai
Kybic, Jan
Lacefield, James
Laine, Andrew
Landman, Bennett
Lango, Thomas
Langs, Georg
Lapeer, Rudy
Laporte, Catherine
Lashkari, Danial
Law, Max W.K.
Le, Yen
Lee, Junghoon
Lee, Su-Lin

Lee, Tim
Lekadir, Karim
Lenglet, Christophe
Lensu, Lasse
Leow, Alex
Lepore, Natasha
Lerch, Jason
Leung, Kelvin
Li, Chunming
Li, Fuhai
Li, Gang
Li, Kaiming
Li, Kang
Li, Ming
Li, Ruijiang
Li, Shuo
Li, Yang
Liang, Jianming
Liao, Hongen
Liao, Rui
Liao, Shu
Lindseth, Frank
Linguraru, Marius George
Linte, Cristian
Liu, Chi
Liu, Huafeng
Liu, Jiamin
Liu, Sheena
Liu, Tianming
Liu, Xiaofeng
Lo, Pechin
Loeckx, Dirk
Loew, Murray
Lori, Nicolas
Lu, Chao
Lu, Le
Luboz, Vincent
Lui, Lok Ming
Ma, YingLiang
Mørup, Morten
Müller, Henning
Machann, Juergen
Machiraju, Raghu
Maddah, Mahnaz
Madore, Bruno
Magee, Derek
Magill, John
Maier-Hein, Lena
Makram-Ebeid, Sherif
Malandain, Gregoire
Manduca, Armando
Mangin, Jean-Francois
Mani, Meenakshi
Manjon, Jose V.
Manniesing, Rashindra
Mansi, Tommaso
Manzke, Robert
Marchal, Maud
Marsland, Stephen
Martí, Robert
Martin-Fernandez, Marcos
Masamune, Ken
Masutani, Yoshitaka
Mattes, Julian
Maurel, Pierre
McClelland, Jamie
Melbourne, Andrew
Menze, Bjoern
Metaxas, Dimitris
Metz, Coert
Meyer, Chuck
Michel, Fabrice
Miga, Michael
Miller, James
Miller, Karol
Mirota, Daniel
Modat, Marc
Modersitzki, Jan
Mohamed, Ashraf
Montillo, Albert
Moore, John
Moradi, Mehdi
Mountney, Peter
Murgasova, Maria
Mylonas, George
Naish, Michael
Nakamura, Ryoichi
Nash, Martyn
Nassiri-Avanaki, Mohammad-Reza
Nichols, Thomas

Nicolau, Stéphane
Niemeijer, Meindert
Niethammer, Marc
Noël, Peter
Noble, Alison
Nolte, Lutz
Ocegueda, Omar
Oda, Masahiro
O'Donnell, Lauren
O'Donnell, Thomas
Ogier, Arnaud
Oguz, Ipek
Okada, Toshiyuki
Olabarriaga, Silvia
Oliver, Arnau
Olmos, Salvador
Orihuela-Espina, Felipe
Orkisz, Maciej
Otake, Yoshito
Ou, Yangming
Oubel, Estanislao
Ozarslan, Evren
Pace, Danielle
Padfield, Dirk
Padoy, Nicolas
Palaniappan, Kannappan
Paniagua, Beatriz
Papadakis, Emanuel
Papademetris, Xenios
Parthasarathy, Vijay
Passat, Nicolas
Patel, Rajni
Patriciu, Alexandru
Paul, Perrine
Paulsen, Rasmus
Pauly, Olivier
Pavlidis, Ioannis
Pearlman, Paul
Peitgen, Heinz-Otto
Pekar, Vladimir
Penney, Graeme
Pernus, Franjo
Petersen, Jens
Peyrat, Jean-Marc
Pham, Dzung
Pitiot, Alain
Pizaine, Guillaume
Platel, Bram
Plewes, Donald
Podder, Tarun
Pohl, Kilian Maria
Poignet, Philippe
Poot, Dirk
Poynton, Clare
Prager, Richard
Prastawa, Marcel
Pratt, Philip
Prima, Sylvain
Prince, Jerry
Punithakumar, Kumaradevan
Qazi, Arish A.
Qian, Xiaoning
Qian, Zhen
Qin, Lei
Quellec, Gwenole
Radeva, Petia
Rajagopal, Vijayaraghavan
Ramamurthi, Krishnakumar
Ramezani, Mahdi
Rasoulian, Abtin
Rathi, Yogesh
Reinertsen, Ingerid
Rettmann, Maryam
Reyes, Mauricio
Reyes-Aldasoro, Constantino
Rhode, Kawal
Ribbens, Annemie
Risholm, Petter
Risser, Laurent
Rit, Simon
Rittscher, Jens
Rivaz, Hassan
Riviere, Cameron
Robb, Richard A.
Roche, Alexis
Rohkohl, Christopher
Rohlfing, Torsten
Rohling, Robert
Rohr, Karl
Ross, James

Rousseau, François
Russakoff, Daniel
Sörensen, Lauge
Sabuncu, Mert Rory
Sadeghi Naini, Ali
Saha, Punam Kumar
Salvado, Olivier
Sanchez, Clara
Savadjiev, Peter
Savinaud, Mickaël
Schaller, Christian
Scherrer, Benoit
Schultz, Thomas
Schweikard, Achim
Sermesant, Maxime
Shah, Mohak
Shah, Shishir
Shamir, Reuben R.
Shams, Ramtin
Shen, Li
Shi, Feng
Shi, Pengcheng
Shi, Yonggang
Shi, Yundi
Shimizu, Akinobu
Siddiqi, Kaleem
Siewerdsen, Jeffrey
Simaan, Nabil
Simpson, Amber
Singh, Nikhil
Sinkus, Ralph
Sjöstrand, Karl
Slabaugh, Greg
Slagmolen, Pieter
Smal, Ihor
Smeets, Dirk
Soler, Luc
Song, Sang-Eun
Song, Xubo
Sotiras, Aristeidis
Sporring, Jon
Staib, Lawrence
Staring, Marius
Stewart, James
Stoianovici, Dan
Stoyanov, Danail
Strother, Stephen
Studholme, Colin
Subramanian, Navneeth
Summers, Ronald
Sun, Bo
Sundar, Hari
Suzuki, Kenji
Syeda-Mahmood, Tanveer
Szczerba, Dominik
Szekely, Gabor
Szilagyi, Laszlo
Sznitman, Raphael
Tahmasebi, Amir
Tanner, Christine
Tao, Xiaodong
Tardif, Christine Lucas
Tasdizen, Tolga
Taylor, Zeike
Thévenaz, Philippe
Thiran, Jean-Philippe
Thiriet, Marc
Thirion, Bertrand
Tiwari, Pallavi
Todd Pokropek, Andrew
Todd, Catherine
Toews, Matthew
Tokuda, Junichi
Toussaint, Nicolas
Tristán-Vega, Antonio
Tsekos, Nikolaos V.
Tsin, Yanghai
Unal, Gozde
Ungi, Tamas
Vaillant, Regis
van Assen, Hans
van der Lijn, Fedde
van Ginneken, Bram
Van Leemput, Koen
van Rikxoort, Eva
van Stralen, Marijn
Vandergheynst, Pierre
Vannier, Michael
Varoquaux, Gael
Vegas-Sánchez-Ferrero, Gonzalo

Venkataraman, Archana
Vercauteren, Tom
Verma, Ragini
Vermandel, Maximilien
Viergever, Max
Vignon, Francois
Villard, Pierre-Frederic
Visentini-Scarzanella, Marco
Vitanovski, Dime
Vitiello, Valentina
von Berg, Jens
Voros, Sandrine
Vos, Pieter
Vosburgh, Kirby
Votta, Emiliano
Vrooman, Henri
Vrtovec, Tomaz
Wachinger, Christian
Waechter-Stehle, Irina
Wahle, Andreas
Wang, Defeng
Wang, Fei
Wang, Hongzhi
Wang, Lejing
Wang, Liansheng
Wang, Peng
Wang, Qian
Wang, Song
Wang, Yang
Wang, Ying
Wassermann, Demian
Weber, Stefan
Weese, Jürgen
Whitaker, Ross
Wiest-Daessle, Nicolas
Wiles, Andrew
Wirtz, Stefan
Wittek, Adam
Wolf, Ivo
Wollny, Gert
Wolz, Robin
Wörz, Stefan
Wu, Guorong
Wu, John Jue
Wu, Xiaodong
Xi, Zhao
Xiao, Gaoyu
Xie, Hua
Xie, Jun, Jun
Xiong, Guanglei
Xu, Jun
Xu, Sheng
Xue, Zhong
Yalamanchili, Raja
Yan, Pingkun
Yang, Xiaoyun
Yao, Jianhua
Yap, Pew-Thian
Ye, Dong Hye
Yeo, B.T. Thomas
Yin, Youbing
Yin, Zhaozheng
Yoo, Terry
Yoshida, Hiro
Young, Stewart
Yushkevich, Paul
Zacharia, Eleni
Zahiri Azar, Reza
Zeng, Wei
Zhan, Liang
Zhang, Daoqiang
Zhang, Dongping
Zhang, Hui
Zhao, Fei
Zheng, Guoyan
Zheng, Yefeng
Zheng, Yuanjie
Zhong, Hua
Zhou, Jinghao
Zhou, Luping
Zhou, S. Kevin
Zhou, X. Sean
Zhu, Hongtu
Zhu, Ying
Zikic, Darko
Zion, Tse
Zwiggelaar, Reyer

Awards Presented at MICCAI 2010, Beijing

MICCAI Society "Enduring Impact Award" Sponsored by Philips.
The Enduring Impact Award is the highest award of the MICCAI Society. It is a career award for continued excellence in the MICCAI research field. The 2010 Enduring Impact Award was presented to Russ Taylor, Johns Hopkins University, USA.

MICCAI Society Fellowships
MICCAI Fellowships are bestowed annually on a small number of senior members of the Society in recognition of substantial scientific contributions to the MICCAI research field and service to the MICCAI community. In 2010, fellowships were awarded to:

James S. Duncan (Yale University, USA)
Stephen M Pizer (University of North Carolina, USA)
Jocelyne Troccaz (CNRS, France)

MedIA-MICCAI Prizes (Split decision)
Jihun Hamm, for the article entitled: "GRAM: A Framework for Geodesic Registration on Anatomical Manifolds," co-authored by: Jihun Hamm, Dong Hye Ye, Ragini Verma, Christos Davatzikos

Samuel Gerber, for the article entitled: "Manifold Modeling for Brain Population Analysis," co-authored by: Tolga Tasdizen, P. Thomas Fletcher, Sarang Joshi, Ross Whitaker

Best Paper in Computer-Assisted Intervention Systems and Medical Robotics, Sponsored by Intuitive Surgical Inc.
Rogerio Richa, for the article entitled: "Robust 3D Visual Tracking for Robotic-Assisted Cardiac Interventions," co-authored by: Rogerio Richa, Antonio P. L. Bo, and Philippe Poignet

MICCAI Young Scientist Awards
The Young Scientist Awards are stimulation prizes awarded for the best first authors of MICCAI contributions in distinct subject areas. The nominees had to be full-time students at a recognized university at, or within, two years prior to submission. The 2010 MICCAI Young Scientist Awards were presented in the following categories to:

Instrument and Patient Localization and Tracking	**Ehsan Dehghan** "Prostate Brachytherapy Seed Reconstruction Using C-Arm Rotation Measurement and Motion Compensation"

Image Reconstruction and Restoration	**Junzhou Huang** "Efficient MR Image Reconstruction for Compressed MR Imaging"
Modelling and Simulation	**Saša Grbić** "Complete Valvular Heart Apparatus Model from 4D Cardiac CT"
Quantitative Image Analysis	**Rémi Cuingnet** "Spatially Regularized SVM for the Detection of Brain Areas Associated with Stroke Outcome"
Functional and Diffusion-Weighted MRI	**Anthony J. Sherbondy** "MicroTrack: An Algorithm for Concurrent Projectome and Microstructure Estimation"

Accepted MICCAI 2011 Papers

by Clinical Theme:

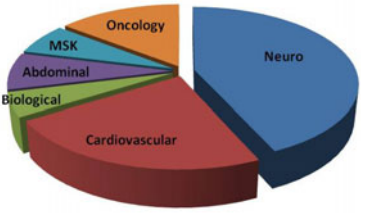

by Technical Theme:

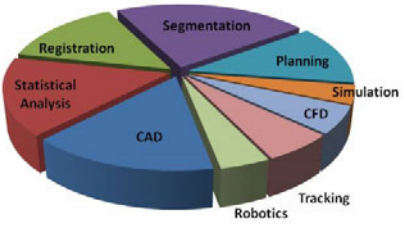

Table of Contents – Part II

Diffusion Weighted Imaging I

Diffusion Weighted Imaging II

fMRI

Statistical Analysis and Shape Modelling I

Statistical Analysis and Shape Modelling II

Registration I

Registration II

Longitudinal Change Detection: Inference on the Diffusion Tensor Along White-Matter Pathways

Antoine Grigis[1,2,⋆], Vincent Noblet[1], Fréderic Blanc[2], Fabrice Heitz[1], Jérome de Seze[2], and Jean-Paul Armspach[2]

[1] University of Strasbourg, LSIIT, UMR 7005, CNRS, France
[2] University of Strasbourg, LINC-IPB, UMR 7237, CNRS, France

Abstract. Diffusion tensor magnetic resonance imaging (DT-MRI) tractography allows to probe brain connections *in vivo*. This paper presents a change detection framework that relies on white-matter pathways with application to neuromyelitis optica (NMO). The objective is to detect global or local fiber diffusion property modifications between two longitudinal DT-MRI acquisitions of a patient. To this end, estimation and testing tools on tensors along the white-matter pathways are considered. Two tests are implemented: a pointwise test that compares at each sampling point of the fiber bundle the tensor populations of the two exams in the cross section of the bundle and a fiberwise test that compares paired tensors along all the fiber bundle. Experiments on both synthetic and real data highlight the benefit of considering fiber based statistical tests compared to the standard voxelwise strategy.

1 Introduction

Diffusion weighted magnetic resonance imaging (DW-MRI) is a non-invasive method for characterizing the diffusion of water molecules in tissues. The automated detection of relevant changes in longitudinal DW-MRI sequences may open promising perspectives for medical diagnosis, follow-up and prognosis. Diffusivity profiles obtained from DW-MRI acquisitions are usually modeled by diffusion tensors of rank 2 (DTs). Some previous work has addressed change detection between DT-derived scalar images, [2], [4], or between DT fields, [3], [6], [5]. In [6], statistical tests have been developed to compare two sets of tensors with application to the comparison of two groups of subjects. Applying these tests in the context of the longitudinal analysis of a given subject requires to extract at each voxel two populations of tensors to be compared. This step is not straightforward since we need to ensure that all the tensors are drawn from the same distribution. A natural idea is to learn tensor distribution at each voxel by considering all the tensors in a surrounding user-defined spatial neighbourhood. This is based on the commonly made assumption of a constant piecewise model, with known limitations, in particular, at the interface between tissues. To circumvent this limitation, a local bootstrap strategy can be used to generate a set

⋆ grigis@unistra.fr

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 1–8, 2011.

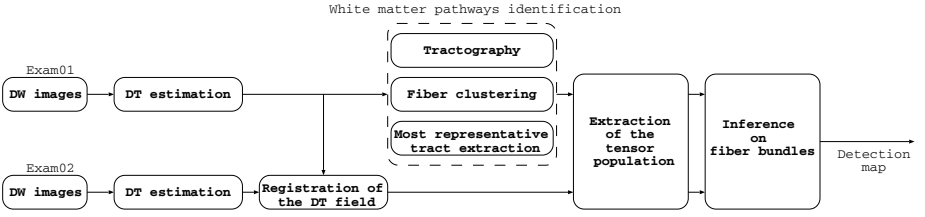

Fig. 1. Flowchart of the proposed framework

of tensors characterizing the variability of each voxel, based on the variability existing in the DW-MRI images [5].

In this paper we suggest a longitudinal DT change detection framework based on tractography features (Fig. 1). The construction of a neighbourhood based on the fiber structures, and the application of statistical tests along the fiber, enable the detection of subtle changes along white-matter pathways. The proposed processing pipeline is as follows. First, eddy current distortions are corrected using FSL[1]. DTs are then estimated from the DW images using a least square algorithm. Then, a fiber tract extraction is performed with Slicer[2] (Label Seeding module) exploiting the first exam information. Fibers of the first exam are clustered to generate homogeneous fiber bundles (section 2.1). The most representative tract of each bundle is extracted and used as a reference tract for the bundle (section 2.2). The DTs of the second exam are warped using a linear interpolation in the Log-eulidean space and the preservation of principal direction (PPD) reorientation strategy to be aligned with the first exam [5]. To this end, an affine and a Bspline based transformation are estimed from the FA images of the two acquisitions [5]. Tensor sets reflecting the local variability are extracted at each sampling point of the reference tract considering the DTs of the two registered exams that belong to the cross section of the bundle (section 3.1). Finally local and global multivariate statistical tests along white-matter principal pathways are applied (section 3.2). Experiments on synthetic data highlight the benefit of considering fiber based statistical tests compared to a standard voxel-wise strategy. Application to neuromyelitis optica (NMO) also demonstrates the clinical relevance of the fiber-based strategy.

2 Identifying White-Matter Fiber Bundles in DT Data

2.1 Fiber Clustering

Tractography methods produce a dense set of curves that bear a close resemblance to known white-matter pathways [7]. Considering a method that automatically clusters and labels these curves into anatomically plausible pathways

[1] http://www.fmrib.ox.ac.uk/fsl/
[2] http://www.slicer.org/

is of great interest for a relevant analysis of white matter properties. To this end, we have implemented a fiber bundle clustering algorithm based on the method described in [7]. Neighboring fiber tracts are grouped using the mean of thresholded closest distance. With this metric the distance between curves Q and R is expressed as follows:

$$d_t(Q, R, t) = mean_{a \in Q, ||a-b|| > t} \min_{b \in R} ||a - b|| \quad (1)$$

where t is the threshold below which the distances are discarded. This threshold enables the separation of two fibers that might run very closely together for a long course and then diverge abruptly for a relatively short course into two different clusters [7]. Since this metric is not symmetric, we consider the longer mean of thresholded closest distances in the clustering algorithm (see [7] for a discussion on the influence of the chosen symmetrized form):

$$d_{Lt}(Q, R, t) = \max(d_t(Q, R, t), d_t(R, Q, t)) \quad (2)$$

With this metric it is possible to identify corresponding tracts with high consistency even when the global shapes of the tracts differ markedly in length or curvature [7]. A single-linkage agglomerative clustering is then performed [7]. The number of clusters is inversely related to a user-defined fusion proximity threshold d_f. The main problem of the tract-clustering algorithm is the computational burden related to the calculation of the pairwise distances between the current fiber and all the previously clustered fibers. To circumvent this issue, pairwise distances are only computed for fibers belonging to clusters whose center of mass are close enough to the center of mass of the current fiber (i.e., specification of a user-defined distance threshold d_{cm}). After the clustering step, bundles with less than N_c fibers are rejected in order to have a minimum number of samples in the statistical population when a pointwise test is performed.

2.2 Bundle Most Representative Tract

We denote each tract belonging to the i^{th} bundle $T_{i,j}$, where n_i is the number of tracts in the bundle i and $j \in \{1, ..., n_i\}$. Then the most representative tract $T_{ref,i}$ of the i^{th} bundle is obtained by selecting the fiber tract that minimizes the sum of distances with other tracts. The metric used here is the average mean of closest distances:

$$d_{MC}(Q, R) = \frac{d_t(Q, R, 0) + d_t(R, Q, 0)}{2} \quad (3)$$

Note that we set $t = 0$ in the thresholded distance d_t since we are dealing with a coherent bundle. Finally, the criterion used to select the most representative tract of a bundle i can be summarized as follows:

$$T_{ref,i} = \arg \min_{T_{i,k}, k \in 1,...,n_i} \left[\sum_{l \in 1,..,n_i} d_{MC}(T_{i,k}, T_{i,l}) \right] \quad (4)$$

Each most representative tract $T_{ref,i}$ is then resampled with N points (N being user-defined) using a cardinal basis spline interpolation.

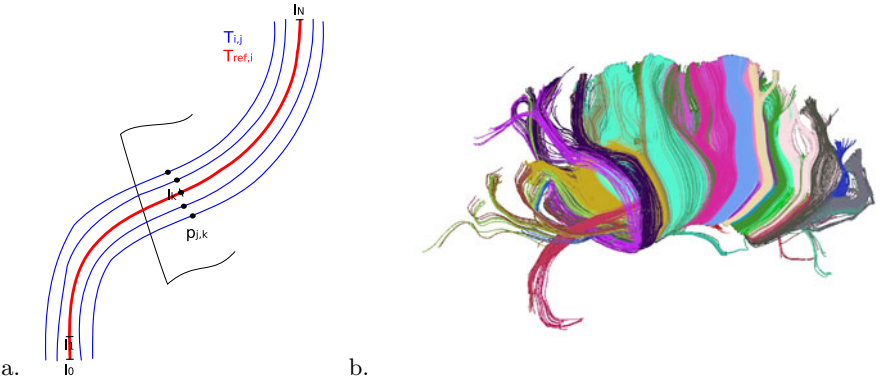

Fig. 2. a. Schematic representation of the extraction of the tensor sets reflecting the local variability: see text, and b. the result of the clustering step on a tractography obtained with Slicer from a ROI in the corpus callosum (parameters: see text).

3 Inference on Fiber Bundles

3.1 Extraction of Tensor Populations Reflecting the Local Variability

Inference on DT is composed of an estimation and a testing step. The estimation of model parameters requires to consider a population of tensors that are all assumed to be drawn from the same distribution. In many situations, considering all the tensors in a given spatial neighbourhood may not satisfy this assumption, especially at the interface between different anatomical structures. To overcome this problem we propose to select a tensor population by taking advantage of the local geometry of the bundles (see Fig. 2-a). For each sampling point l_k of the reference tract $T_{ref,i}$, a population of tensors is extracted by considering the tensors located at the intersections $p_{j,k}$ between the plane perpendicular to $T_{ref,i}$ and each fiber $T_{i,j}$ of the bundle. Since points $p_{j,k}$ may not lie on the grid, corresponding tensors are obtained using linear interpolation in the Log-euclidean space. In some cases, the plane might intersect the bundle in several points (*e.g.* for U-shaped bundles). Consequently, only the closest intersection is considered. Besides to remove outliers, intersections $p_{j,k}$ whose distance to l_k is larger than $3\sigma_i$ are discarded, σ_i being the median absolute distance of the $p_{j,k}$'s to their corresponding sampling point l_k over the i^{th} bundle.

3.2 Multivariate Statistical Estimation and Testing on Tensors

Many statistical tests rely on the normal distribution. Considering the multivariate normal distribution for the symmetric positive definite (SPD) matrices has the drawback to associate matrices with negative or null eigenvalues with a non null probability. To circumvent this limitation, Schwartzman suggests to model the matrix logarithms with the multivariate normal distribution, which

comes to model the SPD matrices with a Log-normal distribution [6]. To this end the Log-euclidean metric is used. Note that the logarithm of an SPD matrix D is obtained as $L = Log(D) = ULog(\Lambda)U^T$, where Λ and U are the matrices derived from the standard spectral decomposition, containing respectively the eigenvalues and the eigenvectors of D. Then, the Log-euclidean distance between two SPD matrices D_1 and D_2 can be defined as the Euclidean distance between their logarithms $d^2(D_1, D_2) = \|Log(D_1) - Log(D_2)\|^2$. According to this metric, an estimator of the mean $\bar{D}$ of a set of n tensors D_i is given by the exponential of the arithmetic Log-tensors mean, i.e. $\bar{D} = exp(\bar{L})$, with:

$$\bar{L} = argmin_{\Sigma} \sum_{i=1}^{n} \|Log(D_i) - \Sigma\|^2 = \frac{1}{n} \sum_{i=1}^{n} Log(D_i) \tag{5}$$

An estimator of the variance is then:

$$s^2 = \frac{1}{6(n-1)} \left[\sum_{i=1}^{n} tr\left(Log(D_i) - \bar{L}\right)^2 \right] \tag{6}$$

Pointwise testing. Based on this model, it is possible to derive a statistical test on tensor's eigenvalues. We consider two populations of n_1 and n_2 tensors respectively ($n = n_1 + n_2$). Under the assumption that the tensor logarithms of the two populations follow the normal distributions $L_1 \sim \mathcal{N}(M_1, \sigma^2 Id)$ and $L_2 \sim \mathcal{N}(M_2, \sigma^2 Id)$, the maximum likelihood estimates of M_1, M_2, and σ^2 are respectively $\bar{L_1}$, $\bar{L_2}$ (Eq.5), and $\hat{\sigma}^2 = [(n_1 - 1)s_1 + (n_2 - 1)s_2]/(n - 2)$ (Eq.6). We consider a test [6] that evaluates whether the two populations of diffusion tensors have similar eigenvalues, but possibly different eigenvectors. Let Λ_1, U_1 and Λ_2, U_2 be the matrices derived from standard spectral decomposition, and containing respectively the eigenvalues and eigenvectors of M_1 and M_2. The test, based on the log-likelihood ratio under hypotheses $H_0 : \Lambda_1 = \Lambda_2$ *vs* $H_1 : \Lambda_1 \neq \Lambda_2$ with $U_1 \neq U_2$ unknown is:

$$T_{pointwise,\sigma} = \frac{n_1 n_2}{3n^2 \hat{\sigma}^2} tr\left[(\bar{\Lambda_1} - \bar{\Lambda_2})^2\right] \tag{7}$$

where $\bar{\Lambda_1}$ and $\bar{\Lambda_2}$ are the eigenvalue matrices of $\bar{L_1}$ and $\bar{L_2}$, respectively. We investigate in the sequel the influence of estimating the variance σ^2 locally for each sampling point l_k or of considering a constant variance $\sigma^2_{b_i}$ estimated over the whole bundle b_i. We denote by $T_{pointwise,\sigma_{b_i}}$ the corresponding test.

Fiberwise testing. It may be of great interest for physicians or neuroscientists to identify global modifications along fiber bundles. To this end, we propose to build a list of putative evolving fibers ordered by decreasing evolution probability. At each of the N sampling points of the most representative tracts $T_{ref,i}$, the difference between the fiber cross section mean Log-tensors are computed. A population of N Log-tensors is thus obtained for each bundle. Under the assumption that the Log-tensor population follows the normal distribution $L \sim \mathcal{N}(M, \sigma^2 Id)$, the maximum likelihood estimates of M, and σ^2 are respectively $\bar{L}$ (Eq.5), and

$\hat{\sigma}^2 = s^2$ (Eq.6). We consider a test that evaluates whether the population of Log-tensors has eigenvalues equal to $\Lambda_0 = [0, 0, 0]$ [6]. Let Λ, U be the matrices derived from standard spectral decomposition, and containing respectively the eigenvalues and eigenvectors of M. The test, based on the log-likelihood ratio under hypotheses $H_0 : \Lambda = \Lambda_0$ *vs* $H_1 : \Lambda \neq \Lambda_0$ is:

$$T_{fiberwise} = \frac{N}{\hat{\sigma}^2} tr\left[(\bar{\Lambda} - \Lambda_0)^2\right] \tag{8}$$

where $\bar{\Lambda}$ is the eigenvalue matrice of $\bar{L}$.

4 Results

Experiments on simulated changes. We consider two repeated DW-MRI acquisitions of the same healthy subject acquired on a 3T SIEMENS MRI scanner with 30 encoding gradients (b-value of 1000 s/mm^2). In this way, the differences between the two scans are only due to the acquisition noise and distortion. The image dimensions are $96 \times 96 \times 55$ and the spatial resolution is $2 \times 2 \times 2mm^3$. The tractography was computed for the first image using the Label Seeding module in Slicer (integration step length of $0.5mm$, minimum tract length of $50mm$, and stopping values on the fractional anisotropy and curvature of 0.1 and 0.8 respectively). For the clustering step, a fusion threshold d_f of $3.5mm$, a cut-off threshold t in the metric of $0.5mm$, and a selection threshold d_{mc} of $4d_f$ were used. Results are presented in Fig. 2-b where each bundle contains at least $N_c = 9$ fibers. A synthetic fiber alteration is simulated in one of the two scans as follows. After the fiber clustering step, a bundle is selected to generate a modification mask. Inside the modification mask, the diffusivity in the principal direction, i.e. the principal eigenvalue, is uniformly modified by applying a multiplicative factor $k \in]1, 1.5]$. The results of the pointwise test are compared with the $3 \times 3 \times 3$ spatial neighbourhood approach (SN) described in [5]. To this end, we use a projection P_1 from the ijk coordinate system to the fiber coordinate system using the nearest neighbour interpolation. The criterion used to compare the different methods is the area under Receiver Operating Characteristic (ROC) curves. A test that allows a perfect discrimination is characterized by ROC plot with an area of one. A resampling of the most representative fiber in $N = 200$ points is done before the bundle cross section point extraction. The Table 1 summarizes the results. The pointwise methods outperform the SN method, thus pointing out the importance of satisfying the constant piecewise assumption. Estimating a constant variance for each bundle slightly increases the performance of the statistical test. The sensitivity of the fiberwise test $T_{fiberwise}$ is also explored (Table 1). Notice that in this particular case the fiberwise test is more appropriate and gives better results since a global modification has been simulated.

Experiments on NMO patients. Neuromyelitis Optica (NMO) is an inflammatory disease of the central nervous system that predominantly affects optic

Table 1. Areas under the ROC curves (diffusion modification: see text)

k	1.01	1.02	1.03	1.04	1.05	1.06	1.07	1.08	1.09	1.1	1.2	1.3	1.4
SN	0.581	0.587	0.622	0.663	0.701	0.736	0.767	0.795	0.822	0.843	0.937	0.957	0.965
$T_{pointwise,\sigma}$	0.698	0.707	0.724	0.746	0.705	0.793	0.815	0.834	0.852	0.868	0.953	0.979	0.989
$T_{pointwise,\sigma_{b_i}}$	0.721	0.727	0.745	0.771	0.801	0.832	0.859	0.884	0.905	0.921	0.986	0.996	0.998
$T_{fiberwise}$	0.605	0.837	0.977	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0

a. b.

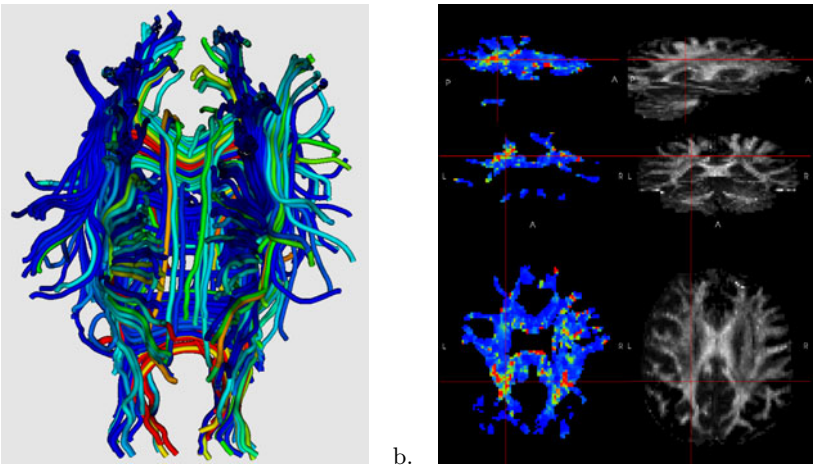

Fig. 3. a. fiberwise test (first patient) and b. projection of the pointwise test from the fiber coordinate system to the ijk coordinate (second patient).

nerves and spinal cord. The frequency of brain lesions is small using standard sequences of MRI (T1, T2 and T2-FLAIR). However, a recent work has demonstrated that the majority of NMO patients have cognitive impairment, which is a subcortical impairment [1]. In the first part of this preliminary study, we compared the cognitive functions outcome of patient 1 (using the French translation of the Brief Repeatable Battery (BRB-N), a cognitive battery with 14 subtests) and brain diffusion MRI outcome (using the fiberwise test) done at the same time at M0 and M18. In the second part, we compared the neurological physical status outcome of patient 2 (using the expanded disability status scale (EDSS)) to brain diffusion MRI outcome (using the pointwise test) done at the same time at M0 and M18. In this latter in order to simplify the interpretation of the result, the fiber detections were transformed back to the voxel space. Images were acquired on a 1.5T SIEMENS MRI scanner with 30 encoding gradients (b-value of $1000s/mm^2$) at 18 months apart. The images dimensions are $128 \times 128 \times 41$ and the spatial resolution is $1.8 \times 1.8 \times 3.5mm^3$. For the first patient the fiberwise test was applied (Fig. 3-a). Two regions stand out: the anterior and posterior parts of the corpus callosum. These modifications are in accordance with cognitive and behavioral status that worsened. Patient 1 had 6 subtests out of 14

inferior to the 5th percentile at M0, and 7 subtests at M18. The second patient had a huge decrease of capacity to walk between M0 and M18. Thus, the EDSS increased from 6.5 (capable to walk with a constant bilateral assistance, such as canes) to 8.5 (essentially restricted to bed much of day, not capable to walk). This aggravation was in accordance with the results of the pointwise test showing modifications in the corticospinal tracts (Fig. 3-b).

5 Conclusions

The proposed multivariate statistical tests along white-matter principal pathways provide complementary information. The global test enables the identification of the fibers involved in the longitudinal evolution of the pathology, while the local strategy enables the detection of more subtle changes. Moreover, we have demonstrated the superiority of such methods compared to the standard voxelwise strategy. Finally, the proposed approach might open promissing perspectives for the follow-up of the NMO pathology, and will give way to further explorations. In the future we want to investigate whether a joint fiber clustering of the two exams would lead to a better mapping of the whitte matter tracts. We also want to study how registration accuracy could affect the statistical tests and investigate the impact of using a symmetric registration.

References

1. Blanc, F., Zéphir, H., Lebrun, C., Labauge, P., Castelnovo, G., Fleury, M., Sellal, F., Tranchant, C., Dujardin, K., Vermersch, P., de Seze, J.: Cognitive functions in neuromyelitis optica. Arch. Neurol. 65(1), 84–88 (2008)
2. Boisgontier, H., Noblet, V., Heitz, F., Rumbach, L., Armspach, J.P.: An automatic method for change detection in serial DTI-derived scalar images. In: Workshop MIAMS - MICCAI (2008)
3. Boisgontier, H., Noblet, V., Heitz, F., Rumbach, L., Armspach, J.P.: Generalized likelihood ratio tests for change detection in diffusion tensor images. In: ISBI, pp. 811–814. IEEE Press, Piscataway (2009)
4. Chung, S., Pelletier, D., Sdika, M., Lu, Y., Berman, J.I., Henry, R.G.: Whole brain voxel-wise analysis of single-subject serial DTI by permutation testing. Neuroimage 39(4), 1693–1705 (2008)
5. Grigis, A., Noblet, V., Renard, F., Heitz, F., Armspach, J.-P., Rumbach, L.: Change detection in diffusion MRI using multivariate statistical testing on tensors. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. 6362, pp. 117–124. Springer, Heidelberg (2010)
6. Schwartzman, A.: Random ellipsoids and false discovery rates: statistics for diffusion tensor imaging data. Ph.D. thesis, Stanford University (2006)
7. Zhang, S., Correia, S., Laidlaw, D.H.: Identifying white-matter fiber bundles in DTI data using an automated proximity-based fiber-clustering method. IEEE Transactions on Visualization and Computer Graphics 14, 1044–1053 (2008)

Voxelwise Multivariate Statistics and Brain-Wide Machine Learning Using the Full Diffusion Tensor

Anne-Laure Fouque[1,2], Pierre Fillard[1,3], Anne Bargiacchi[2], Arnaud Cachia[4], Monica Zilbovicius[2], Benjamin Thyreau[1], Edith Le Floch[1,2], Philippe Ciuciu[1], and Edouard Duchesnay[1,2]

[1] CEA, Neurospin, LNAO, Saclay, France
[2] INSERM-CEA U.1000 Imaging and Psychiatry, Orsay, France
[3] INRIA Saclay-Île-de-France, Parietal, Saclay, France
[4] UMR 894 INSERM - Paris Descartes University, Laboratory of Pathophysiology of Psychiatric Diseases, Sainte-Anne Hospital, Paris, France

Abstract. In this paper, we propose to use the full diffusion tensor to perform brain-wide score prediction on diffusion tensor imaging (DTI) using the log-Euclidean framework., rather than the commonly used fractional anisotropy (FA). Indeed, scalar values such as the FA do not capture all the information contained in the diffusion tensor. Additionally, full tensor information is included in every step of the pre-processing pipeline: registration, smoothing and feature selection using voxelwise multivariate regression analysis. This approach was tested on data obtained from 30 children and adolescents with autism spectrum disorder and showed some improvement over the FA-only analysis.

1 Introduction

Our aim is to generalize to tensors some of the most commonly used diffusion-MRI processing tools in order to exploit the full information of the tensor in every step of the processing pipeline and ultimately perform brain-wide machine learning on the full tensor. We achieve this goal by making use of Log-Euclidean (LE) metrics [4], a simple and fast way to perform computations in the tensor space. Classical tensor-based analysis consists first in computing scalar images from the dataset, then in working only on the resulting scalar maps. Commonly used scalar features include the fractional anisotropy (FA) and the apparent diffusion coefficient (ADC) for diffusion tensor imaging (DTI) studies, and the Jacobian determinant in tensor-based morphometry (TBM) studies. However, these scalar values are computed only from the eigenvalues of the tensors, and therefore do not capture all the information available. Some information (*e.g.* tensor orientation) is thus lost in the process. Recently, several groups performed full-tensor voxelwise analysis of tensor fields with LE metrics, both in the DTI context [6,17] and in the TBM context [10,5]. However, to the best of our knowledge the LE framework has not yet been applied to voxelwise regression nor to brain-wide machine learning. The present paper aims at filling this gap.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 9–16, 2011.

The remainder of this paper is organized as follows. We first present the Log-Euclidean framework and how to use it to perform all the steps necessary to brain-wide machine learning (pre-processings, feature selection using voxelwise regression analysis and finally brain-wide score prediction). We then describe the dataset involved in our experiments and compare the results of the full-tensor approach with those of the classical FA-only approach.

2 Methods

In this section, we present the methods used in this study. Section 2 describes the Log-Euclidean framework. Next sections expose the different steps in the pipeline: registration, feature selection using voxelwise regression analysis and brain-wide score prediction. The whole pipeline is summarized in Fig. 1.

Log-Euclidean Metrics. The tensor space is the space of 3×3 symmetric positive-definite matrices $\mathcal{S}ym_3^+$. The usual Euclidean operations on the space of 3×3 matrices suffer from many defects when applied to tensors [4]: positive-definiteness is not always preserved, and there is tensor swelling effect. To address these problems, affine-invariant metrics have been proposed [13], but they induce a huge numerical complexity. The Log-Euclidean framework [4] is a much simpler and faster way to make computations on tensors, while in practice yielding similar results as the affine-invariant framework. The LE approach relies on a vector space structure defined on tensor space $\mathcal{S}ym_3^+$ with a logarithmic multiplication $S_1 \odot S_2 \overset{def}{=} \exp(\log S_1 + \log S_2)$ and a logarithmic scalar multiplication

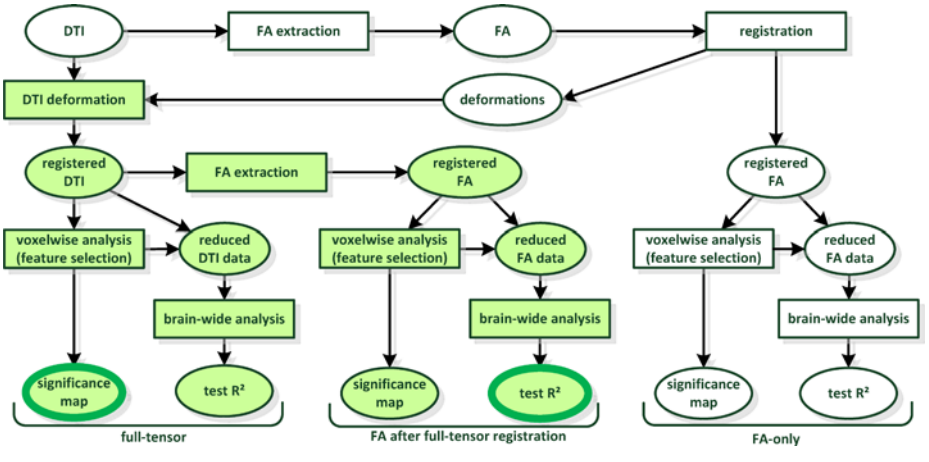

Fig. 1. Flowchart of the whole pipeline – in white classical steps, in green new steps to account for the full-tensor information – thick green border marks the most efficient methods for voxelwise and brain-wide analysis respectively – some pre-processings have been omitted

$\lambda \circledast S \stackrel{def}{=} \exp(\lambda \cdot \log S)$. Computations on the tensors can thus be converted into Euclidean computations on the tensor logarithms. Moreover, this yields a one-to-one mapping between tensor space and $\mathbb{R}^6$ with the operator $V(\boldsymbol{S}) = (\log(\boldsymbol{S})_{1,1}, \sqrt{2}\cdot\log(\boldsymbol{S})_{2,1}, \log(\boldsymbol{S})_{2,2}, \sqrt{2}\cdot\log(\boldsymbol{S})_{3,1}, \sqrt{2}\cdot\log(\boldsymbol{S})_{3,2}, \log(\boldsymbol{S})_{3,3})^T$ ([4], [13]). Computations are therefore further simplified by working on the resulting vectors.

Registration and Smoothing. The classical pipeline involves: (1) extracting scalar values (FA), (2) constructing a study-specific scalar template, (3) nonlinear registering the individual scalar maps to the scalar template, (4) smoothing the resulting scalar maps to improve the power of any further statistical analysis.

To retain the full tensor information in our pipeline, we chose to use the deformations computed from FA registration in step 3, and apply them to the tensor fields (*cf* Fig. 1). During this warping, each tensor must be reoriented to remain consistent with the anatomy [2]. Moreover, the interpolation is performed in the Log-Euclidean domain. Smoothing is performed in the LE domain by applying an isotropic Gaussian convolution kernel to the vector representation of the tensor logarithms.

Feature Selection using Voxelwise Statistical Analysis. Brain-wide methods suffer from the *curse of dimensionality*: as neuroimaging data provides a very large number of features and a low number of subjects, there is a very high risk of overfitting the training data resulting in poor generalization. It is therefore necessary to select a smaller set of relevant features to train the machine learning system. To perform this feature selection step, we use voxel-based regression analysis which allows us to rank voxels according to their statistical significance. A regression is carried out in each voxel yielding an F-value. In the full-tensor case, the regression is multivariate and the F-value is computed via the Pillai-Bartlett's trace [8]. To visualize the results, the p-values maps are thresholded to show only the most statistically significant regions. To use as feature selection, we retain the best m voxels for the rest of the analysis.

Brain-Wide Machine Learning. Some complex brain patterns where several regions interact cannot be detected by voxel-based analysis. Multivariate methods that analyze voxels jointly are thus required to detect these patterns. In this paper, we call these multivariate methods *brain-wide methods* to avoid confusion with multivariate voxel-based methods that work with multivariate objects on only one voxel. The regression method used in this paper is kernel ridge regression (see *e.g.* [9]).

Based on a dataset made of N subjects, we denote by $\boldsymbol{y}$ the $N \times 1$ vector of responses. $(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N)$ are the N subject images (typically $\boldsymbol{x}_i$ is of size $m \times 1$ for FA-only, $m \times 6$ for tensors, where m is the number of selected voxels). $k(\boldsymbol{x}_i, \boldsymbol{x}_j)$ is the kernel function that transforms the data from the data space to the feature space and $\boldsymbol{K}$ is the $N \times N$ kernel matrix where $\boldsymbol{K}_{ij} = k(\boldsymbol{x}_i, \boldsymbol{x}_j)$. The most simple kernel function is the linear kernel, which corresponds to the inner product of the data space. The goal of the kernel ridge regression is to minimize the penalized residual sum of squares: $\hat{f} = \arg\min_{f \in \mathcal{H}_k} \sum_{i=1}^{N} (y_i - f(\boldsymbol{x}_i))^2 + \lambda ||f||^2_{\mathcal{H}_k}$.

As $\hat{f}$ can be expanded into $\hat{f}(\boldsymbol{x}) = \sum_{i=1}^{N} \alpha_i K(\boldsymbol{x}_i, \boldsymbol{x})$, this is equivalent to $\hat{\boldsymbol{\alpha}} = \arg\min_{\boldsymbol{\alpha} \in \mathbb{R}^N} (K\boldsymbol{\alpha} - \boldsymbol{y})^T (K\boldsymbol{\alpha} - \boldsymbol{y}) + \lambda \boldsymbol{\alpha}^T \boldsymbol{K} \boldsymbol{\alpha}$. The ridge regression solution reads: $\hat{\boldsymbol{\alpha}} = (\boldsymbol{K} + \lambda N \boldsymbol{I})^{-1} \boldsymbol{y}$. To validate the method, we perform a leave-one-out (LOO) cross-validation. Prediction accuracy is assessed by the explained variance $R^2 = 1 - \frac{\sum_{i=1}^{N} (y_i - \hat{y}_i)^2}{\sum_{i=1}^{N} (y_i - \overline{y})^2}$ where $\overline{y} = \frac{1}{N} \sum_{i=1}^{N} y_i$ is the mean of the true values and $\hat{y}_i = \hat{f}(\boldsymbol{x}_i)$. $R^2 = 1$ corresponds to a perfect prediction, while we might get $R^2 < 0$ if the prediction error is high.

Fig. 1 presents the two pipelines. In the FA-only pipeline, the data space is simply $\mathbb{R}^m$ where m is the number of voxels selected to predict the score. We use the linear kernel and the regression is thus equivalent to a classical linear ridge regression. In the full-tensor pipeline, the data space is the space of vectors of tensors $(\mathcal{S}ym_3^+)^m$ where m is the number of voxels selected to predict the score: each subject is represented by a vector of m tensors. We construct an inner product on this space using the vector representation V of the tensors : with $\boldsymbol{x}_i = (\boldsymbol{x}_{i,1}, \ldots, \boldsymbol{x}_{i,m})^T$ and $\boldsymbol{x_j} = (\boldsymbol{x}_{j,1}, \ldots, \boldsymbol{x}_{j,m})^T$ members of $(\mathcal{S}ym_3^+)^m$, $\langle \boldsymbol{x}_i, \boldsymbol{x}_j \rangle = \sum_{l=1}^{m} V(\boldsymbol{x}_{i,l})^T V(\boldsymbol{x}_{j,l})$ is an inner product on $(\mathcal{S}ym_3^+)^m$. We can thus compute a linear kernel for the kernel ridge regression.

3 Application to Autism Spectrum Disorders (ASD)

In this section we present the dataset used and compare the results obtained by both approaches for pre-processings, voxelwise regression analysis and brain-wide machine learning.

Data Acquisition. The dataset consisted of diffusion MRI scans and autism severity evaluation for thirty children and adolescents with ASD, and lacks control subjects. The autistic syndrome diagnosis was based on DSM IV-TR criteria (Diagnostic and Statistical Manual of Mental Disorders [3]) and its severity was evaluated with the Autism Diagnostic Interview-Revised (ADI-R) algorithm [11] (a high global ADI-R score means severe autism). For all regression analyses, we used a modified version of the global ADI-R score (mADI-R) previously used by Gendry Meresse et al. in [7].

All diffusion MRI scans were acquired on a GE Signa 1.5 MRI system (General Electric) using a birdcage head coil. The sequence was a dual spin echo echo-planar imaging sequence (echo times $TE = 70$ ms; repetition times $TR = 8800$ ms; 60 axial slices, 2 mm slice thickness, in plane resolution 1.875×1.875 mm, matrix 128×128). For each slice, five images without diffusion weighting ($b = 0$), and 41 images with diffusion gradients ($b = 1500$ s/mm^2) applied along 41 non-collinear directions were acquired. We used the diffusion model pipeline of BrainVISA 3.1[1] to perform a correction of the diffusion images for echo-planar distortions and the evaluation of the diffusion tensor from these corrected images.

Registration. We extracted FA values from the tensor images with the diffusion model pipeline of BrainVISA 3.1[1]. All FA images were nonlinearly registered to

[1] BrainVISA, http://www.brainvisa.info

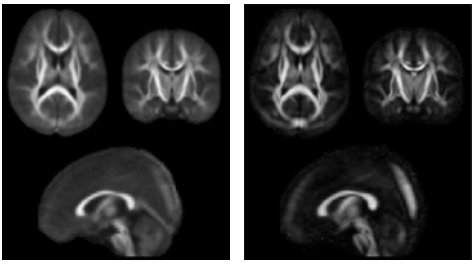

FA only FA + full tensor

Fig. 2. Left: mean of the 30 FA images after FA-only registration – Right: FA of the Log-Euclidean mean of the 30 tensor images after FA + full tensor registration

a study-specific FA template with SPM5[2] running on Matlab 7.6. Before further analysis, all FA maps were smoothed with an 8-mm isotropic Gaussian filter. To obtain registered full-tensor images, we used TTK[3] to apply the deformations computed during the FA registration to the tensor fields (*cf* Fig. 1). This warping was done with a log-Euclidean interpolation scheme and a finite-strain reorientation strategy [2]. All maps were smoothed in the log-Euclidean domain with an 8-mm isotropic Gaussian filter. Using the same level of smoothing in the two approaches enables comparison of the results.

To assess the accuracy of the two methods, we computed the FA mean of the registered images before smoothing that we compared to the FA extracted from the full-tensor mean image (Fig. 2). The FA map obtained from the full-tensor registration is more contrasted than the one obtained from the FA-only registration. This means that the individual images are better aligned to one another, even though the same deformation field is used. This result confirms the importance of the interpolation scheme and the reorientation of tensors.

Voxelwise Regression Analysis. Analyses were performed on the whole brain, with an uncorrected voxelwise threshold of $p < 0.005$. The full-tensor approach finds all of the regions found by the classical FA-only approach and known to be involved in ASD such as superior temporal regions and inferior frontal regions (Fig. 3, top). Moreover, some significant regions highlighted by the full-tensor approach were not found by the FA-only method (Fig. 3, bottom). In particular, the corpus callosum and the anterior commissure have been mentioned in other studies of autism: [1,12,14,15]. The fact that these regions were detected with the full-tensor approach but not with the FA-only approach confirms that the relevant information lies in the orientation of the tensor as we hypothesized, and not only in its eigenvalues.

Brain-Wide Score Prediction. We tested the prediction accuracy with a LOO cross validation. The accuracy was computed for several values of m

[2] Statistical Parametric Mapping, http://www.fil.ion.ucl.ac.uk/spm/
[3] The Tensor Toolkit, https://gforge.inria.fr/projects/ttk/

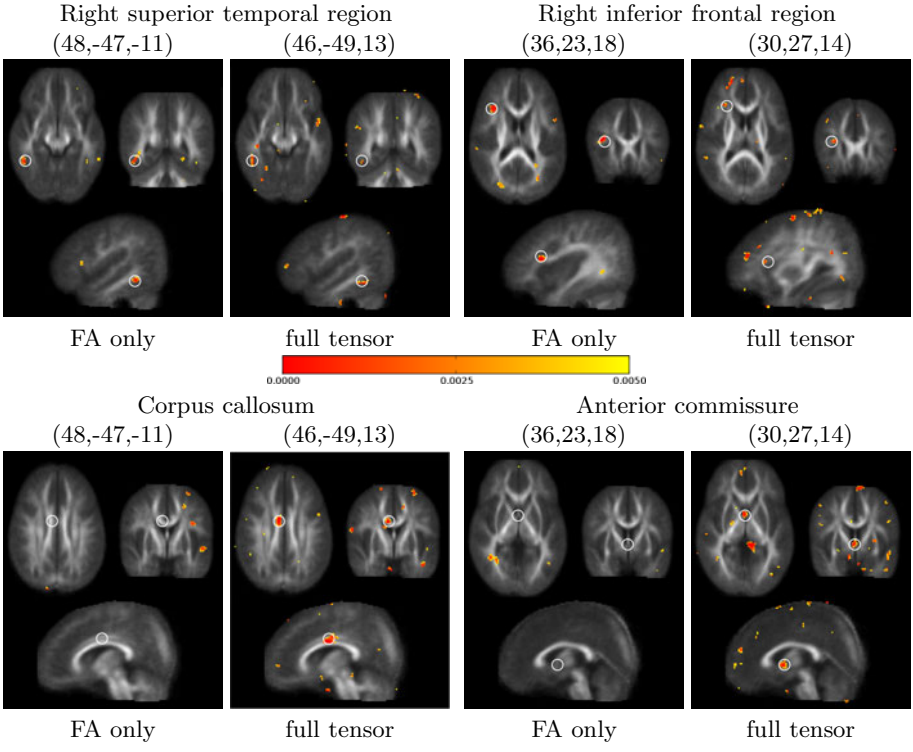

Fig. 3. Thresholded p-values maps for some regions of the brain ($p < 0.005$) – coordinates of the most significant voxel are in Talairach space – Top row: consistent findings of the two approaches in the right superior temporal and inferior frontal regions – Bottom row: the full-DTI approach identified differences in the corpus callosum and the anterior commissure, that were missed by the FA-only analysis.

(number of selected voxels) and the p-values of the resulting explained variances were corrected across the multiple experiments with a maxT procedure [16].

Fig. 4 (left) shows that the FA-only approach does not yield a satisfying result: the cross-validated R^2 is always negative. Taking the full tensor into account does reach an explained variance of 0.50 with $p < 0.05$. However, Fig. 4 (right) shows that the full-tensor approach is more prone to overfitting than the FA-only one. When using the full tensors the dimension of the data space is 6 times the number of voxels selected, so the method overfits as soon as $6m > 29$ (number of training subjects). Therefore we tried an intermediary alternative between the two approaches: we applied deformations to the tensors but extracted the FA for further analysis. This approach achieves a fair trade-off: the brain-wide kernel regression yields a cross-validated R^2 of 0.64 (with a significance of $p < 0.005$). Extracting a scalar value seems to counteract the overfitting problem. Moreover, extracting the FA to perform the brain-wide prediction can be seen as a way to construct a kernel from the tensor values.

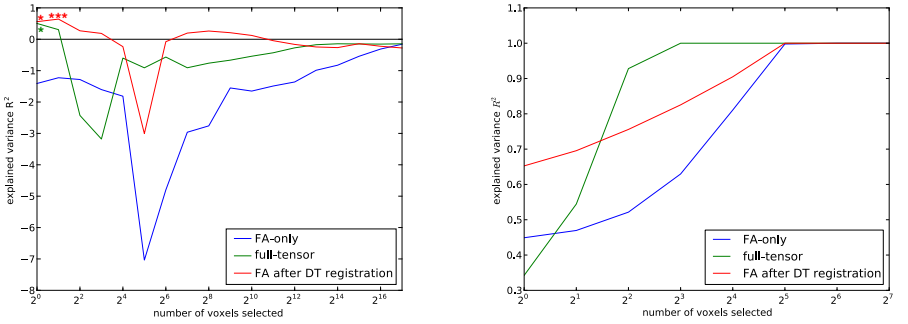

Fig. 4. Test R^2 (left) and train R^2 (right) for different numbers of voxels selected – Significance codes: 0 *** 0.005 ** 0.01 * 0.05 1 – Full-tensor approach overfits rapidly (100% of train explained variance is reached with only $2^3 = 8$ voxels). This contrasts with FA-only that needed 32 voxels to obtain a perfect fit. Finally, the FA after applying deformations to DTI showed an intermediate behavior.

4 Conclusion

The main contribution of this paper was to adapt the DTI analysis pipeline to use the full tensor information in a log-Euclidean framework, with a view to performing brain-wide machine learning on tensor images. We confirmed already observed results: the information contained in the full tensor enhances the correspondence between images and full-tensor voxelwise analysis is more sensitive than the FA-only analysis. We showed that the full-tensor brain-wide score prediction was more prone to overfitting than scalar-based prediction, but that computing the FA after applying deformations to the full tensors yielded better results than the traditional FA-only pipeline.

References

1. Alexander, A.L., Lee, J.E., Bigler, E.D., DuBray, M.B., Froehlich, A., Lange, N., Fletcher, T.P., Chung, M.K., Lainhart, J.E.: White Matter is Diffusely Affected in Autism. In: Proc. 17th ISMRM, vol. 23, p. 639 (2009)
2. Alexander, D.C., Pierpaoli, C., Basser, P.J., Gee, J.C.: Spatial transformations of diffusion tensor magnetic resonance images. IEEE TMI 20(11), 1131–1319 (2001)
3. APA: Diagnostic and Statistical Manual of Mental Disorders, 4th edn. (2000)
4. Arsigny, V., Fillard, P., Pennec, X., Ayache, N.: Log-Euclidean metrics for fast and simple calculus on diffusion tensors. Magnetic Resonance in Medicine 56(2), 411–421 (2006)
5. Brun, C., Nicolson, R., Leporé, N., Chou, Y.-Y., Vidal, C.N., DeVito, T.J., Drost, D.J., Williamson, P.C., Rajakumar, N., Toga, A.W., Thompson, P.M.: Mapping brain abnormalities in boys with autism. Human Brain Mapping 30(12), 3887–3900 (2009)

6. Commowick, O., Fillard, P., Clatz, O., Warfield, S.K.: Detection of DTI white matter abnormalities in multiple sclerosis patients. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part I. LNCS, vol. 5241, pp. 975–982. Springer, Heidelberg (2008)
7. Gendry Meresse, I., Zilbovicius, M., Boddaert, N., Robel, L., Philippe, A., Sfaello, I., Laurier, L., Brunelle, F., Samson, Y., Mouren, M.C., Chabane, N.: Autism severity and temporal lobe functional abnormalities. Annals of Neurology 58(3), 466–469 (2005)
8. Hand, D.J., Taylor, C.C.: Multivariate Analysis of Variance and Repeated Measures: A Practical Approach for Behavioural Scientists. Chapman and Hall, Boca Raton (1987)
9. Hastie, T., Tibshirani, R., Friedman, J.H.: The Elements of Statistical Learning. Springer Series in Statistics. Springer, New York (2001)
10. Leporé, N., Brun, C., Chou, Y.Y., Chiang, M.C., Dutton, R.A., Hayashi, K.M., Luders, E., Lopez, O.L., Aizenstein, H.J., Toga, A.W., Becker, J.T., Thompson, P.M.: Generalized tensor-based morphometry of HIV/AIDS using multivariate statistics on deformation tensors. IEEE TMI 27(1), 129–141 (2008)
11. Lord, C., Rutter, M., Le Couteur, A.: Autism Diagnostic Interview-Revised: a revised version of a diagnostic interview for caregivers of individuals with possible pervasive developmental disorders. J. Autism. Dev. Disord. 24(5), 659–685 (1994)
12. Noriuchi, M., Kikuchi, Y., Yoshiura, T., Kira, R., Shigeto, H., Hara, T., Tobimatsu, S., Kamio, Y.: Altered white matter fractional anisotropy and social impairment in children with autism spectrum disorder. Brain Research, 1–9 (2010)
13. Pennec, X., Fillard, P., Ayache, N.: A riemannian framework for tensor computing. International Journal of Computer Vision 66, 41–66 (2006)
14. Stanfield, A.C., McIntosh, A.M., Spencer, M.D., Philip, R., Gaur, S., Lawrie, S.M.: Towards a neuroanatomy of autism: a systematic review and meta-analysis of structural magnetic resonance imaging studies. European Psychiatry 23(4), 289–299 (2008)
15. Thomas, C., Humphreys, K., Jung, K.-J., Minshew, N., Behrmann, M.: The anatomy of the callosal and visual-association pathways in high-functioning autism: A DTI tractography study. Cortex, 1–11 (2010)
16. Westfall, P.H., Young, S.S.: Resampling-based multiple testing: Examples and methods for p-value adjustment. Wiley, Chichester (1993)
17. Whitcher, B., Wisco, J.J., Hadjikhani, N., Tuch, D.S.: Statistical group comparison of diffusion tensors via multivariate hypothesis testing. Magnetic Resonance in Medicine 57(6), 1065–1074 (2007)

Fiber Modeling and Clustering Based on Neuroanatomical Features

Qian Wang[1,2], Pew-Thian Yap[2], Guorong Wu[2], and Dinggang Shen[2]

[1] Department of Computer Science, University of North Carolina at Chapel Hill
qianwang@cs.unc.edu
[2] Department of Radiology and BRIC, University of North Carolina at Chapel Hill
{ptyap,grwu,dgshen}@med.unc.edu

Abstract. DTI tractography allows unprecedented understanding of brain neural connectivity in-vivo by capturing water diffusion patterns in brain white-matter microstructures. However, tractography algorithms often output hundreds of thousands of fibers, rendering the computation needed for subsequent data analysis intractable. A remedy is to group the fibers into bundles using fiber clustering techniques. Most existing fiber clustering methods, however, rely on fiber geometrical information only by viewing fibers as curves in the 3D Euclidean space. The important neuroanatomical aspect of the fibers is mostly ignored. In this paper, neuroanatomical information is encapsulated in a feature vector called the associativity vector, which functions as the "fingerprint" for each fiber and depicts the connectivity of the fiber with respect to individual anatomies. Using the associativity vectors of fibers, we model the fibers as observations sampled from multivariate Gaussian mixtures in the feature space. An expectation-maximization clustering approach is then employed to group the fibers into 16 major bundles. Experimental results indicate that the proposed method groups the fibers into anatomically meaningful bundles, which are highly consistent across subjects.

1 Introduction

Diffusion Tensor Imaging (DTI) has become a popular imaging modality in exploring brain circuitry in-vivo by capturing water diffusion patterns in brain tissues. Water molecules are more likely to diffuse parallel along neural pathways, since myelin sheaths of axons act as barriers and restrict the mobility of water molecules along directions perpendicular to the neural pathways. In a DT image, each voxel records the anisotropic water diffusion pattern at a specific location of the brain using a second-order tensor. The principal eigenvector of the tensor, corresponding to the maximal eigenvalue, indicates the major direction along which water molecules are diffusing. By tracing along these directions, neural tracts can thus be delineated via a process called tractography [1].

Several fiber tractography algorithms have been proposed in the literature. An intuitive and straightforward streamline-based tractography approach begins fiber tracing from a given seed point and grows smoothly along the principal eigenvectors of the underlying tensor field. This tracking procedure continues until either the fiber enters regions with low fractional anisotropy (FA), where directional information is

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 17–24, 2011.

no longer reliable, or when the curvature of the fiber pathway is too large. With whole brain seeding, a typical fiber tractography algorithm can yield fibers numbering in the order of 10^3-10^6.

Fiber tractography provides an effective approach in visualizing and analyzing brain connectivity, and thus is important in clinical applications. However, the task of analyzing fibers yielded by tractography is non-trivial. The massive amount of fibers often renders subsequent analyses difficult, and makes information provided by the fibers not immediately decipherable. One possible solution is to automatically partition fibers into dozens of bundles, each of which contains fibers that are characterized by similar structural and functional behaviors. Analyses can then be performed on the bundles, instead of the individual fibers.

Existing fiber clustering methods are mostly geometry based approaches. Fibers are typically viewed as a set of curves in the 3D space. In [2], for example, point-to-point correspondences are first established between fibers. The similarity between a pair of fibers is then defined as the ratio of the length of the corresponding segment over the entire fiber length. The correspondence ratio between fibers is maximized when the two fibers are identical, and vanishes if pairwise correspondence is minimal. Similar ideas based on point-to-point correspondences between fibers are applied in [3, 4]. Fiber similarity can also be calculated by counting the number of voxels that the two fibers share [5]. Other approaches involve extracting features from the spatial distributions or parametric representations of fibers [6-9]. The Hausdorff distance, as well as its variants, has also been widely applied to fiber clustering [10-14].

Geometry based methods, however, ignores the neuroanatomical characteristics of fibers and results in bundles which are not clear for interpretations from an anatomical and connectivity point of view. It is therefore important to incorporate anatomical information when performing fiber clustering. For this purpose, Maddah et. al. [7] report an atlas where individual bundles have already been delineated and labeled by experts. Based on this atlas, fibers are clustered and labeled with the known bundles according to geometric similarities. Alternative approaches can also be found in [15, 16], where bundles are tagged according to the regions they connect.

In this paper, we propose a novel fiber clustering scheme that is solely based on the neuroanatomical features of fibers. After performing whole brain tractography, we parcellate the brain with a number of ROIs and diffuse the ROIs according to the underlying fibers. ROI diffusion helps more robustly determine the relationship between the anatomical regions (as delineated by the ROIs) and the fibers by using this fuzzy formulation. We are then able to acquire the *associativity vector* for each fiber to describe its connectivity pattern with respect to all ROIs of different anatomies. Each entry of the associativity vector represents the likelihood of the fiber being connected to a particular ROI. We further model each bundle as a multivariate Gaussian mixture based on the associativity vectors of the fibers belonging to the bundle. An expectation-maximization (EM) approach is thus employed to group the fibers as belonging to one of the bundles or as outliers. Experimental results show that the proposed method can achieve consistent clustering results across individual subjects, implying its potential usage in analyzing a population of DTI data.

2 Method

We will first define the associativity vector (Section 2.1) and then introduce the multivariate-Gaussian-mixture-based EM clustering approach (Section 2.2).

2.1 Neuroanatomical Features of Fibers

Fibers connect different anatomical regions of the brain, forming a connectivity pattern that describes brain circuitry. For annotating different regions of the brain, we warp the "Eve" atlas [17] (called *atlas* for brevity in the rest of the paper) to the native space of each individual subject via non-rigid registration. The atlas contains T1 image, tensor field, and a set of manually delineated ROIs for important anatomical structures of the brain.

Suppose that there are M ROIs, we then define the *associativity vector* $L_i = (\ell_{i1}, \ell_{i2}, \cdots, \ell_{i\mathrm{M}})$ for fiber i, where the entry ℓ_{ij} measures the relationship between fiber i and the j-th ROI. Entry ℓ_{ij} can be set to 1 if any segment of fiber i lies within ROI j, and left unset otherwise. However, this binary formulation, though simple, would result in sparse associativity vectors, pose challenges in estimating fiber distances, and increase the tendency of the clustering algorithms to be trapped in local minima. Moreover, fibers that are prematurely terminated a little short of reaching the ROI due to imaging noise might also be penalized.

In view of this, we propose a fuzzy ROI spatial confidence map by allowing the ROI to grow, or to diffuse, along directions indicated by the fibers. For this purpose, we adopt the fast marching (FM) approach [18]. As illustrated in Fig. 1(a), the initial ROI (dark red area) starts diffusing since time t_s and terminates at t_e. The initial ROI surface (dashed green curve) moves outward following the underlying fibers from upper-left to lower-right. At location x on the surface, diffusion proceeds along the surface normal direction. The contribution of a fiber to the ROI diffusion at x is defined as the inner product of the tangential direction of the fiber and the surface normal. Summing up contributions from all fibers traversing x, we obtain the overall diffusion velocity $v(\mathrm{x})$. By iteratively solving the Eikonal equation $v\|\nabla_{\mathrm{x}}\tau\|=1$ in FM, we then acquire the time value when a certain location is traversed by the ROI surface. The location geodesically closer to the initial ROI will be traversed by the surface at an earlier time. We thus invert all recorded time values and rescale them to the range of [0, 1], giving a spatial confidence map of the ROI. In Fig. 1(a), for example, dark red indicates locations near the initial ROI and dark blue indicates locations far away.

We have also provided a real example in Fig. 1(b)-(d). In Fig. 1(b), the bundle connecting the left superior-frontal gyrus (SPG-L) to the right superior-frontal gyrus (SPG-R) is overlaid in red on the FA map. The ROI of SPG-L diffuses accordingly, and results in the spatial confidence map in panel (c) where the dark red area is the initial ROI. The diffusion pattern of SPG-R is similar and displayed in Fig. 1(d).

All ROIs are diffused according to the fibers given by whole-brain tractography, and resulting in their own spatial confidence maps. We allow the same amount of time for diffusion of all ROIs, in order to assure that the re-scaled spatial confidence maps are comparable. For the associativity vector L_i of fiber i, the entry ℓ_{ij} is

defined as the maximal spatial confidence value that fiber i comes upon the spatial confidence map of ROI j. The atlas consists of 130 ROIs (usually at the scale of gyrus/sulcus), but we only need to detect 16 major fiber bundles as in the following. To reduce the redundancy between entries of the associativity vector, we integrate smaller ROIs into larger ones. In particular, the ROIs we have used in this work include the left/right frontal lobes, the left/right central areas, the left/right parietal lobes, the left/right occipital lobes, the left/right temporal lobes, the left/right sub-cortical areas, the brainstem, and the cerebellum.

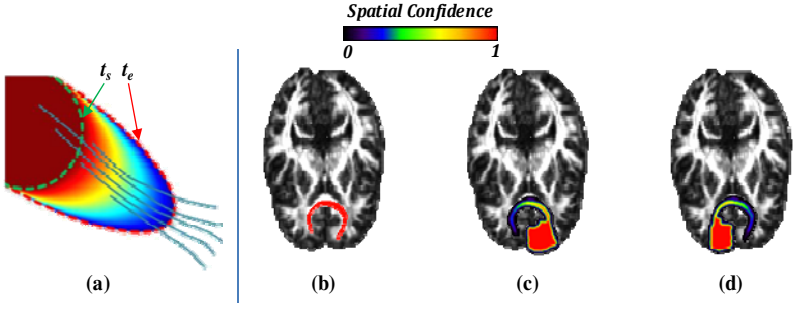

Fig. 1. In (a), the ROI is diffused along the underlying bundle. The diffusion starts from the green dashed curve (at t_s) and terminates at the red dashed curve (at t_e). The spatial confidence map of the ROI can thus be calculated using the traversing time of the diffusion surface. In (b), the bundle connecting SPG-L and SPG-R is overlaid in red on the FA map. The spatial confidence maps of SPG-L and SPG-R after diffusion according to the bundle in (b), are shown in (c) and (d), respectively.

2.2 Bundle Modeling and Clustering

Fibers in the same bundle typically share similar neural pathways and generate highly correlated associativity vectors. As the result, we can model the distribution of fibers in a bundle using their associativity vectors with a multivariate Gaussian mixture. An expectation-maximization (EM) approach based on parametric bundle models can then be used for fiber clustering. For initialization of the EM based clustering, we have manually reproduced 16 major bundles (listed in Table 1) in the atlas space following [19]. Parameters estimated from these bundles will work as constraints to guide fiber clustering in each subject. The fibers of the subject will be either grouped into one of the 16 bundles or tagged as outliers. The parameters for the multivariate Gaussian mixture representing each bundle k $(1 \le k \le K)$ include the mean associativity vector μ_k and the covariance matrix S_k. For fiber i $(1 \le i \le N)$, we denote its associativity vector as L_i and its membership to bundle k as ω_{ik}.

E-step:
According to Bayes rule, the membership of fiber i to bundle k can be estimated as:

$$\omega_{ik} \leftarrow \frac{\alpha_k \cdot N_k(L_i|\mu_k, S_k)}{\sum_{m=1}^{K} \alpha_m \cdot N_m(L_i|\mu_m, S_m)} \tag{1}$$

where $N_k(L_i|\mu_k, S_k)$ is the multivariate Gaussian distribution centered at μ_k with covariance S_k. For the first E-step, the parameters (μ_k, S_k) are directly introduced from the manually delineated bundles in the atlas, while α_k denotes the number of fibers in bundle k.

Table 1. The list of the 16 target bundles

Abbr.	Description	Abbr.	Description
Fminor	Forceps minor	**SLF-R**	Superior longitudinal fasciculus (right)
Fmajor	Forceps major	**SLF-L**	Superior longitudinal fasciculus (left)
CST-R	Corticospinal tract (right)	**ILF-R**	Inferior longitudinal fasciculus (right)
CST-L	Corticospinal tract (left)	**ILF-L**	Inferior longitudinal fasciculus (left)
ATR-R	Anterior thalamic radiation (right)	**UNC-R**	Uncinate fasciculus (right)
ATR-L	Anterior thalamic radiation (left)	**UNC-L**	Uncinate fasciculus (left)
IFO-R	Inferior fronto-occipital fasciculus (right)	**CB-R**	Cingulum (right)
IFO-L	Inferior fronto-occipital fasciculus (left)	**CB-L**	Cingulum (left)

M-step:
The parameters (μ_k, S_k) are updated in the M-step. Based on the membership values estimated in the E-step, we first compute the mixture weighting factor for each bundle $\alpha_k \leftarrow \sum_{i=1}^{N} \omega_{ik}/N$ and then the bundle parameters:

$$\mu_k \leftarrow \lambda\mu_k + \left(\frac{1-\lambda}{\sum_{i=1}^{N} \omega_{ik}}\right) \sum_{i=1}^{N} \omega_{ik} L_i \tag{2}$$

$$S_k \leftarrow \left(\frac{1}{\sum_{i=1}^{N} \omega_{ik}}\right) \sum_{i=1}^{N} \omega_{ik}(L_i - \mu_k)(L_i - \mu_k)^t \tag{3}$$

We set λ of (2) to 0.95 to prevent drastic shifts of μ_k, which might occur due to large fiber variation. Subsequent E-steps and M-steps will be iteratively executed until convergence. Fibers whose highest memberships are lower than a predefined threshold are regarded as outliers.

3 Experimental Results

A total of 15 healthy subjects were used to evaluate the proposed fiber clustering method. The diffusion weighted data were acquired using a Siemens Allegra scanner (b=1000s/mm^2, flip angle 90°, TR/TE=13,640/82ms, matrix 128×128, FoV 256×256mm^2, slice thickness 2mm, 80 contiguous slices). In Fig. 2(a), the manually delineated bundles in the atlas space, following protocols in [19], are annotated in different colors. It is worth noting that all right hemisphere bundles are excluded from Fig. 2(a) and only 9 target bundles are shown for visualization. In panels (b)-(e), we show the clustering results of 4 randomly selected subjects. In each panel, callosal fibers (Fminor and Fmajor) are only shown in the top image and excluded in the bottom image for better visualization of other bundles (i.e., CB). In (b)-(e), bundles in the right hemispheres are also removed for visualization clarity, as we have done in (a). Overall, for all 4 subjects shown in Fig. 2, the clustering results are consistent with the atlas and across subjects.

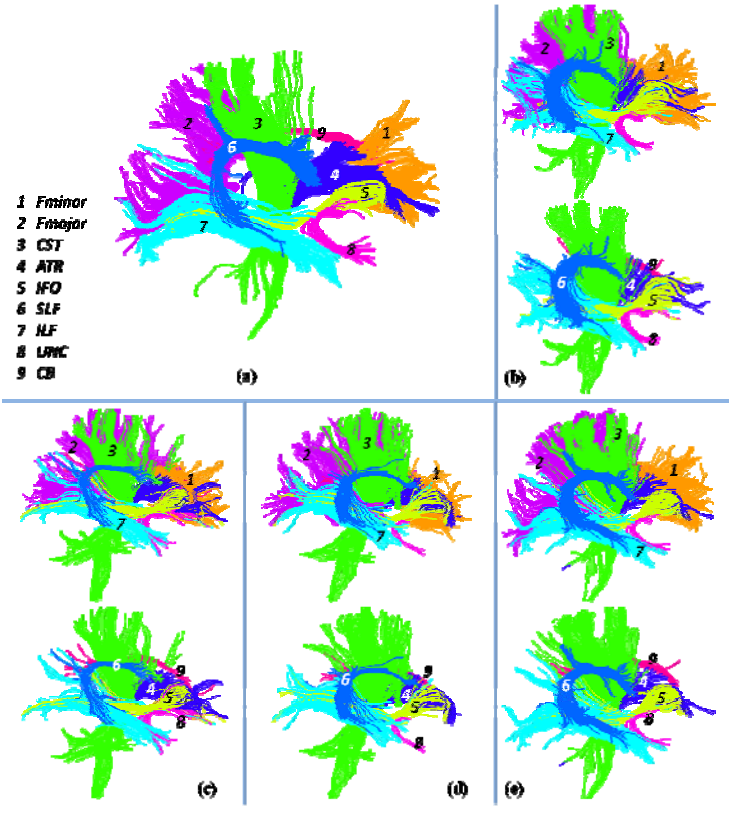

Fig. 2. Target bundles in the atlas are shown and annotated in (a). Results of 4 randomly selected subjects are shown in (b)-(e). In each panel, bundles in the right hemisphere are removed for clarity. Callosal fibers are shown only in the top images of (b)-(e), and are removed from the bottom images for better visualization.

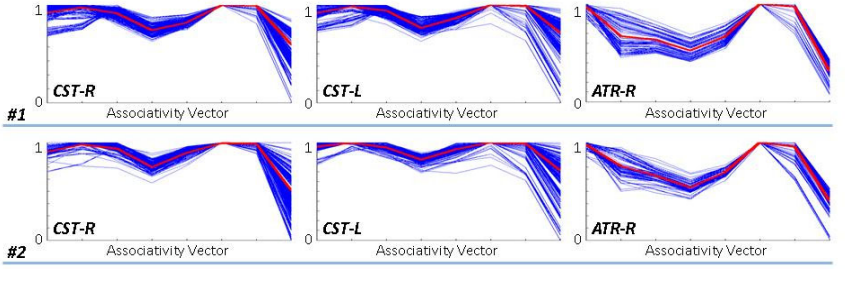

Fig. 3. The associativity vectors for CST-R, CST-L, and ATR-R from two different subjects. The red curves are the mean associativity vectors.

In Fig. 3, we plot associativity vectors of fibers in three bundles (CST-R, CST-L, ATR-R) from two different subjects, respectively. Since the entries of the associativity vectors corresponding to the opposite hemisphere are mostly zeros, we remove them from the associativity vector before plotting. The red curves are the mean associativity vectors. For both subjects, patterns between corresponding bundles are similar to each other, indicating that clustering results are consistent across subjects. We can also observe reflectional symmetry between CST-R and CST-L. Moreover, the associativity vectors for ATR-R, which is neighboring to CST-R, show a clearly different pattern with that of CST-R, suggesting that they can be sufficiently differentiated by the bundle models during fiber clustering.

For corresponding bundles of two individual subjects, we measured their similarity using the mean of the Pearson correlation coefficients of associativity vectors between all fiber pairs. The average correlations of 16 bundles across 15 subjects, as well as standard deviations, are listed in Table 2. In most bundles, the very high correlation scores underline that clustering results are consistent across subjects, which is important to fiber clustering.

Table 2. The mean correlation of associativity vectors, as well as the standard deviation, across all 15 subjects

Bundle	Correlation	Bundle	Correlation	Bundle	Correlation
Fminor	0.891±0.016	**IFO-R**	0.924±0.008	**UNC-R**	0.969±0.009
Fmajor	0.663±0.048	**IFO-L**	0.935±0.032	**UNC-L**	0.975±0.013
CST-R	0.867±0.021	**SLF-R**	0.976±0.009	**CB-R**	0.946±0.020
CST-L	0.906±0.020	**SLF-L**	0.953±0.010	**CB-L**	0.924±0.034
ATR-R	0.898±0.022	**IFL-R**	0.943±0.015		
ATR-L	0.926±0.019	**IFL-L**	0.939±0.030		

4 Discussion

In this paper, we have applied a fuzzy formulation of the associativity vector for fiber representation and modeling. Based on the associativity vectors, we model fiber bundles using multivariate Gaussian mixtures. An EM clustering scheme is employed to group the fibers into 16 major bundles. Compared with geometric methods, the proposed method relies on fiber connectivity patterns only but not geometric similarity, as complex fiber structures (fanning, branching, etc.) are challenging for modeling using geometric information. Moreover, the 16 target bundles are selected due to their easy reproducibility, though no limitation of bundles is assumed. Experimental results show that clustering results are consistent with the atlas and across subjects, which is an important property in applying this method to analysis of population data. More extensive assessment of the algorithm, however, will need to be performed. For this purpose, we will use histological data obtained from excised brains of canine subjects, which will soon be available to us from one of our projects.

References

1. Mori, S., Zijl, P.C.M.v.: Fiber Tracking: Principles and Strategies - A Technical Review. NMR in Biomedicine 15, 468–480 (2002)
2. Ding, Z., Gore, J.C., Anderson, A.W.: Classification and Quantification of Neuronal Fiber Pathways Using Diffusion Tensor MRI. Magnetic Resonance in Medicine 49, 716–721 (2003)
3. Maddah, M., Grimson, W.E.L., Warfield, S.K., Wells, W.M.: A Unified Framework for Clustering and Quantitative Analysis of White Matter Fiber Tracts. Medical Image Analysis 12, 191–202 (2008)
4. Maddah, M., Wells III, W.M., Warfield, S.K., Westin, C.-F., Grimson, W.E.L.: Probabilistic Clustering and Quantitative Analysis of White Matter Fiber Tracts. In: Karssemeijer, N., Lelieveldt, B. (eds.) IPMI 2007. LNCS, vol. 4584, pp. 372–383. Springer, Heidelberg (2007)
5. Klein, J., Bittihn, P., Ledochowitsch, P., Hahn, H.K., Konrad, O., Rexilius, J., Peitgen, H.-O.: Grid-based Spectral Fiber Clustering. In: SPIE Medical Imaging (2007)
6. Maddah, M., Crimson, W.E.L., Warfield, S.K.: Statistical modeling and EM clustering of white matter fiber tracts. In: ISBI (2006)
7. Maddah, M., Mewes, A.U.J., Haker, S., Grimson, W.E.L., Warfield, S.K.: Automated Atlas-Based Clustering of White Matter Fiber Tracts from DTMRI. In: Duncan, J.S., Gerig, G. (eds.) MICCAI 2005. LNCS, vol. 3749, pp. 188–195. Springer, Heidelberg (2005)
8. Brun, A., Knutsson, H., Park, H.-J., Shenton, M.E., Westin, C.-F.: Clustering Fiber Traces Using Normalized Cuts. In: Barillot, C., Haynor, D.R., Hellier, P. (eds.) MICCAI 2004. LNCS, vol. 3216, pp. 368–375. Springer, Heidelberg (2004)
9. Klein, J., Stuke, H., Stieltjes, B., Konrad, O., Hahn, H.K., Peitgen, H.-O.: Efficient fiber clustering using parameterized polynomials. In: SPIE Medical Imaging (2008)
10. Corouge, I., Gouttard, S., Gerig, G.: Towards a Shape Model of White Matter Fiber Bundles Using Diffusion Tensor MRI. In: ISBI (2004)
11. Gerig, G., Gouttard, S., Corouge, I.: Analysis of Brain White Matter via Fiber Tract Modeling. In: IEEE EMBS (2004)
12. O'Donnell, L.J., Westin, C.F.: Automatic Tractography Segmentation Using a High-Dimensional White Matter Atlas. IEEE Trans. Medical Imaging 26, 1562–1575 (2007)
13. Zhang, S., Correia, S., Laidlaw, D.H.: Identifying White-Matter Fiber Bundles in DTI Data Using an Automated Proximity-Based Fiber-Clustering Method. IEEE Trans. Visualization and Computer Graphics 14, 1044–1053 (2008)
14. Maddah, M., Zollei, L., Grimson, W.E.L., Wells, W.M.: Modeling of anatomical information in clustering of white matter fiber trajectories using Dirichlet distribution. In: MMBIA (2008)
15. Xia, Y., Turken, A.U., Whitfield-Gabrieli, S.L., Gabrieli, J.D.: Knowledge-Based Classification of Neuronal Fibers in Entire Brain. In: Duncan, J.S., Gerig, G. (eds.) MICCAI 2005. LNCS, vol. 3749, pp. 205–212. Springer, Heidelberg (2005)
16. Li, H., Xue, Z., Guo, L., Liu, T., Hunter, J., Wong, S.T.C.: A Hybrid Approach to Automatic Clustering of White Matter Fibers. NeuroImage 49, 1249–1258 (2010)
17. Wakana, S., Jiang, H., Nagae-Poetscher, L.M., van Zijl, P.C.M., Mori, S.: Fiber Tract-based Atlas of Human White Matter Anatomy. Radiology 230, 77–87 (2004)
18. Sethian, J.A.: A Fast Marching Level Set Method for Monotonically Advancing Fronts. PNAS 93, 1591–1595 (1996)
19. Wakana, S., Caprihan, A., Panzenboeck, M.M., Fallon, J.H., Perry, M., Gollub, R.L., Hua, K., Zhang, J., Jiang, H., Dubey, P., Blitz, A., van Zijl, P., Mori, S.: Reproducibility of Quantitative Tractography Methods Applied to Cerebral White Matter. NeuroImage 36, 630–644 (2007)

Probabilistic Clustering and Shape Modelling of White Matter Fibre Bundles Using Regression Mixtures

Nagulan Ratnarajah[1], Andy Simmons[2], and Ali Hojjatoleslami[1]

[1] Medical Image Computing, School of Biosciences, University of Kent, U.K.
[2] Neuroimaging Department, Institute of Psychiatry, King's College London, U.K.

Abstract. We present a novel approach for probabilistic clustering of white matter fibre pathways using curve-based regression mixture modelling techniques in 3D curve space. The clustering algorithm is based on a principled method for probabilistic modelling of a set of fibre trajectories as individual sequences of points generated from a finite mixture model consisting of multivariate polynomial regression model components. Unsupervised learning is carried out using maximum likelihood principles. Specifically, conditional mixture is used together with an EM algorithm to estimate cluster membership. The result of clustering is a probabilistic assignment of fibre trajectories to each cluster and an estimate of cluster parameters. A statistical shape model is calculated for each clustered fibre bundle using fitted parameters of the probabilistic clustering. We illustrate the potential of our clustering approach on synthetic and real data.

Keywords: Probabilistic Clustering, Regression Mixture, Fibre Tractography, Shape Model.

1 Introduction

White matter (WM) fibre clustering is becoming an important field of clinical neuroscience research since it facilitates insights about anatomical structures in health and disease, allows clear visualizations of fibre tracts and enables the calculation of relevant statistics across subjects. A number of algorithms have been developed for clustering and labelling WM fibre bundles in DTI. Deterministic clustering algorithms [1-3] assign each trajectory to only one cluster, which may lead to biased estimators of cluster parameters if the clusters overlap. Probabilistic clustering algorithms [4], on the contrary, deal with the inherent uncertainty in assigning the trajectories to clusters. Quantitative parameters can be estimated by a weighted average over cluster members and thus more robust results may be obtained, which are less sensitive to the presence of outliers. Maddah et al. [4] proposed a probabilistic approach using a gamma mixture model and a distance map. This method assumes that the number of clusters is known and the approach requires manual user initialisation of the cluster centres. A problem for this approach was establishing correspondence between points.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 25–32, 2011.

In this paper, we propose a new geometrical framework to automatically cluster WM fibres into biologically meaningful neuro-tracts probabilistically. We are interested in starting with given fibre trajectories and determining whether these trajectories can be naturally clustered into groups. We investigate the model-based clustering of fibre trajectories, where each cluster is modelled as a prototype function with some variability around that prototype. A distinct feature of this model-based approach to clustering is that it produces a distinct model for each cluster. Since we are estimating smooth functions from noisy data it will be natural to use a probabilistic framework. Specifically we use mixtures of polynomial regression models as the basis of clustering. Multivariate clustering technique is used to describe the three dimensional propagations of the fibre trajectories which vary in length. We use a conditional mixture approach as it naturally allows for curves of variable length with unique measurement intervals and missing observations. Polynomial fits also take advantage of smoothness information present in the data. A regression model for each fibre bundle is constructed after performing probabilistic clustering. The probabilistic clustering algorithm is also capable of handling outliers in a principled way.

2 Probabilistic Model for White Matter Trajectories

2.1 Basic Definitions

Let V be a set of M 3-D fibre trajectories, where each trajectory v_i is an $n_i \times 3$ matrix containing a sequence of n_i 3-D points (x, y, z) in $\Re$. The associated $n_i \times 1$ vector u_i of ordered points from 0 to $n_i - 1$ correspond to points of v_i and set $U = \{u_1, u_2, \ldots, u_M\}$. In the standard mixture model framework, probability density function (PDF) for a d-dimensional vector v, is modelled as a function of model parameters φ, by the mixture density

$$p(v|\varphi) = \sum_{k}^{K} \alpha_k p_k(v|\theta_k), \tag{1}$$

in which $\varphi = \{\alpha_k; \theta_k, k = 1 \ldots K\}, \alpha_k(\sum_k^K \alpha_k = 1)$ is the k-th component weight and p_k is the k-th component density with parameter vector θ_k.

In this manner a finite mixture model is a PDF composed of a weighted average of component density functions. Each trajectory v_i is generated by one of the components, but the identity of the generating component is not observed. The parameters of each density component $p_k(v|\theta_k)$, as well as the corresponding weights α_k, can be estimated from the data using the EM algorithm. The estimated component models, $p_k(v|\theta_k)$ are interpreted as K clusters, where each cluster is defined by a PDF. The set of trajectories is clustered to a number of subsets by assigning a membership probability, w_{ik}, to each trajectory, v_i, to denote its membership of the kth cluster. The number of clusters, K, is defined by the user. Finally, each trajectory v_i is assigned to the cluster k with the highest membership probability.

2.2 Model Definition

We model the X directional position (similarly Y and Z) with a p-th order multivariate polynomial regression model in which the order u_i is the independent variable, which is assumed with an additive Gaussian error term. The three regression equations can be defined succinctly in terms of the matrix v_i. The form of the regression equation for v_i is

$$v_i = U_i\beta + \epsilon_i, \quad \epsilon_i \sim N(0, \Sigma) \tag{2}$$

where U_i is the standard $n_i \times (p+1)$ Vandermonde regression matrix associated with vector u_i, β is a $(p+1) \times 3$ matrix of regression coefficients for X, Y, and Z direction and ϵ_i is an $n_i \times 3$ zero-mean matrix multivariate normal error term with a covariance matrix Σ. For simplicity, we assume that $\Sigma = diag(\sigma_x^2, \sigma_y^2, \sigma_z^2)$, so that the X, Y, and Z measurement noise terms are treated as conditionally independent given the model.

The conditional density for the ith trajectory f is a multivariate Gaussian with mean $U_i\beta$ and covariance Σ. The parameter set $\theta = \{\beta, \Sigma\}$.

$$p(v_i|u_i, \theta) = f(v_i|U_i\beta, \Sigma) \tag{3}$$

We can derive regression mixtures for the trajectories by a substitution of Eq. (1) with the conditional density components $p_k(v|u, \theta_k)$, as defined in Eq. (3).

$$p(v_i|u_i, \varphi) = \sum_k^K \alpha_k f_k(v_i|U_i\beta_k, \Sigma_k) \tag{4}$$

Note that in this model each fibre trajectory is assumed to be generated by one of K different regression models. Each model has its own shape parameters $\theta_k = \{\beta_k, \Sigma_k\}$.

The full probability density V given U, $p(V|U, \varphi)$, is also known as the conditional likelihood of the parameter φ given the data set both V and U to be written as

$$L(\varphi|V, U) = p(V|U, \varphi) = \prod_i^M \sum_k^K \alpha_k f_k(v_i|U_i\beta_k, \Sigma_k) \tag{5}$$

The model can handle trajectories of variable length in a natural fashion, since the likelihood equation (Eq. (5)) does not require the number of data points. The product form in Eq. (5) follows from assuming conditional independence of v_i's, given both u_i's and the mixture model, since the fibre trajectories do not influence each other.

2.3 EM Algorithm for Mixture of Regression:

E-Step: In the E-step, we estimate the hidden cluster memberships by forming the ratio of the likelihood of trajectory v_i under cluster k, to the sum-total likelihood of trajectory v_i under all clusters:

$$w_{ik} = \frac{\alpha_k f_k(v_i|U_i\beta_k, \Sigma_k)}{\sum_{j=1}^K \alpha_j f_j(v_i|U_i\beta_j, \Sigma_j)} \tag{6}$$

These w_{ik} give the probabilities that the ith trajectory was generated from cluster k.

M-Step: In the M-step, the expected cluster memberships from the E-step are used to form the weighted log-likelihood function:

$$L(\varphi|V,U) = \sum_i \sum_k w_{ik}\ log\alpha_k f_k(v_i|U_i\beta_k, \Sigma_k) \tag{7}$$

The membership probabilities weight the contribution that the kth density component adds to the overall likelihood. The weighted log-likelihood is then maximized with respect to the parameter set φ.

Let $w_{ik} = w_{ik}I_{n_i}$ and let $W_k = diag(w_{1k}^{'}, w_{2k}^{'}, \ldots, w_{nk}^{'})$ be an $N \times N$ diagonal matrix, where $N = \Sigma_i^M n_i$. Then, we use W_k to calculate the mixture parameters

$$\hat{\beta}_k = (U^{'}W_kU)^{-1}U^{'}W_kV, \quad \hat{\Sigma}_k = \frac{(V-U\hat{\beta}_k)^{'}W_k(V-U\hat{\beta}_k)}{\sum_i^N w_{ik}}, \text{ and } \hat{\alpha}_k = \frac{1}{N}\sum_i w_{ik} \tag{8}$$

for $k = 1, \ldots, K$ where V is an $N \times 3$ matrix containing all the v_i measurements, one trajectory after another, and U is an $N \times (p+1)$ Vandermonde regression matrix corresponding to Y values.

3 Methods

3.1 Implementation of Clustering Algorithm

Algorithm: Consider a set V of M 3-dimensional fibre trajectories (v_i) in the X, Y and Z directions and the associated set U of u_i, which contains ordered values corresponding points of v_i. The number of clusters K and the order of regression p are input parameters.

Step 1. Randomly initialize the membership probabilities w_{ik}
Step 2. Calculate new estimates for parameters $\hat{\beta}_k, \hat{\Sigma}_k$ of the cluster model and mixing weights $\hat{\alpha}_k$ from Eq. (8) using the current w_{ik}.
Step 3. Compute the membership probabilities w_ik using Eq. (6).
Step 4. Loop to step 2 until convergence.
Step 5. Return the final parameter estimates (including mixing proportions) and cluster probabilities. Outliers are deleted from the set of trajectories using a threshold t.

Handling Outliers: It is assumed that each trajectory is assigned a membership probability w_{ik} for each cluster k. There may be trajectories resulting from the tractography which do not resemble any of the regression equations or are not valid due to inaccuracies at the tractography stage. An outlier is identified by imposing a threshold on the membership probabilities. If the membership probability of a given trajectory in all clusters is less than the specified threshold t, that trajectory will be removed as for Maddah et al. [4].

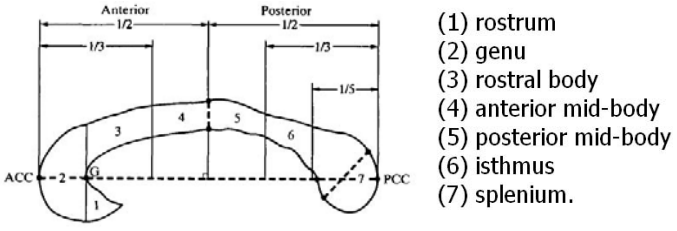

Fig. 1. A schematic of Witelson corpus callosum subdivisions [6] based on the midsaggital slice. ACC and PCC indicate the anteriormost and posteriormost points of the callosum.

3.2 Synthetic Data

We have used PISTE [http://cubric.psych.cf.ac.uk/commondti] synthetic data set (diffusion encoding directions = 30, b-value = 1000 s/mm^2 and voxel resolution: $1 \times 1 \times 1\ mm^3$) to demonstrate some of the basic features of our clustering algorithm, specifically, its ability to cluster a 3D data set into multiple bundles accurately.

Here we consider three example noise free and noisy (SNR=15) data sets: a branching fibre with individual FA in each branch, two orthogonally crossing fibres with individual FA on each fibre and two straight crossing fibres. The noisy synthetic example is intended to demonstrate the robustness of our clustering algorithm in a more hostile environment-one corrupted by additive noise, with complicated fibre structures, and having varying fibre tract lengths. For the three dimensional tract reconstruction, the single-tensor and two-tensor 4th order Runge-Kutta method deterministic tractography were used for branching data and two crossing data respectively. The generated tracts were then clustered into the subdivisions using appropriate K value.

3.3 *In Vivo* Data

Data: 1.5 T DW data were acquired from four healthy adults with an image matrix of 128x128, 60 slice locations covering the whole brain, $1.875 \times 1.875 \times 2.0$ mm^3 spatial resolution, b = 700 s/mm^2 and 41 diffusion directions. To correct for eddy currents and motion, each DW volume was registered to the non-DW volume of the first subject.

Corpus Callosum Clustering: Subdividing the corpus callosum (CC) into anatomically distinct regions is not well defined but is of much importance, especially in studying normal development and in understanding psychiatric and neurodegenerative disorders. Witelson [5] proposed a schematic for seven subdivisions of the CC as shown in Fig. 1. We further divide the splenium into its upper and lower parts to give a finer model.

The ROIs for the CC were outlined by an expert based on information from FA maps for all four subjects. Fibre trajectories were reconstructed using the

4th order Runge-Kutta method for the four subjects and were normalized to a common template (128×128×60 matrix size and voxel size 1×1×1 unit). The CC tracts were then clustered into K=8 subdivisions.

Model Selection: It is important to make decisions about the optimal order of the fibre regression models, the most suitable type of trajectory pre-processing, and the number of clusters that best describes each fibre tract dataset for our method. We fitted regression mixture models with different orders of polynomial to randomly selected training sets of CC fibre trajectories. The experimental results were reported. The choice to use third-order polynomials for the regression models as opposed to other order polynomials was made for two reasons: (a) visual inspection supports this as a sufficient choice and (b) cross-validation also confirms third-order as the optimal choice in this case. We modelled the X position with a cubic polynomial regression model in which u is the independent variable, $x = \beta_3 u^3 + \beta_2 u^2 + \beta_1 u + \beta_0$, and likewise for the Y and Z directions.

4 Results and Discussion

Synthetic Data: The Synthetic data results (Fig. 2) demonstrate the clustering algorithm's ability to accurately separate fibre tracts into meaningful bundles. In our component regression models for the synthetic data a cubic polynomial was used (K=2). This choice is based on the visual inspection of fitted-versus-actual trajectory data. The noise-free synthetic data results in complicated fibre tract structures demonstrating that our clustering algorithm is able to cluster a 3D data set into multiple bundles accurately. The noisy synthetic example results demonstrate the robustness of our clustering algorithm in a noisier environment.

***In Vivo* Data:** Fig. 3 shows the results of clustering approximately 700 trajectories from the corpus callosum into 8 bundles for three subjects. The membership probability of the trajectories for each cluster is obtained and the trajectories in Fig. 3 are coloured based on their maximum membership probabilities. Results showed that our clustering method automatically differentiates CC subdivision fibre bundles consistently across subjects. As a product of the proposed clustering method, regression models of each fibre bundles are obtained in the X, Y,

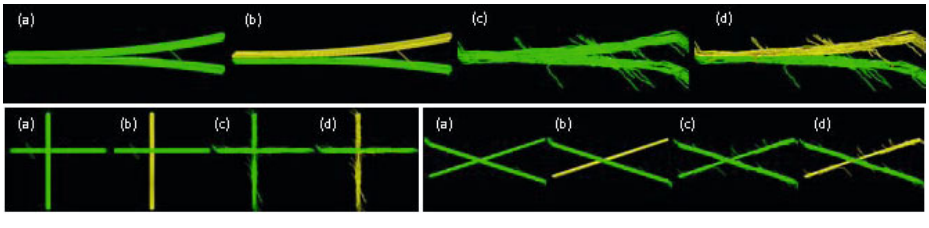

Fig. 2. (a) Synthetic data, (b) clustered trajectories, (c) noisy data and (d) noisy data clustered trajectories of three selected fibre geometries

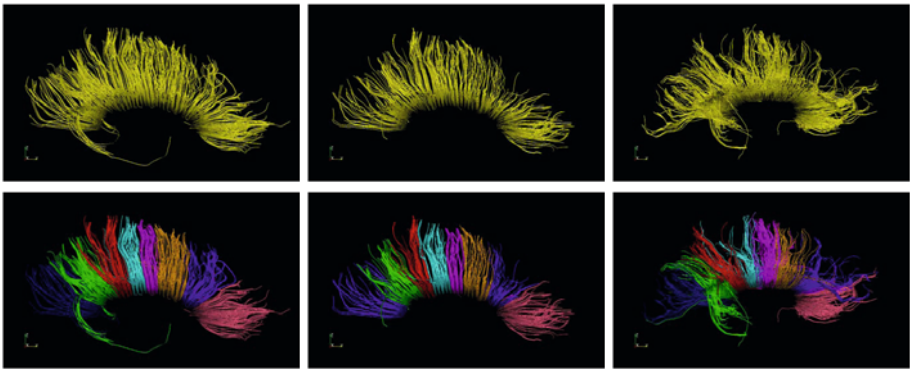

Fig. 3. Clustering of the CC from the first three subjects viewed from a sagittal orientation. Top row: the original fibre tracts and Bottom row: clustered into bundles.

Table 1. Cluster-wise average parameter measures for the sub-divided CC

		Rostrum	Genu	Rostral body	Anterior mid body	Posterior mid body	Isthmus	Upper splenium	Lower splenium
X	β_3	3.09e-4	4.43e-4	3.46e-4	3.74e-4	4.08e-4	3.81e-4	4.08e-4	4.31e-4
	β_2	-0.0348	-0.0422	-0.0367	-0.0393	-0.0363	-0.0368	-0.0350	-0.3908
	β_1	0.7618	0.8090	0.8246	0.9103	0.6254	0.7645	0.4964	0.6804
	β_0	68.034	66.389	65.139	63.545	65.336	63.948	66.064	64.994
Y	β_3	-8.8e-5	9.3e-5	4.99e-5	-8.6e-6	3.26e-5	-1.7e-6	1.00e-4	-2.2e-4
	β_2	-0.0025	-0.0176	-0.0098	-0.0021	-0.0036	0.00053	0.0171	0.0338
	β_1	0.6171	0.8134	0.4960	0.1942	0.1275	-0.0585	-0.6694	-1.335
	β_0	38.215	45.839	53.604	61.118	67.807	74.985	87.812	100.66
Z	β_3	-3.2e-5	8.41e-5	1.35e-4	1.59e-4	-2.6e-4	4.84e-6	8.61e-5	-3.8e-5
	β_2	0.0093	0.00330	3.38e-4	5.17e-5	0.0392	0.0141	0.0138	0.00234
	β_1	-0.5317	-0.4854	-0.6009	-0.6694	-1.4661	-0.9183	-0.5589	-0.0372
	β_0	38.970	44.231	51.163	54.037	54.161	51.672	40.328	28.931

and Z directions. Averages of these quantities are then computed over each cluster for the four subjects. The characteristics (parameters of the cubic regression equation) of each cluster are illustrated in Table 1.

Fig. 4 top row show the X, Y and Z versus order U profiles for all of the tracks with mean curves for subject 1. The cluster groups are colour-coded (the same colour is used as the corresponding cluster in Fig. 3), and the mean curves for each group are highlighted in bold. Mean curves were calculated up to U=70. The mean curve results in each direction show the fibre trajectory points, and how they each differ strongly with direction, especially the Y direction in this case. The mean curve results differ not only in shape but also in location. Fig. 4 bottom row show the cubic polynomial regression models (dotted) fitted to the eight CC subdivision cluster trajectories. The results illustrate that the cubic polynomials provide the best fits among the regression models we considered.

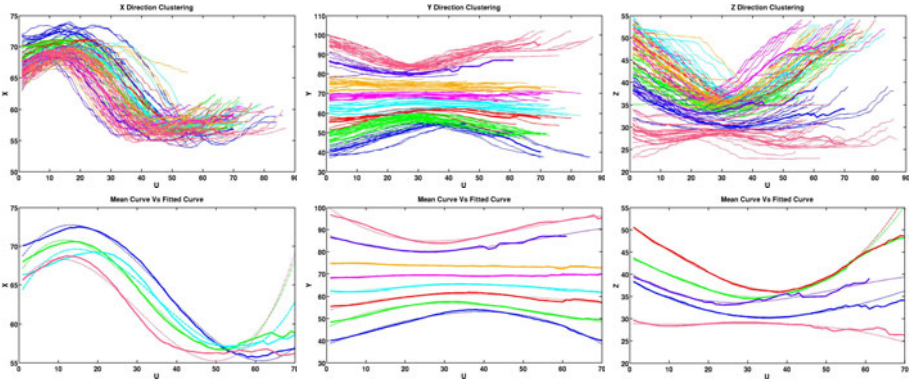

Fig. 4. Top row: all the tracts, and Bottom row: the mean curves and fitted curves for the X, Y and Z directions respectively for subject 1

We presented new techniques for clustering 3D curves into bundles, to remove outlier curves and to develop a technique for shape description of these bundles. Curve-based regression mixture models were used to perform probabilistic clustering of fibre trajectories in 3D space. The number of data points is not required for clustering as the modelling can handle curves with variable lengths. The preliminary results for the synthetic data and in vivo data demonstrate that the new clustering process is quite efficient for bundling sets of curves into anatomically meaningful fibre tracts. Cubic polynomials were found to provide the best fits for CC clustering and modelling among the regression models considered. We have estimated cubic regression equations for each cluster fibre bundle and the equations depending on the coordinate system and image matrix, which we used. Some of the WM trajectories are relatively small, and a successful clustering of them is heavily influenced by such factors as image quality, tractography method, and fibre tracking parameter. In the future, we will investigate how different tractography algorithms such as probabilistic tracking methods and HARDI methods affect the WM fibre clustering procedures.

References

1. Zhang, S., Correia, S., Laidlaw, D.H.: Identifying white-matter fiber bundles in DTI data using an automated proximity-based fiber-clustering method. IEEE Trans Vis. Comput. Graph. 14(5), 1044–1053 (2008)
2. O'Donnell, L.J., Westin, C.-F.: Automatic tractography segmentation using a high-dimensional white matter atlas. IEEE Tr. Med. Im. 26(11), 1562–1575 (2007)
3. Li, H., Xue, Z., Guo, L., Liu, T., Hunter, J., Wong, S.T.: A hybrid approach to automatic clustering of white matter fibers. Neuroimage 49(2), 1249–1258 (2010)
4. Maddah, M., Grimson, W.L., Warfield, S.K.: A unified framework for clustering and quantitative analysis of whitematter fiber tracts. Med. Im. An. 12(2), 191–202 (2008)
5. Witelson, S.F.: Hand and sex differences in the isthmus and genu of the human corpus callosum. Brain 112, 799–835 (1989)

Integrated Parcellation and Normalization Using DTI Fasciculography

Hon Pong Ho[1], Fei Wang[4], Xenophon Papademetris[1,3], Hilary P. Blumberg[3,4], and Lawrence H. Staib[1,2,3]

[1] Departments of Biomedical Engineering
[2] Departments of Electrical Engineering
[3] Departments of Diagnostic Radiology
[4] Departments of Psychiatry,
Yale University, New Haven, CT, USA

Abstract. Existing methods for fiber tracking, interactive bundling and editing from Diffusion Magnetic Resonance Images (DMRI) reconstruct white matter fascicles using groups of virtual pathways. Classical numerical fibers suffer from image noise and cumulative tracking errors. 3D visualization of bundles of fibers reveals structural connectivity of the brain; however, extensive human intervention, tracking variations and errors in fiber sampling make quantitative fascicle comparison difficult. To simplify the process and offer standardized white matter samples for analysis, we propose a new integrated fascicle parcellation and normalization method that combines a generic parametrized volumetric tract model with orientation information from diffusion images. The new technique offers a tract-derived spatial parameter for each voxel within the model. Cross-subject statistics of tract data can be compared easily based on these parameters. Our implementation demonstrated interactive speed and is available to the public in a packaged application.

1 Introduction

Local white matter fiber orientation can be estimated from a set of diffusion weighted images (DWI) [1] by fitting to a diffusion distribution model. The classical second order diffusion model is represented using an ellipsoid, which is mathematically denoted as a symmetric positive definite tensor. The 3D field of ellipsoids becomes the Diffusion Tensor Image (DTI) [2]. Since neuronal fiber diameter is much smaller than a voxel, the measurements represent averaged orientations over a small region [3]. With an end goal of understanding and characterizing brain tracts *in vivo*, existing white matter parcellation methods from diffusion images can be summarized by two main steps: tracking and bundling.

Existing Fiber Tracking and Bundling Paradigm

Streamline tracking is the process of recovering line structures in a vector field [4,5,6]. The vector field $\mathbf{F}(\mathbf{x})$ is provided by orientations from DTI or either

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 33–41, 2011.

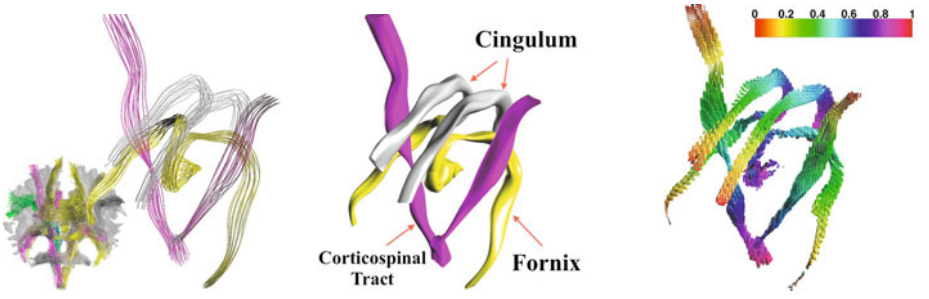

Fig. 1. Fasciculography of major white matter tracts from a human brain DTI. Fasciculography offers both the fibers view (left) and the tract view (middle) without additional processing. All tracts are intrinsically normalized, parametrized (right) and ready for cross-subject analysis.

orientation distribution models. The tracking problem can be written as an initial value problem of the partial differential equation [7]:

$$\frac{dL(s)}{ds} = \mathbf{F}\left(L(s)\right) \qquad , L(s_0) = \mathbf{x}_0 \tag{1}$$

where $s \in \mathbb{R}$ is the arc-length parameter of the streamline $L(s) : \mathbb{R} \rightarrow \mathbb{R}^3$ in the 3D vector field $\mathbf{F}(\mathbf{x}) : \mathbb{R}^3 \rightarrow \mathbb{R}^3$, where $\mathbf{x} \in \mathbb{R}^3$ is a 3D point in the DTI. The best choice of seed points varies across tract types and subjects, and is critical for accurate tract segmentation.

A specific brain tract can be extracted and visualized by properly seeding, tracking and processing those virtual fiber pathways. The selection process has been performed manually or automatically using pattern recognition algorithms (classification/clustering). Interactive parcellation [8,9] usually involves manual post-editing of tracking results, either using whole brain seeding or individual region of interest (ROI) seeding. Automatic bundling [10,11,12] can be done with the help of proper prior models. However, the availability and appropriateness of a training set poses limitations to methods using prior models. Most recently, a generalized cylinder approach has been proposed [13], but their tract solutions are limited to circular cross-sectional areas which is typically not a valid assumption for neural tracts. Moreover, their approach is a tubular-constrained level set surface evolution where the parameterization *within* the tract is not carefully considered and therefore tract normalization is still an issue.

Cross-subject analysis of fiber tracts is also hampered by the definition of point correspondence among bundles with varying fiber lengths and number of fibers. Various Tract-based Morphometry (TBM) methods [14,15,12] tackle the point correspondence issue mainly by generating a common parametrization through post-processing and registering fibers. However, tract parcellation based on fiber bundling relies on good fiber tracking results, while local orientation error from diffusion signals will significantly alter their trajectories.

Motivation of the New Volumetric Approach: Our idea is a fast method that performs streamline tracking and anatomical segmentation simultaneously using a parametric volumetric representation, with less dependence on seed point locations and local orientation errors. Once a tract volume is delineated, we label each image voxel according to the axial parameters $\alpha \in [0, \pi]$ (similar to the arc-length parameter of a curve) of the tubular model. Then, the statistics of the tract image data can be computed based on those labels. For example, we can compare data of the anterior portion (first 20%) of multiple fascicles using $\alpha \in [0 - 0.2\pi]$. In [16], we proposed a fast volumetric tract extraction method using a pair of seed points. We here significantly reduce the dependancy of the seed pairs to one seed and tackle more complex tract geometry by a tract optimization algorithm.

2 Methodology

Our goal is to find a parametric volume $V(\alpha, \beta, \gamma) : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ where $\alpha, \beta, \gamma \in [0, 2\pi]$ and α, γ are symmetric, which best represents an oriented structure in a vector field. We intend to maximize the length of the tract axis through iterative filtering (**Step 3**) and streamline extension (**Step 5**) of a tract volume. Fibers for visualization can be generated by sampling the final $V(\alpha, \beta, \gamma)$ along the axial direction. Throughout this section, α, β and γ refer to the axial, circumferential and radial parameters of a tract volume, respectively.

2.1 Cross-Sectional Tractness Measure

We would like to quantify how relevant a nearby voxel $\mathbf{x} \in \mathbb{R}^3$ is to a streamline $L(s)$ based on its orientation. The coherence of an oriented voxel given a streamline is defined as the inner product of the tensor orientation and the normalized tangent of the given streamline [16]:

$$C(\mathbf{x}; L(s)) = \left(F(\mathbf{x}) \cdot \frac{dL(s)}{ds} \right) \left\| \frac{dL(s)}{ds} \right\|^{-1} \tag{2}$$

Note that this coherence measurement is dependent on the streamline tangent at s. We cannot simply apply a threshold or perform region growing on this measurement without knowing the streamline orientation at s for each x.

Let $P_{Y(\alpha)}(\mathbf{p}) \in \mathbb{R}^2 \rightarrow \mathbb{R}^3$ be the orthogonal plane of streamline $Y(\alpha)$, where $\mathbf{p}$ is the plane parameter. We find the cross-sectional tractness $A_\alpha(\beta, \gamma)$ by the intersection of the thresholded 2D coherence map and FA map computed based on $P_{Y(\alpha)}(\mathbf{p})$ and Equation (2):

$$A_\alpha(\beta, \gamma) = \{C(P_{Y(\alpha)}(\mathbf{p}); Y(\alpha)) \geq \tau_{dp}\} \cap \{F_{fa}(P_{Y(\alpha)}(\mathbf{p})) \geq \tau_{fa}\} \tag{3}$$

Specifically, the tract axis is $A_\alpha(\beta, 0) = Y(\alpha)\ \forall\ \beta$. The tract surface $A_\alpha(\beta, \pi)$ can be found by marching along the thresholded 2D coherence and FA map.

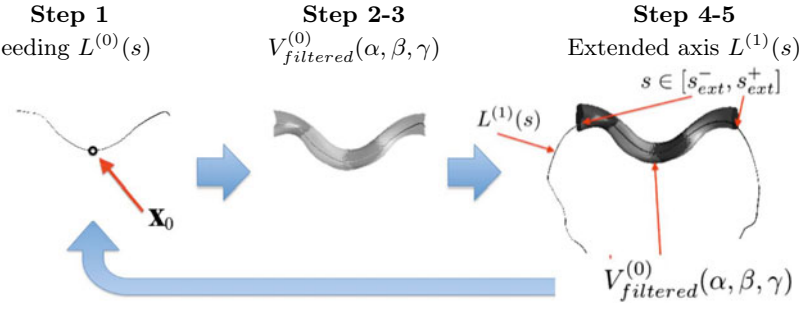

Fig. 2. The first iteration of the fasciculography algorithm in Section 2.2 for the reconstruction of the Tapetum portion of the corpus callosum from real DTI

2.2 Tract Optimization

Since all geometric entities are unknown at the beginning, the search is implemented in an iterative fashion. To find a maximum length tract-axis streamline $Y(\alpha)$ under the influence of the tract surface $S(\alpha,\beta)$ and a generic model $V(\alpha,\beta,\gamma)$ and vice versa, we first begin with a manually seeded streamline $L^{(0)}(s)$ anywhere within the tract. The superscript $^{(t)}$ indicates the iteration number. There is a re-parametrization from s to α since streamline tracking can move in both directions from the seed while we define α to be $[0,\pi]$. Since s and α are a one-to-one mapping, we assume $s = s(\alpha)$ for simplicity. The steps of the algorithm are graphically illustrated in Figure 2.

Fasciculography (FASC) Algorithm

Step 1: Given a seed point $\mathbf{x}_0$ on a vector field $\mathbf{F}$, we compute a streamline $L^{(0)}(s)$ using any numerical method solving Equation (1). Take this line to be the initial tract axis:

$$V^{(0)}(\alpha,\beta,0) = L^{(0)}(s) \qquad \forall\beta \tag{4}$$

Step 2: Compute the surface using Equation (3) and form the volume:

$$V^{(0)}(\alpha,\beta,\pi) = S(\alpha,\beta) = A_\alpha(\beta,\pi) \tag{5}$$

Step 3: Filter the volume with $\mathcal{P}_\mathcal{K}$ in the Fourier domain [17] removing frequencies higher than a tract length-dependent value $\mathcal{K}$:

$$V^{(0)}_{filtered}(\alpha,\beta,\gamma) = \mathcal{F}^{-1}\left\{\mathcal{P}_\mathcal{K}\mathcal{F}\left\{V^{(0)}(\alpha,\beta,\gamma\right\}\right\} \tag{6}$$

Step 4: Estimate the new tract axis :

$$Y^{(0)}(\alpha) = \frac{\int_\beta V^{(0)}_{filtered}(\alpha,\beta_i,\pi)d\beta_i}{\int_\beta d\beta_i} \tag{7}$$

Note that typically $Y^{(0)} \neq L^{(0)}$. The result of each iteration will be an improved axis that has taken tract smoothness and tensor coherence into account. However, it may be shorter than the optimum.

Step 5: Extend $Y^{(0)}(\alpha)$ to form $L^{(1)}(\alpha)$ by re-seeding and streamline tracking from the two end points of $Y^{(0)}(\alpha)$. Let $[s_{ext}^-, s_{ext}^+]$ be the range of parameters indicating the portions of the new $L^{(1)}$ which is copied from Step 4. We solve for $L^{(1)}(s)$ according to:

$$L^{(1)}(s) = Y^{(0)}(\alpha_1) \quad \text{if } s \in [s_{ext}^-, s_{ext}^+] \tag{8}$$

$$\frac{dL^{(1)}(s)}{ds} = \mathbf{F}(L^{(1)}(s)) \quad \text{otherwise} \tag{9}$$

where re-parametrization of the new tract axis is $\alpha_1(s) = \pi \left(\frac{s - s_{ext}^-}{s_{ext}^+ - s_{ext}^-} \right)$.

Repeat Step 1 to 5 until no further extensions can be found, i.e. for some $t > 0$, we stop iterating when $\left|L^{(t+1)}\right| = \left|Y^{(t)}\right|$. This algorithm always converges as it only allows increasing tract length and $\{\mathbf{x} \in \mathbb{R}^3 | F_{fa}(\mathbf{x}) \geq \tau_{fa}\}$ is finite.

3 Comparison Using Synthetic DTI Data

We verify our method using a set of DTI synthetic phantom images with ground truth, PISTE [18]. Each tract in PISTE has a different geometry posing its own tracking challenge. We compare our method against FACT [4] in DtiStudio [8], TEND [5] in TrackVis [9] and Front Propagation [19,20]. The last spiral-cross tract is designed to reflect a more realistic situation where simple thresholding and shortest-path approaches in general would not trivially succeed. To summarize, we ran four methods on 12 synthetic DTIs (4 tract types, 3 noisy DTIs each) to get 48 tracking results. All methods use the mean FA value from the ground truth as the FA threshold and software default 0.5mm for the step size. FACT/TEND use default 30° as the maximum permissible turning angles.

Qualitative Comparison (Phantoms)

Figure 3 shows the 3D appearance of the tracking results using the tested methods in low Signal-to-Noise Ratio (SNR) environments. For the helix and spiral-cross tracts, low or missing fiber coverage is observed between or distant from ROIs. Fibers are diverted outside the expected tract. This result shows the impact of premature termination of fibers due to accumulation of orientation errors. We demonstrated the intrinsic weakness of shortest-path approaches in 3(d). Tracks initiated far away from the source seed often diverted to adjacent tracts.

Quantitative Comparison (Phantoms)

Dice coefficients are computed for the different tract types in different SNR environments. A perfect segmentation will score 100. Figure 4 shows our method

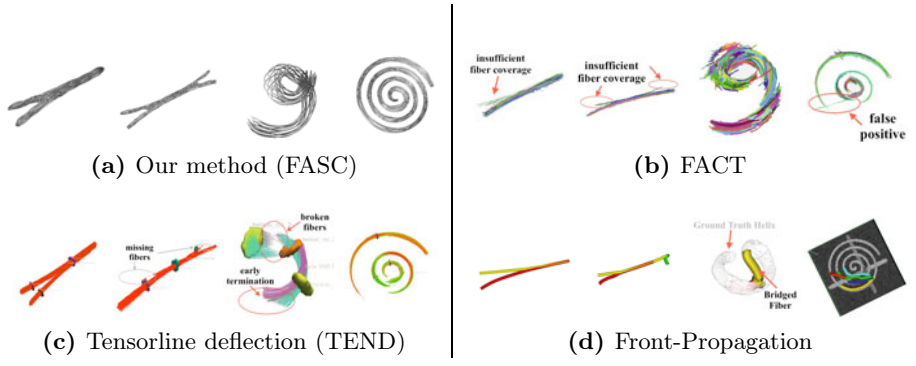

(a) Our method (FASC) **(b)** FACT

(c) Tensorline deflection (TEND) **(d)** Front-Propagation

Fig. 3. Comparing tracking methods in synthetic DTI with increasing tract difficulty (left to right) with SNR 5:1. ROIs for FACT/TEND are derived from ground truth. In a noisy environment, only FASC succeeds in all cases.

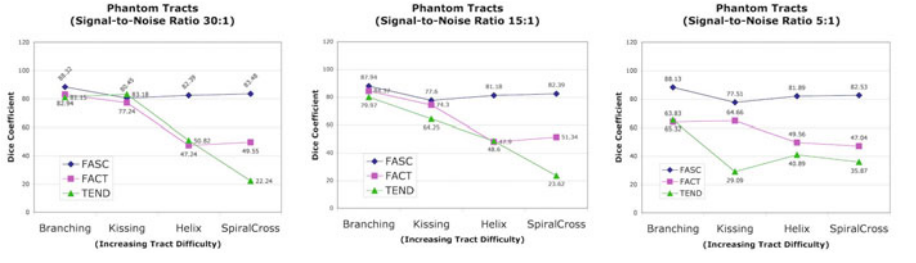

Fig. 4. Quantitative comparison of the tracking performance among Fasciculography (FASC), FACT and tensorline deflection (TEND) using synthetic DTI images. Each plot shows the tracking performance under the same Signal-to-Noise Ratio with increasing tract difficulty (data points from left to right).

has the highest score in 11 out of 12 cases. All FASC lines are relatively flat, as a function of tract difficulty, indicating the robustness of the method regarding difficult tract types, while the others exhibit steeper downward slopes. FASC lines are all above 75% showing our method is robust in noisy environments, while the other lines are much lower (below 70%) at low SNR cases.

4 Experiments Using Real DTI

We acheive similar levels of accuracy (compared with Figure 4) using our FASC method on DTI of 18 cingulum tracts (in terms of averaged Dice coefficient) : 68.6 (FASC), 41.8 (FACT) and 57.1 (TEND). Figure 1(lower-left) shows fibers for more than 30 tracts from a human DTI. In Figure 5, we show that our method is robust against a range of FASC parameters. Reasonable variation of the thresholds or seeding locations within a tract do not significantly reduce

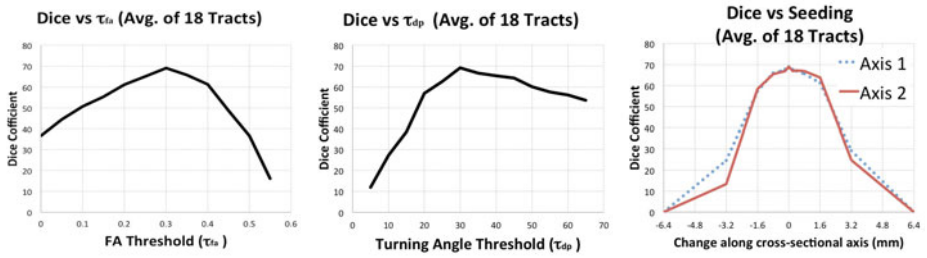

Fig. 5. Sensitivity to major FASC parameters. The Dice cofficients are computed based on 18 manual cingulum segmentations over FA slices.

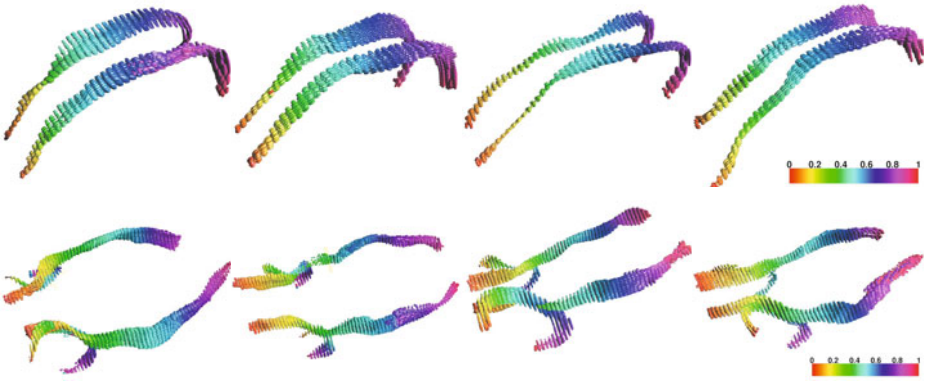

Fig. 6. Normalized cingulum (first row), inferior fronto-occipital and uncinate fasciculus voxels (second row) of 4 subjects from FASC. Voxels are colored by α parameter.

the accuracy of parcellation. Figure 6 shows the normalized tract voxels of the cingulum, inferior fronto-occipital (IFO) and uncinate (UNC) fasciculus. We choose the cingulum and IFO because they run approximately parallel to the sagittal and axial slices respectively which facilitates evaluation of the tracking results. On the other hand, cingulum and IFO bundles are challenging because they are long and thin. The UNC is relatively shorter but has higher curvature and appears to have lower FA (0.2 to 0.3 instead of > 0.4 for cingulum and IFO) which makes it difficult to locate based on FA slices.

5 Conclusion

We have verified our method using third-party synthetic data [18] and real DTI. We have packaged and extensively tested our interactive implementation in a

standalone, *open to the public* application with a convenient user-interface[1]. The application can run on multiple OS platforms.

References

1. Basser, P.J., Mattiello, J., LeBihan, D.: Estimation of the effective self-diffusion tensor from the NMR spin echo. J. Magn. Reson B 103, 247–254 (1994)
2. Lori, N., Akbudak, E., Shimony, J., Cull, T., Snyder, A., Guillory, R., Conturo, T.: Diffusion tensor fiber tracking of human brain connectivity: aquisition methods, reliability analysis and biological results. NMR In Biomedicine 15, 493–515 (2002)
3. Alexander, A.L., Hasan, K.M., Lazar, M., Tsuruda, J.S., Parker, D.L.: Analysis of partial volume effects in diffusion-tensor MRI. Magn. Reson Med. 45, 770–780 (2001)
4. Mori, S., Crain, B.J., Chacko, V.P., van Zijl, P.C.: Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. Ann. Neurol. 45, 265–269 (1999)
5. Lazar, M., Weinstein, D.M., Tsuruda, J.S., Hasan, K.M., Arfanakis, K., Meyerand, M.E., Badie, B., Rowley, H.A., Haughton, V., Field, A., Alexander, A.L.: White matter tractography using diffusion tensor deflection. Hum. Brain Mapp. 18, 306–321 (2003)
6. Catani, M., Howard, R.J., Pajevic, S., Jones, D.K.: Virtual in vivo interactive dissection of white matter fasciculi in the human brain. Neuroimage 17, 77–94 (2002)
7. Faires, J.D., Burden, R.L.: Numerical methods, 3rd edn. Thomson/Brooks/Cole, Pacific Grove, CA (2003)
8. Jiang, H., van Zijl, P.C.M., Kim, J., Pearlson, G.D., Mori, S.: DtiStudio: resource program for diffusion tensor computation and fiber bundle tracking. Comput. Methods Programs Biomed. 81, 106–116 (2006)
9. Wang, R., Benner, T., Sorensen, A.G., Wedeen, V.J.: Diffusion toolkit: A software package for diffusion imaging data processing and tractography. Int'l Society of Magnetic Resonance in Medicine 15, 3720 (2007)
10. O'Donnell, L.J., Westin, C.-F.: Automatic tractography segmentation using a high-dimensional white matter atlas. IEEE Trans. Med. Imaging 26, 1562–1575 (2007)
11. Yushkevich, P.A., Zhang, H., Simon, T.J., Gee, J.C.: Structure-specific statistical mapping of white matter tracts. Neuroimage 41, 448–461 (2008)
12. O'Donnell, L.J., Westin, C.-F., Golby, A.J.: Tract-based morphometry for white matter group analysis. Neuroimage 45, 832–844 (2009)
13. Mohan, V., Sundaramoorthi, G., Tannenbaum, A.: Tubular surface segmentation for extracting anatomical structures from medical imagery. IEEE Trans. Med. Imaging 29, 1945–1958 (2010)
14. Smith, S.M., Jenkinson, M., Johansen-Berg, H., Rueckert, D., Nichols, T.E., Mackay, C.E., Watkins, K.E., Ciccarelli, O., Cader, M.Z., Matthews, P.M., Behrens, T.E.J.: Tract-based spatial statistics: voxelwise analysis of multi-subject diffusion data. Neuroimage 31, 1487–1505 (2006)

[1] Download @ `http://www.fasciculography.com`, `http://noodle.med.yale.edu/~ho`

15. Corouge, I., Fletcher, P.T., Joshi, S., Gouttard, S., Gerig, G.: Fiber tract-oriented statistics for quantitative diffusion tensor MRI analysis. Med. Image Anal. 10, 786–798 (2006)
16. Ho, H.P., Wang, F., Blumberg, H.P., Staib, L.H.: Fast parametrized volumetric DTI tract parcellation. In: IEEE Int'l Symposium on Biomedical Imaging (2011)
17. Oppenheim, A.V., Willsky, A.S., Nawab, S.H.: Signals and systems, 2nd edn. Prentice-Hall, Upper Saddle River (1997)
18. Deoni, A.: Phantom Images for Simulating Tractography Errors, http://cubric.psych.cf.ac.uk/commondti
19. Jackowski, M., Kao, C., Qiu, M., Constable, R., Staib, L.: White matter tractography by anisotropic wavefront evolution and diffusion tensor imaging. Medical Image Analysis 9, 427–440 (2005)
20. Papademetris, X., Jackowski, M., Rajeevan, N., Okuda, H., Constable, R., Staib, L.: Bioimage Suite: An integrated medical image analysis suite, Section of Bioimaging Sciences, Dept. of Diagnostic Radiology, Yale School of Medicine

Predicting Functional Brain ROIs via Fiber Shape Models

Tuo Zhang[1,2,*], Lei Guo[1], Kaiming Li[1,2], Dajing Zhu[2], Guangbin Cui[3], and Tianming Liu[2]

[1] School of Automation, Northwestern Polytechnical University, Xi'an, China
[2] Department of Computer Science and Bioimaging Research Center The University of Georgia, Athens, GA, USA
[3] Department of Radiology, Tangdu Hospital, Xi'an, China

Abstract. Study of structural and functional connectivities of the human brain has received significant interest and effort recently. A fundamental question arises when attempting to measure the structural and/or functional connectivities of specific brain networks: how to best identify possible Regions of Interests (ROIs)? In this paper, we present a novel ROI prediction framework that localizes ROIs in individual brains based on learned fiber shape models from multimodal task-based fMRI and diffusion tensor imaging (DTI) data. In the training stage, ROIs are identified as activation peaks in task-based fMRI data. Then, shape models of white matter fibers emanating from these functional ROIs are learned. In addition, ROIs' location distribution model is learned to be used as an anatomical constraint. In the prediction stage, functional ROIs are predicted in individual brains based on DTI data. The ROI prediction is formulated and solved as an energy minimization problem, in which the two learned models are used as energy terms. Our experiment results show that the average ROI prediction error is 3.45 mm, in comparison with the benchmark data provided by working memory task-based fMRI. Promising results were also obtained on the ADNI-2 longitudinal DTI dataset.

Keywords: diffusion tensor imaging, fMRI, working memory, fiber shape model.

1 Introduction

Mapping of structural and functional connectivities of the human brain via neuroimaging offers an exciting and unique opportunity to understand the brain architecture, and thus has received significant interest [1-5]. However, a fundamental question arises when attempting to map structural and functional connectivities: how to define the best possible Regions of Interests (ROIs) for the connectivity mapping? Essentially, ROIs provide the structural substrates for measuring connectivities within individual brains and for pooling data across population groups. Thus, identification of reliable, reproducible and accurate ROIs is critically important for the success of

[*] Tuo Zhang was supported by Doctorate Foundation of Northwestern Polytechnical University.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 42–49, 2011.

connectivity mapping. However, in our view, this task is challenging due to the unclear boundaries between cortical regions, remarkable variability of brain anatomy, and nonlinearity of ROIs. For instance, a slight change of the shape, size or location of a ROI might dramatically alter its structural and functional connectivity profiles [6].

In this paper, we present a novel framework that learns fiber shape model and anatomical constraint model of functional ROIs based on multimodal task-based fMRI and DTI data in the training stage, and apply the predictive models to localize functional ROIs in testing samples based only on DTI data. The major advantages of this framework and our contributions lie in the following two aspects. 1) In the training stage, the activated brain regions derived from task-based fMRI data provide the benchmark ROI data for learning prior shape models of white matter fibers emanating from these ROIs. Recent literature using working memory task-based fMRI data [6] demonstrated that the fiber connection patterns of corresponding functional ROIs are quite consistent, providing direct evidence that white matter fiber connection pattern is a good predictor of functional landmark [7]. This is similar to the "connectional fingerprint" concept presented in [8]. 2) In the prediction stage, only DTI data is needed to accurately localize the functional ROIs based on our predictive models, without relying on the availability of task-based fMRI data. Typically, a DTI scan needs less than ten minutes, is much less demanding, and is widely available. Therefore, functional brain ROIs prediction based on DTI has wide applications in brain imaging.

2 Materials and Methods

The flowchart of our ROI prediction framework is illustrated in Fig. 1. It consists of two stages: model training and ROI prediction. The model training (the purple box in Fig.1) is conducted on the training dataset in which each subject has both task-based fMRI and DTI data. Then, descriptive shape features of fiber bundles were extracted, and fiber shape models and ROIs location distribution models are learnt. In the prediction stage (the green box in Fig. 1), the learned predictive models are applied on subjects with DTI data by minimizing an energy function to obtain the predicted ROIs. Finally, we use task-based fMRI data to validate the prediction result (black dashed box).

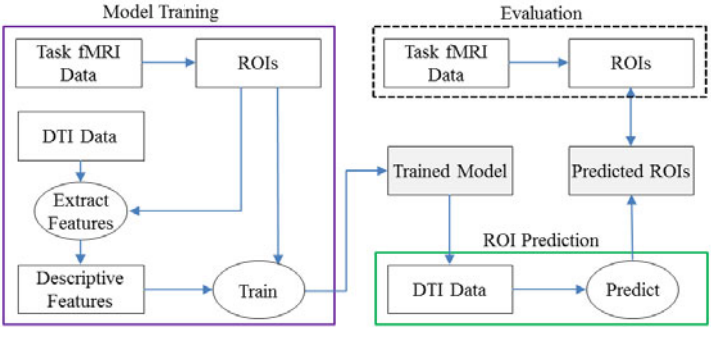

Fig. 1. The flowchart of the model training and ROI prediction framework

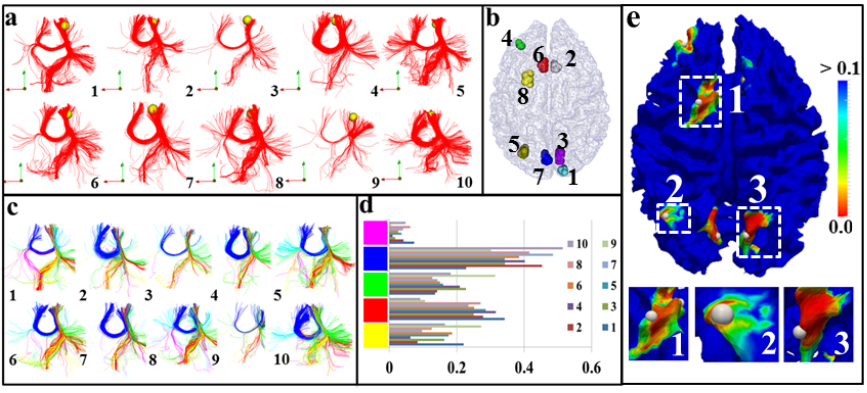

Fig. 2. (a) Fiber bundles extracted for ROI #8 of the training dataset; (b) Examples of consistent ROI locations with ROI IDs labeled. The corresponding ROIs are coded with the same color; (c) Clustered fiber bundles in (a); (d) Shape histograms of fiber bundles in (c). Shape pattern on the vertical axis is color-coded the same way as that in (c); (e) Convergence basins of the energy function Eq. (1) in Section 2.3. More details are referred to the text.

2.1 Data Acquisition and Pre-processing

Fifteen healthy young adults voluntarily participated in this study. Each volunteer performed a modified version of the OSPAN working memory task (3 blocks: OSPAN, Arithmetic and Baseline) while fMRI data was acquired. DTI scans were also obtained for each volunteer. Briefly, both fMRI and DTI scans were acquired on a 3T GE Signa scanner with acquisition parameters as follows: fMRI: 64×64 matrix, 4mm slice thickness, 220mm FOV, 30 slices, TR = 1.5s, TE = 25ms, ASSET = 2; DTI: 128×128 matrix, 2mm slice thickness, 256mm FOV, 60 slices, TR = 15100ms, ASSET = 2, 30 gradient directions, b-value = 1000. The pre-processing of the DTI data includes brain skull removal and motion correction. Afterwards, the diffusion tensor was computed and the fiber was tracked via the MEDINRIA toolkit. The fMRI data of the OSPAN task was analyzed using the FSL FEAT. Individual activation map reflecting the OSPAN (complex span) contrast was identified. For each subject, the DTI space is used as the standard space, in which the gray matter (GM)/white matter (WM) surface is generated via an approach similar to that in [9]. The cortical surface was used as the ROI definition and prediction space. Co-registration between DTI and fMRI data is performed using FSL FLIRT. Currently, 8 most consistently activated working memory ROIs including 1. Left Occipital Pole, 2. Left Paracingulate Gyrus, 3. Left Precuneus, 4. Right Dorsolateral Prefrontal Cortex, 5. Right Lateral Occipital Gyrus, 6. Right Paracingulate Gyrus, 7. Right Precuneus, and 8. Right Superior Frontal Gyrus, were used for model training and ROI prediction. We divided the dataset into two subsets: the training dataset containing ten subjects, and the testing datasets containing the rest five subjects. A randomly selected subject from the training dataset was used as the template onto which all the other subjects were linearly registered.

2.2 Model Training

The main idea of model training is to extract descriptive structural and anatomic features from the activated ROIs, and then learn predictive models based on these features. The structural features are the fiber connection profiles of the activated ROIs, and we propose a fiber shape pattern histogram PCA (principal component analysis) model to represent the fiber connection information for each ROI. The descriptive anatomical features are the spatial location distribution of the activated ROIs, and we propose an ROI location distribution PCA model.

Fiber Shape Pattern Histogram PCA Model. The fiber shape pattern histogram PCA model is constructed for each ROI separately across the training subjects to embed group-wise fiber shape information. Taking ROI i for example, we define R_j^i as the local region centered at location of ROI i on subject j's surface, and fiber bundle penetrating R_j^i was extracted and denoted as F_j^i. We defined $T_i = \{F_j^i | j = 1, \cdots, m\}$ as the fiber bundle set of ROI i, where m is the number of the training subjects. As an example, Fig. 2a shows fiber bundles of T_8, from which consistence of connection pattern of fiber bundles can be observed. As fiber shape is considered as an effective descriptive feature of structural connectivity [10], we constructed a fiber shape pattern histogram PCA model H_i for ROI i on T_i as follows. First, we conducted a fiber shape clustering via the method in [10] on fiber tracts of the entire training dataset and clustered all fiber tracts into five shape patterns. Then for fiber bundle F_j^i, we define the normalized shape pattern histogram h_j^i, in which each bin denotes the ratio of fiber tract number of each shape pattern to total fiber tract number in F_j^i. In Fig. 2c, we illustrate the clustered fiber bundles of Fig. 2a, and their corresponding normalized shape pattern histograms are shown in Fig. 2d, from which the consistency of $\{\mathrm{h}_j^i | j = 1, \cdots, m\}$ can be visually observed. Afterwards, by taking h_j^i as the descriptive feature of F_j^i, we embed the shape information across the training datasets by applying PCA on the feature matrix $[h_1^i, h_2^i, \cdots, h_m^i]^{\mathrm{T}}$. The obtained mean feature h_{mean}^i and eigen vectors, PCA transformation matrix, E_h^i compose the fiber shape pattern model for ROI i, denoted as $H_i = \{h_{mean}^i, E_h^i\}$. It is noteworthy that, for each ROI, we computed the ratio between the first eigen value and the sum of all eigen values. The ratio is as high as 0.88±0.13, suggesting the consistency in T_i across subjects and the effectiveness of feature h_j^i and model H_i.

ROI Location Distribution PCA Model. Despite its variation, the spatial distribution patterns of the ROIs in the template space have certain degree of consistency, reflecting the existence of common human brain anatomical architecture (Fig. 2b), and therefore they were captured and modeled across training subjects. For subject j, we define ROI coordinates vector $d_j = [\mathrm{x}_1, \mathrm{y}_1, z_1, \cdots, \mathrm{x}_n, \mathrm{y}_n, z_n]^T$ to be the feature describing the ROI location distribution, where n denotes the number of ROIs. Similarly, a PCA model $D = \{d_{mean}, E_d\}$ was computed on feature matrix $[d_1, \cdots, d_m]^T$ and the first eigenvalue ratio is 0.72, suggesting the consistency of d_j and the effectiveness of PCA model. It is noted that fiber shape pattern histogram PCA model H_i represent structural connectivity feature for each ROI individually, while ROI location PCA model D represent the ROI spatial distribution in the entire training dataset.

2.3 ROI Prediction Framework

ROI prediction was conducted only based on DTI data of an individual subject and the trained models. The prediction was formulated as an energy function minimization problem by maximally mapping the trained models onto the DTI data of the subject being predicted. The energy function is defined as:

$$E = \lambda E_{int} + (1 - \lambda) E_{ext} \tag{1}$$

where E_{ext} denotes the mapping of fiber shape pattern histogram PCA model, while E_{int} is the ROI location PCA model constraint, and ë trades off between them.

Similar to section 2.2, let $\tilde{d} = [\mathrm{x}_1, \mathrm{y}_1, z_1, \cdots, \mathrm{x}_n, \mathrm{y}_n, z_n]^T$ be the location of ROIs being predicted on the cortical surface. For each ROI, fiber bundles $\tilde{F}^i$were extracted from local region R^i, and the corresponding shape histogram $\tilde{h}^i$were obtained.

To derive E_{ext}, we mapped $\tilde{h}^i$ into PCA model$H_i = \{h^i_{mean}, E^i_h\}$ and reconstructed a new histogram as follows:

$$\tilde{h}^i_{rec} = h^i_{mean} + {E^i_h}^T E^i_h (\tilde{h}^i_{rec} - h^i_{mean}) \tag{2}$$

If $\tilde{h}^i$is consistent with those $\{\mathrm{h}^i_j | j = 1, \cdots, m\}$ in the training dataset, then the error between it and $\tilde{h}^i_{rec}$ is expected to be small. Therefore, we define E_{ext} to be the sum of the errors of all ROIs:

$$E_{ext} = \frac{1}{n} \sum_{i=1}^{n} \chi^2 (\tilde{h}^i, \tilde{h}^i_{rec}) \tag{3}$$

where $\div^2$ denotes $\div^2$-test measuring the similarity between two histograms if we let l be its bin number:

$$\chi^2(\tilde{h}^i, \tilde{h}^i_{rec}) = \sum_{k=1}^{l} (\tilde{h}^i(k) - \tilde{h}^i_{rec}(k))^2 / ((\tilde{h}^i(k) + \tilde{h}^i_{rec}(k)) \tag{4}$$

E_{int} can be defined as the Euclidean distance between location $\tilde{d}$ and the its reconstructed version $\tilde{d}_{rec}$ from model D:

$$E_{int} = dist(\tilde{d}, \tilde{d}_{rec}) \tag{5}$$

Global search was conducted to minimize the energy function. The search starts from the mean ROIs location of in training dataset and ends when the location becomes stable. It is noted that the two energy terms are normalized into [0, 1] prior to global search.

3 Experimental Results

At current stage, we did not take ROI size into consideration, therefore R^i defined as uniform 3-ring vertices neighborhood on the cortical surface was used for all ROIs to extract fiber bundles. λ was empirically assigned 0.3.

3.1 Evaluation and Validation via Task-Based fMRI Data

The activated ROI locations detected from task-based fMRI data of the five subjects in prediction dataset were used as benchmark to evaluate the prediction results. We randomly selected a subject from the prediction dataset to illustrate the prediction result.

First of all, we illustrate the effectiveness of the prediction framework in Fig. 2e where the 8 benchmark ROI locations are highlighted by white bubbles, and we computed the energy E defined in section 2.3 between the benchmark ROI location and the neighbor vertices on surface around them. Energy values mapped onto the surface intuitively illustrate the convergence basins around the benchmark ROIs (see three zoomed-in basins in Fig. 2e), suggesting the convergence of our prediction framework.

Then, the prediction results of the 8 ROIs were shown in Fig. 3. It is evident that the fiber bundles emanating from the predicted ROIs are quite similar to those from the benchmark ROIs, indicating the effectiveness and accuracy of our ROI prediction framework. In addition, fiber bundles emanating from the predicted ROIs and benchmark ROIs are quite similar to those from the corresponding ROIs in the training dataset. Quantitatively, Table 1 shows the mean prediction errors, the Euclidian distances between the predicted ROIs and the benchmark ones, of the five subjects in the prediction dataset. On average, the prediction error for 8 eight ROIs over 5 subjects is 3.45 mm. For most ROIs, the prediction errors are approximately 2~4 mm. Notably, the less accurate result on ROI #1 reveals that the energy function may be trapped in a local minimum.

For the purpose of comparison, we used linear (via FSL FLIRT) and nonlinear registration (via the HAMMER software package [11]) methods to warp the training subjects onto the subjects to be predicted, and their functional ROIs were correspondingly warped and used as prediction results. The mean prediction errors between warped ROIs and benchmark ones in the prediction dataset are also shown in Table 1. On average, the prediction errors by FSL FLIRT and HAMMER are 5.72 mm and 5.53 mm, respectively. As can be seen, our method (3.45 mm) significantly outperforms both of them.

3.2 Application on ADNI-2 longitudinal DTI Data

We applied our method on the ADNI-2 longitudinal DTI dataset under the premise that the white matter of working memory system in ADNI-2 mild cognitive impairment (MCI) patient is not distinctively different from the normal controls. Two scans of 10 MCI patients' data were obtained from the ADNI-2 project (http://adni.loni.ucla.edu/). The time interval between the two scans was around 3 months. As an example, the prediction results of one patient are shown in Fig. 4. The average distance between the predicted ROI locations of two scans is (2.54±1.53mm). This relatively small distance and the high similarity of the fiber bundles (Fig. 4b) is a strong evidence of the reproducibility of our ROI prediction method. This result also suggests that there is no distinctive change in the white matter of MCI subject, at least in the working memory system.

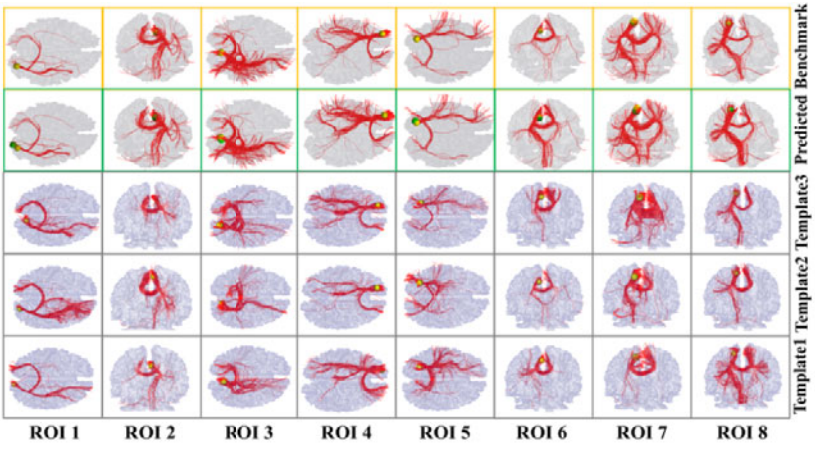

Fig. 3. Visualization of prediction results. Top yellow frame: fiber bundles emanating from benchmark ROIs (yellow bubbles); Green frame: fiber bundles emanating from predicted ROIs (green bubbles), and benchmark ROIs in the yellow frame are also displayed for comparison; Gray frames: the fiber bundles emanating from corresponding ROIs (yellow bubbles) of three subjects in the training dataset illustrated as templates for comparison.

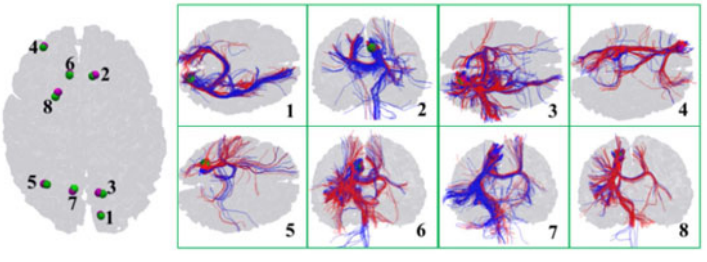

Fig. 4. Prediction results of 2 scans of ADNI-2 subject. The predicted ROI locations of 2 scans are illustrated on the left side with ROI ID labeled. Fiber tracts extracted of the 2 scans are overlapped on the right side. (The 1st scan: green bubbles, blue fibers; the 2nd scan: purple bubbles, red fibers).

Table 1. Distances between the predicted ROIs and the benchmark ROIs (mm)

ROIs	ROI 1	ROI 2	ROI 3	ROI 4	ROI 5	ROI 6	ROI 7	ROI 8	Mean
Our method	7.56	1.21	4.01	1.76	2.67	1.88	5.71	2.77	**3.45**
FSL FLIRT	7.48	4.95	6.24	6.18	5.41	4.51	6.46	4.54	**5.72**
HAMMER	6.18	4.92	6.40	7.50	5.15	3.06	6.08	4.95	**5.53**

4 Conclusion

We presented a novel framework for functional ROI prediction using working memory network as a testing bed. Our work has demonstrated that fiber shape models

of functional ROIs have remarkable prediction capability, providing direct support to the connectional fingerprint concept [8]. Our future work will include application and evaluation of this framework in other brain networks, and further application of this ROI prediction framework to clinical datasets such as the DTI data of ADNI-2 subjects.

References

1. Biswal, B.B., Mennes, M., Zuo, X.N., Gohel, S., Kelly, C., Smith, S.M., Beckmann, C.F., Adelstein, J.S., Buckner, R.L., Colcombe, S., Dogonowski, A.M., Ernst, M., Fair, D., Hampson, M., Hoptman, M.J., Hyde, J.S., Kiviniemi, V.J., Kötter, R., Li, S.J., Lin, C.P., Lowe, M.J., Mackay, C., Madden, D.J., Madsen, K.H., Margulies, D.S., Mayberg, H.S., McMahon, K., Monk, C.S., Mostofsky, S.H., Nagel, B.J., Pekar, J.J., Peltier, S.J., Petersen, S.E., Riedl, V., Rombouts, S.A., Rypma, B., Schlaggar, B.L., Schmidt, S., Seidler, R.D., Siegle, G.J., Sorg, C., Teng, G.J., Veijola, J., Villringer, A., Walter, M., Wang, L., Weng, X.C., Whitfield-Gabrieli, S., Williamson, P., Windischberger, C., Zang, Y.F., Zhang, H.Y., Castellanos, F.X., Milham, M.P.: Toward discovery science of human brain function. PNAS. 107(10), 4734--4739 (2010)
2. Sporns, O., Tononi, G., Kötter, R.: The human connectome: A structural description of the human brain. PLoS Comput Biol. 1(4), e42 (2005)
3. Van, Dijk. K.R., Hedden, T., Venkataraman, A., Evans, K.C., Lazar, S.W., Buckner, R.L.: Intrinsic functional connectivity as a tool for human connectomics: theory, properties, and optimization. J Neurophysiol. 103(1): 297--321 (2010)
4. Hagmann, P., Cammoun, L., Gigandet, X., Gerhard, S., Grant, P.E., Wedeen, V., Meuli, R., Thiran, J.P., Honey, C.J., Sporns, O.: MR connectomics: Principles and challenges. J Neurosci Methods. 194(1), 34--45 (2010)
5. Human Connectome Project, http://www.humanconnectomeproject.org/overview/
6. Li, K., Guo, L., Faraco, C., Zhu, D., Deng F., Zhang, T., Jiang, X., Zhang, D., Chen, H., Hu, X., Miller, S., Liu, T.: Individualized ROI Optimization via Maximization of Group-wise Consistency of Structural and Functional Profiles. NIPS. (2010)
7. Honey, C.J., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J.P., Meuli, R., Hagmann, P.: Predicting human resting-state functional connectivity from structural connectivity. PNAS. 106(6), 2035--2040 (2009)
8. Passingham, R.E., Stephan, K.E., Kötter, R.: The anatomical basis of functional localization in the cortex. Nat Rev Neurosci. 3(8), 606--616 (2002)
9. Liu, T., Li, H., Wong, K., Tarokh, A., Guo, L., Wong, S.T.: Brain Tissue Segmentation Based on DTI Data. NeuroImage. 38(1), 114--123 (2007)
10. Hu, X., Guo, L., Zhang, T., Li, G., Nie, J., Jiang, X., Zhang, D., Liu, T.: Joint analysis of fiber shape and cortical folding patterns. ISBI. (2010)
11. Shen, D., Davatzikos, C.: HAMMER: hierarchical attribute matching mechanism for elastic registration. IEEE Trans Med Imaging. 21(11), 1421--1439 (2002)

Predictive Modeling of Cardiac Fiber Orientation Using the Knutsson Mapping

Karim Lekadir, Babak Ghafaryasl, Emma Muñoz-Moreno, Constantine Butakoff, Corné Hoogendoorn, and Alejandro F. Frangi

Center for Computational Imaging & Simulation Technologies in Biomedicine
Universitat Pompeu Fabra and CIBER-BBN, Barcelona, Spain
karim.lekadir@upf.edu

Abstract. The construction of realistic subject-specific models of the myocardial fiber architecture is relevant to the understanding and simulation of the electro-mechanical behavior of the heart. This paper presents a statistical approach for the prediction of fiber orientation from myocardial morphology based on the Knutsson mapping. In this space, the orientation of each fiber is represented in a continuous and distance preserving manner, thus allowing for consistent statistical analysis of the data. Furthermore, the directions in the shape space which correlate most with the myocardial fiber orientations are extracted and used for subsequent prediction. With this approach and unlike existing models, all shape information is taken into account in the analysis and the obtained latent variables are statistically optimal to predict fiber orientation in new datasets. The proposed technique is validated based on a sample of canine Diffusion Tensor Imaging (DTI) datasets and the results demonstrate marked improvement in cardiac fiber orientation modeling and prediction.

1 Introduction

The study of fiber structure in the myocardium is of considerable importance for the understanding of its electro-mechanical behavior and associated cardiac pathologies. Over the years, there has been an increasing interest in the use of Diffusion Tensor Imaging (DTI) for the measurement of cardiac fiber orientation [1]. However, the modality is currently infeasible *in vivo* due to its great sensitivity to cardiac motion. On the other hand, myocardial fibers tend to be arranged according to consistent patterns across individuals yet they are crucial in determining electrical propagation and contraction patterns. As a result, computational techniques are required to construct accurate subject-specific models of myocardial fiber orientation from *ex vivo* datasets so as to enable realistic simulation of the electro-mechanical function of the heart. Existing research emphasizes the need for modeling techniques that can encode the relationship between myocardial shape and fiber orientation, as well as that can take into account the inter-subject variability in fiber structure [1]. Amongst existing methods, a synthetic fiber model has been used in [2], which provides only the overall trend of the structure and thus lacks accuracy in subject-specific modeling. The work in [3] estimates fiber structure by warping the data to a predefined single template but the method is not based on a sample population and thus does not take

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 50–57, 2011.

into account the variability found in fiber orientation and shape. The Streeter model [4], on the other hand, is built from a database of histological datasets and thus far remains the most established model in cardiac simulation. The technique fits to the training data parametric equations relating the position in the myocardium with a selected set of global and local shape information, such as the left ventricular axis and position relative to the myocardial walls. It can be argued, however, that a significant amount of shape information is ignored, thus producing a model that might not approximate well enough the complex distribution of the fibers.

This paper presents an alternative approach through which the relationship between myocardial morphology and fiber orientation is statistically encoded based on a training sample population. To this end, the proposed technique firstly introduces a representation of the fibers in a 5D space that is continuous and distance preserving, thus allowing consistent manipulation of the orientation data. Additionally, the technique makes use of all the available shape information and extracts the latent directions in the shape space that correlate most with the fiber orientation. A non-linear regression technique is then used to accurately estimate fiber orientation in new datasets based on the myocardial morphology. The proposed predictive model is validated based on a sample of normal canine cardiac DTI datasets.

2 Methods

2.1 Dataset Description and Preprocessing

To construct a predictive model of the myocardial fiber orientations with the proposed approach, we use a database of seven normal *ex vivo* DTI datasets of canine hearts (Center of Cardiovascular Bioinformatics and Modeling, at John Hopkins University). The images were produced using a 1.5T GE CV/I MRI scanner with an enhanced gradient system, 40 mT/m maximum gradient amplitude and a 150 T/m/s slew rate. The first eigenvector of the obtained DTI tensors was calculated to estimate the 3D vector describing the fiber orientation at each voxel. Furthermore, the predictive model proposed in this paper is aimed at encoding the statistical relationship between myocardial shape and fiber orientation. As a result, it is required to build surface meshes of the myocardial shapes and volumetric meshes of the myocardial fibers that correspond between the training datasets. To this end, the left ventricular walls are firstly segmented in each of the DTI volumes over the maps of Fractional Anisotropy (FA) using a semi-automatic tool [5]. The point correspondence between the obtained surface meshes is established based on the position of key anatomical features (*i.e.*, mitral valve, apex, and aorta). The mean shape is subsequently computed and used to derive a volumetric mesh that contains the fiber orientation at uniformly distributed nodes within the myocardium. Finally, to obtain the fiber information at corresponding locations of the myocardium, the mean volumetric mesh is then mapped onto all image volumes by using the Thin Plate Spline (TPS) technique [6].

2.2 Knutsson Mapping of Fiber Orientation

To enable a consistent manipulation of the data within the proposed predictive model, it is important to use a suitable representation of the fiber orientations. The 3D vectors

obtained from the diffusion tensors are non-directional and thus opposite vectors can represent the same fiber orientation. This introduces a fundamental ambiguity that must be specifically handled to avoid potential errors in the statistical analysis. In 2D, this is resolved easily by using the double angle trick, *i.e.*, the new orientation coordinates are the cosine and sinus of twice the angle between the line and a fixed coordinate axis. In 3D, however, the situation is more complicated and the solution presented in [7] involves a mapping to a 5D space with interesting mathematical properties. Given a unit vector characterized by the angles θ, ϕ of its spherical representation, the obtained Knutsson function $M_K : R^3 \rightarrow R^5$ can be derived as [7]:

$$\begin{aligned} s &= \sin^2\theta \cos 2\phi \\ t &= \sin^2\theta \sin 2\phi \\ u &= \sin 2\theta \cos\phi \\ v &= \sin 2\theta \sin\phi \\ w &= \sqrt{3}(\cos^2\theta - \frac{1}{3}) \end{aligned} \tag{1}$$

The Knutsson coordinates s, t, u, v, w essentially map half of a sphere onto a 5D space such that the surface is uniformly stretched in all directions and all pairs of opposite vectors are mapped onto the same location in the 5D space. The mapping described in Eq. (1) is therefore unique and removes the ambiguity found in directional vectors. More importantly, it can be shown that the Knutsson mapping induces a continuous representation that preserves distance [8] (thus the importance of the normalizing constants in the definition of the coordinate w in Eq. (1)). Applications of the Knutsson mapping in the literature include curvature estimation in 3D images [8] and the analysis of high angular diffusion tensors of the brain [9].

With this representation, operations on the cardiac fiber orientations can be carried out consistently and the results need to be mapped back to the original 3D space, for example for visualization and simulation purposes. This can be carried out by first approximating the product $\mathbf{x}\mathbf{x}^T$ from the mapping M_K as follows:

$$\mathbf{x}\mathbf{x}^T = \frac{1}{2}\begin{pmatrix} \frac{2-\sqrt{3}w}{3} + s & t & u \\ t & \frac{2-\sqrt{3}w}{3} - s & v \\ u & v & \frac{1+\sqrt{3}w}{3} \end{pmatrix} \tag{2}$$

Subsequently, it can be then demonstrated that the 3D orientation can be derived from the principal eigenvector of $\mathbf{x}\mathbf{x}^T$ [10].

2.3 Nonlinear Predictive Modeling

Existing research on cardiac fiber architecture suggests that a strong relationship exists between myocardial morphology and fiber structure. It is not trivial, however, to define for each region of the myocardium the shape information that is most suitable to estimate the corresponding fiber orientations. In the Streeter model [4], for example, it is assumed that the fibers are better described with respect to a number of preselected variables, such as the left ventricular axis and relative position to the myocardial walls. As a result, a substantial amount of available shape information is simply ignored with such approach, which can potentially limit the accuracy of the model. In this paper, the aim is to use all available shape information and automatically extract the variables in the shape space that are most relevant to the estimation of fiber orientation based on statistical criteria. In other words, the axes of shape variation that correlate the most with the fiber orientation maps need to be extracted and used to optimize the prediction capability of the obtained model.

Partial Least Squares (PLS) regression [11] is a dimensionality reduction technique with several advantages that make it ideal for the predictive modeling task of this paper. It can produce efficient regression models in situations where the number of predictors is very large and at the same time the training sample is small (thus introducing singular data matrices). This is beneficial for predictive modeling of fiber orientation due to the generally limited availability of *ex vivo* DTI datasets. In the following, let $\mathbf{X}$ be the data matrix representing the aligned myocardial shapes in the training set and $\mathbf{Y}$ the matrix of the 5D Knutsson data representing the fiber structures inside the myocardium. The aim of partial least squares regression is to perform a simultaneous decomposition of $\mathbf{X}$ and $\mathbf{Y}$ such that the score vectors obtained along the new representation axes of both the input and output matrices correlate most. One solution to the problem can be obtained through the NIPALS algorithm [11], which iteratively estimates new vectors $\mathbf{a}$ and $\mathbf{b}$ as:

$$[\mathrm{cov}(\mathbf{Xa}, \mathbf{Yb})]^2 = \max_{\|\mathbf{c}\|=\|\mathbf{d}\|=1} [\mathrm{cov}(\mathbf{Xc}, \mathbf{Yd})]^2 \tag{3}$$

At each iteration, the data matrices $\mathbf{X}$ and $\mathbf{Y}$ are deflated by subtracting the contribution along the latent axes found at previous step. At convergence, the predictors and outputs are decomposed as $\mathbf{X} = \mathbf{TP}^T + \mathbf{E}$ and $\mathbf{Y} = \mathbf{UQ}^T + \mathbf{F}$, respectively, where $\mathbf{P}$ and $\mathbf{Q}$ are the extracted latent vectors. The final step is to estimate the regression coefficient matrix $\mathbf{D}$ and the predictive model is obtained as:

$$\hat{\mathbf{Y}} = \mathbf{TDQ}^T \tag{4}$$

It is important to note that only a subset of the latent vectors are used in the prediction and these are selected using cross validation such that they minimize the reconstruction of the output data in the training.

The mapping introduced in this paper for the representation of fiber orientation, as well as the nature of the data, can introduce a nonlinear interdependency between the matrices $\mathbf{X}$ and $\mathbf{Y}$. As a result, it is more appropriate to use a nonlinear implementation of the partial least square algorithm. Using the kernel version developed in [12], this can be achieved with a kernel transformation of the input data

followed by the application of the linear PLS technique described above. The kernel function Φ used for this purpose is typically defined as Gaussian or polynomial. The kernel gram matrix $\mathbf{K} = \Phi\Phi^T$ of the cross product between all input data points is computed and used to obtain the final nonlinear predictive model as follows:

$$\hat{\mathbf{Y}} = \mathbf{K}_t\mathbf{U}(\mathbf{T}^T\mathbf{K}\mathbf{U})^{-1}\mathbf{T}^T\mathbf{Y} \tag{5}$$

where $\mathbf{K}_t$ is the kernel Gram test matrix estimated from the new input data. More details regarding the kernel PLS implementation can be found in [11].

3 Results

The proposed predictive model is validated using the seven DTI datasets of normal canine hearts described in Section 2.1 on a leave-one-out basis. The absolute value of the dot product between the predicted fiber orientations and the original data is calculated to measure the extent of agreement and the accuracy of the proposed technique. A value close to 0 indicates strong disagreement, while a dot product close to 1 indicates high prediction accuracy. Firstly, we applied a template warping approach between all pairs of datasets similarly to the work in [3] and using TPS warping. We found the results to be inconsistent depending on the choice of the subject to serve as a template, with a similarity measure equal to 0.47 for subject 5 versus subject 1 and equal to 0.85 for subject 4 versus subject 1. This justifies well the use of a statistical approach as presented in this paper, where the variability in fiber orientation between the datasets is specifically modeled.

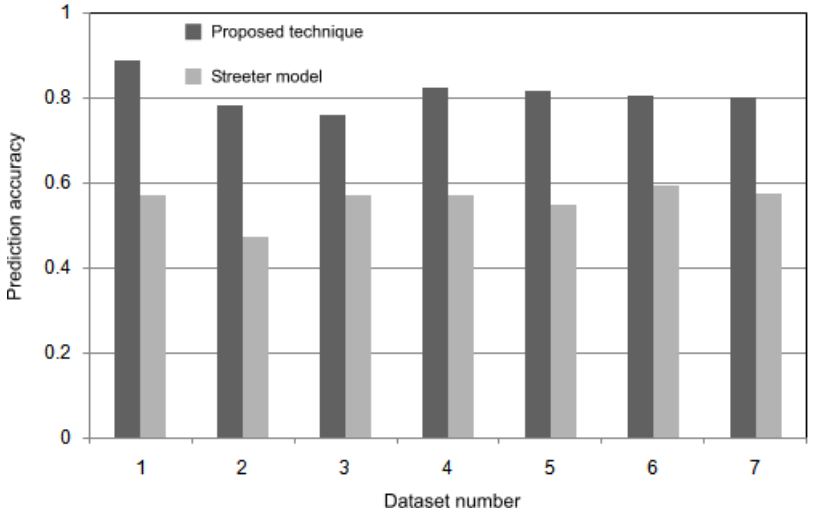

Fig. 1. Dot product accuracy results for the seven canine DTI images using the Streeter model and the proposed technique, showing significant improvement throughout all datasets

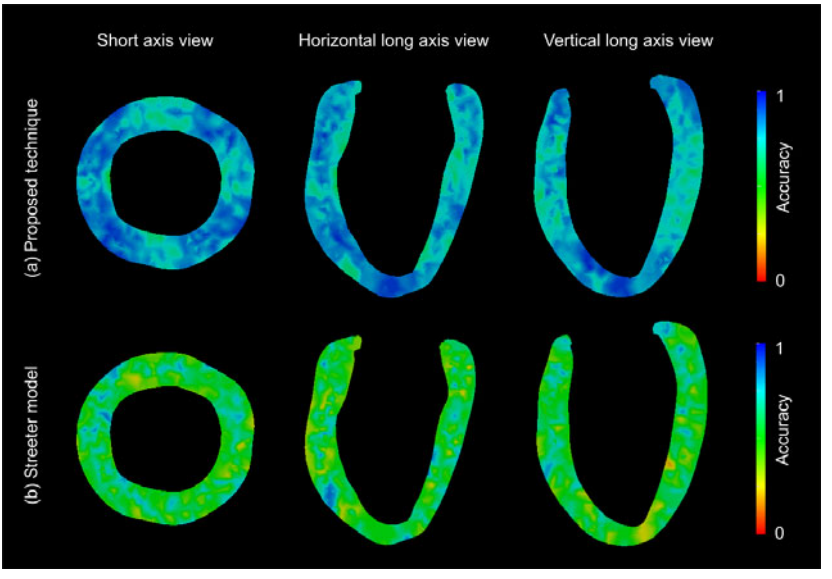

Fig. 2. Average accuracy maps for the proposed technique and the existing Streeter model

Furthermore, the Streeter model [4] was implemented and tested against the accuracy obtained with the approach presented in this paper. The results for both techniques are plotted in Fig. 1, where it can be seen that the proposed method improves the prediction of fiber orientation for all datasets. These results are significant particularly given the fact that the predictive models were built from a small training set, which shows the performance capability of the non-linear PLS method in this work. Table I summarizes the average, standard deviation and minimal accuracy results, where it can be observed that the proposed predictive model is not only accurate but also consistent throughout all datasets, with an average accuracy equal to 0.81 and a minimal accuracy of 0.75. It is also worth noting from Table I how the introduction of the Knutsson mapping for the representation of fiber orientation (third column) and the use of the nonlinear kernel technique for prediction (fourth column of the table) increase the accuracy of the model.

Table 1. Prediction accuracy using the proposed technique

Accuracy	Streeter model	L-PLSR in 3D space	L-PLSR in Knutsson space	NL-PLSR in Knutsson space
Average	0.55	0.71	0.75	0.81
Std	0.06	0.08	0.06	0.04
Minimum	0.47	0.65	0.68	0.75

For detailed visualization of the performance of the proposed technique, the average prediction accuracy for each fiber location in the myocardium is displayed in Fig. 2 for different long- and short-axis views. It can be seen that the fiber orientation prediction is accurate consistently for all regions of the myocardium, unlike the Streeter model which displays less accurate results. This is also due to the fact that the Streeter technique uses the same parametric model for all regions of the myocardium. On the contrary, the proposed technique extracts the latent structures most relevant for the prediction of each individual fiber orientation, which is a natural approach since the relationship between shape and fiber is expected to vary for different regions of the myocardium. Finally, an illustration of the predicted fiber orientations using the proposed technique is given in Fig. 3, where significant agreement with the original DTI data can be observed in all views and regions.

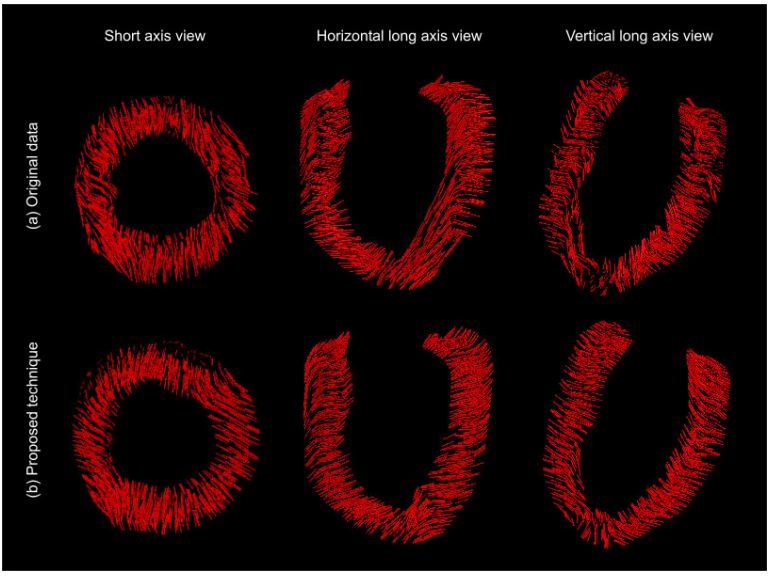

Fig. 3. Illustration of the fiber structure obtained by the proposed technique (b) as compared to the ground truth (a) derived from the original DTI data

4 Conclusions

Our conclusions following the work presented in this paper are twofold: firstly, the relationship between myocardial shape and fiber orientation is significant yet it is not trivial. Therefore it must be defined automatically and statistically by extracting the optimal latent structures from a training sample population. Secondly, the representation of orientation data in general, and particularly for the study of fiber structure, requires continuous and distance preserving mappings in order to allow for consistent statistical analysis, such as by using the Knutsson mapping described in this work. The initial results reported in this paper show the potential of the technique for fiber structure modeling and prediction.

Acknowledgement. This work was funded by the CENIT program of the CDTI, the Industrial and Technological Development Centre of Spain, under the research project cvREMOD (ref. CEN20091044).

References

1. Peyrat, J.-M., Sermesant, M., Pennec, X., Delingette, H., Xu, C., McVeigh, E.R., Ayache, N.: Towards a statistical atlas of cardiac fiber structure. In: Larsen, R., Nielsen, M., Sporring, J. (eds.) MICCAI 2006. LNCS, vol. 4190, pp. 297–304. Springer, Heidelberg (2006)
2. Sermesant, M., Rhode, K., Sanchez-Ortiz, G.I., Camara, O., Andriantsimiavona, R., Hegde, S., Rueckert, D., Lambiase, P., Bucknall, C., Rosenthal, E., Delingette, H., Hill, D.L.G., Ayache, N., Razavi, R.: Simulation of cardiac pathologies using an electromechanical biventricular model and XMR interventional imaging. Medical Image Analysis 9(5), 467–480 (2005)
3. Sundar, H., Shen, D., Biros, G., Litt, H., Davatzikos, C.: Estimating myocardial fiber orientation by template warping. In: IEEE International Symposium on Biomedical Imaging (2006)
4. Streeter, D.: The cardiovascular system: gross morphology and fiber geometry of the heart, in Handbook of physiology. Williams & Wilkins, Baltimore (1979)
5. Larrabide, I., Omedas, P., Martelli, Y., Planes, X., Nieber, M., Moya, J.A., Butakoff, C., Sebastián, R., Camara, O., Craene, M.D., Bijnens, B., Frangi, A.F.: GIMIAS: an open source framework for efficient development of research tools and clinical prototypes. In: Functional Imaging and Modeling of the Heart (2009)
6. Bookstein, F.L.: Principal warps: thin-plate splines and the decomposition of deformations. IEEE Transactions on Pattern Analysis and Machine Intelligence 11(6), 567–585 (1989)
7. Knutsson, H.: Producing a continuous and distance preserving 5-D vector representation of 3-D orientation. In: IEEE Computer Society Workshop on Computer Architecture for Pattern Analysis and Image Database Management, Miami Beach, Florida (1985)
8. Rieger, B., Timmermans, F.J., Vliet, L.J.v., Verbeek, P.W.: Curvature estimation of surfaces in 3D grey-value images. In: International Conference on Pattern Recognition, ICPR (2002)
9. Bhalerao, A., Westin, C.-F.: Hyperspherical von Mises-Fisher Mixture (HvMF) Modelling of High Angular Resolution Diffusion MRI. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) MICCAI 2007, Part I. LNCS, vol. 4791, pp. 236–243. Springer, Heidelberg (2007)
10. Knutsson, H.: Representing local structure using tensor. In: Scandinavian Conference in Image Analysis, Oulu, Finland (1989)
11. Abdi, H.: Partial least squares regression (PLS-regression), in Encyclopedia for research methods for the social sciences. Sage, Thousand Oaks (2003)
12. Rosipal, R., Trejo, L.J.: Kernel patial least squares regression in reproducing kernel Hilbert space. Journal of Machine Learning Research 2, 97–123 (2002)

Sparse Multi-Shell Diffusion Imaging

Yogesh Rathi[1], O. Michailovich[2], K. Setsompop[3], S. Bouix[1], M.E. Shenton[1], and C.-F. Westin[1]

[1] Brigham and Women's Hospital, Harvard Medical School, Boston
[2] Department of Electrical Engg., University of Waterloo, Canada
[3] Massachusetts General Hospital, Harvard Medical School, Boston

Abstract. Diffusion magnetic resonance imaging (dMRI) is an important tool that allows non-invasive investigation of neural architecture of the brain. The data obtained from these *in-vivo* scans provides important information about the integrity and connectivity of neural fiber bundles in the brain. A multi-shell imaging (MSI) scan can be of great value in the study of several psychiatric and neurological disorders, yet its usability has been limited due to the long acquisition times required. A typical MSI scan involves acquiring a large number of gradient directions for the 2 (or more) spherical shells (several b-values), making the acquisition time significantly long for clinical application. In this work, we propose to use results from the theory of compressive sampling and determine the minimum number of gradient directions required to attain signal reconstruction similar to a traditional MSI scan. In particular, we propose a generalization of the single shell *spherical ridgelets* basis for sparse representation of multi shell signals. We demonstrate its efficacy on several synthetic and *in-vivo* data sets and perform quantitative comparisons with solid spherical harmonics based representation. Our preliminary results show that around 20-24 directions per shell are enough for robustly recovering the diffusion propagator.

1 Introduction

A popular dMRI acquisition technique is High Angular Resolution Diffusion Imaging (HARDI), which involves acquiring diffusion information for a single b-value (single shell) in several gradient directions uniformly spread on a sphere [1]. While this protocol allows for resolving the angular structure of the neural fibers, it does not provide information about the radial signal decay, which is sensitive to white matter anomalies.

To obtain accurate information about the neural architecture, diffusion spectrum imaging (DSI) was proposed by [2]. This high resolution technique requires upwards of 512 gradient directions and more than an hour to scan each subject (spatial resolution of $2mm^3$), which makes it impractical to use in clinical settings. A few works have attempted to reduce the scan time using compressed sensing for DSI [3,4], however, the acquisition time is still too long for clinical applications. Consequently, other imaging schemes have been proposed, namely, Hybrid Diffusion Imaging (HYDI) [5], Diffusion Propagator Imaging (DPI) [6],

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 58–65, 2011.

Diffusion Kurtosis Imaging (DKI) [7], spherical polar Fourier basis [8] and high-order tensor models [9,10], all of which fall under the category of multi-shell imaging (MSI). Each of these techniques captures different aspect of the underlying tissue geometry. They acquire important information about the neural tissues, which are missed by HARDI methods, yet, they are seldom used in clinical studies due to their long scan times. In general, it takes around 30-50 minutes to scan each subject with: $2mm^3$ spatial resolution, 60 gradient directions per shell, 2-4 shells with b-values of 1000 to 6000. The acquisition time is directly proportional to the number of gradient direction acquisitions. Thus, if we can reduce the number of gradient acquisitions by half (without sacrificing the quality), the scan time reduces by 50%. This fact is the main motivation behind the proposed work.

In order to recover the MSI signal from very few measurements N, we propose to use the theory of *compressive sampling* (CS). Formalizing such a reconstruction approach in the context of MSI constitutes the main contribution of this work. In particular, we propose to generalize the spherical ridgelets (SR) basis of [11] for the case of multi-shell signals. We demonstrate that the MSI signal representation in the proposed basis is indeed sparse and compute error statistics on synthetic data sets. Further, the representation allows for a direct quantification of the signal decay as a function of b-values, which can be a useful measure in neuroimaging studies.

2 Compressive Sampling

The diffusion signal $S(\mathbf{q})$ is a real-valued function, which determines the value of S at location $\mathbf{q}$ in q-space. The scalar q is given by $q = \|\mathbf{q}\|$, with $\mathbf{q} = q\mathbf{u}$, where $\mathbf{u} \in \mathbb{S}^2$. In the context of MSI, the signal $S(\mathbf{q})$ is measured along N discrete orientations $\{\mathbf{u}_k\}_{k=1}^{N}$ for several different q values (Q shells). Thus, for each q value, measurements are made along N directions uniformly spread on a sphere, giving the measurements a multi-shell (also referred to as multi-b or multi-q) structure. In such a case, all experimental information on $S(\mathbf{q})$ is represented by its NQ values $\{S_k\}_{k=1}^{NQ}$ corresponding to each of the $\mathbf{q}$ values. The most fundamental question in this regard is: what is the minimum number of diffusion directions N (on each shell) required to unambiguously represent the signal $S(\mathbf{q})$ in terms of its discrete values S_k?

A particularly important answer to the above question is offered by the theory of CS [12,13]. In particular, the theory specifies conditions under which the original signal $S(\mathbf{q})$ can be perfectly recovered from a much smaller number of its samples than what would be required by the classical sampling theory. Since diffusion measurements are linear, the discrete values S_k can be expressed in the form of inner products $S_k = \langle S(\mathbf{q}), \varphi_{j_k}(\mathbf{q})\rangle$, with $\{\varphi_{j_k}\}_{k=1}^{N}$ being a subset of a *Dirac sampling basis* $\{\varphi_i\}_{i\in\mathcal{I}}$. Moreover, let $\{\psi_j\}_{j\in\mathcal{J}}$ be *another* basis in the signal space, which we will use for *representation* of $S(\mathbf{q})$. In particular, we are interested in representing the diffusion (MSI) signal $S(\mathbf{q})$ in the form of a linear combination $S(\mathbf{q}) = \sum_{j\in\mathcal{J}} c_j\, \psi_j(\mathbf{q})$, where $\mathcal{J}$ denotes the set of indices

over which the basis functions ψ_j are counted. Note that, in a more general setup, the set $\{\psi_j\}_{j\in\mathcal{J}}$ may be *overcomplete*, but finite, with its total number of elements being equal to M. Then, the theory of CS proves that an accurate approximation of S is possible from only $\mathcal{O}\left(\mu^2 \log(M) L\right)$ of its measurements, if the following conditions are satisfied:

(a) S is assumed to be *sparsely representable* by $\{\psi_j\}_{j\in\mathcal{J}}$, which implies that the number L of non-zero coefficients c_j is significantly less than M.
(b) The bases $\{\varphi_i\}_{i\in\mathcal{I}}$ and $\{\psi_j\}_{j\in\mathcal{J}}$ are *incoherent*, implying that the value of $\mu = \sup_{i,j} |\langle \varphi_i(\mathbf{q}), \psi_j(\mathbf{q})\rangle|$ is relatively small.

The above considerations suggest that the applicability of CS to MSI depends on the availability of a basis $\{\psi_j\}_{j\in\mathcal{J}}$ for which the assumptions (a) and (b) above would be valid. Such a basis was introduced for HARDI data in [11], where it is called a basis of *spherical ridgelets*. In the next section, we demonstrate how to generalize this basis for sparse representation of MSI signals. Just as in the case of HARDI, the energy of the proposed basis of spherical ridgelets is distributed alongside the great circles of $\mathbb{S}^2$ (for every shell), which is very incoherent with respect to the Dirac sampling basis $\{\varphi_i\}_{i\in\mathcal{I}}$. Thus, the amount of incoherence μ for this basis is 0.56. Note that, this value is computed by normalizing each of the basis elements $\{\psi_j\}_{j\in\mathcal{J}}$ to unit norm.

3 Methods

Spherical ridgelets (SR) were proposed in [14,11] following the theory of multiresolution analysis on the sphere. For the case of HARDI data, the SR basis is given by: $\mathbb{F} := \left\{\Psi_{j,\mathbf{v}} \mid \mathbf{v} \in \mathbb{S}^2, j \in \mathbb{N}\cup\{-1\}\right\}$, where

$$\begin{aligned}
\Psi_{j,\mathbf{v}} &= \frac{1}{2\pi}\begin{cases}\mathcal{K}_{0,\mathbf{v}}, & \text{if } j=-1,\\ \mathcal{K}_{j+1,\mathbf{v}} - \mathcal{K}_{j,\mathbf{v}}, & \text{if } j\in\mathbb{N}, \quad \text{with,}\end{cases}\\
\mathcal{K}_{j,\mathbf{v}}(\mathbf{u}) &= \sum_{n=0}^{\infty} \frac{2n+1}{4\pi}\,\lambda_n\,\kappa_j(n)\,P_n(\mathbf{u}\cdot\mathbf{v}), \text{ where, } \kappa_j(n) = \kappa(\rho, 2^{-j}n)\\
\kappa(\rho,x) &= \exp\{-\rho\,x\,(x+1)\}, \quad \lambda_n = \begin{cases}2\pi(-1)^{n/2}\frac{1\cdot 3\cdots(n-1)}{2\cdot 4\cdots n}, & \text{if } n \text{ is even}\\ 0, & \text{if } n \text{ is odd.}\end{cases}
\end{aligned} \tag{1}$$

where P_n is the legendre polynomial of order n, $\mathbf{u}\in\mathbb{S}^2$ and ρ is a user-defined parameter. Notice that, the basis functions $\Psi_{j,\mathbf{v}}$ can only represent signals defined on a single spherical shell.

3.1 From Single to Multiple Shells

The structure and magnitude of the signal varies significantly as the b-value (or q-value) increases. This is evident from Figure 1, where in the top row we show the frequency spectrum (in spherical harmonic basis) required to represent signals with increasing b-values. Notice that, a). higher frequencies are required

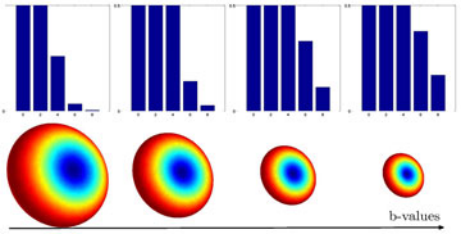

Fig. 1. Top row (bar graph) shows the spherical harmonic frequency spectrum (upto order 8) required to represent each of the signals at the bottom. Also notice the decay in amplitude of the signal as the b-values increase.

to represent sharper diffusion signals, and b). the signal decays in magnitude with b-values. We propose to use these two important facts in designing the MSI SR basis:

$$\mathcal{K}_{j,\mathbf{v}}(\mathbf{q}) = \underbrace{\exp(-\alpha q^2)\{1 - \sum_{m=0}^{h} a_m H_m(q)\}}_{I_1} \underbrace{\sum_{n=0}^{\infty} \frac{2n+1}{4\pi} \lambda_n \, \kappa_j(q,n) \, P_n(\mathbf{u}\cdot\mathbf{v})}_{I_2}, \quad (2)$$

where H_m are hermite polynomials of order m, and $\kappa_j(q,n)$ is now a function of q given by: $\kappa(q,n) = \exp\left(-\beta W(q) n(n+1)\right)$, with $W(q)$ given by the Weibull distribution function: $W(q) = \frac{k}{l}\left(\frac{q}{l}\right)^{k-1} \exp(-(q/l)^k)$ where we set $k = 0.8$, $l = 2$ in this work and β is a constant. We now describe the rationale behind choosing such a function $\mathcal{K}$ by individually examining the terms I_1 and I_2.

The term I_1 essentially models the decay rate of the signal magnitude. Some studies have shown a biexponential decay of the signal [15], while others have seen diffraction patterns with high q-values [16]. The expression $I_1 = \exp(-\alpha q^2)\{1 - \sum_{m=0}^{h} a_m H_m(q)\}$ can model both these phenomena, with the term $1 - \sum_{m=0}^{h} a_m H_m(q)$ modeling the departure from an exponential decay. Note that, the Hermite polynomials are capable of modeling diffraction patterns as shown in [16] and hence we use it in our model.

The term I_2 models the frequency component (sharpness) of the signal. As n increases in the summation in I_2, the legendre polynomial P_n incorporates higher frequencies. The desired behavior of the combined term $\kappa(q,n)P_n$ is to use only low frequencies for lower q-values and incrementally allow higher frequencies with increasing q. Thus, the term $\kappa(q,n)$ should act as a bandpass filter by selectively adding high frequency components of P_n for the appropriate q-values. This behavior can be modeled by choosing $W(q)$ to be a Weibull distribution function with parameters $l = 2$, $k = 0.8$. This function has a heavy tail, preventing the value of $W(q)$ from converging to zero too quickly as in an exponential function.

Combining the effects of increasing frequency and decreasing signal magnitude results in the desired behavior for sparse representation of multi-shell signals. The free parameters of this model are α and $\{a_m\}_{m=1}^{h}$ along with the sparse set of weights $\mathbf{c}$ for the overcomplete basis. Note that, if we set $\alpha = 0$ and $\{a_m\}_{m=1}^{h} = \{0\}$, then the expression for $\mathcal{K}_{j,\mathbf{v}}(\mathbf{q})$ in (2) reduces to that for a single shell HARDI (1), with the parameter ρ in (1) determined by evaluating

$W(q)$ for a specific q-value. Thus, the proposed basis is a generalization of the spherical ridgelets of [14]. We should note that the basis is still given by $\mathbb{F} := \left\{\Psi_{j,\mathbf{v}} \mid \mathbf{v} \in \mathbb{S}^2, j \in \mathbb{N} \cup \{-1\}\right\}$, as defined before, albeit with the modified form of the function $\mathcal{K}_{j,\mathbf{v}}(\mathbf{q})$.

3.2 Sparse Estimation

Given the measurements $\{S(\mathbf{q}_i)\}_{i=1}^{NQ}$, one can use (2) to compute the values of the spherical ridgelet basis for all $\mathbf{q}_i$ (Q is the total number of shells and N is the number of gradient directions per shell). The resulting values can be then stored into an $NQ \times M$ matrix $\mathcal{A}$, where M is the number of elements in the overcomplete basis $\mathcal{A}$ as defined in [11]. We use 3 discrete resolution levels $\{-1, 0, 1\}$ for each shell and set $h = 3$ in equation (2). Subsequently, if $\mathbf{c} \in \mathbb{R}^M$ is defined to be a (column) vector of ridgelet coefficients and $y := [S(\mathbf{q}_1), S(\mathbf{q}_2), \ldots, S(\mathbf{q}_{NQ})]^T$, then the measurement model can be formally expressed as $\mathcal{A}\,\mathbf{c} = y + e$, where e is an error vector that accounts for both measurement and model noises. From the theory of CS, a sparse estimate of coefficients $\mathbf{c}$ can be found by solving

$$\mathbf{c} = \arg\min_{\mathbf{c}} \|\mathbf{c}\|_1 \text{ subject to } \|\mathcal{A}\,\mathbf{c} - y\|_2 \leq \eta., \tag{3}$$

where η depends on the level of noise expected in the signal. Note that (3) is a convex optimization problem, and can be solved using the L_1 homotopy algorithm of [17][1].

In the present scenario, we not only have to estimate the sparse vector $\mathbf{c}$, but also the parameters for the radial decay. We do this by the method of *coordinate descent*, wherein the coefficients $\mathbf{c}$ and $\{\alpha, a_m\}$ are estimated alternately. The estimation framework is as follows:

Algorithm 1. Algorithm for sparse estimation of MSI data

1: Initialize $\alpha = \alpha_0$ and $a_m = \{0\}$.
2: **while** $\|\mathbf{c}\|_1$ is decreasing **do**
3: Estimate $\mathbf{c}$ using equation (3).
4: Update α and $\{a_m\}_{m=1}^{h}$ using gradient descent of $\|\mathcal{A}\,\mathbf{c} - y\|^2$.
5: **end while**

The gradient with respect to a_m is given by $< [.. \bigtriangledown_{a_m} \Psi_{j,\mathbf{v}} ..]\mathbf{c}, \mathcal{A}\,\mathbf{c} - y >$, where $< ., . >$ is the euclidean inner product. The gradient only affects the radial part of $\Psi_{j,\mathbf{v}}$, resulting in $\bigtriangledown_{a_m} \Psi_{j,\mathbf{v}} = -2e^{-\alpha q^2} H_m(q) I_2$, where I_2 is as defined in (2). A similar expression for gradient of α can be easily deduced from (2).

4 Experiments

Synthetic Data: We tested the proposed algorithm on synthetic data and quantitatively compared it with the solid spherical harmonics (SH) based method of

[1] http://users.ece.gatech.edu/~sasif/homotopy/

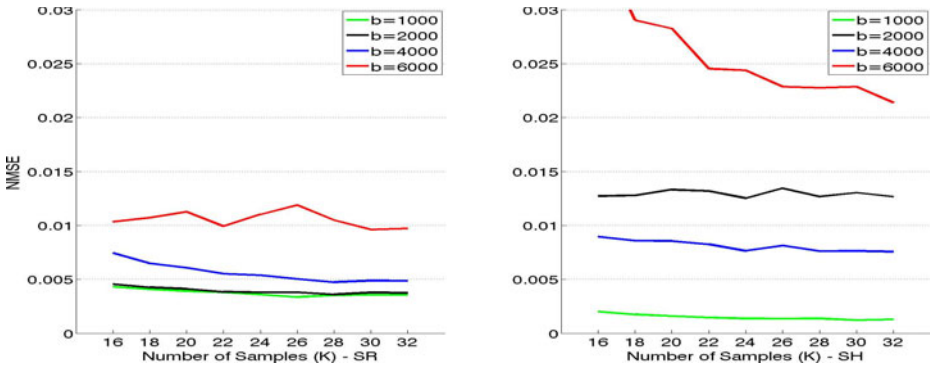

(a) NMSE

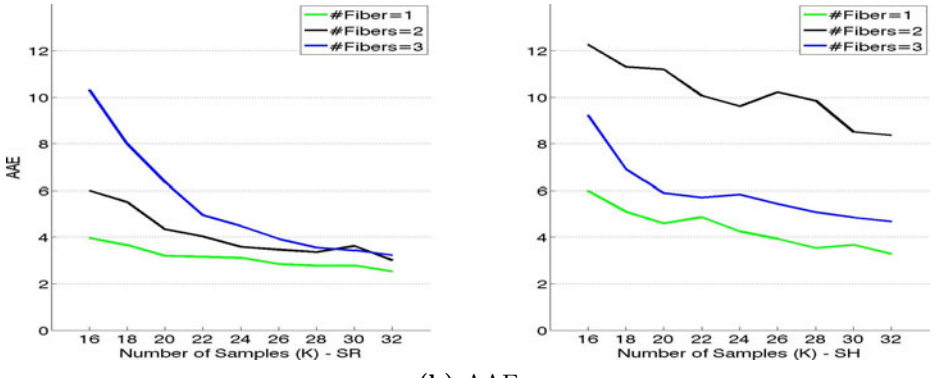

(b) AAE

Fig. 2. The left two figures (a) show the NMSE for different values of $K = [16, 32]$ (x-axis) and for b=$\{1000, 2000, 4000, 6000\}$. First figure is with SR basis and second with SH basis. (b) shows AAE in degrees (x-axis is K) for different number of crossing fibers (1,2,3). 3^{rd} figure is with the proposed SR basis and 4^{th} with the SH basis of [6].

[6]. Synthetic data sets were generated using a mixture of biexponential models: $S(\mathbf{q}) = \sum_i f_i G_i(\mathbf{q})$, where f_i are weight fractions set to $1/F$, where F is the number of fibers and $G_i(\mathbf{q}) = 0.7\exp(-q^2\mathbf{u}D\mathbf{u}^T) + 0.3\exp(-\theta q^2\mathbf{u}D\mathbf{u}^T)$, where $D = diag\{0.0017, 0.0003, 0.0003\}$ and $\theta = 1/3$. Four different b-values were used $b = \{1000, 2000, 4000, 6000\}$ to obtain 4 shells. Rician noise was added to each shell so that the SNR for each of the shells was: $\{5, 4, 3, 2\}$ respectively, where SNR is defined by σ_s/σ_n, with σ_s being the standard deviation of signal and σ_n the standard deviation of noise. 1000 random samples were generated by randomly choosing the number of fiber crossings (maximum of 3) and random orientation between the fibers. Two different error metrics were used to determine the quality of fit using the proposed SR basis and the SH basis of [6]: *(i)* Normalized Mean Squared Error (NMSE) given by $\frac{\|S(\mathbf{q})-\hat{S}(\mathbf{q})\|^2}{\|S(\mathbf{q})\|^2}$, where $\hat{S}(\mathbf{q})$ is the estimated signal. *(ii)* Average Angular Error (AAE) was computed between principal diffusion

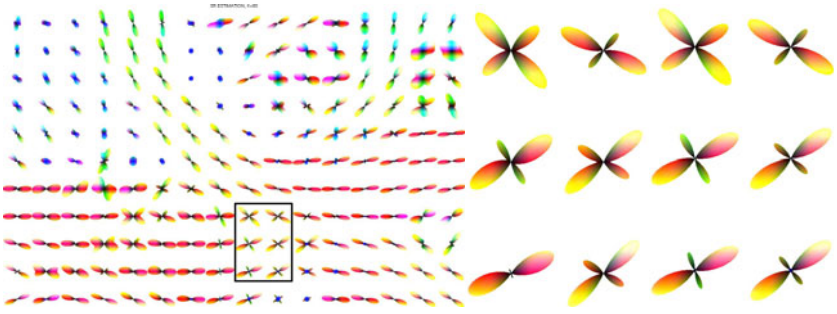

(a) K=60 (all samples) **(b)** K=16 **(c)** K=28

Fig. 3. Top left (a) shows the propagator for radius $R_0 = 20\mu m$ with $K = 60$. The remaining figures show a zoomed-in version (the black box) of the propagator as computed by estimating the signal with $K = 16, 28$ per shell. Note that the values of the propagator were scaled to uniform size for better visualization.

directions of the known ground truth and the ones estimated from the diffusion propagator. For the SR basis, the propagator was computed numerically as in [5]. Error statistics were computed by estimating the signal using different number of measurements $K = [16, 32]$ per shell. Thus, with $K = 16$, only a total of 64 measurements (for 4 shells) were used in estimating the signal.

Figure 3(a) shows NMSE as computed for each of the 4 shells for various K. Also shown is the AAE for various K and different number of crossing fibers.

As seen in the figures, the proposed SR basis better fits the signal with lower errors for higher b-values. $K = [20, 24]$ per shell seems to be sufficient to model the MSI signal without significant errors. The SH basis on the other hand tends to oversmooth the signal (as seen by higher NMSE and AAE). *We should note that, having 4 shells is not a necessity for our method, but we used it to demonstrate the accuracy with which the proposed SR basis fits the signal for all b-values.*

In-vivo Data: Our in-vivo data consisted of a human brain scanned on a 3T Siemens scanner with the following parameters: $2.5mm^3$ spatial resolution, $b = \{900, 2000, 5600\}$ and 60 gradient directions per shell. For each shell, the signal was subsampled to obtain the desired measurements for different $K = [16, 32]$ and the proposed SR basis was used to fit the data. Figure 4 shows the diffusion propagator for various K and radius $R_0 = 20\mu m$.

5 Conclusion

In this paper, we introduced a generalization of the spherical ridgelet basis for modeling multi-shell diffusion signal. The representation in this basis is sparse and as such allows one to faithfully recover the signal with few gradient directions (around 20-24 per shell). The proposed representation models the radial decay and the frequency components separately.Thus, one can compute a measure of the "overall" signal decay at each voxel. Future work involves examining its

correlation with the underlying tissue properties. This work also showed results on in-vivo data with 3 shells and the same number of directions for each shell. Figuring out the optimum number of shells, the corresponding b-value and the distribution of gradient directions is still a topic of active research. We hope that by recovering the signal from sparse measurements on each shell and with a few (2-3) shells, we will be able to reduce the scan time of MSI significantly and thus make it clinically feasible.

References

1. Tuch, D., Reese, T., Wiegell, M., Wedeen, V.: Diffusion MRI of complex neural architecture. Neuron 40, 885–895 (2003)
2. Wedeen, V.J., Hagmann, P., Tseng, W.Y.I., Reese, T.G., Weisskoff, R.M.: Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. Magnetic Resonance in Medicine 54, 1377–1386 (2005)
3. Landman, B.A., Wan, H., Bogovic, J.A., van Zijl, P.C., Bazin, P.L., Prince, J.L.: Accelerated Compressed Sensing of Diffusion-Inferred Intra-Voxel Structure through Adaptive Refinement. In: ISMRM (2010)
4. Merlet, S., Deriche, R.: Compressed Sensing for Accelerated EAP Recovery in Diffusion MRI. In: CDMRI Workshop, MICCAI, p. 14 (2010)
5. Wu, Y.C., Alexander, A.L.: Hybrid diffusion imaging. NeuroImage 36, 617–629 (2007)
6. Descoteaux, M., Deriche, R., Bihan, D.L., Mangin, J.F., Poupon, C.: Multiple q-shell diffusion propagator imaging. Medical Image Analysis (2010)
7. Jensen, J.H., Helpern, J.A., Ramani, A., Lu, H., Kaczynski, K.: Diffusional kurtosis imaging: The quantification of non-gaussian water diffusion by means of magnetic resonance imaging. Magnetic Resonance in Medicine 53, 1432–1440 (2005)
8. Assemlal, H.-E., Tschumperlé, D., Brun, L.: Efficient computation of PDF-based characteristics from diffusion MR signal. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part II. LNCS, vol. 5242, pp. 70–78. Springer, Heidelberg (2008)
9. Barmpoutis, A., Vemuri, B., Forder, J.: Fast displacement probability profile approximation from hardi using 4th-order tensors. In: ISBI, pp. 911–914 (2008)
10. Ghosh, A., Deriche, R.: Fast and closed-form ensemble-average-propagator approximation from the 4th-order diffusion tensor. In: ISBI, pp. 1105–1108 (2010)
11. Michailovich, O., Rathi, Y.: Fast and accurate reconstruction of HARDI data using compressed sensing. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. 6361, pp. 607–614. Springer, Heidelberg (2010)
12. Donoho, D.L.: Compressed sensing. IEEE Transactions on Information Theory 52, 1289–1306 (2006)
13. Candès, E.J., Romberg, J., Tao, T.: Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. IEEE Transactions on information theory 52, 489–509 (2006)
14. Michailovich, O., Rathi, Y.: On approximation of orientation distributions by means of spherical ridgelets. IEEE Trans. on Image Processing 19, 1–17 (2010)
15. Cohen, Y., Assaf, Y.: High b-value q-space analyzed diffusion-weighted MRS and MRI in neuronal tissues–a technical review. NMR in Biomedicine 15, 516–542 (2002)
16. Ozarslan, E., Koay, C., Basser, P.: Simple harmonic oscillator based estimation and reconstruction for one-dimensional q-space mr. In: ISMRM, vol. 16
17. Asif, M.S., Romberg, J.K.: Dynamic updating for l1 minimization. CoRR abs/0903.1443 (2009)

Longitudinal Tractography with Application to Neuronal Fiber Trajectory Reconstruction in Neonates

Pew-Thian Yap[1], John H. Gilmore[2], Weili Lin[1], and Dinggang Shen[1]

[1] BRIC, Department of Radiology and
[2] Department of Pyschiatry
University of North Carolina at Chapel Hill, NC
{ptyap,jgilmore,weili_lin,dgshen}@med.unc.edu

Abstract. This paper presents a novel tractography algorithm for more accurate reconstruction of fiber trajectories in low SNR diffusion-weighted images, such as neonatal scans. We leverage information from a later-time-point longitudinal scan to obtain more reliable estimates of local fiber orientations. Specifically, we determine the orientation posterior probability at each voxel location by utilizing prior information given by the longitudinal scan, and with the likelihood function formulated based on the Watson distribution. We incorporate this Bayesian model of local orientations into a state-space model for particle-filtering-based probabilistic tracking, catering for the possibility of crossing fibers by modeling multiple orientations per voxel. Regularity of fibers is enforced by encouraging smooth transitions of orientations in subsequent locations traversed by the fiber. Experimental results performed on neonatal scans indicate that fiber reconstruction is significantly improved with less stray fibers and is closer to what one would expect anatomically.

1 Introduction

The human brain is a complex system that is capable of integrating massive amount of information with startling efficiency. A comprehensive description of the architecture of the anatomical connectivity patterns is therefore fundamentally important in cognitive neuroscience and neuropsychology, as it reveals how functional brain states emerge from their underlying structural substrates and provides new mechanistic insights into the association of brain functional deficits with the underlying structural disruption [1].

The neonatal brain provides a window for insights into perhaps the most important phase of human brain development. To this end, diffusion-weighted imaging (DWI) plays an indispensable role in the *in vivo* characterization of brain structural circuity, which relates different functional regions. DWI-based connectivity analysis of neonates, however, is often impeded by the unreliability of neuronal fiber trajectory reconstruction due to the often lower quality of the diffusion-weighted scans. Noisy estimates of local fiber orientations can propagate and accummultate in the course of trajectory reconstruction, especially if

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 66–73, 2011.

the tractography algorithm is greedy in nature, rendering the validity of subsequent tract-based analysis questionable. We present in this paper a remedy to this problem by leveraging prior information from longitudinal scans to improve the accuracy of local fiber orientation estimates for more accurate trajectory reconstruction of the neuronal fibers.

Cook et al. [2] proposed an atlas based approach for better tractography outcome by modifying the stochastic white matter tractography algorithm developed by Friman et al. [3]. An atlas was first generated by computing for each voxel the dyadic tensor of the principal directions from a set of diffusion tensor images. The atlas encapsulates the mean local fiber orientations as well as their degrees of orientation dispersion, serving as prior information to the Bayesian stochastic tractography framework in [3]. This approach, while effective, has two shortcomings: 1) Its formulation is limited to one orientation per voxel, seriously limiting the ability of the algorithm in accounting for fiber crossing, branching or kissing, and 2) The atlas is not subject-specific, causing loss of fiber tracts in regions where the subject and the atlas disagree.

The contributions of this paper are: 1) The formulation and evaluation of a fiber tractography algorithm that is guided by longitudinal prior information for more robust fiber trajectory reconstruction, especially in images with lower SNR, and 2) The evaluation of how modeling multiple local fiber orientations improves tractography outcome in regions with complex diffusion architecture.

2 Approach

2.1 Modeling Local Fiber Orientations

A white matter fiber can be modeled as a finite-length path parameterized by a train of unit length vectors. We use the following notation for such a path: $\mathbf{v}_{(1:T)} = \{\mathbf{v}_{(1)}, \ldots, \mathbf{v}_{(T)}\}$. We further assume that a fiber path can be traced by tracking the trajectory of a particle traveling in an orientation field. Each particle is endowed with an initial speed in an appropriate direction. It then moves with constant speed to position $\mathbf{x}_{(t)}$ according to

$$\mathbf{x}_{(t+1)} = \mathbf{x}_{(t)} + s\mathbf{v}_{(t)} \tag{1}$$

where t is the time index and s is the step length. To reconstruct the fiber trajectories, we need to determine the probability density function (PDF) of the local fiber orientation $f(\mathbf{v}_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}_{(t)}, \boldsymbol{\theta}'_{(t)}, D_{(t)})$. Variable $D_{(t)} = 1, \ldots, \Omega$ is the orientation index; each voxel may contain up to Ω orientations, as in the case of high angular resolution diffusion imaging (HARDI). Sets $\boldsymbol{\theta}_{(t)}$ and $\boldsymbol{\theta}'_{(t)}$ are collections of orientations and their strengths in the neighborhood of $\mathbf{x}_{(t)}$ of the neonatal and longitudinal scans, respectively. Specifically, we define a neighborhood $\mathcal{N}(\mathbf{x}_{(t)})$ (e.g., a $3 \times 3 \times 3$ neighborhood) in the vicinity of $\mathbf{x}_{(t)}$ and, for each voxel i in the neighborhood, collect all corresponding orientations $\mathbf{v}_{(t),i}^{[D_{(t)}]}$ and their strengths $\rho_{(t),i}^{[D_{(t)}]}$ (e.g., magnitudes of the orientation distribution functions at corresponding orientations), spatially weighted by a Gaussian kernel

so that voxels further away from the neighborhood center are deemphasized. The orientations are first sorted to avoid the bias discussed in [4]. Applying Bayes' theorem, we have

$$f(\mathbf{v}_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}_{(t)}, \boldsymbol{\theta}'_{(t)}, D_{(t)}) = \frac{f(\boldsymbol{\theta}_{(t)}|\mathbf{v}_{(t)}, \mathbf{v}_{(t-1)}, \boldsymbol{\theta}'_{(t)}, D_{(t)}) f(\mathbf{v}_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}'_{(t)}, D_{(t)})}{f(\boldsymbol{\theta}_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}'_{(t)}, D_{(t)})}. \tag{2}$$

Since the orientations in both images are estimated independently in a voxel-wise fashion, $f(\boldsymbol{\theta}_{(t)}|\mathbf{v}_{(t)}, \mathbf{v}_{(t-1)}, \boldsymbol{\theta}'_{(t)}, D_{(t)}) = f(\boldsymbol{\theta}_{(t)}|\mathbf{v}_{(t)}, D_{(t)})$ and $f(\boldsymbol{\theta}_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}'_{(t)}, D_{(t)}) = f(\boldsymbol{\theta}_{(t)}|D_{(t)})$. The equation can then be written as

$$f(\mathbf{v}_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}_{(t)}, \boldsymbol{\theta}'_{(t)}, D_{(t)}) = \frac{f(\boldsymbol{\theta}_{(t)}|\mathbf{v}_{(t)}, D_{(t)}) f(\mathbf{v}_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}'_{(t)}, D_{(t)})}{f(\boldsymbol{\theta}_{(t)}|D_{(t)})}. \tag{3}$$

The factor $f(\boldsymbol{\theta}_{(t)}|D_{(t)})$ normalizes the posterior probability function to have a unit volume and can thus be written as the integral of the numerator

$$f(\boldsymbol{\theta}_{(t)}|D_{(t)}) = \int_{\mathbf{v}_{(t)}, \mathbf{v}_{(t-1)}, \boldsymbol{\theta}'_{(t)}} f(\boldsymbol{\theta}_{(t)}|\mathbf{v}_{(t)}, D_{(t)}) f(\mathbf{v}_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}'_{(t)}, D_{(t)}) d\mathbf{v}_{(t)} d\mathbf{v}_{(t-1)} d\boldsymbol{\theta}'_{(t)}. \tag{4}$$

In what follows, we will discuss how the likelihood $f(\boldsymbol{\theta}_{(t)}|\mathbf{v}_{(t)}, D_{(t)})$ and prior $f(\mathbf{v}_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}'_{(t)}, D_{(t)})$ can be computed.

Likelihood: We assume that the orientations observed in the neighborhood of $\mathbf{x}_{(t)}$ can be regarded as noisy observations of $\mathbf{v}_{(t)}$, characterized by the Watson distribution with probability density function (PDF) [5]

$$f(\mathbf{w}|\boldsymbol{\mu}, \kappa) = C(\kappa) \mathrm{e}^{\kappa(\mathbf{w}^{\mathrm{T}}\boldsymbol{\mu})^2}. \tag{5}$$

The parameter $\boldsymbol{\mu}$ is a unit vector called the *mean orientation* and κ is a positive constant called the *concentration parameter*. The squared exponential in (5) ensures that the distribution is antipodal symmetric. The density has maxima at $\pm\boldsymbol{\mu}$ and becomes more concentrated around $\pm\boldsymbol{\mu}$ as κ increases. The density is also rotationally invariant around $\pm\boldsymbol{\mu}$. $C(\kappa)$ is a normalizing constant to ensure that the density function integrates to unity over the unit sphere. By letting $\boldsymbol{\mu} = \mathbf{v}_{(t)}$, the joint distribution, or the likelihood, of the observed orientations $\boldsymbol{\theta}_{(t)}$ can be written as

$$f(\boldsymbol{\theta}_{(t)}|\mathbf{v}_{(t)}, D_{(t)}) = \prod_{(\mathbf{w},\rho)\in\boldsymbol{\theta}_{(t)}^{[D_{(t)}]}} C(\kappa) e^{\rho\kappa(\mathbf{w}^{\mathrm{T}}\mathbf{v}_{(t)})^2} \tag{6}$$

where $\boldsymbol{\theta}_{(t)}^{[D_{(t)}]}$ is a subset of $\boldsymbol{\theta}_{(t)}$ consisting only of the group of orientations specified by $D_{(t)}$. This equation is substituted into (3) to obtain the posterior distribution.

Priors: Via the probability function $f(\mathbf{v}_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}'_{(t)}, D_{(t)})$, we encode our prior knowledge about fiber regularity and about orientation information from the longitudinal scan. We define the prior probability function as

$$f(\mathbf{v}_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}'_{(t)}, D_{(t)}) \propto \begin{cases} \left[\mathbf{v}_{(t)}^{\mathrm{T}}\mathbf{v}_{(t-1)}\right]^2 \left[C\left(\kappa'^{[D_{(t)}]}_{(t)}\right) \mathrm{e}^{\kappa'^{[D_{(t)}]}_{(t)}(\mathbf{v}_{(t)}^{\mathrm{T}}\boldsymbol{\mu}'^{[D_{(t)}]}_{(t)})^2}\right], & \mathbf{v}_{(t)}^{\mathrm{T}}\mathbf{v}_{(t-1)} > 0 \\ 0, & \text{otherwise.} \end{cases} \tag{7}$$

The first term on the right enforces the regularity constraint during fiber reconstruction. The second term transfuses orientation information from the longitudinal scan. The Watson distribution is used to represent the distribution of the orientations in $\boldsymbol{\theta}'^{[D_{(t)}]}_{(t)}$, with mean orientation $\boldsymbol{\mu}'^{[D_{(t)}]}_{(t)}$ and concentration parameter $\kappa'^{[D_{(t)}]}_{(t)}$. Maximum likelihood estimates of the these parameters can be obtained using the method described in [5].

2.2 State-Space Model

Tractography is assumed to be a stochastic process that can be represented using a state-space model with the local fiber orientation $\mathbf{v}_{(t)}$ as the observation, and the orientation index $D_{(t)}$ as the state. Given the previous orientation $\mathbf{v}_{(t-1)}$ and the current state $D_{(t)}$, the observation probability is defined according to (3). Noting that the posterior probability for selection of a particular value of the orientation index $D_{(t)}$ is

$$f(D_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}_{(t)}, \boldsymbol{\theta}'_{(t)}) = \frac{f(\boldsymbol{\theta}_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}'_{(t)}, D_{(t)}) f(D_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}'_{(t)})}{f(\boldsymbol{\theta}_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}'_{(t)})} \tag{8}$$

and by letting

$$f(\boldsymbol{\theta}_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}'_{(t)}, D_{(t)}) \propto \bar{\rho}^{[D_{(t)}]}_{(t)}, \quad f(D_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}'_{(t)}) \propto \left(\mathbf{v}_{(t-1)}^{\mathrm{T}}\boldsymbol{\mu}'^{[D_{(t)}]}_{(t)}\right)^2, \tag{9}$$

the transition probability can be defined as

$$f(D_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}_{(t)}, \boldsymbol{\theta}'_{(t)}) \propto \bar{\rho}^{[D_{(t)}]}_{(t)} \left(\mathbf{v}_{(t-1)}^{\mathrm{T}}\boldsymbol{\mu}'^{[D_{(t)}]}_{(t)}\right)^2. \tag{10}$$

where $\bar{\rho}^{[D_{(t)}]}_{(t)}$ denotes the average strength of orientations in $\boldsymbol{\theta}^{[D_{(t)}]}_{(t)}$. For sequential sampling of fiber paths, we need to draw random samples of local fiber orientations from the observation probability (3) and transition probability (10). For drawing samples from complicated and high dimensional PDFs, one can always resort to Markov Chain Monte Carlo (MCMC) techniques. The probability of a path of a given length T is

$$f(\mathbf{v}_{(1:T)}|\boldsymbol{\theta}_{(1:T)}, \boldsymbol{\theta}'_{(1:T)}) = f(D_{(1)}) f(\mathbf{v}_{(1)}|\mathbf{v}_{(1)}, \boldsymbol{\theta}_{(1)}, \boldsymbol{\theta}'_{(1)}, D_{(1)}) \times \Pi_{t=2}^{T} f(D_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}_{(t)}, \boldsymbol{\theta}'_{(t)}) f(\mathbf{v}_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}_{(t)}, \boldsymbol{\theta}'_{(t)}, D_{(t)}). \tag{11}$$

Tracking is stopped if the trajectory reaches a voxel with *orientation coherence*

$$\beta^{[D_{(t)}]} = 1 - \frac{1}{2}\left\{\sqrt{\frac{\lambda_2^{[D_{(t)}]} + \lambda_3^{[D_{(t)}]}}{2\lambda_1^{[D_{(t)}]}}} + \sqrt{\frac{\lambda'^{[D_{(t)}]}_2 + \lambda'^{[D_{(t)}]}_3}{2\lambda'^{[D_{(t)}]}_1}}\right\}, \quad \beta^{[D_{(t)}]} \in [0,1], \tag{12}$$

falling below a predefined threshold β_0, or simply when the brain boundary is encountered. The λ's are the eigenvalues of the dyadic tensors computed from $\boldsymbol{\theta}_{(t)}^{[D_{(t)}]}$ and $\boldsymbol{\theta}'^{[D_{(t)}]}_{(t)}$ [2,4]. Perfect alignment of the orientations indicated by $D_{(t)}$ results in $\beta^{[D_{(t)}]} = 1$ and an uniform distribution of orientations results in $\beta^{[D_{(t)}]} = 0$.

3 Results

3.1 Materials

Diffusion-weighted images of 10 infants were acquired at two time points: one month and one year after birth. Diffusion gradients were applied in 42 non-collinear directions with diffusion weighting $b = 1000$ s/mm^2, repetition time (TR) = 7,680 ms and echo time (TE) = 82 ms. The scans covered the whole brain with a resolution of $2 \times 2 \times 2$ mm^3. Data post-processing includes brain skull removal, motion correction and eddy current correction using algorithms developed and distributed as part of the FMRIB Software Library (FSL) package. Each neonatal scan was co-registered with their respective longitudinal scan so that they reside in a common space.

3.2 Tractography

To evaluate the effectiveness of the proposed method, we performed tractography based on seeds placed at the points where the midline crosses the splenium of the corpus callosum. 3000 trajectories were initiated from each seed point. The step size was fixed at 1 mm. The maximum allowable orientation coherence was set to 0.1. Various configurations were used:

(a) Tractography on the neonatal scan alone with no prior information from the longitudinal scan.
(b) Tractography on the neonatal scan guided by prior information from the longitudinal scan.
(c) Tractograph on the longitudinal scan.
(d) Tractography on the neonatal scan, using only the first principal orientation with no prior information.
(e) Tractography on the neonatal scan guided by longitudinal prior information from the longitudinal scan, using only the first principal orientations.
(f) Tractography on the longitudinal scan using only the first principal orientations.

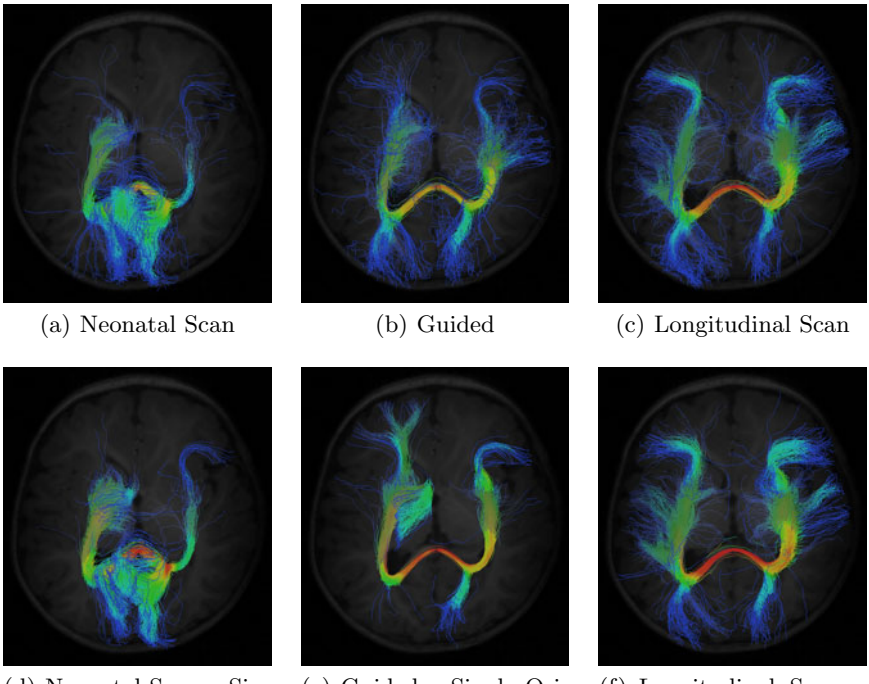

(a) Neonatal Scan (b) Guided (c) Longitudinal Scan

(d) Neonatal Scan – Single Orientation (e) Guided – Single Orientation (f) Longitudinal Scan – Single Orientation

Fig. 1. Comparison of different tractography schemes. (a) Neonatal scan alone, (b) Neonatal scan + longitudinal scan, (c) Longitudinal scan alone, (d) Neonatal scan alone (single orientation), (e) Neonatal scan + longitudinal scan (single orientation), and (f) Longitudinal scan alone (single orientation).

The respective results, shown in Fig. 1, indicate that the proposed method (Fig. 1(b)) gives reasonable results. Tractography, when performed based on orientation information given by the neonatal scan alone, results in noisy fiber tracts (Fig. 1(a) and Fig. 1(d)). This is not surprising since neonatal scans typically suffer from lower SNR. The results were improved remarkably by employing prior information given by the longitudinal scan (Fig. 1(c) and Fig. 1(f)). The reconstructed trajectories are generally cleaner with less stray fibers and are in higher agreement with our anatomical understanding of the fibers. Note that by allowing only one orientation per voxel, major fibers are lost. Fig. 1(e) shows that, despite with longitudinal guidance, fibers connecting one of the occipital lobes cannot be correctly preserved and reconstructed.

3.3 Single- and Multi-orientation Schemes

For quantification of the effect of single-orientation and multi-orientation schemes on the reconstructed trajectories, various measures were employed. First, we

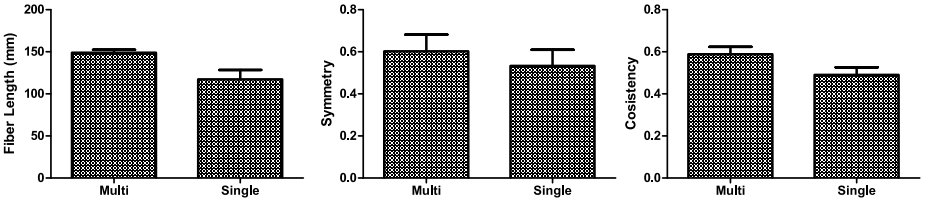

(a) Fiber Length (b) Symmetry (c) Consistency

Fig. 2. Comparison between the multi- and single-orientation schemes using different statistics. Each bar indicates the mean value, and the error bar indicates the corresponding standard error.

compared the average length of the reconstructed trajectories. A longer length average generally indicates that fiber reconstruction is less likely to be terminated prematurely, possibly due to noisy local fiber orientations. Fig. 2(a) shows that the proposed method results in on average longer trajectories, hinting that by allowing more than one orientation per voxel results in less chances of premature termination at fiber crossings.

Anatomical scans (e.g., T1-weighted and T2-weighted images) of the brain indicate that the brain is mostly symmetrical between the hemispheres. Therefore, one would expect that, to be anatomically sound, the reconstructed trajectories should exhibit some degree of symmetry. For our purpose, we evaluated the degree of symmetry of the reconstructed trajectories by computing the degree of correlation of the connectivity matrices derived from the fibers between the left and right hemispheres. Specifically, this was done by parcellating the brain into 116 regions according to the Automated Anatomical Labelling (AAL) atlas [6], computing the connectivity matrix for each hemisphere based only on the $(116/2 = 58)$ ROIs in that particular hemisphere and then computing the normalized scalar product of the connectivity matrices of both hemispheres. Each element of the (58×58) connectivity matrix records the number of fibers connecting a pair of ROIs. The results, shown in Fig. 2(b), indicates the multi-orientation scheme yields higher inter-hemispheric consistency compared with the single-orientation scheme, again validating that modeling complex orientation information can be conducive to more accurate reconstruction of fiber trajectories.

Major fiber bundles known to exist in all subjects can be used to test fiber tracking consistency across subjects. Since the seeds used in our evaluation were placed on a major white major structure, we expect the reconstructed trajectories to be relatively consistent across subjects. We tested this by using a similar methodology described in the previous paragraph. But instead of generating the connectivity matrix for each hemisphere separately, we generated a whole brain connectivitiy matrix by considering both hemispheres at the same time. We then computed the mean connectivity matrix over all the 10 subjects. The connectivity matrix of each subject was then compared with this mean matrix (via normalized scalar product) as an indicator of consistency. The higher the

similarity of all individual connectivity matrices with the mean matrix, the greater is the consistency. Fig. 2(c) indicates that greater consistency can be achieved using the multi-orientation scheme, again validating the effectivenes of the proposed method.

4 Conclusion

We have presented a tractography algorihtm that caters especially for low SNR diffusion-weighted images by leveraging prior information from the respective longitudinal scans. Experimental results performed using neonatal scans indicate that the proposed method yields fiber trajectories that are more consistent with our anatomical knowledge. The ability of taking into account multi-orientation information gives further improvement over previous methods by allowing more accurate modeling of complex white matter architecture involving crossing fibers.

Acknowledgment. This work was supported in part by NIH grants: EB006733, EB008374, EB009634, MH088520, HD05300, MH064065, and NS055754.

References

1. Sporns, O., Tononi, G., Kötter, R.: The human connectome: A structural description of the human brains. PLoS Computational Biology 1(4), e42 (2005)
2. Cook, P.A., Zhang, H., Awate, S.P., Gee, J.C.: Atlas-guided probabilistic diffusion-tensor fiber tractography. In: IEEE International Symposium on Biomedical Imaging (ISBI 2008), pp. 951–954 (2008)
3. Friman, O., Farnebäck, G., Westin, C.F.: A bayesian approach for stochastic white matter tractography. IEEE Transactions on Medical Imaging 25, 965–977 (2006)
4. Basser, P.J., Pajevic, S.: Statistical artifacts in diffusion tensor MRI (DT-MRI) caused by background noise. Magnetic Resonance in Medicine 44, 41–50 (2000)
5. Schwartzman, A., Dougherty, R.F., Taylor, J.E.: False discovery rate analysis of brain diffusion direction maps. Annals of Applied Statistics 2(1), 153–175 (2008)
6. Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., Joliot, M.: Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. Neuroimage 15, 273–289 (2002)

Quantitative Body DW-MRI Biomarkers Uncertainty Estimation Using Unscented Wild-Bootstrap*

M. Freiman[1], S.D. Voss[2], R.V. Mulkern[2], J.M. Perez-Rossello[2], and S.K. Warfield[1]

[1] Computational Radiology Laboratory, Childrens Hospital, Harvard Medical School, Boston, USA
[2] Department of Radiology, Children's Hospital, Harvard Medical School, Boston, USA

Abstract. We present a new method for the uncertainty estimation of diffusion parameters for quantitative body DW-MRI assessment. Diffusion parameters uncertainty estimation from DW-MRI is necessary for clinical applications that use these parameters to assess pathology. However, uncertainty estimation using traditional techniques requires repeated acquisitions, which is undesirable in routine clinical use. Model-based bootstrap techniques, for example, assume an underlying linear model for residuals rescaling and cannot be utilized directly for body diffusion parameters uncertainty estimation due to the non-linearity of the body diffusion model. To offset this limitation, our method uses the Unscented transform to compute the residuals rescaling parameters from the non-linear body diffusion model, and then applies the wild-bootstrap method to infer the body diffusion parameters uncertainty. Validation through phantom and human subject experiments shows that our method identify the regions with higher uncertainty in body DWI-MRI model parameters correctly with realtive error of ∼36% in the uncertainty values.

1 Introduction

Diffusion Weighted MRI (DW-MRI) is a non-invasive imaging technique that is rapidly evolving as a surrogate imaging biomarker for extra-cranial applications, including organ functionality evaluation [8,17], oncological disease classification [9], and early assessment of body tumor response to therapy [5,6].

DW-MRI is particularly appealing for pediatric patients undergoing multiple, repeat imaging examinations because a) it does not involve ionizing radiation associated with traditional imaging methods such as CT, MIBG, or FDG-PET; and b) it is performed without administration of intravenous contrast agents, which, in some patients, may be contraindicated [14].

* This investigation was supported in part by NIH grants R01 RR021885, R01 EB008015, R03 EB008680 and R01 LM010033.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 74–81, 2011.

In-vivo measurement of the body tissue diffusivity involves acquisition of a sequence of DW-MRI with multiple b-values, and estimation of the Apparent Diffusion Coefficient (ADC) from the sequence. Traditionally, the ADC value is computed for each voxel using the Stejskal-Tanner exponential model: $S_b = S_0 e^{-b*ADC}$ [13]. However, in living tissue, physiological motions unrelated to diffusion (e.g., the intra-voxel micro-capillary perfusion effect) can mimic diffusion processes and confound *in-vivo* measurements.Hence, accurate estimation of the ADC values requires fitting a multi-exponential model that describes both diffusion and perfusion effects [12,1].

Clinical applications of DW-MRI typically involve the comparison of mean or median ADC values in particular, Regions of Interest (ROI) [8]; or may use voxel-wise analysis [10] to visualize the change in the ADC over time. Estimation of ADC from multiple b-value images has an intrinsic uncertainty, however, due in part, to noise in the image acquisition process; and in part, to artifacts such as those caused by patient motion. The uncertainty that characterizes ADC value estimation is therefore needed to achieve quantitative ADC assessment in clinical applications.

Direct estimation of the ADC measurement uncertainty using traditional asymptotic estimators or the bootstrap technique [7] requires repeated acquisitions, however. These techniques are therefore impracticable in routine clinical use due to lengthy scan times per patient, and to resulting increases in artifacts. Moreover, model-based bootstrap techniques require homogeneous leverages of the different sample points. Although it is possible to rescale the heterogeneous leverages of different points using an underlying log-linear model derived from the Stejskal-Tanner exponential diffusion model in intra-cranial Diffusion Tensor MRI (DT-MRI) [3]; body DW-MRI exhibits a multi-exponential model [12,1] that cannot be linearized through a log transform. As a result, model-based bootstrap techniques cannot be utilized directly for the uncertainty estimation of body DW-MRI ADC.

We present a new model-based bootstrap technique to infer the measurement uncertainty in ADC estimation from body DW-MRI *without* acquiring additional images, thus eliminating the underlying linear model requirement. First, a multi-exponential model is fitted to the acquired data. Next, the residuals between the observed data and the estimated model are computed. Heterogeneous leverages scaling parameters are computed using the Unscented transform [11], and the residuals are scaled accordingly. Last, the wild-bootstrap resampling technique [4] is used to estimate the ADC distribution from which the model uncertainty can be derived. The estimated uncertainty values can then be used for further analysis of the change in the ADC values in response to the therapy. Our Unscented wild-bootstrap technique does not assume any linearity in the underlying model and thus can be utilized directly in other cases exhibiting underlying non-linear models, such as crossing fibers in DT-MRI.

We compared the uncertainty estimation using our method to that which was obtained through repeated acquisitions of both phantom and human subject experiments. In all experiments, our method identified the regions with higher

uncertainty in body DWI-MRI model parameters correctly with realtive error of ∼36% in the uncertainty values. Thus, our method can provide the necessary information for reliable quantitative DW-MRI biomarkers assessment, without repeated acquisitions.

2 Method

The goal is to infer the uncertainty in ADC estimation from DW-MRI. Our proposed approach utilizes a body DW-MRI dataset consisting of six or more images acquired with different b-values in the range of 0-800 s/mm^2. First, a multi-exponential diffusion model is fitted to the data and the raw residuals between the obtained fit and the observed images are computed. Next, the Unscented Transform (UT) [11] is used to compute the scaling parameters required to account for heterogeneous leverages in the different sampling points, and the residuals are scaled accordingly to satisfy the wild-bootstrap method [4] conditions. Last, wild-bootstrap resampling is used to infer the uncertainty in the diffusion model parameters estimation. In the next section, we describe each step in detail.

2.1 Diffusion Model Fitting

The DW-MRI measurement of tissue diffusivity is affected by both thermal-driven water molecules diffusion and intra-voxel, randomly oriented, microcapillary blood flow. Thus, the DW-MRI measurements should be modeled using a multi-exponential diffusion model [12,1]:

$$S(b) = S_0 \left(f \cdot exp(-b \cdot PER) + (1-f) \cdot exp(-b \cdot (ADC + PER))\right) \quad (1)$$

where b is the b-value used to acquire the images, ADC is the Apparent Diffusion Coefficient (ADC) describing the signal attenuation due to the real diffusion effect; PER is the perfusion coefficient describing the signal attenuation due to the blood perfusion effect; and f is the fractional volume of the perfusion effect.

Assuming a Rician noise model with scaling parameter σ_R, the Probability Distribution Function (PDF) of the DW-MRI signal is given by:

$$P(M|S,\sigma_R) = \frac{M}{\sigma_R^2} exp\left(-\frac{M^2+S^2}{2\sigma_R^2}\right) I_0\left(\frac{MS}{\sigma_R^2}\right) \quad (2)$$

where M is the observed value; S is the "true" underlying signal; and $I_0(\cdot)$ is the zero-ordered modified Bessel function of the first kind.

The likelihood of the diffusion model parameters $\Omega = \{ADC, PER, f, S_0\}$ given the observed data M and σ_R is:

$$L(\Omega; M, \sigma_R) = \prod_{i=1}^{N} P(M_i|\Omega, \sigma_R) \quad (3)$$

where N is the number of the observations.

The Maximum Likelihood (ML) estimation of the model parameters Ω is computed by maximizing the following log-likelihood function:

$$\widehat{\Omega} = \underset{\Omega}{\operatorname{argmax}} \sum_{i=1}^{N} I_0 \left(\frac{S(b_i;\Omega) M_i}{\sigma_R^2} \right) - \sum_{i=1}^{N} \frac{S(b_i;\Omega)^2}{2\sigma_R^2} \tag{4}$$

where $S(b_i;\Omega)$ is the signal estimation at b-value b_i given the model parameters Ω computed using Eq. 1. The estimation of σ_R is computed by ML fitting of a Rayleigh distribution to a pre-defined background region on each image.

Initial estimation of the model parameters is obtained as in [2], and the maximization is done using the BOBYQA non linear optimization algorithm [15]. The raw residuals between the obtained fit and the observed images are then given by: $\epsilon_i = M_i - S(b_i)$

2.2 Heterogeneous Leverages Scaling

The wild-bootstrap resampling technique assumes homogeneous leverage of the different sampling points. However, the estimation of model fitting uncertainty usually involves heterogeneous leverage. Appropriate rescaling of the raw residuals ϵ_i is required to satisfy this condition. In linear models, these rescaling parameters can be obtained directly from the model equation as in [3,4]. However, there is no direct way to compute the heterogeneous leverage of different points in the non-linear model underlying body DW-MRI.

To offset this limitation, we use the Unscented Transform (UT) [11] to estimate the heterogeneous leverage of different points by propagating a set normally distributed test points with $\mu = \widehat{\Omega}$ and $\sigma = 1$ throughout the non-linear model, and computing the variance of the propagated points.

Given the multi-exponential diffusion model (Eq. 1) and the estimated parameters Ω with n as the model parameters number; we define a set of $2n$ test points p as in [3]. A new set of model parameters with $\mu = \widehat{\Omega}$ and $\sigma = 1$ is then generated by adding each test point p_j to the estimated model parameters:

$$\widetilde{\Omega_j} = \widehat{\Omega} + p_j \tag{5}$$

Next, the DW-MRI signal values $S(b_i;\widetilde{\Omega_j})$ are computed for each b_i with the new model parameters $\widetilde{\Omega_j}$ using Eq. 1, and the variance of the propagated errors is given for each sampling point i by:

$$h_i = \frac{1}{2n} \sum_{j=1}^{2n} \left(S(b_i;\widetilde{\Omega_j}) - \overline{S(b_i;\widetilde{\Omega_j})} \right)^2 \tag{6}$$

Finally the scaled residuals $\widehat{\epsilon_i}$ are computed from the raw residuals ϵi by:

$$\widehat{\epsilon_i} = \frac{\epsilon_i}{\sqrt{1 - h_i^2}} \tag{7}$$

2.3 Wild-Bootstrap Estimation of the Model Uncertainty

Following the method described in [4], the wild-bootstrap resampling is defined as:

$$S^*(b_i;\widehat{\Omega})_k = S(b_i;\widehat{\Omega}) + t_k\widehat{\epsilon_i} \tag{8}$$

where $S^*(b_i;\widehat{\Omega})_k$ is the resampled measure at the b-value b_i and t_k is a two-point Rademacher distributed random variable with $P(t=1)=0.5$ and $P(t=-1)=0.5$ defined for each b-value separately.

The wild-bootstrap resamples are generated for the entire set of b_i values and a multi-exponential diffusion model is fitted using Eq. 4. Resampling and model fitting are repeated for a large number of fixed repetitions to obtain a large set of diffusion models. The estimation of the model uncertainty is then obtained by computing the standard deviation of the estimated model parameters from the set of bootstrap diffusion models.

3 Experimental Results

We evaluated the performance of our Unscented wild-bootstrap uncertainty estimation method through DW-MRI scans of an American College of Radiology (ACR) MRI phantom and a human subject. In the following section, we describe each experiment in detail.

3.1 ACR Phantom Experiment

Ten DW-MRI datasets of an ACR MRI phantom were acquired using the gradient encoding scheme [16] on a 1.5T MRI machine (Magnetom Avanto, Siemens Medical Solutions, Erlangen, Germany). Each dataset consisted of diffusion images with values $= [5, 50, 100, 200, 270, 400, 600, 800]$s/mm^2. One out of the ten datasets was acquired during slight vibrations of the phantom to introduce motion artifacts and enlarge the model fit uncertainty. The additional 9 datasets were acquired with the stabilized phantom. A multi-exponential diffusion model was fitted to each dataset using Eq. 4. The performance of our method was evaluated using the Leave-One-Out (LOO) methodology, where the "ground-truth" uncertainty in model estimation was defined as the standard deviation of the ADC measure at each voxel over 8 out of the 9 stabilized acquisitions, and the ADC uncertainty was estimated from the additional dataset solely using our technique with a varying number of bootstrap resampling iterations between 50 and 1000. Finally, the Relative error between the estimated uncertainty and the "ground-truth" uncertainty was computed for each dataset. In addition, we compared the ADC uncertainty estimated from the vibrated phantom dataset to a "ground-truth" uncertainty obtained from the 9 stabilized phantom datasets.

Fig. 1 presents the DW-MRI image used in the experiment, the "ground-truth" uncertainty map, the estimated uncertainty map computed from one of the stabilized phantom datasets and the estimated uncertainty computed

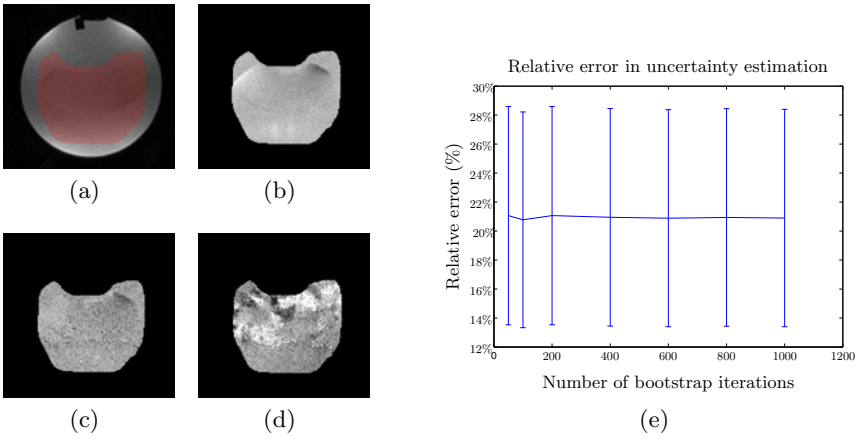

Fig. 1. ACR phantom experiment: (a) DW-MRI image of the ACR phantom with b-value=50s/mm^2; pixels used for evaluation are marked in red. (b) The "ground-truth" ADC uncertainty. (c-d) Estimated ADC uncertainty using our method from a (c) stabilized phantom dataset and (d) vibrated phantom dataset. (e) RMS error between our estimation and the "ground-truth" using a varying number of bootstrap resampling iterations.

from the vibrated phantom dataset. As expected, our method estimated higher uncertainty in the vibrated phantom dataset (41%) compared to the stabilized phantom datasets (∼21%).

Fig. 1e presents the relative error between the estimated and the ground-truth uncertainty maps. Our method estimated the uncertainty with an average (std) error of 21% (6.5%) of the "ground-truth" value. Increasing the number of bootstrap resampling iterations, however, did not yield any significant improvement in the uncertainty estimation accuracy.

3.2 Human Subject Experiment

Five (5) body DW-MRI datasets of a healthy human subject were acquired as in the ACR phantom experiment (Sec. 3.1) during free-breathing of the subject. A multi-exponential diffusion model was fitted to each dataset using Eq. 4. The performance of our method was evaluated using same methodology as in the ACR phantom experiment (Sec. 3.1).

Fig. 2 presents the DW-MRI image acquired with b-value=200s/mm^2, the ADC map computed using one of the datasets, the "ground-truth" uncertainty map, and the estimated uncertainty map. Our method identify the organ boundaries as the regions with the highest uncertainty, due to the respiratory motion effect on these regions. The uncertainty in the liver is higher than that in the kidneys and the spleen due to the larger effect of respiratory motion on the liver.

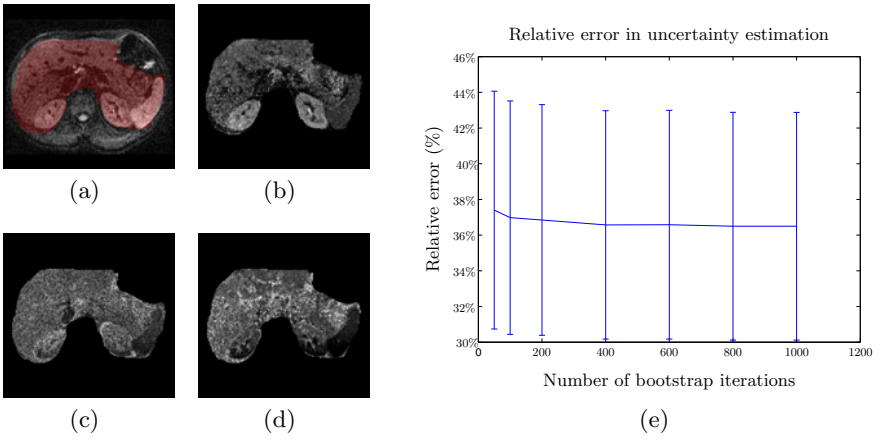

Fig. 2. Human subject experiment: (a) Body DW-MRI image with b-value=200s/mm^2; pixels used for evaluation are marked in red. (b) ADC map obtained by fitting the multi-exponential model to one of the datasets. (c) The "ground-truth" ADC uncertainty. (d) Estimated ADC uncertainty using our method from one of the datasets. (e) Relative error between our estimation and the "ground-truth" using a varying number of bootstrap resampling iterations.

Fig. 2e presents the relative error between the estimated uncertainty and the ground-truth uncertainty. Our method estimated the uncertainty accurately with an average (std) error of 36.8% (6.2%) of the ground-truth value. Increasing the number of bootstrap resampling iterations did not, however, yield significant improvement in the uncertainty estimation accuracy. Similar relative error was obtained for the estimation of the uncertainty in the perfusion comparatment (PER). Higher uncertainty values and larger estimation errors were observed in the human imaging study; these numbers are primarily driven by motion artifact (i.e., breathing) that typically takes place during any acquisition with DW-MRI.

4 Conclusions

We have presented a new Unscented wild-bootstrap method for the estimation of diffusion model parameters uncertainty in body MR-DWI imaging. The estimation of the model parameters uncertainty is necessary to achieve quantitative clinical assessment of body DW-MRI biomarkers. The main advantage of our method is two-fold: First, it does not require repeated acquisitions, which is undesirable in routine clinical applications; and second, it is not limited to underlying linear models that are less suitable in body DW-MRI. Evaluation of our method against uncertainty estimations obtained by repeated acquisitions of both phantom and human subject DW-MRI data shows that our method can estimate the uncertainty associated with a multi-exponential body diffusion model fit accurately without repeated acquisitions.

References

1. Bihan, D.L., Breton, E., Lallemand, D., Aubin, M.L., Vignaud, J., Laval-Jeantet, M.: Separation of diffusion and perfusion in intravoxel incoherent motion MR imaging. Radiology 168(2), 497–505 (1988)
2. Chandarana, H., Lee, V., Hecht, E., Taouli, B., Sigmund, E.: Comparison of Biexponential and Monoexponential Model of Diffusion Weighted Imaging in Evaluation of Renal Lesions: Preliminary Experience. Invest Radiol. 46(5), 285–291 (2010)
3. Chung, S., Lu, Y., Henry, R.G.: Comparison of bootstrap approaches for estimation of uncertainties of DTI parameters. NeuroImage 33(2), 531–541 (2006)
4. Davidson, R., Flachaire, E.: The wild bootstrap, tamed at last. J. of Econometrics 146(1), 162–169 (2008)
5. Dudeck, O., Zeile, M., Pink, D., Pech, M., Tunn, P., Reichardt, P., Ludwig, W., Hamm, B.: Diffusion-weighted magnetic resonance imaging allows monitoring of anticancer treatment effects in patients with soft-tissue sarcomas. J. Magn. Reson. Imaging 27(5), 1109–1113 (2008)
6. Eccles, C., Haider, E., Haider, M., Fung, S., Lockwood, G., Dawson, L.: Change in diffusion weighted mri during liver cancer radiotherapy: preliminary observations. Acta Oncol. 48(7), 1034–1043 (2009)
7. Efron, B., Tibshirani, R.J.: An Introduction to the Bootstrap. Chapman & Hall, New York (1993)
8. Fujimoto, K., Tonan, T., Azuma, S., Kage, M., Nakashima, O., Johkoh, T., Hayabuchi, N., Okuda, K., Kawaguchi, T., Sata, M., Qayyum, A.: Evaluation of the mean and entropy of apparent diffusion coefficient values in chronic hepatitis C: correlation with pathologic fibrosis stage and inflammatory activity grade. Radiology 258(3), 739–748 (2011)
9. Gahr, N., Darge, K., Hahn, G., Kreher, B., von Buiren, M., Uhl, M.: Diffusion-weighted MRI for differentiation of neuroblastoma and ganglioneuroblastoma/ganglioneuroma. Eur. J. Radiol (2010) (in press)
10. Galbán, C., Chenevert, T., Meyer, C., Tsien, C., Lawrence, T., Hamstra, D., Junck, L., Sundgren, P., Johnson, T., Ross, D., Rehemtulla, A., Ross, B.: The parametric response map is an imaging biomarker for early cancer treatment outcome. Nat. Med. 15(5), 572–576 (2009)
11. Julier, S., Uhlmann, J.: Unscented filtering and nonlinear estimation. Proc. of the IEEE 92(3), 401–422 (2004)
12. Koh, D.M., Collins, D.J., Orton, M.R.: Intravoxel incoherent motion in body diffusion-weighted mri: Reality and challenges. AJR Am. J. Roentgenol. 196(6), 1351–1361 (2011)
13. Koh, D.M., Thoeny, H.C., Chenevert, T.L.: Principles of Diffusion-Weighted Imaging (DW-MRI) as Applied to Body Imaging. In: Diffusion-Weighted MR Imaging. In: Medical Radiology. Springer, Heidelberg (2010)
14. Perazella, M.A.: Current status of gadolinium toxicity in patients with kidney disease. Clin. J. Am. Soc. Nephrol. 4(2), 461–469 (2009)
15. Powell, M.: The BOBYQA algorithm for bound constrained optimization without derivatives. technical report NA2009/06, Dep. App. Math. and Th. Physics, Cambridge, England (2009)
16. Scherrer, B., Warfield, S.: Toward an accurate multi-fiber assessment strategy for clinical practice. In: ISBI 2011 (2011)
17. Xu, Y., Wang, X., Jiang, X.: Relationship between the renal apparent diffusion coefficient and glomerular filtration rate: preliminary experience. J. Magn. Reson. Imaging 26(3), 678–681 (2007)

Axon Diameter Mapping in Crossing Fibers with Diffusion MRI

Hui Zhang[1], Tim B. Dyrby[2], and Daniel C. Alexander[1]

[1] Microstructure Imaging Group, Department of Computer Science, University College London, London WC1E 6BT, United Kingdom
[2] Danish Research Center for Magnetic Resonance, Copenhagen University Hospital, Hvidovre, Denmark

Abstract. This paper proposes a technique for a previously unaddressed problem, namely, mapping axon diameter in crossing fiber regions, using diffusion MRI. Direct measurement of tissue microstructure of this kind using diffusion MRI offers a new class of biomarkers that give more specific information about tissue than measures derived from diffusion tensor imaging. Most existing techniques for axon diameter mapping assume a single axon orientation in the tissue model, which limits their application to only the most coherently oriented brain white matter, such as the corpus callosum, where the single orientation assumption is a reasonable one. However, fiber crossings and other complex configurations are widespread in the brain. In such areas, the existing techniques will fail to provide useful axon diameter indices for any of the individual fiber populations. We propose a novel crossing fiber tissue model to enable axon diameter mapping in voxels with crossing fibers. We show in simulation that the technique can provide robust axon diameter estimates in a two-fiber crossing with the crossing angle as small as 45^o. Using *ex vivo* imaging data, we further demonstrate the feasibility of the technique by establishing reasonable axon diameter indices in the crossing region at the interface of the cingulum and the corpus callosum.

1 Introduction

Axon diameter mapping using diffusion MRI offers exciting new possibilities for investigating white matter in health and disease beyond diffusion-tensor imaging. Information about axon diameter and density informs the role and performance of white matter pathways [1,2]. Specific changes in axon diameter have also been linked to diseases such as multiple sclerosis [3] and amyotrophic lateral sclerosis [4]. Direct measurement of such features can therefore shed new light into white matter development and disease mechanisms.

Most techniques for axon diameter mapping adopt the model-based strategy in which a geometric model of the tissue predicts the MR signal from water diffusing within. Earlier methods [5,6] assume a single and known orientation of axons in the tissue model, which limits their application to nerve tissue samples or small regions of brain with uniform orientation. More recently, estimating

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 82–89, 2011.

axon diameters of unknown orientation on clinical scanners has been shown to be feasible, first in simulation [7], and later in live human brains [8]. However, the models in [7,8] still assume a single albeit unknown fiber orientation. Most recently, Zhang et al [9] relax this assumption to some extent by modeling the axonal-orientation distribution as a Watson distribution. The Watson model accommodates the presence of axonal-orientation dispersion and extends axon diameter mapping beyond the most coherent white matter regions like the corpus callosum to a much wider subset of the white matter. Nevertheless, its model of axonal-orientation distribution remains unimodal. Like the earlier models, it is inappropriate in crossings or partial volume between tracts with significantly different orientations.

Fiber crossings occur in many areas of the brain, resulting in the observation of two or more distinct fiber populations in a significant number of voxels. In such voxels, considerable success has been achieved in resolving the precise underlying crossing fiber configurations. Since the earliest breakthroughs in crossing fiber resolution [10,11], major progress has produced very effective modern techniques (See [12] for a review). Despite that, there has been no attempt at direct measurement of additional microstructure features in these locations. A solution to this problem is prerequisite to realizing whole-brain axon diameter mapping.

This paper describes a technique that addresses the combined problem of crossing fiber resolution and microstructure imaging. We propose a crossing fiber tissue model that enables the simultaneous estimation of crossing configurations and microstructure features. We demonstrate the technique both in simulation and in brain data. The rest of the paper is organized as follows: Section 2 describes the proposed crossing-fiber tissue model and the data fitting procedure; Section 3 details the design of the simulation and brain data experiments for validating the proposed technique and reports the findings; Section 4 summarizes the contribution and discusses future work.

2 Crossing-Fiber Tissue Model

The proposed model generalizes the minimal model of white matter diffusion (MMWMD) in [7,8] to accommodate the presence of multiple fiber populations in a single voxel. The MMWMD represents white matter as an ensemble of impermeable cylindrical axons, with both a single diameter and a single orientation, embedded in a homogeneous medium. Water molecules diffuse with an identical intrinsic diffusivity both inside and outside the axons but without exchanges between the compartments. By assuming a single axon diameter rather than a distribution as in [6,13], the MMWMD enables orientation-invariant estimation of axon diameter for the first time, providing a single summary statistic of the axon diameter distribution, called the axon diameter index [8]. The axon diameter index is simpler to estimate than models of the full distribution in [6,13], but still discriminates naturally occurring axon diameter distributions [8].

To model crossing fibers, we instead represent white matter as a set of $N \geq 1$ distinct fiber populations embedded in a common homogeneous medium. Each

fiber population is individually modeled as prescribed by the MMWMD. Assuming that no exchange occurs between the fiber populations and their common surrounding medium, as well as between the fiber populations themselves, the normalized (MR) signal from white matter, A_{wm}, is thus $\nu_{ic}\sum_{i=1}^{N} f_{ic}^{i} A_{ic}^{i} + (1-\nu_{ic})A_{ec}$, where $\nu_{ic} \in [0,1]$ is the intra-cellular volume fraction within the white matter, A_{ic}^{i} the normalized signal from the i-th fiber population with a volume fraction $f_{ic}^{i} \in [0,1]$ relative to the whole population of fibers, and A_{ec} the normalized signal from the extra-cellular medium.

The MMWMD contains two additional compartments: one models isotropically free Gaussian diffusion to capture the partial volume with cerebrospinal fluid (CSF); the other models isotropically restricted diffusion to capture observed restrictions parallel to axons in fixed tissue, potentially due to water trapped within glial cells [5]. Here we include only the latter, because the white matter area considered in our brain data experiment is several voxels away from the boundary with the ventricles and thus void of CSF contamination. The full normalized signal model, A, is therefore $(1-\nu_{ir})A_{wm} + \nu_{ir}A_{ir}$, where A_{ir} is the normalized signal from the isotropically restricted compartment with a volume fraction ν_{ir}. We set $A_{ir} = 1$, following the stationary water assumption in [8], i.e., the compartment signal remains unattenuated by diffusion weighting. The subsequent sections detail the modeling of A_{ic}^{i} and A_{ec} and the fitting procedure.

Intra-Cellular Model. Water diffusion in this compartment is cylindrically restricted. The intra-cellular signal from the i-th fiber population, A_{ic}^{i}, depends on the axon diameter and orientation of the population, denoted by a_i and $\boldsymbol{n}_i$ respectively, the intrinsic diffusivity of water, denoted by d, as well as the imaging protocol. For the pulsed-gradient spin-echo (PGSE) sequence, as used in our experimental evaluation, we compute A^{ic} using the Gaussian phase distribution approximation [14] as in [7,8] to model the signal from restricted diffusion.

Extra-Cellular Model. Water diffusion in this compartment is hindered. The extra-cellular signal, A_{ec}, is modeled with simple (Gaussian) anisotropic diffusion using an (apparent) diffusion tensor as in [15]. We model the diffusion tensor, $\mathbb{D}_h$, as $\sum_{i=1}^{N} f_{ic}^{i}\mathbb{D}_{cyl}(\nu_{ic}, \boldsymbol{n}_i, d)$, where $\mathbb{D}_{cyl}(\nu_{ic}, \boldsymbol{n}_i, d)$ is the diffusion tensor representing the hindered diffusion in the i-th fiber population, defined according to the MMWMD. This follows from the common medium assumption, i.e., water exchanges freely within different regions of the extra-cellular space.

Model Fitting. We fit the proposed model to data with the three-stage routine described in [8]. It provides robust estimates of the model parameters with the Rician Markov Chain Monte Carlo (MCMC) procedure in [7], after an initial grid search and then gradient descent to determine the maximum likelihood (ML) estimates of the parameters.

The full set of model parameters are the N axon diameters (a_1, a_2,..., a_N) and orientations ($\boldsymbol{n}_1$,$\boldsymbol{n}_2$, ..., $\boldsymbol{n}_N$), the $N-1$ relative volume fractions (f_{ic}^{1}, f_{ic}^{2}, ..., f_{ic}^{N-1}), the other two volume fractions ν_{ic} and ν_{ir}, and the intrinsic diffusivity

d. Note that, because $\sum_{i=1}^{N} f_{ic}^{i} = 1$, only $N-1$ of the relative volume fractions are independent. Throughout, we fix d to 0.6×10^{-9} $\mathrm{m^2 s^{-1}}$, its expected value in the *ex vivo* data, as in [8].

3 Experiments and Results

Ex Vivo Imaging of Monkey Data Set. *Ex vivo* diffusion weighted imaging (DWI) of a 32-month perfusion-fixed Vervet monkey was acquired on a 4.7T Varian system with maximum gradient strength $|\boldsymbol{G}| = 400$ mT/m. (See [16] for brain preparation.) A total of 360 images were collected using a PGSE DWI sequence with single-line spin-echo readout (TE = 36ms, TR = 2500ms). Each has isotropic 0.5x0.5x0.5 mm^3 voxels and 30 sagittal slices centered on the mid-sagittal plane of the corpus callosum with in-plane matrix size of 128x256. The protocol, determined using the experimental design optimization [7] adapted for fixed tissue, consists of three HARDI shells with the number of diffusion gradients [103, 106, 80], $|\boldsymbol{G}| = [300, 210, 300]$ mT/m, $\delta = [5.6, 7.0, 10.5]$ ms, and $\Delta = [12, 20, 17]$ ms, corresponding to b-values = [2084, 3084, 9550] $\mathrm{s/mm^2}$. Ethical rules concerning the handling and care of live animals were followed.

Synthetic Data Experiment. We synthesize diffusion MR signal from a broad set of two-fiber crossing substrates using the publicly available Monte-Carlo diffusion simulator in Camino [17] with the *ex vivo* imaging protocol described above. Synthetic Rician noise with $\sigma = 0.05$ is added to match the SNR of the *ex vivo* data (around 20). For each substrate, the proposed model with $N = 2$ is then fitted to its synthetic data and the parameter estimates are compared against the known ground-truth settings of the substrate.

The synthetic substrates assume reasonable crossing configurations and microstructure features and are constructed as follows: We consider six representative axon diameter combinations: $\{a_1, a_2\} \in \{\{2,2\},\{2,4\},\{2,6\},\{4,4\},\{4,6\},\{6,6\}\}$ μm. For each combination, we test three relative volume fractions, $f_{ic}^{1} \in \{0.3, 0.5, 0.7\}$, and a broad range of crossing angles, varying from 30^o to 90^o in 15^o increment. Similar to [7], ν_{ic} is set to 0.7, its typical value in brain white matter. Since the diffusion simulator does not model isotropically restricted compartment, ν_{ir} is 0. This results in a total of 90 different crossing fiber substrates. Finally, to avoid possible orientation dependence, 20 different instances of each substrate are created by applying random 3-D rotations to the initial configuration. Independent Rician noise described above are added to each instance.

For each substrate, we compute the mean and the standard deviation of the parameter estimates for its 20 random instances. This accounts for the effect of noise and the dependence to orientation. We report our findings from the assessment of all 90 substrates. Due to limit in space, the findings are illustrated with only one axon diameter combination ($\{2,6\}\mu m$) and $f_{ic}^{1} = 0.3$ in Fig. 1.

The relative volume fraction f_{ic}^{1} can be estimated accurately for all axon diameter combinations and for the crossing angles larger than 45^o; for the crossing angle of 30^o, it is consistently over-estimated. The intra-cellular volume fraction

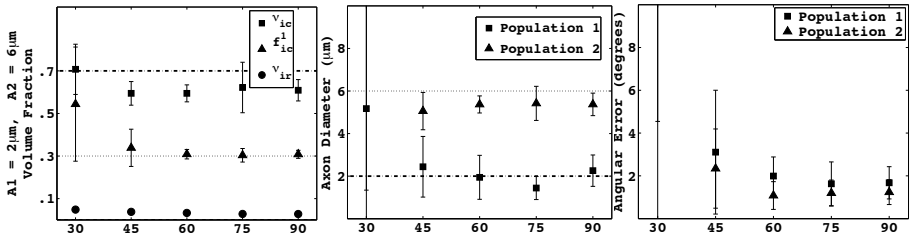

Fig. 1. The parameter estimates for the substrates: $\{2, 6\}\mu m$, $f_{ic}^1 = 0.3$, and all the crossing angles (the horizontal axis)

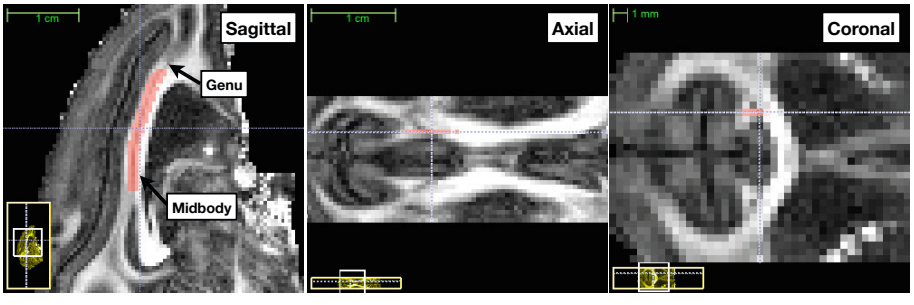

Fig. 2. The manually defined crossing fiber ROI (in red) overlaid on the FA map in three orthogonal views

ν_{ic} is consistently under-estimated for all substrates, by about 10%. The volume fraction of the isotropically restricted compartment ν_{ir} can be consistently estimated for all substrates, with a slight over-estimation.

For axon diameters, $2\mu m$ and $4\mu m$ are difficult to differentiate from one another for all values of f_{ic}^1 and crossing angles. However, both can be differentiated from $6\mu m$ for all values of f_{ic}^1 and crossing angles larger than 45^o.

Orientations can be estimated accurately for all axon diameter combinations, all values of f_{ic}^1, and the crossing angles larger than 45^o. The error in the orientation estimates increases as the crossing angle decreases and as the axon diameter increases. The error in the estimate of the fiber population with the lower relative volume fraction is higher than that of the other population.

In summary, under practical imaging protocol, the proposed model is able to both resolve the crossing fiber configuration and estimate microstructure features for crossing angles larger than 45^o.

Monkey Data Experiment. This experiment uses the *ex vivo* monkey data set described earlier to demonstrate the efficacy of the proposed model for mapping axon diameters in brain tissue with crossing fibers. We manually delineate a well-defined two-fiber crossing region-of-interest (ROI) between the corpus callosum and the cingulum bundle as shown in Fig. 2. Note the dark band in the center

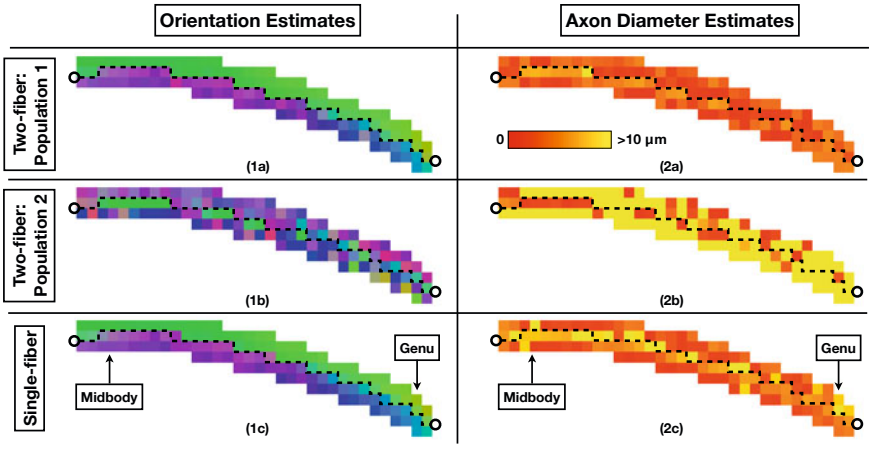

Fig. 3. The orientation and axon diameter estimates for both the two-fiber and single-fiber models. The orientation estimates are represented using the standard RGB-encoding (Red: medial-lateral, Blue: inferior-superior, Green: anterior-posterior) [18].

of the ROI. The appearance of such voxels flanked by the ones with higher FA is characteristic of crossing fibers due to partial volume between two tracts with substantially different orientations. We fit both the proposed two-fiber model and the MMWMD (single-fiber) to each voxel in this region, and compare their respective estimates.

Fig. 3 shows the orientation and axon diameter estimates for both the two-fiber and single-fiber models. For the two-fiber model, the fiber population 1 corresponds to the population initialized, during model fitting, with the primary eigenvector of the diffusion tensor estimate at each voxel. The well-defined trajectory of the cingulum, which traverses from anterior to posterior, allows us to identify the voxels colored in green in Fig. 3(1a) as the ones with the fiber population 1 being the cingulum. The dashed line separates these voxels from the ones below them, which are primarily the corpus callosum fibers.

The most striking observation is that, despite the orientation estimates of the single-fiber model (Fig. 3(1c)) being largely consistent with those of the fiber population 1 of the two-fiber model (Fig. 3(1a)), their respective axon diameter estimates are distinctly different. In particular, among the voxels immediately adjacent to the dashed line, the putative interface between the cingulum and the corpus callosum, the axon diameter estimates of many, from the single-fiber model (Fig. 3(2c)), are close to or higher than $10\mu m$, much higher than the values determined from histology (about $2\mu m$) [2]. In contrast, their corresponding estimates for the fiber population 1 from the two-fiber model (Fig. 3(2a)) are in the same range as the histologically estimated values, highlighting a key benefit of using the two-fiber model in such crossing fiber voxels.

The orientation and axon diameter estimates for the fiber population 2 from the two-fiber model require more care in their interpretation. The most apparent feature of the axon diameter estimates (Fig. 3(2b)) is that a large majority of them are higher than $10\mu m$. Further inspection confirms that the values for almost all these voxels are close or equal to 40 μm, the maximum value allowed in our fitting routine to indicate the negligible presence of the corresponding fiber population, in this case, the fiber population 2. On the other hand, a number of voxels have axon diameter estimates with values comparable to those from histology. They can be divided into two groups according to their spatial locations and orientations. In the first group are the voxels that are immediately adjacent to the putative cingulum and corpus callosum interface and have their orientation estimates for the fiber population 2 consistent with being part of the cingulum, i.e., colored in green in Fig. 3(1b). For these voxels, the two-fiber model is able to both resolve the crossing configuration and estimate the axon diameters of individual fiber populations. In the second group are the voxels that are away from the interface and inside the cingulum region. The fiber population 2 of these voxels may correspond to the outward-projecting cingulum fibers.

4 Discussion

We have described a technique for joint estimation of crossing fiber configuration and microstructure features using a new crossing-fiber white matter model that includes individual axon diameter and volume fraction parameters for each fiber population. The results from simulation and brain data demonstrate promising possibilities for extending axon diameter mapping to the whole brain. In particular, the brain data findings suggest that the crossing-fiber model can resolve crossing configurations and estimate axon diameters simultaneously under some circumstances; when it fails to do so, it can still provide more sensible axon diameter estimates for the dominant fiber population. Given that it is challenging to estimate axon diameter in voxels with only a single orientation, it is remarkable that we can make progress at crossing. Even improving the estimate of the dominant popualtion is itself useful and a significant step forward.

Nevertheless, the proposed crossing fiber model can be extended in several ways: First, the current model assumes that axons are strictly parallel within each fiber population. This assumption can be relaxed to account for axon spreading in each distinct population using the orientation dispersion model in [9]. Second, the current model assumes that the extra-cellular space is a common homegeneous medium. A more general model may require the modeling of the extra-cellular space of each fiber population individually. Lastly, we will examine other crossing regions, such as the pons and the crossing regions between the corpus callosum and the corona radiata in centrum semiovale.

Acknowledgement. We thank Prof Maurice Ptito, University of Montreal, for providing the monkey brain. The future and emerging technologies program of the EU FP7 framework funds the CONNECT consortium (brain-connect.eu), which supports this work. EPSRC fund DCA under grant EP/E007748.

References

1. Ritchie, J.M.: On the relation between fibre diameter and conduction velocity in myelinated nerve fibres. Proc. R. Soc. Lond. B 217, 29–35 (1982)
2. Lamantia, A.S., Rakic, P.: Cytological and quantitative characteristics of four cerebral commissures in the rhesus monkey. J. Comp. Neurol. 291, 520–537 (1990)
3. Shintaku, M., Hirano, A., Llena, J.F.: Increased diameter of demyelinated axons in chronic multiple sclerosis of the spinal cord. Neuropathol. Appl. Neurobiol. 14, 505–510 (1988)
4. Sasaki, S., Maruyama, S.: Increase in diameter of the axonal initial segment is an early change in amyotrophic lateral sclerosis. J. Neurological. Sci. 110, 114–120 (1992)
5. Stanisz, G.J., Szafer, A., Wright, G.A., Henkelman, M.: An analytical model of restricted diffusion in bovine optic nerve. Magn. Reson Med. 37, 103–111 (1997)
6. Assaf, Y., Blumenfeld-Katzir, T., Yovel, Y., Basser, P.J.: AxCaliber: a method for measuring axon diameter distribution from diffusion MRI. Magn. Reson Med. 59, 1347–1354 (2008)
7. Alexander, D.C.: A general framework for experiment design in diffusion MRI and its application in measuring direct tissue-microstructure features. Magn. Reson Med. 60, 439–448 (2008)
8. Alexander, D.C., Hubbard, P.L., Hall, M.G., Moore, E.A., Ptito, M., Parker, G.J.M., Dyrby, T.B.: Orientationally invariant indices of axon diameter and density from diffusion MRI. NeuroImage, 1374–1389 (2010)
9. Zhang, H., Hubbard, P., Parker, G.J.M., Alexander, D.C.: Axon diameter mapping in the presence of orientation dispersion with diffusion MRI. NeuroImage 56, 1301–1315 (2011)
10. Tuch, D.S., Reese, T.G., Wiegell, M.R., Makris, N., Belliveau, J.W., Wedeen, V.J.: High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. Magn. Reson Med. 48, 577–582 (2002)
11. Jansons, K.M., Alexander, D.C.: Persistent angular structure: new insights from diffusion magnetic resonance imaging data. Inverse Problems 19, 1031–1046 (2003)
12. Seunarine, K.K., Alexander, D.C.: Multiple fibres: beyond the diffusion tensor. In: Behrens, T.E.B., Johansen-Berg, H. (eds.) Diffusion MRI, pp. 55–72. Elsevier, Amsterdam (2009)
13. Barazany, D., Basser, P.J., Assaf, Y.: In-vivo measurement of the axon diameter distribution in the corpus callosum of a rat brain. Brain 132, 1210–1220 (2009)
14. Murday, J.S., Cotts, R.M.: Self-diffusion coefficient of liquid lithium. J. Chem. Phys. 48, 4938–4945 (1968)
15. Assaf, Y., Freidlin, R.Z., Rhode, G.K., Basser, P.J.: New modeling and experimental framework to characterize hindered and restricted water diffusion in brain white matter. Magn. Reson Med. 52, 965–978 (2004)
16. Dyrby, T.B., Baaré, W.F.C., Alexander, D.C., Jelsing, J., Garde, E., Søgaard, L.V.: An ex vivo imaging pipeline for producing high-quality and high-resolution diffusion-weighted imaging datasets. Hum. Brain Mapp. 32, 544–563 (2010)
17. Hall, M.G., Alexander, D.C.: Convergence and parameter choice for Monte-Carlo simulations of diffusion MRI. IEEE Trans. Med. Imaging 28, 1354–1364 (2009)
18. Pajevic, S., Pierpaoli, C.: Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: application to white matter fiber tract mapping in the human brain. Magn. Reson Med. 42, 526–540 (1999)

Detecting Structure in Diffusion Tensor MR Images

K. Krishna Nand[1], Rafeef Abugharbieh[1], Brian G. Booth[2], and Ghassan Hamarneh[2]

[1] Biomedical Signal and Image Computing Lab, University of British Columbia
[2] Medical Image Analysis Lab, School of Computing Science, Simon Fraser University
{kkrishna,rafeef}@ece.ubc.ca, {bgb2,hamarneh}@sfu.ca

Abstract. We derive herein first and second-order differential operators for detecting structure in diffusion tensor MRI (DTI). Unlike existing methods, we are able to generate full first and second-order differentials without dimensionality reduction and while respecting the underlying manifold of the data. Further, we extend corner and curvature feature detectors to DTI using our differential operators. Results using the feature detectors on diffusion tensor MR images show the ability to highlight structure within the image that existing methods cannot.

1 Introduction

Feature detection is a core component in most image processing and analysis tasks with ubiquitous applications ranging from data registration to object segmentation, classification, and recognition. Of the multitude of feature detectors available today, those that are widely used typically involve first and second order spatial differential operators [9,18]. We desire the ability to use such feature detectors to highlight anatomical structure in diffusion tensor images (DTI). To date, only a few attempts have been made to develop differential operators for such manifold-valued image data. Table 1 summarizes the current state-of-the-art and highlights a conspicuous gap in the current methodology: the ability of current differential operators to respect the manifold-valued nature of diffusion tensor data.

Table 1. Summary of existing and proposed work in differential-based feature detection for scalar and manifold-valued data. The manifold of symmetric second-order positive definite tensors is represented as SPD(3). Methods listed for $\mathbb{R}^N$ also apply to $\mathbb{R}^3$. Note our contributions in the lower-right of the table. They are derived in Equations (1 - 3).

Data Dimensionality	First-Order Differentials		Second-Order Differentials	
	Magnitude Only	**Full Differential**	**Magnitude Only**	**Full Differential**
$\mathbb{R}$	Edge Detection [5] Corner Detection [18]		Frangi *et al.* Curvature [9] Curvature Magnitude [12]	
$\mathbb{R}^3$	MPT Projection [20]		Quaternions [19]	
$\mathbb{R}^N$	Channel Normalization [8] Third-Order Tensors [17] Structure Tensor [6]		Directional Derivative of Gradient [1]	Weighted Channel Averaging [13]
SPD(3)	Distance Metrics [2]	*Our Contribution: Equations (1 - 2)*	*Our Contribution: Equation (3)*	

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 90–97, 2011.

In this paper, we derive first and second order differential operators for diffusion tensor fields that respect the manifold of second-order tensors. Using these differential operators, we extend the Harris and Shi-Tomasi corner detectors to DTI. We further construct curvature based features detectors to detect *tube-like* and *sheet-like* structures in DTI. These features allow us to capture structural information within a DT image that has previously not been explored. We note here that our approach to detecting these structural features is different from other approaches [11,15,21] in that we work with more fundamental operators of low-level computer vision as opposed to higher level characterization of the data.

2 Methods

2.1 First Order Operators

We begin our exposition by deriving the first-order differential operator of a tensor field. To ensure that we respect the manifold of second-order tensors, we employ the Log-Euclidean (LE) mapping proposed in [2] to map the space of symmetric second order positive definite tensors, SPD(3), into a vector space. Let $\boldsymbol{D}(\boldsymbol{x})$ be a 3D DT image indexed by $\boldsymbol{x} = [x_1 \;\; x_2 \;\; x_3]$. Then the LE mapping gives $\boldsymbol{L}(\boldsymbol{x}) = \text{logm}(\boldsymbol{D}(\boldsymbol{x}))$, where $\boldsymbol{L}(\boldsymbol{x})$ is the LE tensor and logm$(\cdot)$ is the matrix logarithm.

Let $\boldsymbol{J}(\boldsymbol{x})$ denote the Jacobian matrix, given by $J_{ij} = \frac{\partial L_i}{\partial x_j}$. $\boldsymbol{J}(\boldsymbol{x})$ generalizes the gradient of a scalar field to the derivatives of a vector field. We define the 3×3 symmetric positive semi-definite matrix $\boldsymbol{S}(\boldsymbol{x}) = \boldsymbol{J}(\boldsymbol{x})^{\boldsymbol{T}}\boldsymbol{J}(\boldsymbol{x})$. The matrix $\boldsymbol{S}(\boldsymbol{x})$, referred to as the structure tensor, measures the directional dependence of total change. Let $\lambda_1^{\boldsymbol{S}} \geq \lambda_2^{\boldsymbol{S}} \geq \lambda_3^{\boldsymbol{S}}$ denote the eigenvalues of $\boldsymbol{S}(\boldsymbol{x})$ and let $\hat{\boldsymbol{e}}_{\boldsymbol{i}}^{\boldsymbol{S}}$ be the eigenvector corresponding to $\lambda_i^{\boldsymbol{S}}$. The squared Euclidean norm of $d\boldsymbol{L}(\boldsymbol{x})$ can be written in terms of $\boldsymbol{S}(\boldsymbol{x})$ as $||d\boldsymbol{L}(\boldsymbol{x})||^2 = (d\boldsymbol{x})^T\boldsymbol{S}(\boldsymbol{x})d\boldsymbol{x}$. This positive definite quadratic form is called the first fundamental form. For a unit vector $\boldsymbol{n}$, $A(\boldsymbol{x}) = \boldsymbol{n}^T\boldsymbol{S}(\boldsymbol{x})\boldsymbol{n}$ measures the rate of change of the image in the direction $\boldsymbol{n}$ and is referred to as the squared local contrast [6]. It is maximum when $\boldsymbol{n}$ is in the direction of the eigenvector corresponding to the largest eigenvalue of $\boldsymbol{S}(\boldsymbol{x})$. As a result, this eigenvector is used as the gradient direction $\hat{\boldsymbol{g}}(\boldsymbol{x})$. To resolve the sign ambiguity of the eigenvector, we use a voting approach across channels. The gradient magnitude $|\boldsymbol{g}(\boldsymbol{x})|$ is set to be the square root of the corresponding eigenvalue of $\boldsymbol{S}(\boldsymbol{x})$. The resulting gradient vector is then given by:

$$\boldsymbol{g}(\boldsymbol{x}) = \sqrt{\lambda_1^{\boldsymbol{S}}}\hat{\boldsymbol{e}}_{\boldsymbol{1}}^{\boldsymbol{S}}. \tag{1}$$

A closed-form solution for (1) in terms of $\boldsymbol{S}(\boldsymbol{x})$ can be derived but is not presented here due to space limitations.

An averaged version of $\boldsymbol{S}$ around a local neighborhood, $\bar{\boldsymbol{S}}$, allows the integration of first order information from the neighborhood and a more stable numerical derivation. We call $\bar{\boldsymbol{S}}$, the log-Euclidean structure tensor and construct it as:

$$\bar{S}_{ik}(\sigma_I, \boldsymbol{x}) = \boldsymbol{\omega}_{\sigma_I}(\boldsymbol{x}) * \sum_j \frac{\partial L_j}{\partial x_i}\frac{\partial L_j}{\partial x_k}, \tag{2}$$

where $\boldsymbol{\omega}_{\sigma_I}$ is a Gaussian window with standard deviation σ_I and L_j is the j^{th} component of $\boldsymbol{L}(\boldsymbol{x})$.

2.2 Second Order Operators

The Hessian matrix is a symmetric matrix consisting of second order partial derivatives. It describes the local second order structure around each pixel in an image. We derive two forms of the Hessian matrix. The first, $\boldsymbol{H_1}$, is based on simple weighted averaging of the Hessian matrices of each component of $\boldsymbol{L}(\boldsymbol{x})$, as done for color images in [13]. Our second formulation is based on the gradient vector we derived earlier in (1). Given $\boldsymbol{g}(\boldsymbol{x}) = [g_{x_1}\ g_{x_2}\ g_{x_3}]^T$ has been computed at every location $\boldsymbol{x}$ of the DT image, the gradient vector field $\boldsymbol{g}$ can be considered as a multivalued function with the mapping $\Omega_g \colon \mathbb{R}^3 \rightarrow \mathbb{R}^3$. Using this context, we compute the derivative of the gradient vector field via the Jacobian matrix, $\boldsymbol{G}(\boldsymbol{x})$, as $G_{ij} = \frac{\partial g_i}{\partial x_j}$. $\boldsymbol{G}(\boldsymbol{x})$ encodes second order information of $\boldsymbol{L}(\boldsymbol{x})$ and, in scalar images, $\boldsymbol{G}(\boldsymbol{x})$ represents the Hessian matrix. However in multivalued images, $\boldsymbol{G}(\boldsymbol{x})$ is generally non-symmetric. Our second formulation for the Hessian matrix $\boldsymbol{H_2}$ in multivalued images is thus given by:

$$\boldsymbol{H_2}(\boldsymbol{x}) = \frac{\boldsymbol{G}(\boldsymbol{x}) + \boldsymbol{G}(\boldsymbol{x})^{\boldsymbol{T}}}{2}, \tag{3}$$

which is the symmetric part of $\boldsymbol{G}(\boldsymbol{x})$ and is the best L_2 approximation to $\boldsymbol{G}(\boldsymbol{x})$ from the set of symmetric matrices [3]. Our first order (1,2) and second order (3) differential operators can easily be extended to different scales.

2.3 Feature Detectors

Our proposed differential operators fit naturally into existing feature detectors and can thus be effectively used to detect structure in DTI. Here we provide four examples from a potentially long list of feature detectors that can incorporate our differential operators. Using our first order differential in (2), we detect corners in DT images using the popular Harris (C_H) [16] and Shi-Tomasi (C_{ST}) [18] corner detectors which we define as:

$$C_H = \frac{\det(\bar{\boldsymbol{S}})}{\mathrm{tr}(\bar{\boldsymbol{S}}) + \epsilon} \tag{4}$$

$$C_{ST} = \min(\lambda_1^{\bar{\boldsymbol{S}}}, \lambda_2^{\bar{\boldsymbol{S}}}, \lambda_3^{\bar{\boldsymbol{S}}}) \tag{5}$$

where $det(\bar{\boldsymbol{S}})$ is the determinant of $\bar{\boldsymbol{S}}$, $tr(\bar{\boldsymbol{S}})$ the trace of $\bar{\boldsymbol{S}}$ and $\lambda_1^{\bar{\boldsymbol{S}}} \geq \lambda_2^{\bar{\boldsymbol{S}}} \geq \lambda_3^{\bar{\boldsymbol{S}}}$ are the eigenvalues of $\bar{\boldsymbol{S}}$.

Curvature cues can also be obtained from our Hessian matrix derived in (3). As examples, we extend the vesselness filter of [9] and the sheetness filter of [7] to detect tube-like and sheet-like fiber tracts in DTI[1]. Our tubular-ness and sheet-ness measures are thus defined as:

$$C_{tube} = \left(1 - \exp\left(-\frac{\mathcal{R}_A^2}{2\alpha^2}\right)\right)\left(\exp\left(-\frac{\mathcal{R}_B^2}{2\beta^2}\right)\right)\left(1 - \exp\left(-\frac{\mathcal{S}^2}{2c^2}\right)\right), \tag{6}$$

$$C_{sheet} = \left(\exp\left(-\frac{\mathcal{R}_A^2}{2\alpha^2}\right)\right)\left(1 - \exp\left(-\frac{\mathcal{R}_D^2}{2\eta^2}\right)\right)\left(1 - \exp\left(-\frac{\mathcal{S}^2}{2c^2}\right)\right), \tag{7}$$

[1] Here, it is important to note that tube-like and sheet-like refer to the structural characteristics of the DT image and not to the diffusion tensors.

where C_{tube} and C_{sheet} measure how tubular and sheet-like the structure is, $\mathcal{R}_A = |\lambda_2^H|/|\lambda_1^H|$, $\mathcal{R}_B = |\lambda_3^H|/\sqrt{|\lambda_1^H \lambda_2^H|}$, $\mathcal{R}_D = |(2|\lambda_1^H| - |\lambda_2^H| - |\lambda_3^H|)|/|\lambda_1^H|$ and $\mathcal{S} = \sqrt{(\lambda_1^H)^2 + (\lambda_2^H)^2 + (\lambda_3^H)^2}$. The measures $\mathcal{R}_A$, $\mathcal{R}_B$ and $\mathcal{S}$ are used to penalize sheet-like structures, blob-like structures and noise, respectively. $\mathcal{R}_D$ is used to penalize tube like and blob like structures. These terms are computed from the eigenvalues $\lambda_1^H \geq \lambda_2^H \geq \lambda_3^H$ of the Hessian matrix in (3). The constants α, β, η and c control the sensitivity of C_{tube} and C_{sheet} to each term.

3 Results

We present results of the proposed differential operators and feature detectors on synthetic and real DT data. The real DT data consisted of 24 images of normal adult brains, of which, 12 were taken from the John Hopkins LBAM [14] database and 12 from the MIDAS database [4]. In our experiments, the constants α, β and η were set to 0.5 (as suggested in [7,9]) and c was set to 0.1.

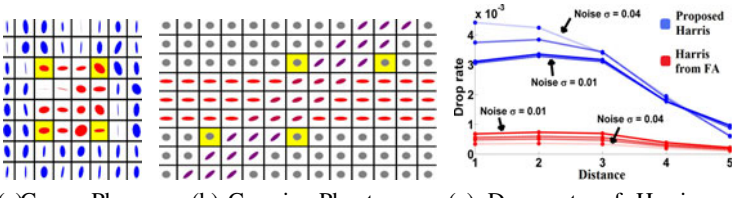

(a) Corner Phantom (b) Crossing Phantom (c) Drop rate of Harris response vs. distance

Fig. 1. Results of the proposed Harris and Shi-Tomasi corner detectors on synthetic data. (a) and (b): two synthetic slices with the ground truth corner locations shown in yellow. (c): the drop rate of the Harris feature response as a function of distance from the ground truth corners for different levels of noise (standard deviation of noise ranging from 0.01 to 0.04). Results are shown for both our method and the approach in [16] on the FA maps. Note the response of our method falls rapidly for all levels of noise, demonstrating localization accuracy.

Figure 1 shows the proposed Harris and Shi-Tomasi corner detector results on synthetic data. The conventional scalar Harris and Shi-Tomasi features obtained from the fractional anisotropy (FA) maps are used for comparison. Figure 1(a) shows a noisy image where the tensors are pointing horizontally in the foreground (colored in red) while tensors point vertically in the background (colored in blue). Figure 1(b) shows an example of two fibers crossing each other. The ground truth corner points are shown in yellow. Varying levels of noise were added to these images using the method in [10]. We show in Figure 1(c) the drop rate in the feature response as a function of distance to the nearest ground truth corner point in the examples in Figures 1(a) and 1(b). Our results demonstrate how the response of the proposed Harris detector falls very quickly for all levels of noise when compared to the response of the scalar Harris detector on FA as we move away from the ground truth corners. We also observed consistent stronger response to corners for our proposed Harris detector compared to the approach based on

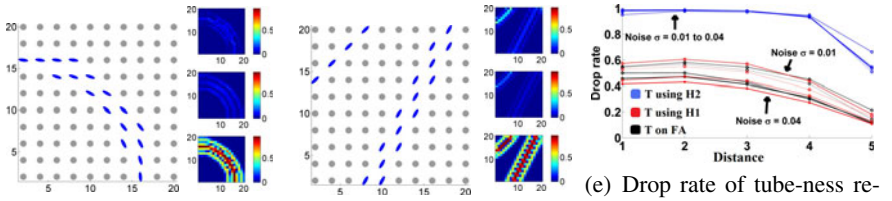

(a) Bent fiber tract (b) (c) Linear fibers (d) (e) Drop rate of tube-ness response vs. distance

Fig. 2. Tubular-ness results for synthetic 2D slices. Figures (a) and (c) show a bent fiber tract and linear fiber tracts oriented in different angles, respectively. Figures (b) and (d) show the corresponding tubular-ness responses obtained using Hessian from FA (top), $\boldsymbol{H_1}$ (middle) and $\boldsymbol{H_2}$ (bottom). Note that tubular-ness obtained using $\boldsymbol{H_2}$ gives the strongest response in both cases. Figure (e) shows the drop of the filter's response as a function of distance to the ground truth medial of the tube. Note that even when noise is high, tubular-ness using $\boldsymbol{H_2}$ decreases more rapidly than tubular-ness using $\boldsymbol{H_1}$ or FA, indicating greater localization ability.

the FA maps. A similar behavior was also observed for the Shi-Tomasi detector (figure not shown due to space constraints). This rapid decrease in response for our method indicates its greater ability to localize the corner points in various noise conditions.

Figure 2 shows results of tubular-ness filter on synthetic data. The tubular-ness obtained from the corresponding FA map using the scalar version of the filter in [9] is included for comparison. Figure 2(a) shows a bent fiber tract and figure 2(c) shows linear fiber tracts oriented along different angles. Figures 2(b) and 2(d) show the corresponding tubular-ness results. We observe that tubular-ness using $\boldsymbol{H_2}$ gives the highest response in both cases. In figure 2(e), we add noise to these images and, as before, we measure the rate of decrease of the filter's response as a function of distance to the ground truth medial of the tube. We observe that tubular-ness using $\boldsymbol{H_2}$ falls more rapidly when compared to the other filters, even when the noise is high. Again, this result illustrates our method's greater localization ability.

To quantify the information captured by the differential operators, we compute various norms of our structure tensor, gradient vector and Hessian matrices for twelve adult brain DTI from the LBAM database. Histograms of these norms are shown in Figure 3 for the proposed methods as well as those obtained from the FA map. Note that the norms of our proposed differential operators are spread over larger values than the norms of the differentials obtained from the FA maps. This result stems directly from avoiding the dimensionality reduction of the tensor data to FA values. Also note that $\boldsymbol{H_2}$ captures a greater amount of differential information than the other two Hessian approaches as $\boldsymbol{H_2}$ works with the whole tensor and not its individual channels or FA.

We applied our corner and tube detectors to 24 real adult brain DTI. Figure 4 shows a representative result for the corner detectors for both our method and the conventional scalar Harris [16] and Shi-Tomasi [18] features obtained from the FA maps. First, we observe that the feature response of our approach is stronger than the feature response from the FA-based approach by three orders of magnitude. Secondly our approach generates a more distinctive response to the corners of both the Genu and Splenium of the

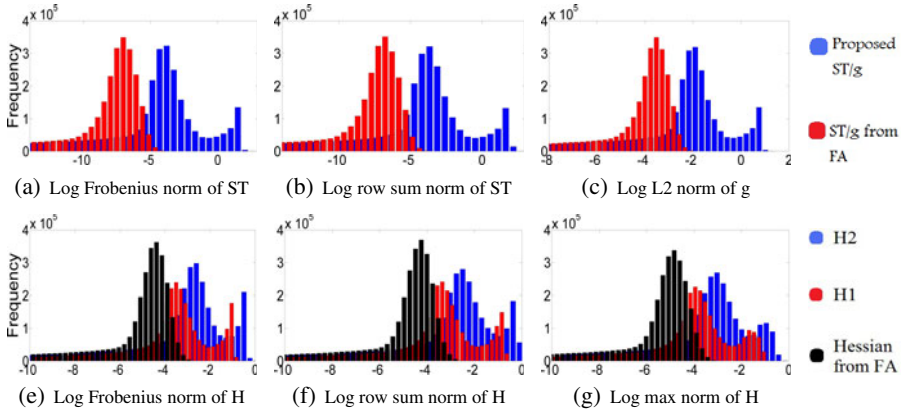

(a) Log Frobenius norm of ST (b) Log row sum norm of ST (c) Log L2 norm of g

(e) Log Frobenius norm of H (f) Log row sum norm of H (g) Log max norm of H

Fig. 3. Histograms of various norms of the structure tensor, gradient vector and Hessian matrix (abbreviated in the figure as ST, g and H, respectively) for twelve adult brain DT images from the LBAM database. Shown are results for the proposed methods and FA-based feature detection. Note the greater norm of our proposed differential operators compared to FA-based approaches, demonstrating that we are capturing more differential information.

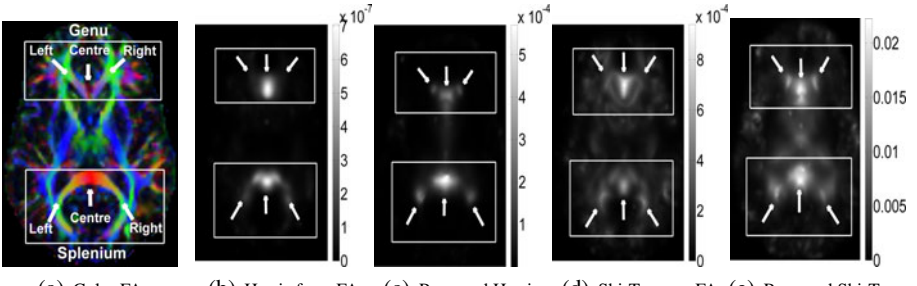

(a) Color FA (b) Harris from FA (c) Proposed Harris (d) Shi-Tom. on FA (e) Proposed Shi-Tom.

Fig. 4. Representative result for the Harris and Shi-Tomasi corner detectors on an adult brain DT image. Shown are (a) the color FA map, (b) Harris features from the FA map, (c) Harris features using our method, (d) Shi-Tomasi features from the FA map, and (e) Shi-Tomasi features using our method. Note that since the responses of the FA-based detectors are extremely weak when compared to the responses of our proposed detectors, they are shown with different intensity scales. Also note the clearer response to the corners of the Genu and Splenium using our method.

Corpus Callosum than the FA-based detectors. These results were consistently observed throughout the datasets.

Finally, Figure 5 shows representative tubular-ness and sheet-ness results. Note that tubular tracts such as the Cingulum and Fornix are well detected using $\boldsymbol{H_2}$. Sheet-like tracts such as the Inferior Longitudinal Fasciculus, Corpus Callosum and Corticospinal Tract are also better detected using $\boldsymbol{H_2}$ compared to $\boldsymbol{H_1}$ and FA-based features. The FA-based features fail to detect structures in most cases. These results were consistently observed throughout the datasets.

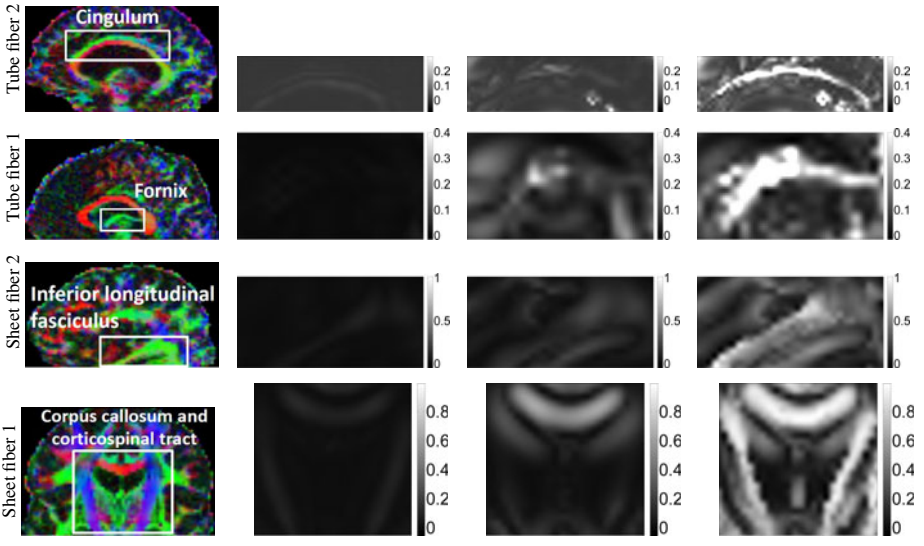

Fig. 5. Representative results for the tubular-ness and sheet-ness filters on adult brain DTI data. Shown from top to bottom are two results for tubular-ness and two results for sheet-ness. Shown from left to right are (a) color FA image, (b) results obtained using [7,9] on FA, (c) results obtained using $\boldsymbol{H_1}$ and (d) results obtained using $\boldsymbol{H_2}$. Note the identified tubular tracts (Cingulum, Fornix) and the identified sheet-like tracts (Inferior Longitudinal Fasciculus, Corpus Callosum, Corticospinal Tract) are better detected using $\boldsymbol{H_2}$ than with $\boldsymbol{H_1}$ and FA based method.

4 Conclusion

We derived novel first and second order differential operators for DTI. Unlike existing state-of-the-art, our operators respect the manifold of symmetric second-order tensors through the use of the log-Euclidean mapping. We further show how our differential operators can be naturally incorporated into various feature detectors in order to find structure in diffusion tensor images that, to date, has not been possible. We extend the Harris and Shi-Tomasi corner detectors to DTI and show that our approach better distinguishes corners in DTI data and is more robust to noise. We also extend the vesselness filter of [9] and the sheetness filter of [7] to detect tube-like and sheet-like structures. We show that our methods better detect these tube-like and sheet-like structures in DTI data and are more robust to noise than existing methods. We believe that our derived low-level operators and image features are very versatile and will be of great advantage to classification, registration, and segmentation of DTI data.

This work was supported in part by NSERC and the Institute for Computing, Information and Cognitive Systems (ICICS) at UBC.

References

1. Alshatti, W., Lambert, P.: Using eigenvectors of a vector field for deriving a second directional derivative operator for color images. In: Chetverikov, D., Kropatsch, W.G. (eds.) CAIP 1993. LNCS, vol. 719, pp. 149–156. Springer, Heidelberg (1993)
2. Arsigny, V., Fillard, P., Pennec, X., Ayache, N.: Log-Euclidean metrics for fast and simple calculus on diffusion tensors. Magnetic Resonance in Medicine 56(2), 411–421 (2006)
3. Bailey, R.A., Gower, J.C.: Approximating a symmetric matrix. Psychometrika 55(4), 665–675 (1990)
4. Bullitt, E., Smith, J., Lin, W.: Designed database of MR brain images of healthy volunteers, http://www.insight-journal.org/midas/community/view/21 (accessed May 2010)
5. Canny, J.: A computational approach to edge detection. IEEE Trans. Pattern Anal. Mach. Intell. 8(6), 679–714 (1986)
6. Cumani, A.: Edge detection in multispectral images. Graphical Models and Image Processing 53(1), 40–51 (1991)
7. Descoteaux, M., Audette, M., Chinzei, K., Siddiqi, K.: Bone enhancement filtering: Application to sinus bone segmentation and simulation of pituitary surgery. In: Duncan, J.S., Gerig, G. (eds.) MICCAI 2005. LNCS, vol. 3749, pp. 9–16. Springer, Heidelberg (2005)
8. Evans, A.N., Liu, X.U.: A morphological gradient approach to color edge detection. IEEE Trans. on Image Processing 15(6), 1454–1463 (2006)
9. Frangi, A.F., Niessen, W.J., Vincken, K.L., Viergever, M.A.: Multiscale vessel enhancement filtering. In: Wells, W.M., Colchester, A.C.F., Delp, S.L. (eds.) MICCAI 1998. LNCS, vol. 1496, pp. 130–137. Springer, Heidelberg (1998)
10. Frindel, C., Robini, M., Croisille, P., Zhu, Y.-M.: Comparison of regularization methods for human cardiac diffusion tensor MRI. Medical Image Analysis 13(3), 405–418 (2009)
11. Goodlett, C.B., Fletcher, P.T., Gilmore, J.H., Gerig, G.: Group analysis of DTI fiber tract statistics with application to neurodevelopment. NeuroImage 45(1), 133–142 (2009)
12. Kindlmann, G., Tricoche, X., Westin, C.-F.: Delineating white matter structure in diffusion tensor MRI with anisotropy creases. Medical Image Analysis 11(5), 492–502 (2007)
13. Ming, A., Ma, H.: A blob detector in color images. In: CIVR 2007, pp. 364–370. ACM, New York (2007)
14. Mori, S.: John Hopkins Medical Institute: Laboratory of Brain Anatomical MRI, in vivo human database, http://lbam.med.jhmi.edu/ (accessed February 2009)
15. Niethammer, M., Zach, C., Melonakos, J., Tannenbaum, A.: Near-tubular fiber bundle segmentation for diffusion weighted imaging: Segmentation through frame reorientation. NeuroImage 45(1), 123–132 (2009)
16. Noble, J.A.: Finding corners. Image and Vision Computing 6(2), 121–128 (1988)
17. Pajevic, S., Aldroubi, A., Basser, P.J.: A continuous tensor field approximation of discrete DT-MRI data for extracting microstructural and architectural features of tissue. Journal of Magnetic Resonance 154(1), 85–100 (2002)
18. Shi, J., Tomasi, C.: Good features to track. In: CVPR 1994, pp. 593–600 (1994)
19. Shi, L., Funt, B., Hamarneh, G.: Quaternion color curvature. In: IS&T CIC 2008, pp. 338–341 (2008)
20. Tsai, P., Chang, C.-C., Hu, Y.-C.: An adaptive two-stage edge detection scheme for digital color images. Real-Time Imaging 8(4), 329–343 (2002)
21. Yushkevich, P.A., Zhang, H., Simon, T.J., Gee, J.C.: Structure-specific statistical mapping of white matter tracts. NeuroImage 41(2), 448–461 (2008)

Diffeomorphism Invariant Riemannian Framework for Ensemble Average Propagator Computing*

Jian Cheng[1,2], Aurobrata Ghosh[2], Tianzi Jiang[1], and Rachid Deriche[2]

[1] Center for Computational Medicine, LIAMA, Institute of Automation, Chinese Academy of Sciences, China
[2] Athena Project Team, INRIA Sophia Antipolis – Méditerranée, France
jiancheng@nlpr.ia.ac.cn

Abstract. Background: In Diffusion Tensor Imaging (DTI), Riemannian framework based on Information Geometry theory has been proposed for processing tensors on estimation, interpolation, smoothing, regularization, segmentation, statistical test and so on. Recently Riemannian framework has been generalized to Orientation Distribution Function (ODF) and it is applicable to any Probability Density Function (PDF) under orthonormal basis representation. Spherical Polar Fourier Imaging (SPFI) was proposed for ODF and Ensemble Average Propagator (EAP) estimation from arbitrary sampled signals without any assumption.
Purpose: Tensors only can represent Gaussian EAP and ODF is the radial integration of EAP, while EAP has full information for diffusion process. To our knowledge, so far there is no work on how to process EAP data. In this paper, we present a Riemannian framework as a mathematical tool for such task.
Method: We propose a state-of-the-art Riemannian framework for EAPs by representing the square root of EAP, called wavefunction based on quantum mechanics, with the Fourier dual Spherical Polar Fourier (dSPF) basis. In this framework, the exponential map, logarithmic map and geodesic have closed forms, and weighted Riemannian mean and median uniquely exist. We analyze theoretically the similarities and differences between Riemannian frameworks for EAPs and for ODFs and tensors. The Riemannian metric for EAPs is diffeomorphism invariant, which is the natural extension of the affine-invariant metric for tensors. We propose Log-Euclidean framework to fast process EAPs, and Geodesic Anisotropy (GA) to measure the anisotropy of EAPs. With this framework, many important data processing operations, such as interpolation, smoothing, atlas estimation, Principal Geodesic Analysis (PGA), can be performed on EAP data.
Results and Conclusions: The proposed Riemannian framework was validated in synthetic data for interpolation, smoothing, PGA and in real data for GA and atlas estimation. Riemannian median is much robust for atlas estimation.

1 Introduction

Diffusion MRI (dMRI) is the unique non-invasive technique to study the white matter in vivo by probing the water diffusion. Ensemble Average Propagator (EAP) $P(\mathbf{R})$ is

* This work was partially granted by the French Government Award Program, the Natural Science Foundation of China (30730035,81000634), the National Key Basic Research and Development Program of China (2007CB512305), the French ANR "Neurological and Psychiatric diseases" NucleiPark and the France-Parkinson Association.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 98–106, 2011.

the Probability Density Function (PDF) in R^3 to describe the diffusion process of water molecules. Under the narrow pulse assumption, $P(\mathbf{R})$ is the Fourier Transform of signal attenuation $E(\mathbf{q})$, i.e. $P(\mathbf{R}) = \int E(\mathbf{q}) \exp(-2\pi i \mathbf{q}^T \mathbf{R}) d\mathbf{q},$. $\mathbf{R} = R\mathbf{r}$ is the displacement vector in $\mathbf{R}$-space, $\mathbf{q} = q\mathbf{u}$ is the reciprocal vector in $\mathbf{q}$-space, $\mathbf{u}$ and $\mathbf{r}$ are unit vectors.

In Diffusion Tensor Imaging (DTI) [4], $P(\mathbf{R})$ is a Gaussian PDF parameterized by its covariance matrix, i.e. tensor. Gaussian distribution is well studied in Information Geometry theory [1]. Riemannian framework [14,13,9] based on Information Geometry was proposed to process the tensor data. Log-Euclidean framework [2] is to approximate Riemannian framework and efficiently process tensors. Riemannian metric for tensors is affine-invariant, while Euclidean metric is not [14]. Riemannian framework and Log-Euclidean framework have proved useful in many works on tensor estimation, interpolation, filtering, segmentation, registration, statistical analysis [14,13,9,2,11].

Since DTI can not deal with complex diffusion process, many methods beyond DTI were proposed to estimate EAP or Orientation Distribution Function (ODF), which is the radial integral of EAP. Recently Spherical Polar Fourier Imaging (SPFI) [3,8] was proposed to analytically and robustly estimate both ODF and EAP by representing $E(\mathbf{q})$ using Spherical Polar Fourier (SPF) basis. ODF and EAP are both PDFs. Thus it is possible to extend the Riemannian framework from tensors to ODFs and EAPs based on Information Geometry theory. [7,6] and [12] recently developed separately and in parallel a nonparametric Riemannian framework to process ODF data. To our knowledge, recent works for EAPs mainly focus on EAP estimation, and there is no work on how to process EAP data.

In this paper, we first propose the Riemannian framework for EAPs based on the theoretical results in [7,6]. We analyze theoretically the similarities and differences between Riemannian frameworks for EAPs and for ODFs and tensors. For instance the isotropic EAP is not unique, which brings a different definition of Geodesic Anisotropy (GA) as well as the Log-Euclidean framework. Then we implement the framework using the orthonormal basis in SPFI [8]. SPFI provides dSPF basis which analytically obtains GA and Log-Euclidean framework. For the application part, we propose GA for measuring the anisotropy of water diffusion, and some Riemannian operations for EAP computing, such as interpolation, PGA, smoothing, atlas estimation.

2 Riemannian Framework for EAPs

In this section we will show Riemannian framework for EAPs and also analyze and compare it with the Riemannian framework for ODFs [8] and highlight the new problems faced and their solutions, since the EAP has different properties from the ODF.

Parametric Family. In quantum mechanics, the square root of the probability of finding the subject at a certain time and position is called as wavefunction. Analogously, the square root of EAP, denoted by $\psi(\mathbf{R})$, is also called as wavefunction. $\psi(\mathbf{R}) = \sqrt{P(\mathbf{R})} \geq 0$. Let $\{B_i(\mathbf{R})\}_{i \in N}$ is a given orthonormal basis function set in R^3 which could sparsely represent $\psi(\mathbf{R})$. then $\psi(\mathbf{R})$ could be represented by finite linear combination of $\{B_i(\mathbf{R})\}$,i.e. $\psi(\mathbf{R}|\boldsymbol{c}) = \sum_{i=1}^{K} c_i B_i(\mathbf{R})$, where $\boldsymbol{c} = (c_1, c_2, ..., c_K)^T$ is called the Riemannian Coordinate [7]. In practice, we always could choose a large enough K to represent $\psi(\mathbf{R})$, so the assumption here is very weak. Then the Parametric Family or EAP

space, called PF_K, could be formulated in (1) [7,6]. Based on PF_K in (1), the Parameter Space (PS), denoted as PS_k, is a subset of sphere S^{K-1}. $PS_K = \{\boldsymbol{c} \mid \|\boldsymbol{c}\| = \sum_{i=1}^{K} c_i^2 = 1,\ \sum_{i=1}^{K} c_i B_i(\mathbf{R}) \geq 0,\ \forall \mathbf{R} \in R^3\}$. In Fig. 1(A), we visualize PS_3 as an example where $\{B_i\}$ were chosen as three piecewise constant orthonormal functions, i.e. three bins for the histogram. Since the formulation in (1) has been well studied in [7,6] and successfully applied for ODFs, *in the following we will list the results in [7,6] and modify them if necessary so that they can be applied to EAPs*.

$$PF_K = \left\{ P(\mathbf{R}|\boldsymbol{c}) = \psi(\mathbf{R}|\boldsymbol{c})^2 :\ \psi(\mathbf{R}|\boldsymbol{c}) = \sum_{i=1}^{K} c_i B_i(\mathbf{R}) \geq 0,\ \int P(\mathbf{R}|\boldsymbol{c})\mathrm{d}\mathbf{R} = \|\boldsymbol{c}\|^2 = 1,\ \forall \mathbf{R} \in R^3 \right\} \quad (1)$$

Riemannian Metric. [7] proved that based on formulation in (1), the Riemannian metric [1] is $g_{ij} = \int \frac{\partial \sqrt{P(\mathbf{R}|\boldsymbol{c})}}{\partial c_i} \frac{\partial \sqrt{P(\mathbf{R}|\boldsymbol{c})}}{\partial c_j} \mathrm{d}\mathbf{R} = 4\delta_{ij}$. The constant 4 could be ignored, then PS_K is a subset of the unit sphere S^{K-1} and the metric is just the Euclidean metric in the sphere. The geodesic distance between two points $P(\mathbf{R}|\boldsymbol{c})$ and $P(\mathbf{R}|\boldsymbol{c}')$ will be the angle between them, i.e. $\arccos(\boldsymbol{c}^T\boldsymbol{c}') = \arccos(\int_{R^3} \psi(\mathbf{R}|\boldsymbol{c})\psi(\mathbf{R}|\boldsymbol{c}')\mathrm{d}\mathbf{R})$. Denote $\boldsymbol{v_c}$ is the tangent vector at $\boldsymbol{c}$ towards $\boldsymbol{c}'$, then the geodesic, exponential map, logarithmic map all have closed forms. See Fig. 1(A).

$$\text{Geodesic: } \gamma(t) : P(\mathbf{R}|\boldsymbol{c}(t)), \quad \text{where } \boldsymbol{c}(t) = Exp_{\boldsymbol{c}}(tLog_{\boldsymbol{c}}(\boldsymbol{c}')) \quad (2)$$

$$\text{Exponential Map: } Exp_{\boldsymbol{c}}(\boldsymbol{v_c}) = \boldsymbol{c}' = \boldsymbol{c}\cos\varphi + \frac{\boldsymbol{v_c}}{\|\boldsymbol{v_c}\|}\sin\varphi, \text{ where } \varphi = \|\boldsymbol{v_c}\| \quad (3)$$

$$\text{Logarithmic Map: } Log_{\boldsymbol{c}}(\boldsymbol{c}') = \boldsymbol{v_c} = \frac{\boldsymbol{c}' - \boldsymbol{c}\cos\varphi}{\|\boldsymbol{c}' - \boldsymbol{c}\cos\varphi\|}\varphi, \text{ where } \varphi = \arccos(\boldsymbol{c}^T\boldsymbol{c}') \quad (4)$$

Please note that if we change the orthonormal basis $\{B_i(\mathbf{R})\}$ the exponential map and logarithmic map are invariant under a change of basis matrix, and the geodesic $\gamma(t)$ is also invariant. Different orthonormal basis will obtain equivalent Riemannian framework. The final results of the following Riemannian operations will be the same if the approximation error is negligible. That is why here we consider the formulation using orthonormal basis in [7] instead of the formulation using histogram in [12].

Properties of Parameter Space. [7] showed that PS_K for EAPs is a closed convex set of S^{K-1} and it is contained in a convex cone with 90^o. See Fig. 1(C). [7] also proved for ODFs that "*the projection of any* $\mathbf{c} \in PS_K$ *on the Riemannian Coordinate* $\mathbf{u}$ *of the isotropic ODF should be more than* $\frac{1}{\sqrt{4\pi}}$, if ODFs are less than 1", i.e. $\boldsymbol{c}^T\boldsymbol{u} = \cos(\boldsymbol{c}, \boldsymbol{u}) = \int_{\chi} \sqrt{p}(x|\boldsymbol{c})\frac{1}{\sqrt{|\chi|}}\mathrm{d}x > \frac{1}{\sqrt{|\chi|}}\int_{\chi} p(x|\boldsymbol{c})\mathrm{d}x = \frac{1}{\sqrt{4\pi}}$. However, It is specific for ODFs. Since the ODF and EAP are both continuous function, it may be more than 1 in some area although the integral in the whole domain is 1. If $P(\mathbf{R}|\boldsymbol{c}) > 1$, then $\sqrt{P(\mathbf{R})} < P(\mathbf{R})$ and the conclusion will be problematic. Normally for typical ODFs the values are always less than 1. However, we found that the values of EAPs are normally much larger than 1 because the diffusion time τ is small. Moreover, the isotropic ODF is unique. While

the isotropic EAP is not unique, because $P(\mathbf{R})$ is isotropic if $P(R\mathbf{r}) = F(R)$, $\forall \mathbf{r} \in S^2$ and $F(R)$ could be any positive definite function.

Geodesic Anisotropy (GA), Log-Euclidean Framework. GA for ODFs is defined as the geodesic distance between the ODF and the isotropic one. Log-Euclidean framework for ODFs could be used to approximate Riemannian framework by projecting the ODFs onto the tangent space of the isotropic ODF [7]. However, the isotropic EAP is not unique as discussed above. Please note that the isotropic tensor is not unique in DTI neither. All tensors with three equal eigenvalues are isotropic, i.e. $D = [\lambda, \lambda, \lambda]$. GA for tensors was defined as the geodesic distance from the nearest isotropic tensor [9]. And the identity tensor was chosen and fixed [2] for Log-Euclidean framework. Analogously we define the GA as the geodesic distance between the EAP and the nearest isotropic EAP. It could be proved that *for any given EAP with coordinate* $\mathbf{c}$*, the Riemannian coordinate of the nearest isotropic EAP with the coordinate* $\mathbf{u}$*, is just the normalized version of the isotropic part of* $\mathbf{c}$. We omit the rigorous proof due to page limitations. Furthermore we can also fix a typical isotropic Gaussian distribution for all EAPs. Then Log-Euclidean framework could be obtained by projecting EAPs onto the tangent space of the fixed isotropic EAP. The projection diffeomorphism is defined as $F(\boldsymbol{c}) = Log_{\boldsymbol{u}}(\boldsymbol{c})$ where $\boldsymbol{u}$ is the Riemannian coordinate for the fixed isotropic EAP [7]. See Fig. 1(B).

Weighted Mean, Weighted Median and Principal Geodesic Analysis (PGA). Given N EAPs $f_1, f_2, ..., f_N$ in PS_K and the weight vector $\boldsymbol{w} = (w_1, w_2, ..., w_N)^T$ with $\sum_{i=1}^{N} w_i = 1$, $w_i \in [0, 1]$, the weighted Riemannian mean $\mu_{\boldsymbol{w}}$ is defined as the minimizer of the weighted sum of squared distances [14,13,9,7,5]. And the weighted Riemannian median $m_{\boldsymbol{w}}$ is defined as the minimizer of the weighted sum of distances [10,6].

$$\mu_{\boldsymbol{w}} = \arg\min_{f \in PS_K} \sum_{i=1}^{N} w_i d(f, f_i)^2 \qquad m_{\boldsymbol{w}} = \arg\min_{f \in PS_K} \sum_{i=1}^{N} w_i d(f, f_i) \tag{5}$$

[7,6] proved that the weighted Riemannian mean and weighted Riemannian median uniquely exist in the manifold PS_K based on the general results in [5,10]. And they can be obtained efficiently from gradient descent method on the manifold. Normalized Euclidean mean is chosen as the initialization, which makes the methods converge fast. See Algorithm 1 and 2. For Log-Euclidean framework, the Riemannian mean has closed form, i.e. $\mu_{\boldsymbol{w}} = F^{-1}\left(\sum_{i=1}^{N} w_i F(f_i)\right)$ [7]. When Riemannian mean μ of $\{f_i\}$ is obtained, we can find some Principal Components (PCs) based on Principal Geodesic Analysis (PGA) [11] by eigen-decomposition of the covariance matrix at μ. If $\boldsymbol{v}$ is one eigenvector with eigenvalue λ_k, then the PC will be $Exp_{\mu}(\alpha_k \boldsymbol{v})$, where $\alpha_k \in R$ is the mode variation. Normally α_k is chosen in $[-3\sqrt{\lambda_k}, 3\sqrt{\lambda_k}]$ [11]. PGA has been proposed for tensor analysis [11] and for ODF analysis [12]. However PGA in [12] decompose a covariance matrix in a very high dimension, which is much inefficient. Actually we can perform PGA on ODFs and EAPs just in a low dimension via orthonormal basis representation, which will get the same final results as we discussed above.

Diffeomorphism Invariant: A Natural Extension. What is the connection between the proposed metric and the previous metric for tensors [14,13,9]? Actually they are all the Fisher Information metric in Information Geometry. The Fisher information metric is actually diffeomorphism invariant. Denote $P_X(x|\boldsymbol{c})$ is a PDF on

Algorithm 1: Weighted Riemannian Mean

Input: $f_1, ..., f_N \in PS^K$, $\boldsymbol{w} = (w_1, ..., w_N)^T$

Output: $\mu_{\boldsymbol{w}}$, the *Weighted Mean*.

Initialization: $\mu_{\boldsymbol{w}}^{(0)} = \frac{\sum_{i=1}^N w_i f_i}{\|\sum_{i=1}^N w_i f_i\|}$, $k = 0$

Do

$v_{\mu_{\boldsymbol{w}}^{(k)}} = \sum_{i=1}^N w_i Log_{\mu_{\boldsymbol{w}}^{(k)}}(f_i)$

$\mu_{\boldsymbol{w}}^{(k+1)} = Exp_{\mu_{\boldsymbol{w}}^{(k)}}(v_{\mu_{\boldsymbol{w}}^{(k)}})$

$k = k + 1$

while $\|v_{\mu_{\boldsymbol{w}}^{(k)}}\| > \varepsilon$

Algorithm 2: Weighted Riemannian Median

Input: $f_1, ..., f_N \in PS^K$, $\boldsymbol{w} = (w_1, ..., w_N)^T$

Output: $\mu_{\boldsymbol{w}}$, the *Weighted Median*.

Initialization: $m_{\boldsymbol{w}}^{(0)} = \frac{\sum_{i=1}^N w_i f_i}{\|\sum_{i=1}^N w_i f_i\|}$, $k = 0$

Do

$v_{m_{\boldsymbol{w}}^{(k)}} = \sum_{i=1}^N \frac{w_i / d(m_{\boldsymbol{w}}^k, f_i)}{\sum_{j=1}^N w_j / d(m_{\boldsymbol{w}}^k, f_j)} Log_{m_{\boldsymbol{w}}^{(k)}}(f_i)$

$m_{\boldsymbol{w}}^{(k+1)} = Exp_{m_{\boldsymbol{w}}^{(k)}}(v_{m_{\boldsymbol{w}}^{(k)}})$

$k = k + 1$

while $\|v_{\mu_{\boldsymbol{w}}^{(k)}}\| > \varepsilon$

domain χ, $h : \chi \mapsto \chi$ is a diffeomorphism. $P_Y(y|\boldsymbol{c})$ is the PDF under g, i.e. $y = h(x)$. Then $P_Y(y) = |\nabla h^{-1}(y)| P_X(h^{-1}(y))$. By considering $dy = |\nabla h(y)| dx$, we have $g_{ij} = \int_\chi \frac{\partial \sqrt{P_Y(y|\boldsymbol{c})}}{\partial c_i} \frac{\partial \sqrt{P_Y(y|\boldsymbol{c})}}{\partial c_j} \mathrm{d}y = \int_\chi \frac{\partial \sqrt{P_X(x|\boldsymbol{c})}}{\partial c_i} \frac{\partial \sqrt{P_X(x|\boldsymbol{c})}}{\partial c_j} \mathrm{d}x$, which means the metric g_{ij} is diffeomorphism invariant. So for any two given ODFs or EAPs, the Riemannian distance between them is diffeomorphism invariant. If we constrain the PDF $P_X(x|\mathbf{D})$ as a Gaussian distribution parameterized by tensor $\mathbf{D}$, then it could be easily proved that $P_Y(y|\mathbf{D})$ is still Gaussian if and only if $h(x)$ is an affine transform, i.e. $h(x) = Ax$, A is a nonsingular matrix. In this sense the proposed Riemannian metric for EAPs is actually a natural extension of previous affine-invariant metric for tensors [14,13,9]. The diffeomorphism invariant metric is probably useful in registration.

3 Implementation via Spherical Polar Fourier Imaging (SPFI)

In theory, the Riemannian framework for EAPs could be implemented by any orthonormal basis. However, so far there is no direct way to estimate the Riemannian Coordinate $\boldsymbol{c}$ from DWI signals $E(\mathbf{q})$. Existing methods only estimate EAPs. SPFI is a fast, regularized, robust method to estimate EAPs without any assumption [3,8]. In SPFI, $E(\mathbf{q})$ is represented by an orthonormal basis $\{B_{nlm}(\mathbf{q})\}$ in formula (6), named SPF basis, where ζ is a fixed scale parameter, $R_n(q)$ is the Gaussian-Laguerre function and $Y_l^m(\mathbf{u})$ is the l order m degree Spherical Harmonic. [8] proved that the EAP $P(\mathbf{R})$ could be analytically obtained in formula (8) from the same coefficients $\{a_{nlm}\}$, where ${}_1F_1$ is the confluent hypergeometric function. $\{D_{nlm}(\mathbf{R})\}$ is actually an orthonormal basis in $\mathbf{R}$-space, called Fourier dual Spherical Polar Fourier (dSPF) basis, because of Parseval's theorem, i.e. $\delta_{nlm}^{n'l'm'} = \int_{R^3} B_{nlm}(\mathbf{q}) B_{n'l'm'}(\mathbf{q}) \mathrm{d}\mathbf{q} = \int_{R^3} D_{nlm}(\mathbf{R}) D_{n'l'm'}(\mathbf{R}) \mathrm{d}\mathbf{R}$. So SPFI actually provides two orthonormal basis. One is $\{B_{nlm}(\mathbf{q})\}$ in $\mathbf{q}$ space for $E(\mathbf{q})$ and the other one is $\{D_{nlm}(\mathbf{R})\}$ in $\mathbf{R}$ space for $P(\mathbf{R})$.

$$E(q\mathbf{u}) = \sum_{n=0}^{N} \sum_{l=0}^{L} \sum_{m=-l}^{l} a_{nlm} R_n(q) Y_l^m(\mathbf{u}) \qquad B_{nlm}(\mathbf{q}) = R_n(q) Y_l^m(\mathbf{u}) \tag{6}$$

$$R_n(q) = \kappa_n(\zeta) \exp\left(-\frac{q^2}{2\zeta}\right) L_n^{1/2}\left(\frac{q^2}{\zeta}\right) \qquad \kappa_n(\zeta) = \left[\frac{2}{\zeta^{3/2}} \frac{n!}{\Gamma(n + 3/2)}\right]^{1/2} \tag{7}$$

$$P(R\mathbf{r}) = \sum_{n=0}^{N} \sum_{l=0}^{L} \sum_{m=-l}^{l} a_{n,l,m} F_{nl}(R) Y_l^m(\mathbf{r}) \qquad D_{nlm}(\mathbf{R}) = F_{nl}(R) Y_l^m(\mathbf{r}) \tag{8}$$

$$F_{nl}(R) = \frac{\zeta^{0.5l+1.5}\pi^{l+1.5}R_0^l\kappa_n(\zeta)}{(-1)^{l/2}\Gamma(l+1.5)}\sum_{i=0}^{n}\binom{n+0.5}{n-i}\frac{(-1)^i}{i!}2^{0.5l+i+1.5}\Gamma(0.5l+i+1.5)_1F_1(\frac{2i+l+3}{2};l+\frac{3}{2};-2\pi^2R^2\zeta) \tag{9}$$

After obtaining the continuous function $P(\mathbf{R})$ in (8) represented by $\{D_{nlm}(\mathbf{R})\}$, we can get many discrete samples $\{\psi(\mathbf{R}_i)\}$ from the wavefunction $\psi(\mathbf{R}) = \sqrt{P(\mathbf{R})}$. Then similarly with [7], a least square fitting using the same basis $\{D_{nlm}(\mathbf{R})\}$ is applied to estimate the Riemannian coordinates $\boldsymbol{c}$ from these samples $\{\psi(\mathbf{R}_i)\}$. After the Riemannian Coordinate $\boldsymbol{c} = \{c_{nlm}\}$ is estimated, GA is the distance between $\boldsymbol{c}$ and the nearest isotropic EAP with coordinate $\boldsymbol{u}$, which is the normalized version of isotropic part of $\boldsymbol{c}$. With the basis $\{D_{nlm}(\mathbf{R})\}$, the isotropic part of $\{c_{nlm}\}$ is analytically obtained as $\{c_{nlm}\delta^{00}_{lm}\}$. The isotropic EAP with coordinate $(1, 0, ..., 0)^T$, which is an isotropic Gaussian based on SPF basis in (6), was chosen and fixed for Log-Euclidean framework. Thus Log-Euclidean framework can be obtained analytically from the given Riemannian Coordinate thanks to dSPF basis. Moreover we can perform many data processing algorithms, such as interpolation, smoothing etc. [7,12,14].

4 Applications and Experiments

Experiments on Diffeomorphism Invariance. Here we give a simple example to demonstrate the diffeomorphism invariance for the Riemannian metric using a affine transform. Mixture of tensors model $P(\mathbf{R}) = w_1P(\mathbf{R}|\mathbf{D}_1) + w_2P(\mathbf{R}|\mathbf{D}_2)$ was used to generate the synthetic data. $\mathbf{D}_1$ and $\mathbf{D}_2$ have the same eigenvalues $[1.7, 0.3, 0.3]\times10^{-3}mm^2/s$ with different directions. We estimated Riemannian Coordinates and computed the Riemannian distance between the EAP in the center and the EAPs in other voxels. Afterwards a given affine transform A was performed on the mixture model, i.e. $(A \circ P)(\mathbf{R}) = w_1P(\mathbf{R}|A^T\mathbf{D}_1A) + w_2P(\mathbf{R}|A^T\mathbf{D}_2A)$. With this transformed model, we calculated again the distance maps. In Fig. 1, (C1) shows the original EAPs at $15\mu m$ and (C2) shows the transformed EAPs. The distance maps were used to color the glyphs and were set as the background. It is clear that the distance is invariant under the transform. This experiment showed that Riemannian metric for EAPs is diffeomorphism invariant, which is a natural extension of affine-invariant metric for tensors.

Interpolation and PGA. We demonstrate the Lagrange interpolation of EAPs in 1D and 2D in Fig. 1, where the weights are coded by the spatial distance [7]. For 1D case, two EAPs were generated from tensors with eigenvalues $[1.7, 0.3, 0.3]\times10^{-3}mm^2/s$ and $[0.3, 0.3, 0.3]\times10^{-3}mm^2/s$. For 2D case, one EAP was generated from $[0.7, 0.7, 0.7]\times10^{-3}mm^2/s$, the other three from $[1.7, 0.3, 0.3]\times10^{-3}mm^2/s$ with 3 orthogonal directions. Fig. 1(D1) and (D3) show the results for Riemannian framework and (D2) and (D4) for Log-Euclidean framework. EAPs were visualized at $15\mu m$ and the glyphs were colored by GA values. It is clear that the interpolation from Riemannian framework and Log-Euclidean framework have very similar results. The first two principal components of EAP field in Fig. 1(D3) based on PGA are shown in (E1) and (E2). The two components showed clearly the three orthogonal directions as well as the mean EAP.

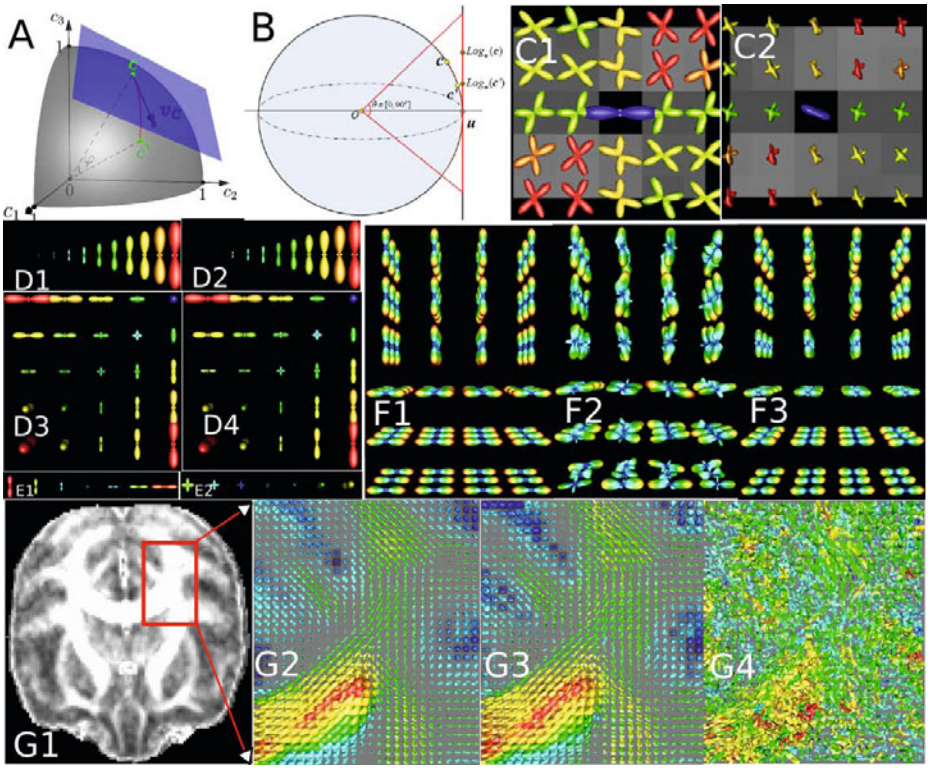

Fig. 1. (A): an example of PS_3; (B): properties of PS_K; (C1,C2): original EAPs and transformed EAPs; (D1,D2): 1D interpolation via Riemannian framework and Log-Euclidean framework; (D3,D4): 2D interpolation via Riemannian framework and Log-Euclidean framework; (E1,E2): the first two PCs in PGA for EAPs in D3; (F1,F2,F3): original EAPs, noisy EAPs, smoothed EAP; (G1): GA map; (G2,G3): original and noisy atlas; (G4): noisy EAPs from one subject;

Smoothing. Actually all filtering algorithms in [14] for tensors and in [6,12] for ODFs could be generalized into EAP field. Here we just demonstrate a simple Gaussian smoothing method. In this method the filtered EAP in every voxel is the weighted mean/median of EAPs in a given neighborhood. The weights were chosen from spatial Gaussian kernel. Fig. 1(F1) shows the ground truth EAP profiles generated from tensor model with eigenvalues $[1.7, 0.3, 0.3] \times 10^{-3} mm^2/s$ EAPs in (F2) were estimated from SPFI from single shell DWI signals with $b = 1500 s/mm^2$ and $SNR = 5$ [8]. The results of Gaussian smooth using Riemannian mean were shown in Fig. 1(F3). The deviation of the Gaussian kernel was set as 1 voxel.

GA and Riemannian Median Atlas Estimation. Riemannian median is more robust than Riemannian mean and Euclidean mean, which makes it more appropriate for atlas estimation [10,6]. We construct an EAP atlas from five monkey data, with three b values ($500/1500/3000 s/mm^2$), 30 directions at each shell. Since so far there is no common registration method for EAP data and it is also not our focus here, we just use a naive

way to align the EAP data. All DWIs from 5 subjects were aligned via affine registration to a template made by non-diffusion weighted images. The affine transformation was used to rotate the gradient directions for each subject through the finite strain method. After that, EAP images were estimated via SPFI from registrated DWIs and reorientated gradient directions. The atlas in every voxel was estimated as the Riemannian median of the five EAP images. And GA map in Fig. 1(G1) was estimated from the atlas. To test the robustness, we add Rician noise with $SNR = 10$ to the DWIs of one subject, then re-estimate the atlas. Fig. 1 showed the noisy EAPs (G4) and the EAPs of the two atlases (G2,G3) in the red region in (G1). The atlas from noisy data is much similar with the one from the real data, which validates the robustness.

5 Conclusion

In this paper, we propose a diffeomorphism invariant Riemannian framework for EAP computing, which is a natural extension of previous Riemannian framework for tensors and ODFs. In the proposed framework, exponential map, logarithmic map and geodesic have closed forms, weighted Riemannian mean and median uniquely exist and could be calculated efficiently. We analyze theoretically the similarities and differences between Riemannian frameworks for EAPs and for ODFs and tensors. The isotropic EAP is not unique, which brings a different definition of GA and the Log-Euclidean framework compare to the Riemannian framework for ODFs. The Riemannian framework for EAPs is more like the Riemannian framework for tensors, since we need to define what is the nearest isotropic EAP. And if we constrain the EAP as a Gaussian, the diffeomorphism invariance property will be the affine-invariance for tensors. The proposed Riemannian framework was implemented via SPFI by representing the wavefunction with dSPF basis, which has closed forms for GA and Log-Euclidean framework. For application part, we propose GA to measure anisotropy, Lagrange interpolation, Gaussian smoothing, PGA and median atlas estimation for EAP computing.

References

1. Amari, S., Nagaoka, H.: Methods of Information Geometry. American Mathematical Society, USA (2000)
2. Arsigny, V., Fillard, P., Pennec, X., Ayache, N.: Log-Euclidean metrics for fast and simple calculus on diffusion tensors. Magnetic Resonance in Medicine 56, 411–421 (2006)
3. Assemlal, H.E., Tschumperlé, D., Brun, L.: Efficient and robust computation of PDF features from diffusion MR signal. Medical Image Analysis 13, 715–729 (2009)
4. Basser, P.J., Mattiello, J., LeBihan, D.: MR diffusion tensor spectroscopy and imaging. Biophysical Journal 66, 259–267 (1994)
5. Buss, S.R., Fillmore, J.P.: Spherical averages and applications to spherical splines and interpolation. ACM Transactions on Graphics 20, 95–126 (2001)
6. Cheng, J., Ghosh, A., Jiang, T., Deriche, R.: Riemannian median and its applications for orientation distribution function computing. In: ISMRM (2010)
7. Cheng, J., Ghosh, A., Jiang, T., Deriche, R.: A riemannian framework for orientation distribution function computing. In: Yang, G.-Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (eds.) MICCAI 2009. LNCS, vol. 5761, pp. 911–918. Springer, Heidelberg (2009)

8. Cheng, J., Ghosh, A., Jiang, T., Deriche, R.: Model-free and analytical EAP reconstruction via spherical polar fourier diffusion MRI. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. 6361, pp. 590–597. Springer, Heidelberg (2010)
9. Fletcher, P.T.: Statistical Variability in Nonlinear Spaces Application to Shape Analysis and DT-MRI. Ph.D. thesis, University of North Carolina (2004)
10. Fletcher, P.T., Venkatasubramanian, S., Joshi, S.: The geometric median on riemannian manifolds with application to robust atlas estimation. NeuroImage 45, S143–S152 (2009)
11. Fletcher, P., Joshi, S.: Riemannian geometry for the statistical analysis of diffusion tensor data. Signal Processing 87, 250–262 (2007)
12. Goh, A., Lenglet, C., Thompson, P., Vidal, R.: A nonparametric Riemannian framework for processing high angular resolution diffusion images and its applications to ODF-based morphometry. NeuroImage (2011)
13. Lenglet, C., Rousson, M., Deriche, R.: Statistics on the manifold of multivariate normal distributions theory and application to diffusion tensor MRI processingy. Journal of Mathematical Imaging and Vision 25(3), 423–444 (2006)
14. Pennec, X., Fillard, P., Ayache, N.: A riemannian framework for tensor computing. International Journal of Computer Vision 66, 41–66 (2006)

Assessment of Bias for MRI Diffusion Tensor Imaging Using SIMEX

Carolyn B. Lauzon[1,2], Andrew J. Asman[1], Ciprian Crainiceanu[3], Brian C. Caffo[3], and Bennett A. Landman[1,2]

[1] Department of Electrical Engineering, Vanderbilt University, Nashville, TN, USA
[2] Institute of Imaging Science, Vanderbilt University, Nashville, TN, USA
[3] Department of Biostatistics, Johns Hopkins University Bloomberg School of Public Health, Baltimore, MD, USA

Abstract. Diffusion Tensor Imaging (DTI) is a Magnetic Resonance Imaging method for measuring water diffusion *in vivo*. One powerful DTI contrast is fractional anisotropy (FA). FA reflects the strength of water's diffusion directional preference and is a primary metric for neuronal fiber tracking. As with other DTI contrasts, FA measurements are obscured by the well established presence of bias. DTI bias has been challenging to assess because it is a multi-variable problem including SNR, six tensor parameters, and the DTI collection and processing method used. SIMEX is a modern statistical technique that estimates bias by tracking measurement error as a function of added noise. Here, we use SIMEX to assess bias in FA measurements and show the method provides; i) accurate FA bias estimates, ii) representation of FA bias that is data set specific and accessible to non-statisticians, and iii) a first time possibility for incorporation of bias into DTI data analysis.

Keywords: DTI, FA, bias, SIMEX, parameter estimation, bias correction, diffusion, tensor, imaging.

1 Introduction

Radiological practice appears to be on the cusp of reliably and routinely using imaging biomarkers as surrogates for traditional tissue biopsy (e.g., virtual biopsy). One promising magnetic resonance imaging (MRI) technique, diffusion tensor imaging (DTI) provides unique, non-invasive contrasts sensitive to tissue micro-architecture, which can be correlated with cellular organization – for example, as seen through histological analysis [1]. However, quantification of observed data is an *estimation* process and numeric values must be viewed in light of their distributional properties. Extensive research has been invested in characterizing the impacts of noise and study design on DTI contrasts[2, 3], and it is well recognized that both lead to changes in precision (*variance*) as well as accuracy (*bias*) of diffusion derived measures. Much progress has been made on variability analysis in DTI using Monte Carlo and boot-strap techniques [4], but quantitative *post hoc* assessment of the bias in diffusion estimated parameters has remained elusive.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 107–115, 2011.

Bias is explicitly defined as the difference between the expectated value of an experimentally observed measure, $\mathbf{E}(\Theta_{obs})$, and the true value of that measure, Θ_{truth}; bias = $\mathbf{E}(\Theta_{obs}) - \Theta_{truth}$. As seen, bias causes the statistical average of a measured DTI contrast to converge on the wrong value. Also, differing propensities for motion and/or underlying anatomy could lead to systematically different tensor bias between disease and healthy subgroups, thus statistically impacting analyses. Bias in DTI is a multi-variable function including imaging noise, field in homogeneity, DTI experimental parameters, and the underlying biological diffusion parameters. No analytic equation for bias in DTI is known and no published method to account for bias in empirical DTI data has been tried. Herein, we present a statistical approach to estimate the expected level of bias in DTI contrasts on a voxel-wise basis.

In this initial evaluation, we focus on quantifying bias in the commonly studied and important DTI contrast fractional anisotropy (FA) as measured using the traditional log-linear minimum mean-squared error (LLMMSE) tensor estimation framework. Significant bias in FA has been well documented both experimentally and through computational studies[3, 5]. Within a clinical and research SNR range (~15:1 – 40:1), FA bias is generally a monotonic decreasing function, with lower FA values tending to have larger bias than higher FA values.

To estimate FA bias, we test the SIMulation EXtrapolation method (SIMEX) which emerged from the statistical literature as a way to compensate for bias due to additive noise within a regression framework[6, 7]. However, it can also be viewed as a general procedure to extrapolate the degree of bias in an estimate given noisy data by studying the sensitivity of the model fitting procedure when *additional synthetic* noise is added. Two key assumptions of SIMEX are (i) the bias as a function of SNR is monotonic and smooth and (ii) the noise level of the observed data is well understood and estimated. Fortunately, (i) is generally holds and for (ii), noise estimation has been well-studied in the DTI literature and several robust methods exist to estimate the noise field given empirical data [8-10].

This manuscript is organized as follows. First, we briefly review the diffusion tensor formulation and describe how SIMEX bias estimation can be applied to bias estimation for FA. Second, we present a simulation based on empirical data which enables knowledge of the true bias level for test comparisons with the SIMEX estimated bias. Third, we present single and multi-voxel demonstrations of SIMEX estimation of bias in FA. We close with a brief critique and discussion of potential applications for this approach and avenues for continuing research.

2 Theory

2.1 Diffusion Tensor Imaging

In DTI, magnetic gradients (**g**) are used to label the relative position of water at the start and end of an experiment. Given the random thermal motions of water, the labels add incoherently and result in an attenuated, diffusion weighted image (dwi). In general, the greater probabilistic movement, the greater the signal attenuation. In the tensor model, the dwi will be a function of the unweighted signal intensity, symbolized as S_o, diffusion direction and magnitude (encoded in diffusion tensor **D),** the experimental timing and gradient strength (encoded in parameter *b*), and the direction of the gradient applied (encoded in **g**). For a single voxel we write

$$\mathrm{dwi_j} = \mathrm{S_o}\exp\left(-b\mathbf{g_j^T}\cdot\mathbf{D}\cdot\mathbf{g_j}\right). \quad (1)$$

Here, *j* indexes the gradient directions. We define the observed data for a single voxel, $\vec{\mathrm{X}}_{\mathrm{obs}}$, as a vector with elements x_i, representing dwi_j for *j =1,2,…m,* and $\mathrm{S_o}$.

$$\vec{\mathrm{X}}_{\mathrm{obs}} = \left\{\mathrm{dwi}\,|_{j=1}^{m}, \mathrm{S_o}\right\} = \left\{\mathrm{x}_i|_{i}^{m+1}\right\}. \quad (2)$$

D is a 3 x 3 symmetric matrix whose Eigenvalues, λ_1, λ_2, and λ_3, describe the three principal axis of diffusion. Herein, we consider the FA of the tensor which is the normalized standard deviation of the Eigenvalues and reflects the relative degree of orientation (i.e., 0: no orientation preference = isotropic diffusion, 1: infinite directional preference = 1-D diffusion).

$$\mathrm{FA} = \sqrt{\frac{3}{2}\frac{(\lambda_1-\bar{\lambda})^2+(\lambda_2-\bar{\lambda})^2(\lambda_3-\bar{\lambda})^2}{\lambda_1^2+\lambda_2^2+\lambda_3^2}}. \quad (3)$$

2.2 SIMEX Applied to DTI

The premise of SIMEX is simple. Bias can be understood by adding noise to data in incremental amounts and measuring the resulting contrast. The trend in the contrast with added noise should enable prediction of the contrast with 'removed' noise. For simplicity, the following description assumes a single DTI experiment is performed (e.g. single subject, one DTI dataset) and describes the single-voxel case, though the analysis can be extended to multiple experiments and multiple voxels.

Let a truth data set be described by zero experimental noise

$$\vec{\mathrm{X}}_{\mathrm{truth}} = \{\mathrm{x}_i\,|_{i=1}^{m+1}\}\; s.t.\; (\sigma_{\mathrm{E}} = 0), \quad (4)$$

where σ_{E} is the standard deviation of the experimental noise. The function **T** mapping $\vec{\mathrm{X}}_{\mathrm{truth}}$ to FA determines the calculated ground truth FA value, $\mathrm{FA_{truth}}$,

$$\mathrm{FA_{truth}} = \mathbf{T}\left(\{\vec{\mathrm{X}}_{\mathrm{truth}}\}\right). \quad (5)$$

The observed data of an actual experiment contains experimental noise and is described by,

$$\vec{\mathrm{X}}_{\mathrm{obs}} = \vec{\mathrm{X}}_{\mathrm{truth}} + \sigma_{\mathrm{E}}\vec{\mathrm{Z}}. \quad (6)$$

Here, $\vec{\mathrm{Z}}$ is a vector of *m+1* random drawings from a standard normal distribution (and noise is independent across voxels when expanded to the multi-voxel case). The resulting observed FA from the experiment is then,

$$\mathrm{FA_{obs}} = \mathbf{T}\left(\{\vec{\mathrm{X}}_{\mathrm{obs}}\}\right). \quad (7)$$

In SIMEX, synthetic normally distributed noise with standard deviation $\omega^{1/2}\sigma_{\mathrm{E}}$ is added to $\vec{\mathrm{X}}_{\mathrm{obs}}$, and the new simulated data is written as,

$$\vec{\mathrm{X}}_{\mathrm{M.C.}}(\omega) = \vec{\mathrm{X}}_{\mathrm{obs}} + \omega^{1/2}\sigma_{\mathrm{E}}\vec{\mathrm{Z}}_{\mathrm{M.C.}}. \quad (8)$$

M.C. stands for 'Monte Carlo' (Note, M.C. replaces 'b' as defined in the SIMEX literature[7]). The simulated FA is then described by,

$$\mathrm{FA}_{\mathrm{M.C.}}(\omega) = \mathbf{T}\left(\{\vec{\mathrm{X}}_{\mathrm{M.C.}}(\omega)\}\right). \tag{9}$$

The variance of the simulated data as a function of ω is

$$\begin{gathered}\mathrm{var}\left(\vec{\mathrm{X}}_{\mathrm{M.C.}}(\omega)\right) = \\ \mathrm{var}\left(\sigma_{\mathrm{E}}\vec{\mathrm{Z}} + \omega^{1/2}\sigma_{\mathrm{E}}\vec{\mathrm{Z}}_{\mathrm{M.C.}}\right) = \\ \sigma_{\mathrm{E}}^2(1+\omega)\,.\end{gathered} \tag{10}$$

A key observation is to note that at the value ω = -1, (which cannot be simulated) the variance goes to 0. For repeated Monte Carlo simulations of $\mathrm{FA}_{\mathrm{M.C.}}(\omega)$, a mean value can be calculated, $\overline{\mathrm{FA}}_{\mathrm{M.C.}}$, and the trend approximated by fitting $\overline{\mathrm{FA}}_{\mathrm{M.C.}}(\omega)$ with a standard non-linear equation, $\widehat{\overline{\mathrm{FA}}}(\omega)$,

$$\widehat{\overline{\mathrm{FA}}}(\omega) = \mathrm{a} + \frac{\mathrm{b}}{1+\omega} \tag{11}$$

The FA at zero noise can then be estimated by solving eq. 11 for ω = -1,

$$\mathrm{FA}_{\mathrm{SIMEX}} = \widehat{\overline{\mathrm{FA}}}(-1)\ . \tag{12}$$

The bias in the data can then be estimated,

$$\text{Estimated Bias} = \mathrm{FA}_{\mathrm{obs}} - \mathrm{FA}_{\mathrm{SIMEX}}\ . \tag{13}$$

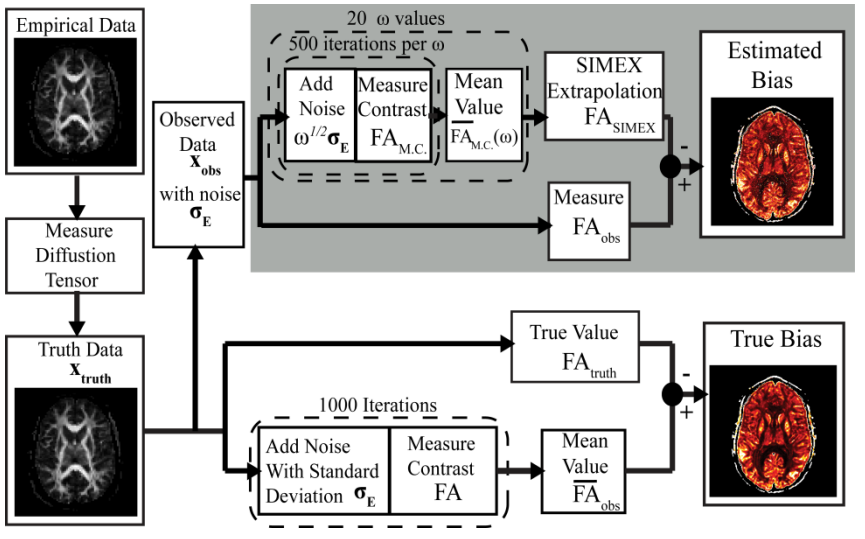

Fig. 1. Flow chart of steps to create SIMEX Estimated Bias and True Bias maps. Variables refer to terms defined explicitly by equations in the Theory section.

The true bias can also be calculated if the truth data is known,

$$\text{True Bias} = \mathbf{E}(\text{FA}_{\text{obs}}) - \text{FA}_{\text{Truth}} . \quad (14)$$

If the true extrapolant function could be used in place of eq. 11, then FA_{SIMEX} would be an asymptotically unbiased estimator of FA_{truth}. Under the condition where the true but unknown extrapolant is smooth w.r.t. ω, the approximating extrapolant, eq. 11, can provide approximately unbiased results.

3 Methods and Results

The steps used here to evaluate SIMEX on FA are summarized in the flow chart of Fig. 1 and exampled by individual voxels in Fig. 2. Unless otherwise stated, all processing and analysis was performed in Matlab 2010 (Mathworks, Natick, MA). Diffusion tensor estimates were calculated by fitting the diffusion model (eq.1) to the data using a simple LLMMSE. In all cases, noise was added to data using a Rician distribution, which is approximately Gaussian for SNR > 5:1[11].

Empirical Data. Empirical data (KKI2009-33-DTI) were downloaded from the Multi-Modal MRI Reproducibility study available online at www.nitrc.org [12]. Full collection details are provided in the reference. Briefly, DTI data was collected using a multi-slice, single-shot, echo planar imaging (EPI) sequence. Thirty two gradient directions were used with b = 700 s/mm^2 (m = 32, b = 700, eq. 1 & 2). Five S_o images were collected and averaged into a single S_o. The resulting images consisted of sixty-five transverse slices with a field of view = 212 x 212 reconstructed to 256 x 256.

Creating Noiseless Ground Truth (eq. 4). The 32 dwi images were registered to the single S_o image using FLIRT affine registration (FMRIB, Oxford, UK). For spatial consistency, the resulting affine transformation matrix was applied to the gradient table. A single axial slice of interest was selected and diffusion tensor estimates (**D**) for that slice were calculated. Using the S_o image, the estimated diffusion tensor, and the given **g** and b as inputs to eq. 1, dwi images were simulated. The S_o image and the 32 noiseless dwi images make the noiseless ground truth dataset, $\vec{\text{X}}_{\text{truth}}$.

Creating Observed Data (eq. 6). To create the observed data, $\vec{\text{X}}_{\text{obs}}$, noise was added to the ground truth data set such that SNR = 35:1. The signal strength was estimated from the average signal intensity of the masked S_o image. Noise, σ_E = signal/35, was added to the ground truth data to create the 32 dwi images and the single S_o image.

Calculating the Estimated Bias (eq. 13). To calculate FA_{SIMEX}, first a plot of $\overline{\text{FA}}_{\text{M.C.}}(\omega)$ needed to be constructed. 500 Monte Carlo simulations for 20 different **ω** values were performed and $\overline{\text{FA}}_{\text{M.C.}}(\omega)$ calculated from the average. Each simulation consisted of a noise adding step to create $\mathbf{X}_{\text{M.C.}}$ (eq. 8), followed by a calculation of $FA_{M.C.}$. Example $\overline{\text{FA}}_{\text{M.C.}}(\omega)$ for three voxels are shown in Fig. 2A. $\widetilde{\text{FA}}(\omega)$ was fit using the empirical non-linear equation (eq. 11). The fit equation was extrapolated to ω = -1, producing FA_{SIMEX} (Fig. 2B). The FA_{obs} from eq. 13 was calculated directly from diffusion tensor fits to the observed data set, $\vec{\text{X}}_{\text{obs}}$. The difference of FA_{SIMEX} and FA_{obs} was taken and the Estimated Bias calculated.

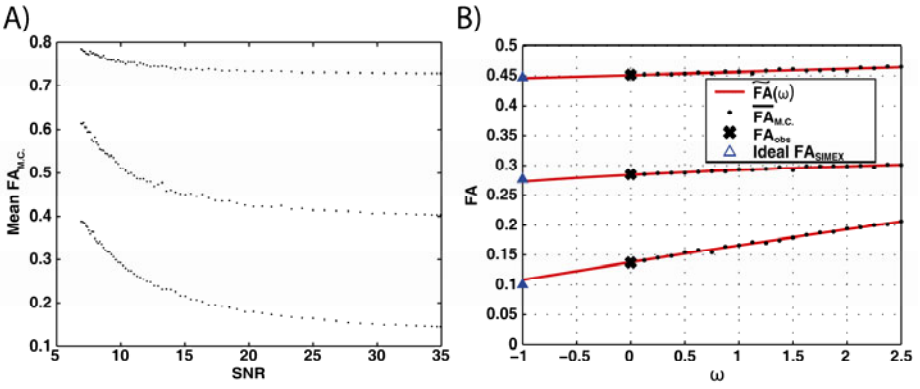

Fig. 2. (A) Exploratory SIMEX plots of simulated $FA_{M.C.}$ as a function of SNR for three voxels from the simulated observed data. 1000 Monte Carlo simulations were averaged for each $\overline{FA}_{M.C.}$. The FA_{obs} is the point at the highest SNR value (SNR = 35:1). An upward bias with decreasing SNR is observed for all three voxels. (B) $FA_{M.C.}(\omega)$ and the extrapolation fit for three voxels in the observed data set. The x-axis for extrapolation is ω ~1/SNR and the fit is extrapolated to ω = -1. The ideal FA_{SIMEX} (blue triangles) represents the extrapolated FA_{SIMEX} value if the true bias was correctly estimated.

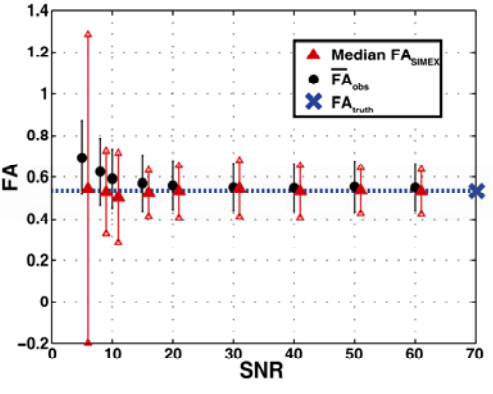

Fig. 3. Simulated outcomes for SIMEX at several SNR values. $\overline{FA}_{obs}$ values for each SNR value (*circles*) were calculated starting from a voxel in the simulated truth data. The $\overline{FA}_{obs}$ values are shown with the standard deviation (*error bars*) from 1000 simulations run at each SNR value. 100 observations at each SNR value were randomly chosen for SIMEX. The median FA_{SIMEX} values (*triangles*) are shown with the standard deviation(*error bars*). The median value was chosen due to the decreased robustness of only 100 simulations.

Calculating True Bias (eq. 14). The average of 1000 iterations was used to estimate the expected value of FA_{obs}, $\overline{FA}_{obs} \sim \mathbf{E}(FA_{obs})$. Each simulation consisted of first adding noise (SNR = 35:1) to $\vec{X}_{truth}$ followed by calculation of FA. FA_{truth} was calculated directly from $\vec{X}_{truth}$. The True Bias was calculated from the difference between $\overline{FA}_{obs}$ and FA_{truth}. Exploratory M.C. simulations applying SIMEX to different SNR levels of $\vec{X}_{obs}$ for a

single ground truth voxel are shown in Fig. 3. The True Bias and Estimated Bias are compared in Fig. 4. The results shown in Fig. 4 were typical for multiple test trials (data not shown).

4 Discussion

Potential systematic differences in distributional properties of tensor-derived contrasts pose serious hazards for statistical analysis. Understanding "statistical artifacts" in DTI has long been an essential challenge in proper interpretation of DTI contrasts [13]. Routine approaches for interpreting contrasts (e.g., regions of interest, voxel-based morphometry, tract-based spatial statistics) hinge upon the assumption of equivalent (or zero) bias between populations. If this assumption is violated then the inferences (and p-values) may be suspect. In any empirical experiment, there will be variations in bias due to subject characteristics, protocols, scanner stability, hardware/software changes, etc. Without a method to estimate bias in measures derived from individual empirical datasets, evaluation of any change requires an extensive simulation study based on a range of potential anatomical models.

With the proposed SIMEX approach, it is possible to estimate bias in contrasts on a voxel-by-voxel level given the empirical data. SIMEX proved a sensitive method for detecting FA bias as evidenced by its success on the high SNR (and therefore relatively low bias) dataset used here (Fig. 4). Although bias may be negligible for some individual cases at high SNR values (such as exampled for high SNR values in Fig. 3), bias may become increasingly important at lower SNR values or in hypothesis test settings where even a small bias can significantly inflate the Type-I error rate. To that end, the individual voxel test case (Fig. 3) suggests SIMEX to be a promising technique for empirical data collected at lower SNR levels (stable median FA_{SIMEX} values) as well as a reliable method for repeated use on individual datasets within grouped data (stable error bars). The instability of the SIMEX estimate for SNR $\leq$ 10:1 (Fig. 3) is likely to be a result of the non-monotonic behavior of bias at extremely low SNR values [5].

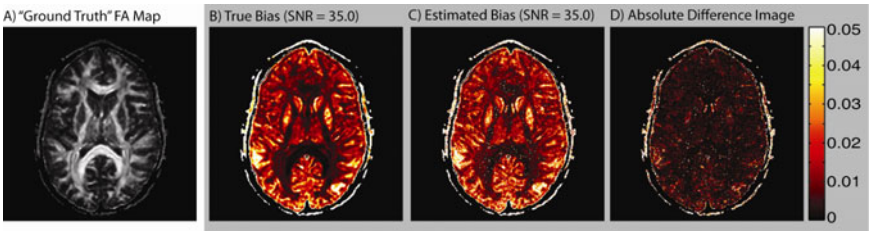

Fig. 4. (A) The 'ground truth' FA map (FA_{truth}) of the slice selected for data analysis. (B) The True Bias map for an SNR of 35:1 is compared to (C) the absolute value of the Estimated Bias at SNR = 35:1. The absolute value difference map was calculated from the raw bias data (not the difference in the absolute values). For a significant majority of voxels the True Bias is positive and the Estimated Bias is greater than the True Bias.

We note that at these early stages, we are not specifically recommending using the estimated bias measures to correct for voxel-wise bias. Optimization of user input SIMEX parameters (e.g. the approximation function in eq. 11 or the number of Monte Carlo iterations) and repeat testing on empirical data of various SNR levels is underway. Refinements may help decrease the variability in the SIMEX estimate (Fig. 3 red error bars) and mediate some of the isolated outliers seen in Fig. 4D. Additionally, methods for constructing error estimates of the measured SIMEX bias exist and have yet to be tested in the context of DTI[6]. With continued future refinement and careful incorporation of robust/advanced tensor estimators it may be possible to produce unbiased contrast estimates without the disadvantages of direct maximum likelihood methods[14].

In summary, a method for direct, empirical assessment of bias in tensor derived contrasts opens many possibilities for useful statistical assessments and improved analysis approaches. Further investigation is warranted into bias estimation with other tensor contrasts, using advanced tensor fitting methods, and examining quantitative tissue models beyond the tensor formulation.

Acknowledgments. This work was supported in part by NIH/NINDS 1R01NS056307 and NIH/NINDS 1R21NS064534.

References

1. Le Bihan, D., van Zijl, P.: From the diffusion coefficient to the diffusion tensor. NMR Biomed 15, 431–434 (2002)
2. Landman, B.A., Farrell, J.A., Jones, C.K., Smith, S.A., Prince, J.L., Mori, S.: Effects of diffusion weighting schemes on the reproducibility of DTI-derived fractional anisotropy, mean diffusivity, and principal eigenvector measurements at 1. 5T. Neuroimage 36, 1123–1138 (2007)
3. Jones, D.K., Basser, P.J.: Squashing peanuts and smashing pumpkins: How noise distorts diffusion-weighted MR data. Magnetic Resonance in Medicine 52, 979–993 (2004)
4. Whitcher, B., Tuch, D.S., Wisco, J.J., Sorensen, A.G., Wang, L.: Using the wild bootstrap to quantify uncertainty in diffusion tensor imaging. Hum Brain Mapp 29, 346–362 (2008)
5. Landman, B.A., Farrell, J.A., Huang, H., Prince, J.L., Mori, S.: Diffusion tensor imaging at low SNR: nonmonotonic behaviors of tensor contrasts. Magn. Reson Imaging (2008)
6. Carroll, R.J.: Measurement error in nonlinear models: a modern perspective. CRC Press, Boca Raton (2006)
7. Cook, J., Stefanski, L.: Simulation-extrapolation estimation in parametric measurement error models. Journal of the American Statistical Association 89 (1994)
8. Landman, B.A., Bazin, P.L., Smith, S.A., Prince, J.L.: Robust estimation of spatially variable noise fields. Magn. Reson Med. 62, 500–509 (2009)
9. Rajan, J., Poot, D., Juntu, J., Sijbers, J.: Noise measurement from magnitude MRI using local estimates of variance and skewness. Phys Med Biol 55, N441-N449 (2010)

10. Aja-Fernandez, S., Tristan-Vega, A., Alberola-Lopez, C.: Noise estimation in single- and multiple-coil magnetic resonance data based on statistical models. Magn. Reson Imaging 27, 1397–1409 (2009)
11. Gudbjartsson, H., Patz, S.: The Rician distribution of noisy MRI data. Magn Reson Med. 34, 910–914 (1995)
12. Landman, B., Huang, A., Gifford, A., Vikram, D., Lim, I., Farrell, J., Bogovic, J., Hua, J., Chen, M., Jarso, S., Smith, S., Joel, S., Mori, S., Pekar, J., Barker, P., Prince, J., van Zijl, P.: Multi-Parametric Neuroimaging Reproducibility: A 3T Resource Study. Neuroimage 4, 2854–2866 (2011)
13. Basser, P.J., Pajevic, S.: Statistical artifacts in diffusion tensor MRI (DT-MRI) caused by background noise. Magn. Reson Med. 44, 41–50 (2000)
14. Anderson, A.W.: Theoretical analysis of the effects of noise on diffusion tensor imaging. Magn Reson Med. 46, 1174–1188 (2001)

Impact of Radial and Angular Sampling on Multiple Shells Acquisition in Diffusion MRI

Sylvain Merlet[1], Emmanuel Caruyer[1], and Rachid Deriche[1]

Athena Project-Team, INRIA Sophia Antipolis - Méditerranée, France

Abstract. We evaluate the impact of radial and angular sampling on multiple shells (MS) acquisition in diffusion MRI. The validation of our results is based on a new and efficient method to accurately reconstruct the Ensemble Average Propagator (EAP) in term of the Spherical Polar Fourier (SPF) basis from very few diffusion weighted magnetic resonance images (DW-MRI). This approach nicely exploits the duality between SPF and a closely related basis in which one can respectively represent the EAP and the diffusion signal using the same coefficients. We efficiently combine this relation to the recent acquisition and reconstruction technique called Compressed Sensing (CS). Based on results of multi-tensors models reconstruction, we show how to construct a robust acquisition scheme for both neural fibre orientation detection and attenuation signal/EAP reconstruction.

Keywords: Diffusion MRI, Compressed sensing, Ensemble Average Propagator recovery, Propagator, Orientation Distribution Function, Spherical Polar Fourier, Multiple Shells Sampling.

1 Introduction

Since the introduction of CS by [5], this method has been used in a large range of domains including image and video compression as well as geophysics and medical imaging. In [8], we apply CS in diffusion magnetic resonance imaging (dMRI) for modelling the EAP, in SPF basis introduced by [1], from highly under sampled diffusion weighted MR images (DW-MRIs).

In [8], the quality of reconstruction is sensitive to the acquisition scheme. Hence, in order to remove the variance of the results due to the random aspect of the sampling scheme, it is necessary to find a robust and efficient way to acquire DW-MRIs. For this purpose, we propose to evaluate and compare several sampling protocols. In this study, we begin in 2 by summarizing the EAP based CS-reconstruction proposed in [8]. The section 3 aims to describe some techniques to build sampling scheme. Finally in section 4 we present several sets of experiments in order to examine the robustness and efficiency of these schemes.

2 EAP Based CS Reconstruction

The method described in [8] enables the modeling of the EAP in the SPF basis, using the recent technique known as Compressed Sensing (CS). The method

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 116–123, 2011.

allows us to analytically reconstruct the propagator at any radius and, also, to derive one of its famous feature: the Orientation Distribution Function (ODF). The CS reconstruction is based on a l_1 minimization scheme promoting the signal sparsity. For a more mathematical definition, suppose our signal of interest is a vector $\mathbf{x} \in \mathbb{R}^m$. Let $\mathbf{y} \in \mathbb{R}^n$, with $n \ll m$, be an observation representative of $\mathbf{x}$ given by the sensor of a given application. $\mathbf{y} = \mathbf{A}\mathbf{x} + \eta$, where $\mathbf{A} \in \mathbb{R}^{n \times m}$ is the measurement matrix, so called the CS matrix, and $\eta \in \mathbb{R}^n$ represents the acquisition noise. Our goal is to find $\mathbf{x}$ given the measurement vector $\mathbf{y}$. Since $\mathbf{y}$ has less entries than $\mathbf{x}$, this ill-posed problem cannot be resolved without any prior knowledge about the signal to recover.

We condider the signal admits a sparse representation with respect to an orthonormal basis B. $\mathbf{c} \in \mathbb{R}^{n_c}$ is the vector of transform coefficients $\{c_i = \langle x, b_i \rangle, b_i \in B, i = 1, ..., n_c\}$. We can constrain most of the transform coefficients to be zero by minimizing the l_1 norm defined by $\|\mathbf{c}\|_1 = \sum_{i=1}^{n_c} |c_i|$ [7].The solution $\mathbf{x}$ of our problem is given by solving the following convex optimization problem :

$$\arg\min_{\mathbf{x}} J(\mathbf{x}) = \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{c}\|_1. \tag{1}$$

The first term is the data consistency constraint, $\|\mathbf{c}\|_1$ is the sparsity constraint. λ is the Lagrange parameter that balances the confidence between the measured signal $\mathbf{y}$ and the sparsity constraint. The data consistency constraint enables the solution to remain close to the raw data acquisition, whereas the minimization of the second term promotes sparsity. In short, this mathematical problem searches for the sparsest solution while remaining close to the acquired data.

In our previous paper [8], we define B as the so called Spherical Polar Fourier (SPF) basis in spherical coordinates (r, θ, φ). This orthonormal basis is the combination of the real spherical harmonics Y_m^l and the Gauss-Laguerre functions R_n. It is expressed as $\Psi_{n\ell m}(r, \theta, \phi) = R_n(r) Y_m^\ell(\theta, \phi)$, where $\ell \in \mathcal{N}$ is the spherical harmonic order, $-l \leq m \leq l$ the SH degree and n the Gauss-Laguerre order. This basis enables a complete description of the diffusion propagator.

We described a method to accurately reconstruct the EAP P from undersampled measurements. In this method P was estimated by a truncated linear combination of the SPF basis functions $\Psi_{n\ell m}$

$$P(r, \theta, \phi) = \sum_{n=0}^{N} \sum_{\ell=0}^{L} \sum_{m=-\ell}^{\ell} c_{n\ell m} \Psi_{n\ell m}(r, \theta, \phi), \tag{2}$$

where $c_{n\ell m} = \langle P, \Psi_{n\ell m} \rangle$ are the SPF transform coefficients.

While modelling the EAP with respect to the SPF basis, we have shown that we can reconstruct the corresponding attenuation signal E by keeping the same coefficients $\{c_{n\ell m}, n = 0, ..., N, l = 0, ..., L, m = -l, ..., l\}$ but by using a new family of functions called the SPF dual basis $\{\Phi_{n\ell m}, n = 0, ..., N, l = 0, ..., L, m = -l, ..., l\}$. E can, thus, be written as

$$E(q, \theta, \varphi) = \sum_{n=0}^{N} \sum_{\ell=0}^{L} \sum_{m=-\ell}^{\ell} c_{n\ell m} \Phi_{n\ell m}(q, \theta, \phi) \tag{3}$$

where $q = |\mathbf{q}|$ is the norm of the effective gradient $\mathbf{q}$ in q-space and θ, φ the direction angles. Suppose n_q is the number of measurement samples, $\mathbf{E} \in \mathbb{R}^{n_q}$ a vector representing the signal attenuation, $\mathbf{c} \in \mathbb{R}^{n_c}$ a vector of the SPFd coefficients $c_{n\ell m}$ and $\mathbf{\Phi} \in \mathbb{R}^{n_q \times n_c}$ the matrix constructed with the SPFd basis functions

$$\mathbf{\Phi} = \begin{pmatrix} \Psi_{n\ell m}(r_1, \theta_1, \phi_1) & \cdots & \Psi_{NLL}(r_1, \theta_1, \phi_1) \\ \vdots & \ddots & \vdots \\ \Psi_{n\ell m}(r_{n_q}, \theta_{n_q}, \phi_{n_q}) & \cdots & \Psi_{NLL}(r_{n_q}, \theta_{n_q}, \phi_{n_q}) \end{pmatrix}, \tag{4}$$

We can write equation (3) as an over determined linear system, $\mathbf{E} = \mathbf{\Phi c}$ Let $\mathbf{\Phi}_u \in \mathbb{R}^{n_u \times n_c}$ be the undersampled version of $\mathbf{\Phi}$ operator and $\mathbf{E}_u \in \mathbb{R}^{n_u}$ the vector of undersampled signal attenuation. We can rewrite the problem described in equation 1

$$\arg\min_{\mathbf{c}} J(\mathbf{c}) = \|\mathbf{\Phi_u c} - \mathbf{E_u}\|_2^2 + \lambda \|\mathbf{c}\|_1. \tag{5}$$

Eq. (5) searches for the EAP coefficients with respect to the SPF basis, i.e. we can compute a continuous version of the true propagator. Using the same coefficients, we can as well model the attenuation signal with respect to the SPFd basis functions. Moreover, in [8] we give an analytical estimate of the ODF in terms of spherical harmonic functions and the coefficients $c_{n\ell m}$.

In [8], we randomly took 80 measurements spread on 3 shells with b values 1000, 2000, 3000 s/mm^2. However, the random aspect of the acquisition process makes the reconstruction very sensitive to the sampling scheme. We selected it as follow: on 100 sampling schemes generated, we kept the one that leads to the best results. In this way, good results were observed while performing the reconstruction on our data set but were not observed in all cases. In the next section, we propose several sampling schemes in order to make the method more robust to this reconstruction phase.

3 Sampling Design

3.1 Jones

References [6,4] give an algorithm to uniformly distribute N points $q_n \in R^3$ on a sphere by considering each point as an antipodal pair of electrical charges. The method involves the minimization of the electrostatic force of repulsion between each couple of charges. The electrostatic repulsion between two points q_i and q_j is given by

$$E(q_i, q_j) = \frac{1}{\|q_i + q_j\|} + \frac{1}{\|q_i - q_j\|} \tag{6}$$

For a set of N points, the energy to minimize becomes

$$J_J = \sum_{i \neq j} E(q_i, q_j) \tag{7}$$

Reference [3] provides Camino, an Open-Source Diffusion-MRI Reconstruction and Processing software. They include several sets of directions, from N=3 to 150 points, computed by electrostatic energy minimization.

3.2 Generalized Jones

This method is proposed by [2] as a generalization of [6] to multi-shells acquisition. It enables the distribution of N points $q_n \in R^3$ on K shells of radius r_k. The points from each shell have staggered directions and follow a near-optimal uniform distribution. Another important point in this method is the possibility to balance the proportion α_k of samples between shells. We will take advantage of this feature in order to test out different spherical distributions.

Firstly, the method consists in minimizing the electrostatic repulsion between every point for each shell independently, that is

$$E_1 = \sum_k r_k \alpha_k \sum_{i \neq j \ s.t \ \|q_i\| = \|q_j\| = r_k} E(q_i, q_j) \tag{8}$$

Then, in order to have staggered directions between shells, [2] introduces a new term that minimizes the electrostatic repulsion of the N points projected on the unit sphere. It comes to minimize

$$E_2 = \sum_{i \neq j} \frac{1}{\left\| \frac{q_i}{\|q_i\|} - \frac{q_j}{\|q_j\|} \right\|} + \frac{1}{\left\| \frac{q_i}{\|q_i\|} + \frac{q_j}{\|q_j\|} \right\|} \tag{9}$$

Finally, the energy to minimize is $J_{GJ} = (1 - \mu)E_1 + \mu E_2$, where μ is a weighting factor.

3.3 5 Sampling Schemes

We perform our experiments on five sampling schemes to evaluate the impact of the angular sampling in MS acquisition : **Regular sampling (RS)** means we take the same directions on each shell. These directions are provided by the application of Jones algorithm. **Uniform Jones sampling (UJS)** uses the generalized Jones algorithm by setting the parameters in such way that the samples are distributed along a spherical uniform law (The number of point on each shell is proportional to the square of its radius). **Constant Jones sampling (CJS)** uses the generalized Jones algorithm by setting the parameters in such way that there is a constant number of samples by shell. **Constant random sampling (CRS)** means we randomly take directions on each shell by setting a constant number of samples by shell. **Uniform random sampling (URS)** means we randomly take directions on each shell in such way that the samples are distributed along a spherical uniform law.

4 Experimental Results

In this section, we review the outcome of angular sampling as well as radial sampling on the CS reconstruction defined in section 2. The performance of each sampling scheme is determined on both attenuation signal reconstruction and neural fibre orientation which is given by the maxima of the estimated ODF estimated as in [8]. We evaluate the maxima extraction by the Percentage of Corrected Number of Detected Maxima (PCNDM) obtained on a predefined number of trial. Each time the number of detected maxima N_m is correct we also compute the Angular Error in degrees $\frac{1}{N_m}\sum_{m=0}^{N_m}\frac{180}{\pi}\arccos(\tilde{u}_m \cdot u_m)$, where $\tilde{u}_m$ is the orientation of the detected maxima and u_m the ground truth. The mean over all the trial gives the Mean Angular Error (MAE).

The quality of the signal attenuation estimation $\tilde{S}$ is given by the Normalized Mean Square Error (NMSE). For N sampling points q_n the NMSE is $\frac{\sum_{i=1}^{N}|S(q_i)-\tilde{S}(q_i)|^2}{\sum_{i=1}^{N}|S(q_i)|^2}$, where S is the ground truth signal attenuation. Then, we average the NMSE obtained in all the trials.

Before initiating the procedure, we set some reconstruction parameters. SNR ζ=20, for average quality data. In eq. 5, we set λ to 0.01 as in [8].

We reconstruct the propagator from a multi-Gaussian model through four scenarios : One fiber, two 60°-crossing fibres, two 70°-crossing fibres, two 90°-crossing fibres. All the results are obtained on 1500 independant trials.

Angular Profile: For each scenario we use three shells with b values equal to 500, 1500, and 3000. We begin by evaluate the five sampling schemes presented in section 3.3 while using only N=60 samples in the reconstruction. Figure 1 presents the results through the computation of the Percentage of Corrected Number of Detected Maxima (PCNDM), the Mean of Angular Error (MAE) in degrees and the Normalized Mean Square Error (NMSE) respectively on line 1, 2 and 3. Each color corresponds to one the five sampling schemes. Each group of five bars correspond to one of the scenarios presented at the beginning of the section (One fiber, two 60°-crossing fibres, two 70°-crossing fibres, two 90°-crossing fibres). We can see that, in term of PCNDM, two schemes stand out: the UJS and RS schemes (green and dark blue). However, due to the bad capacity of the RS scheme to resolve 60°-crossing fibres and most probably low degrees as well, we cannot rely upon it for maxima extraction. This is verified by looking at the MAE (second line) of RS scheme. Once again UJS gives the best results, i.e. the lowest MAE. Overall, orientation detection is better performed when a spherical uniform distribution is applied compared to a constant distribution. Constant distribution means we take the same number of samples on each shell. Random angular sampling confirms this point. Concerning the NMSE (third line), three sampling schemes are quite equivalents: the RS, the UJS and CJS schemes. Even if the regular one is slightly better than the other, it is difficult to distinguish one scheme from another. Here again random sampling does not meet the CS expectations.

We illustrate, as well, the evolution of the PCNDM, MAE and NMSE (from left to right in fig. 2) against the number of samples taken in the reconstruction.

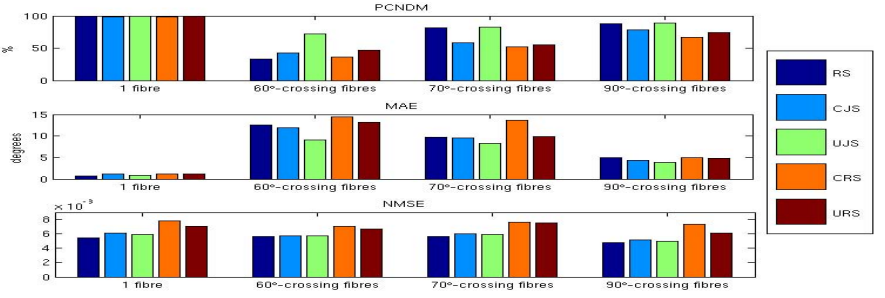

Fig. 1. Reconstruction results while using only N=60 samples in the reconstruction. Five angular sampling schemes are examined : Regular sampling (dark blue bars), Constant Jones sampling (light blue bars), Uniform Jones sampling (green bars) Constant random sampling (orange bars), Uniform random sampling (red bars). From top to bottom the three lines respectively represent the Percentage of Corrected Number of Detected Maxima (PCNDM), the Mean of Angular Error (MAE) in degrees and the Normalized Mean Square Error (NMSE). Each group of five bars correspond to one of the following scenarios : One fiber, two 60°-crossing fibres, two 70°-crossing fibres, two 90°-crossing fibres (from left to right).

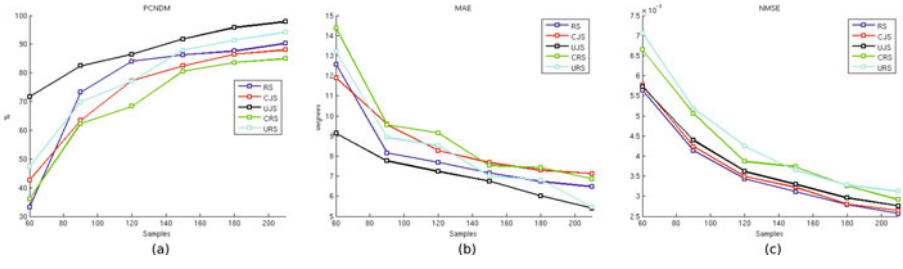

Fig. 2. Evolution of the reconstruction results against the number of samples for the 60°-crossing fibres. Five angular sampling schemes are examined : Regular sampling (dark blue), Constant Jones sampling (red), Uniform Jones sampling (black) Constant random sampling (green bars), Uniform random sampling (light blue). (a) represents the Percentage of Corrected Number of Detected Maxima (PCNDM), (b) the Mean of Angular Error (MAE) in degrees and (c) the Normalized Mean Square Error (NMSE).

Only the results for 60°-crossing fibres are represented. For maxima extraction (see PCNDM (a) and MAE (b)), the UJS overcomes the other schemes for every number of samples. In term of NMSE, the figure 2.c shows that the gap between the RS, the UJS and CJS is not sufficiently large in order to distinguish one scheme from another. Random sampling gives less stable curves than the others. It shows the problem of sensibility for this kind of scheme in our reconstruction.

In conclusion with respect to the angular sampling, UJS is a robust and efficient way to build MS schemes.

Radial Profile: Now we want to examine the influence of the radial sampling. Due to the robustness of UJS, we decide to keep it while changing the number of shells from 1 to 12.

Let us examine the figure 3(a) and (b). These results show that, above 3 shells, the MAE and the PCNDM do not vary a lot. It means the uniform sampling scheme is not sensitive to the number of shells used for the acquisition when fibre orientation detection is required. With 1 and 2 shells, we get a higher PCNDM than previously. However this improvement is done to the price of an increasing MAE. Moreover, when few shells are taken for sampling, it is more difficult to catch significant information of the radial profile. The b values have to be cautiously chosen if few shells are used. By regularly taking more and more shells in a MS process, we cover more precisely the radial profile of our signal and we get rid of the choice of b-values.

Concerning the NMSE, figure 3(c) shows that increasing the number of sample decreases the NMSE until the number of shells exceeds a specific value. This limit may due to the fact that a too big increase of the radial resolution lead to a fall of the angular resolution. Figure 3(c) shows a limit of 9 shells in our case. However, we dont need to reach this point. Indeed the quasi-flatness of the curves before 9 shells allows us to use less shells for sampling while keeping the advantage of MS sampling.

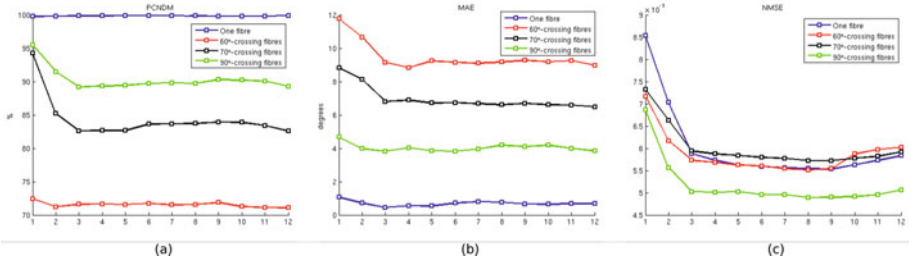

Fig. 3. Evolution of the reconstruction results against the number of shells. Uniform Jones sampling is used while changing the number of shells from 1 to 12. Four scenarios are examined : One fiber (blue curve), two 60°-crossing fibres (red curve), two 70°-crossing fibres (black curve), two 90°-crossing fibres (green curve). (a) represents the Percentage of Corrected Number of Detected Maxima (PCNDM), (b) the Mean of Angular Error (MAE) in degrees and (c) the Normalized Mean Square Error (NMSE).

Conclusion of the Experiments: With respect to the angular profile, one sampling scheme stand out : the UJS scheme. It allows us to correctly detect the orientation of neural fibres and especially for low number of samples (N=60). It also gives good results in terms of NMSE. The results show again that the

signal attenuation reconstruction and maxima extraction are very sensitive to the sampling scheme. Even if random sampling (CRS, URS) can work better than near-regular sampling (RS, CJS, UJS), we cannot ensure that it will work in every case. Fibre orientation detection is not sensitive while changing the number of shells of the UJS scheme. A great advantage of MS sampling, compare to one shell sampling, lies in the quasi non-sensibility in the choice of b-values. Indeed we just need to acquire our signal on shells regularly spaced. On the other hand, due to the small number of point used, a compromise has to be done between radial and angular resolution.

5 Conclusion

The main contribution of this paper is to evaluate different way to sample q-space in dMRI. We showed that multiple shells acquisition is of great interest when dealing with fibres orientation detection and attenuation signal reconstruction. Hence this method generalizes the Q-ball imaging while being able to reconstruct the EAP and signal attenuation at any radius. In our study, we also showed that generalized Jones algorithm is a good way to build robust multiple shells sampling schemes and that the use of a spherical uniform distribution improves the results.

References

1. Assemlal, H.-E., Tschumperlé, D., Brun, L.: Efficient computation of PDF-based characteristics from diffusion MR signal. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part II. LNCS, vol. 5242, pp. 70–78. Springer, Heidelberg (2008)
2. Caruyer, E., Lenglet, C., Sapiro, G., Deriche, R.: Incremental gradient table for multiple q-shells diffusion mri. In: HBM 17th Annual Meeting (June 2011)
3. Cook, P.A., Bai, Y., Nedjati-Gilani, S., Seunarine, K.K., Hall, M.G., Parker, G.J., Alexander, D.C.: Camino: Open-source diffusion-mri reconstruction and processing. In: 14th ISMRM, Seattle, USA (2006)
4. Deriche, R., Calder, J., Descoteaux, M.: Optimal real-time q-ball imaging using regularized kalman filtering with incremental orientation sets. Medical Image Analysis 13(4), 564–579 (2009)
5. Donoho, D.L.: Compressed sensing. IEEE Trans. on Information Theory 52(4), 1 (2004)
6. Jones, D., Horsfield, M., Simmons, A.: Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. Magnetic Resonance in Medicine 42(39), 515–525 (1999)
7. Lustig, M., Donoho, D., Pauly, J.: Sparse mri: The application of compressed sensing for rapid mr imaging. Magnetic Resonance in Medicine 58(6), 1182–1195 (2007)
8. Merlet, S., Cheng, J., Ghosh, A., Deriche, R.: Spherical polar fourier eap and odf reconstruction via compressed sensing in diffusion mri. In: Proceedings of ISBI (2011)

Super-Resolution in Diffusion-Weighted Imaging

Benoit Scherrer, Ali Gholipour, and Simon K. Warfield

Computational Radiology Laboratory, Department of Radiology
Children's Hospital Boston, 300 Longwood Avenue, Boston, MA, 02115, USA

Abstract. Diffusion-weighted imaging (DWI) enables non-invasive investigation and characterization of the white-matter but suffers from a relatively poor resolution. In this work we propose a super-resolution reconstruction (SRR) technique based on the acquisition of multiple anisotropic orthogonal DWI scans. We address the problem of patient motions by aligning the volumes both in space and in q-space. The SRR is formulated as a maximum *a posteriori* (MAP) problem. It relies on a volume acquisition model which describes the generation of the acquired scans from the unknown high-resolution image. It enables the introduction of image priors that exploit spatial homogeneity and enables regularized solutions. We detail our resulting SRR optimization procedure and report various experiments including numerical simulations, synthetic SRR scenario and real world SRR scenario. Super-resolution reconstruction in DWI may enable DWI to be performed with unprecedented resolution.

Keywords: Diffusion imaging, super-resolution, orthogonal acquisitions.

1 Introduction

Diffusion-weighted imaging (DWI) is a key imaging technique to investigate and characterize the white-matter architecture and microstructure. It is based on the acquisition of several diffusion-sensitized images, probing the water diffusion in various directions and at various diffusion scales. DWI is, however, strongly limited by the relatively low resolution achievable by today's imaging techniques: while individual axon diameter is on the order of 1-30μm, typically achievable DWI resolution is on the order of 2x2x2mm^3. Anisotropic acquisitions with a better *in-plane* resolution (up to 1x1mm^2) can be performed on modern scanners but lead to a lower signal to noise ratio (SNR) and are not adapted to further perform tractography. Consequently, due to strong partial volume effect, DWI has been limited to the investigation of the major fiber "highways" in the brain. Increasing the resolution of DWI acquisitions holds out the potential to allow investigation of novel fiber structures and will enable a more accurate white-matter and brain connectivity assessment. However, increasing the resolution is challenging in DWI. First, the common anatomical imaging resolution enhancement techniques, based on the modification of the pulse sequence, cannot be employed in DWI. The data of a same slice cannot be acquired over many excitations due to phase inconsistencies resulting from even minimal physiological

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 124–132, 2011.

motion during the application of the sensitizing gradients. Second, the SNR is directly proportional to the voxel size, and proportional to the square root of the number of averages. Consequently, 64 averages are necessary to increase the resolution from 2x2x2mm^3 to 1x1x1mm^3 while ensuring a similar SNR. A 5 minute acquisition would become a 5 hour scan, which is not realistic.
Solutions to achieve higher resolution include improvements of the MRI scanner hardware itself, such as employing higher magnetic fields (7 Tesla, 11 Tesla) or stronger and faster gradients. Another solution is to consider adapted acquisitions and post-processing algorithms. Super-resolution (SR) approaches were originally developed for the reconstruction of high-resolution (HR) images from a set of low-resolution (LR) images in video sequences [6]. SR techniques have also been applied to anatomical magnetic resonance imaging (MRI) [4]. Calamante *et al.* [1] have explored the application of interpolation of fiber tracts inside voxels but do not increase the resolution of the imaging data. To our knowledge, only [9] have used SR in DWI. They proposed to employ the Irani-Peleg SR technique [6] from a set of spatially in-plane subpixel-shifted scans. However, MRI being a Fourier acquisition technique, employing in-plane shifting has been shown to be equivalent to a global phase shift in k-space [3]. Such a technique does not enable any resolution enhancement in MRI but is equivalent to interpolation by zero-padding of the raw data in the temporal domain. Recently, sub-voxel spatial shifts in the *slice-select* dimension have been shown to enable SR in anatomical MRI [4]. Scattered data interpolation has been used to combine multiple DWI images of moving subjects [7]. Other techniques using multiple and *orthogonal* fast slice scans have enabled the SR reconstruction of moving subjects in fetal imaging [2].

In this work we propose to investigate a novel super-resolution reconstruction (SRR) approach for DWI. It is based on the acquisition of *multiple anisotropic orthogonal* DWI scans (see Fig.1a). First, we propose a technique to align each volume both in space and in q-space. Second, we formulate the super-resolution reconstruction from multiple scans as a maximum *a posteriori* estimation problem. Inspired by recent developments in fetus anatomical imaging [2], our approach relies on an image acquisition model. It describes the generation of the acquired volumes from the unknown HR volumes we aim to recover. Our formulation enables introduction of image priors that exploit spatial homogeneity and provide regularized solutions. We report various experiments including numerical simulations, synthetic SRR scenario and real world SRR scenario. The results indicate resolution enhancement in DWI through SRR.

2 Material and Methods

DW Signal Smoothness Hypothesis and Interpolation in q-space. We consider K *orthogonal* DWI LR acquisitions containing G sensitizing-gradients each. The KG volumes should all be properly aligned to enable the super-resolution reconstruction. First we perform an alignment in space: we register each volume to a reference volume, chosen as the $B = 0\text{s/mm}^2$ volume of the first

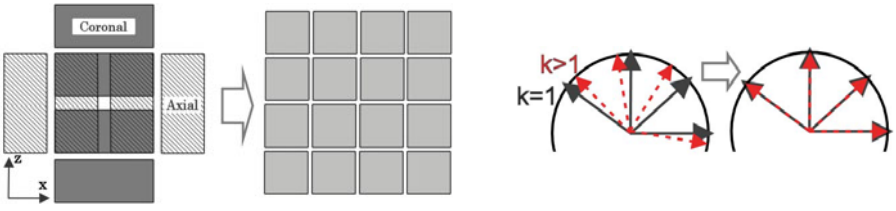

(a) (b)

Fig. 1. (a) Scheme illustrating the super-resolution reconstruction from the acquisition of two orthogonal thick slices. (b) Alignment in q-space: the gradient images of each acquisition $k > 1$ are resampled so that its gradient directions $\mathbf{g}^k$ (red dots arrows) correspond to the reference gradient directions of the first acquisition $\tilde{\mathbf{g}}$ (grey plain arrows). At each voxel, we compute novel intensities corresponding to the gradients $\tilde{\mathbf{g}}$ by interpolation in q-space from the observed intensities corresponding to $\mathbf{g}^k$.

DWI LR acquisition. Each gradient orientation is compensated for the rotation component of the transformation, providing a gradient set $\mathbf{g}^k = (g_1^k, \ldots, g_G^k)$ for each acquisition k.

Because of possible patient motions between the scans, the resulting G gradient images may correspond to slightly modified sensitizing gradients across the K DWI acquisitions. However, it is *essential* that the images combined by the SRR technique correspond to the same "scene", namely that they correspond to the exact same gradient direction and show identical diffusion-attenuation patterns. Consequently, we propose to perform an alignment of the volumes in q-space (see Fig.1b). We consider that the DW-signal varies smoothly in q-space and propose to resample the gradient images. We consider the gradients of the first DWI LR acquisition as the reference gradients $\tilde{\mathbf{g}} = \mathbf{g}^1$. We align in q-space each other $k \geq 2$ DWI acquisition so that their gradients exactly match $\tilde{\mathbf{g}}$. This is done by using interpolation in q-space. At each voxel, we consider the G intensity values corresponding to the gradients $\mathbf{g}^k$. We then compute the new intensity values corresponding to the gradients $\tilde{\mathbf{g}}$. The interpolation is performed via Kriging [8], a general and efficient statistical interpolation framework originally introduced for geology and mining applications. This enables us to easily perform scattered data interpolation. It determines the weights of the contribution of each observed data via the resolution of a simple linear system. In the absence of motion (*i.e.* a gradient g_j exactly matches a gradient $\tilde{g}_{j'}$), the interpolated intensity match exactly the observed intensity. As a result, the K LR acquisitions are all aligned in space and represent the same sensitizing-gradient set $\tilde{\mathbf{g}}$.

Super-Resolution Model-Based Reconstruction. In this section each LR volume is considered to be aligned in space and in q-space. The SRR technique we now address is performed for each gradient image separately. For each gradient $\tilde{g}_j$, we aim to recover the HR image $\mathbf{x}_j$ underlying the K LR images $\mathbf{y}_j = (\mathbf{y}_{1,j}, \ldots, \mathbf{y}_{K,j})$. By omitting the gradient dependency to simplify the notations,

we consider the K LR volumes $\mathbf{y} = (\mathbf{y}_1, \ldots, \mathbf{y}_K)$ to be the degraded version of the same unknown HR volume $\mathbf{x}$. $\mathbf{x}$ is estimated according to the maximum *a posteriori* principle, by maximizing:

$$\widehat{\mathbf{x}}_{\mathrm{MAP}} = \arg\max_{\mathbf{x}} p(\mathbf{x}|\mathbf{y}) = \arg\max_{\mathbf{x}} p(\mathbf{y}|\mathbf{x})p(\mathbf{x}) = \arg\max_{\mathbf{x}} \left[\ln p(\mathbf{y}|\mathbf{x}) + \ln p(\mathbf{x})\right] , \quad (1)$$

which decomposes into a likelihood term and a prior term.

Likelihood term $p(\mathbf{y}|\mathbf{x})$: we consider a model describing how the LR volumes are obtained from the unknown HR volume. In a trade-off between a realistic model and a feasible solution, we consider the following acquisition model:

$$\mathbf{y}_k = \mathbf{D}_k\mathbf{B}_k\mathbf{M}_k\mathbf{x} + \epsilon_k , \quad (2)$$

where the volumes $\mathbf{y}_k$ and $\mathbf{x}$ are expressed as column vectors by a lexicographical reordering of the pixels, $\mathbf{D}_k$ is a down-sampling matrix, $\mathbf{M}_k$ is the warping matrix that maps the HR volume $\mathbf{x}$ to the LR volume $\mathbf{y}_k$ and ϵ_k is the residual noise vector. $\mathbf{B}_k$ is the blur, or point spread function (PSF) of the MRI signal acquisition process. It is constructed from the imaging parameters. The PSF can be separated into three components corresponding to the slice-selection direction and the phase- and frequency-encoding directions. As in [3], we currently consider a PSF in the slice-selection direction only, which describes the slice selection profile. We consider a Gaussian slice selection profile of variance σ^2_{PSF}. Consequently, on the basis of Eq.2, the unknown HR volume $\mathbf{x}$ goes through geometric and signal operations, including motion, signal averaging, and resampling, to generate the acquired LR volume $\mathbf{y}_k$. Assuming a Gaussian noise with zero-mean and variance σ_k for ϵ_k, the likelihood of the LR volume $\mathbf{y}_k$ under the model in Eq.2 can be written as:

$$p(\mathbf{y}_k|\mathbf{x}, \sigma_k) = \frac{1}{\sigma_k\sqrt{2\pi}} \exp\left(-\frac{||\mathbf{y}_k - \mathbf{D}_k\mathbf{B}_k\mathbf{M}_k\mathbf{x}||^2}{2\sigma_k^2}\right) . \quad (3)$$

Assuming statistical independence of the noise between the acquisitions, we have $p(\mathbf{y}|\mathbf{x}, \sigma) = \prod_{k=1}^{K} p(\mathbf{y}_k|\mathbf{x}, \sigma_k)$ with $\sigma = (\sigma_1, \ldots, \sigma_K)$.

Prior term $p(\mathbf{x})$: the term $p(\mathbf{x})$ in Eq.1 enables us to incorporate a prior knowledge on $\mathbf{x}$. In this work we consider a regularization prior that exploits spatial homogeneity. More precisely, we favour smoothness of $\mathbf{x}$ by setting $p(\mathbf{x}|\lambda) = \exp(-\lambda||\mathbf{Q}\mathbf{x}||^2)$ where the matrix $\mathbf{Q}$ is symmetric definite positive and represents a linear high-pass operation. The parameter λ controls the regularization strength. In this work, $\mathbf{Q}$ is chosen as the 3-D discrete Laplacian corresponding to the following approximation of the partial derivative for a 3-D image $\mathbf{I}$ indexed by $\mathbf{u} \in I\!N^3$, and for a direction $\mathbf{u}_m \in I\!N^3$ ($m \in \{1, 2, 3\}$):

$$\partial_m\mathbf{I}(\mathbf{u}) \approx \Big(\mathbf{I}(\mathbf{u} + \mathbf{u}_m) - 2\mathbf{I}(\mathbf{u}) + \mathbf{I}(\mathbf{u} - \mathbf{u}_m)\Big)/(2||\mathbf{u}_m||). \quad (4)$$

Ultimately, by considering the same σ_k across acquisitions, maximization of the posterior distribution in Eq.1 leads to the following minimization:

$$\widehat{\mathbf{x}} = \arg\min_{\mathbf{x}} \sum_{k=1}^{K} ||\mathbf{y}_k - \mathbf{D}_k\mathbf{B}_k\mathbf{M}_k\mathbf{x}||^2 + \lambda||\mathbf{Q}\mathbf{x}||^2 . \quad (5)$$

DWI-SRR Optimization Procedure. The matrix $\mathbf{D}_k\mathbf{B}_k\mathbf{M}_k$ is especially large and the classical solution through pseudo-inverse is prohibitive. Instead we use a steepest descent iterative minimization approach. Differentiation of Eq.5 leads to the following update at each step:

$$\widehat{\mathbf{x}}^{n+1} = \widehat{\mathbf{x}}^{n} - \alpha \left[\sum_{k=1}^{K} \mathbf{M}_k^T \mathbf{B}_k^T \mathbf{D}_k^T \left(\mathbf{D}_k \mathbf{B}_k \mathbf{M}_k \widehat{\mathbf{x}}^n - \mathbf{y}_k \right) - \lambda Q^T Q \widehat{\mathbf{x}}^n \right], \qquad (6)$$

where α is the step size and $\mathbf{M}_k^T$ denotes the transpose of $\mathbf{M}_k$. The iterative algorithm is initialized by setting $\widehat{\mathbf{x}}^0$ equal to the mean of the aligned LR volumes. The iterative minimization is stopped when $||\widehat{\mathbf{x}}^{n+1} - \widehat{\mathbf{x}}^n||_1 < \mathcal{T}_{\text{STOP}}$. The DWI-SRR optimization procedure is synthesized by the following *pseudo-code*:

```
FOR each LR study k
  FOR each gradient image of k
    M_k^T ← Register to the reference volume
    Apply the transform to the gradient
  ENDFOR
ENDFOR
FOR each output gradient g_g
  FOR each LR study k
    y_k ← Compute the gradient image for g_g (q-space interpolation)
  ENDFOR
  Compute x̂^0 ← Mean of the M_k^T D_k^T y_k
  WHILE ||x̂^n − x̂^(n−1)||_1 ≥ T_STOP
    x̂^(n+1) ← Update with Eq. 6
  ENDWHILE
ENDFOR
```

3 Results

The SRR procedure was implemented in C++ and optimized with various techniques to reduce the processing burden. The $\mathbf{M}_k^T\mathbf{B}_k^T\mathbf{D}_k^T\mathbf{D}_k\mathbf{B}_k\mathbf{M}_k$ and $\mathbf{M}_k^T\mathbf{B}_k^T\mathbf{D}_k^T\mathbf{y}_k$ matrices were precomputed, and the derivative of the Laplacian corresponding to the finite difference scheme in Eq.4 computed analytically. To accelerate the convergence, the steepest descent algorithm was implemented with a variable per-voxel step-size α which incorporates inertia: the step-size is multiplied by 1.1 when the sign of two consecutive computed gradient does not change, and divided by two otherwise. α is initialized to 0.01 and constrained to lie in $\left[0.1, 10^{-6}\right]$. The FWHM for the Gaussian slice model was set to half the slice thickness, by setting $\sigma_{\text{PSF}} = (\text{slice thickness})/(4\sqrt{2\ln 2})$. Other parameters were set to $\lambda = 0.001$ and $\mathcal{T}_{\text{STOP}} = 10^{-5}$.

Numerical Simulations. We first performed numerical simulations. The DW-signal for tensors (see Fig.2a) was simulated with a b-value of 1000s/mm^2 for 15 directions, and corrupted by Rician noise (SNR of 30dB for the b=0s/mm^2 image). Linear down-sampling with a factor of 4 was applied to each gradient

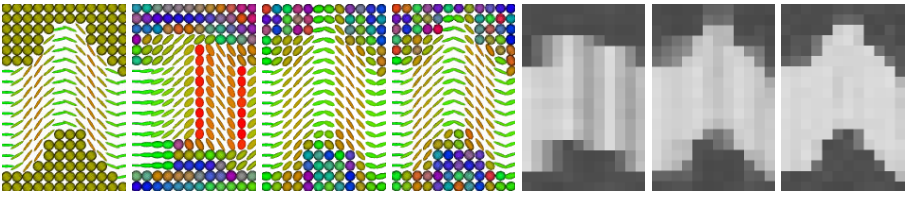

Fig. 2. Numerical simulations from the tensors of Fig.a. Fig.b-d: Tensors estimated resp. from a single LR acquisition, from the mean of the LR acquisitions and from the SRR. Fig.e-g: Corresponding tensor fractional anisotropy. It shows the tensor directions to be well estimated from the mean (Fig.c). However, the SRR provides a much more accurate reconstruction of the complete tensor (see the better FA uniformity in Fig.g).

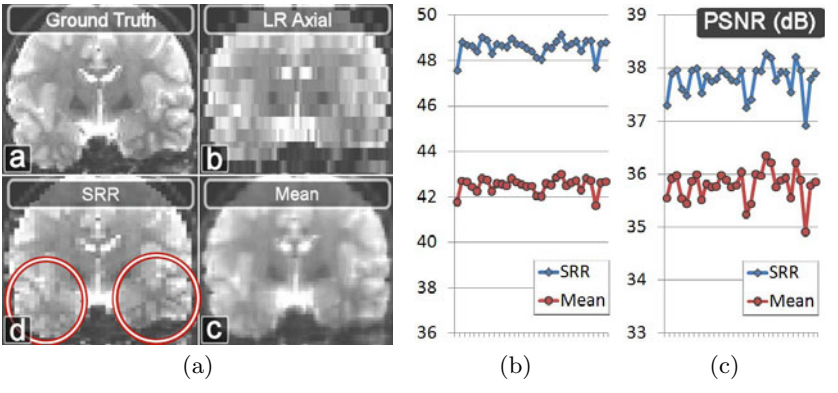

Fig. 3. Fig.a: Synthetic SRR scenario from a real acquisition. a.a: b=0 image. a.b: Axial down-sampled b=0 image with a factor of 4. a.c: Mean of the b=0 images of the LR acquisitions. a.d: SRR of the b=0 image. The SRR is better contrasted and is less blurry than the mean. Fig.b and Fig.c: Quantitative evaluation of the reconstruction accuracy in term of PSNR for the /2 and /4 down-sampling, for *each* gradient direction.

image in each of the three directions, providing three simulated LR acquisitions. Various tensor estimation were performed (see Fig.2b-d), and the corresponding fractional anisotropy (FA) computed (Fig.2e-g) (ground truth FA=0.8).

Synthetic SRR Scenario. We then simulated a SRR scenario by down-sampling in each of the three directions a real DWI acquisition (Siemens 3T Trio, 32 channel head coil, 68 slices, FOV=220mm, matrix=128x128, resolution=1.7x1.7x2mm^3, TE=86ms/TR=8800ms, 30 directions at B=1000s/mm^2, 5 B=0s/mm^2). A down-sampling of factor 2 and 4 were considered. The SRR at the original resolution was estimated and qualitatively compared to the original image (see Fig.3a). The SRR estimation time was approximately 2 hours on a

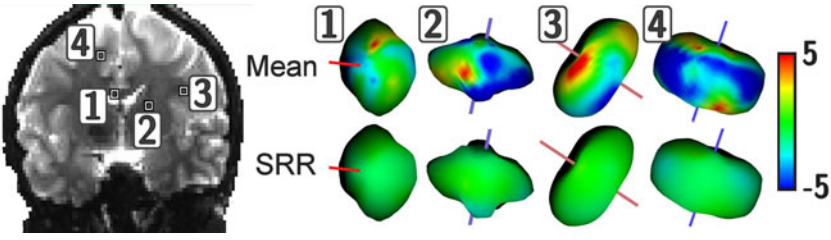

Fig. 4. 3-Dimensional angular reconstructions of the diffusion signal at four voxels whose position is shown on the b=0s/mm^2 image (left image). The voxels were chosen to have a high FA ($FA > 0.9$). The obtained 3-D shapes are proportional to the apparent diffusion coefficient (ADC). Comparison between the 3-D reconstruction performed from the mean image (first line) and from the SRR estimate (second line). The stick indicates the major fiber direction estimated by a single-tensor model. The color encodes for the error with the ground-truth (difference in image intensity). It shows the SRR estimate to provide a much better reconstruction for each gradient image.

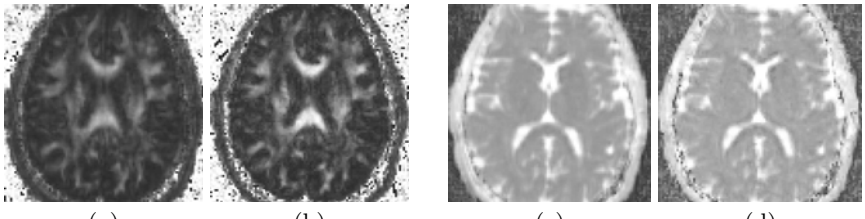

Fig. 5. Real SRR scenario. Fig.a-b: FA computed from the mean of the LR acquisition (a) and from the SRR (b). Fig.c-d: idem for MD. It shows that the SRR leads to more contrasted and less blurry FA and MD estimates.

3Ghz Intel Xeon (3 to 4min per gradient image). Fig.3b-c report, for each gradient direction, the Peak Signal to Noise Ratio (PSNR) with the original acquisition. The PSNR is defined by $20\log_{10}(\text{MAX}/\sqrt{\text{MSE}})$ with MAX the maximum intensity and MSE the mean square error. Fig.4 synthesizes, for four voxels, the error with the ground-truth for all the gradient directions via a 3-D angular reconstruction.

Real SRR scenario. Finally, we performed the acquisition of three anisotropic DWI scans (same parameters as before, except: 1.6x1.6mm^2 in-plane res., 5mm slice thickness, 38 slices, TE/TR=87/4700ms) and achieved the reconstruction at 1.6x1.6x 2.5mm^3. Fig.5 shows the FA and the mean-diffusivity (MD) computed from the mean of the LR acquisitions and from our SRR technique. As in Fig.2, the images are less blurry and more contrasted when employing our SRR technique.

4 Discussion

We have proposed a novel SRR technique for DWI based on the acquisition of orthogonal anisotropic DWI scans. To our knowledge, it is the first attempt to perform SRR in DWI in the last decade, since [9]. In contrast to [9], we take into account possible patient motions by aligning the volumes in both space *and* q-space. In addition, we formulate the SRR as a MAP estimation problem. For each gradient, we estimate the underlying unknown HR volume given the acquired LR DWI scans. Our formulation enables us to integrate an image acquisition model and to integrate image priors. We have shown that the SRR estimate outperforms the mean of the LR acquisitions, providing a better contrast and less blurry results for the gradient images (Fig.3a) and for diffusion parameters such as the FA or the MD (Fig.2 and Fig.5). The quantitative evaluation showed an increase of PSNR on the order of 6dB and 2dB for our two synthetic down-sampling scenarios (Fig.3b-c). In future work, we will evaluate the benefits of correcting for the geometric distortions. Indeed, the acquisitions show locally very different geometric distortion patterns due to different phase-encoding directions. As a result, accurate alignment of the images is difficult, which can locally perturb the SRR. We will investigate the effectiveness of employing a distortion correction technique by acquisition of a magnetic field map [10] or by the acquisition of two DWI scans with reversed phase directions [5]. Finally, we will investigate the introduction of novel priors. Particularly, incorporating the brain anatomy description provided by a HR T2-w scan may enable improved SRR in DWI.

SRR techniques are of great interest for medical imaging because they enable us to go beyond the limits dictated by the hardware. With today's scanners, they may enable routine HR investigation of the brain white-matter in clinically compatible scan time. Combined with future MRI hardware improvements, they may enable DWI to be performed with unprecedented resolution.

Acknowledgement. This investigation was supported in part by NIH grants R01 RR021885, R01 EB008015, R03 EB008680, R01 LM010033 and UL1 RR025758-03.

References

1. Calamante, F., Tournier, J.-D., Jackson, G.D., Connelly, A.: Track-density imaging (TDI): super-resolution white matter imaging using whole-brain track-density mapping. NeuroImage 53(4), 1233–1243 (2010)
2. Gholipour, A., Estroff, J.A., Warfield, S.K.: Robust super-resolution volume reconstruction from slice acquisitions: application to fetal brain MRI. IEEE Trans. on Med. Imag. 29(10), 1739–1758 (2010)
3. Greenspan, H.: MRI inter-slice reconstruction using super-resolution. Magn. Res. Imag. 20(5), 437–446 (2002)
4. Greenspan, H.: Super-Resolution in Medical Imaging. Comput. J. 52(1), 43–63 (2009)

5. Holland, D., Kuperman, J.M., Dale, A.M.: Efficient correction of inhomogeneous static magnetic field-induced distortion in echo planar imaging. NeuroImage 50(1), 175–183 (2010)
6. Irani, M., Peleg, S.: Motion Analysis for Image Enhancement: Resolution, Occlusion, and Transparency. J. Vis. Com. and Image Rep. 4(4), 324–335 (1993)
7. Jiang, S., Xue, H., Counsell, S., Anjari, M., Allsop, J., Rutherford, M., Rueckert, D., Hajnal, J.V.: Diffusion tensor imaging (DTI) of the brain in moving subjects: application to in-utero fetal and ex-utero studies.. Magn. Res. in Med. 62(3), 645–655 (2009)
8. Matheron, G.: Principles of geostatistics. Economic Geology 58(8), 1246–1266 (1963)
9. Peled, S., Yeshurun, Y.: Superresolution in MRI: application to human white matter fiber tract visualization by diffusion tensor imaging.. Magn. Res. in Med. 45(1), 29–35 (2001)
10. Windischberger, C., Robinson, S., Rauscher, A., Barth, M., Moser, E.: Robust field map generation using a triple-echo acquisition. J. Magn. Reson. Imaging 20(4), 730–734 (2004)

Reconstruction of Fiber Trajectories via Population-Based Estimation of Local Orientations

Pew-Thian Yap[1], John H. Gilmore[2], Weili Lin[1], and Dinggang Shen[1]

[1] BRIC, Department of Radiology and
[2] Department of Pyschiatry
University of North Carolina at Chapel Hill, NC
{ptyap,jgilmore,weili_lin,dgshen}@med.unc.edu

Abstract. White matter fiber tractography plays a key role in the *in vivo* understanding of brain circuitry. For tract-based comparison of a population of images, a common approach is to first generate an atlas by averaging, after spatial normalization, all images in the population, and then perform tractography using the constructed atlas. The reconstructed fiber trajectories form a common geometry onto which diffusion properties of each individual subject can be projected based on the corresponding locations in the subject native space. However, in the case of High Angular Resolution Diffusion Imaging (HARDI), where modeling fiber crossings is an important goal, the above-mentioned averaging method for generating an atlas results in significant error in the estimation of local fiber orientations and causes a major loss of fiber crossings. These limitatitons have significant impact on the accuracy of the reconstructed fiber trajectories and jeopardize subsequent tract-based analysis. As a remedy, we present in this paper a more effective means of performing tractography at a population level. Our method entails determining a bipolar Watson distribution at each voxel location based on information given by all images in the population, giving us not only the local principal orientations of the fiber pathways, but also confidence levels of how reliable these orientations are across subjects. The distribution field is then fed as an input to a probabilistic tractography framework for reconstructing a set of fiber trajectories that are consistent across all images in the population. We observe that the proposed method, called PopTract, results in significantly better preservation of fiber crossings, and hence yields better trajectory reconstruction in the atlas space.

1 Introduction

Diffusion Tensor Imaging (DTI) is a powerful imaging modality that allows probing into the intricate micro-architecture of white matter. It plays an indispensable role in characterizing neural pathways *in vivo* by means of fiber tractography, which entails reconstructing the trajectories of fiber paths by tracing the direction of maximal water diffusion. Such estimated fiber paths can subsequently be used, for instance, to investigate brain connectivity alterations in

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 133–140, 2011.

mental and neurological disorders. Tract pathology can be evident, as in the case of brain tumors where tracts are grossly displaced, or subtle, as in the case of neuropsychiatric disorders, such as schizophrenia where disruptions are manifested as change of diffusion properties within the tracts.

The core assumption of DTI — Gaussianity of water diffusion — however, does not always hold true. Modeling water diffusion using the single tensor formulation ignores this complexity and results in loss of information. One of the methods proposed to remedy the shortcomings of DTI is High Angular Resolution Diffusion Imaging (HARDI) [1], where diffusion signals are acquired along a significantly larger number of directions than is normally employed in DTI. This allows modeling of the complex non-Gaussian diffusion process and construction of spherical functions with multiple local maxima which are potentially aligned with the underlying fiber bundle orientations. Accordingly, one naturally wants to extend existing DT-based tractography algorithms to work with HARDI data to better deal with fiber crossings and to reconstruct fiber trajectories that resemble more closely the anatomical fibers connecting different functional regions of the brain.

In this paper, we detail a tractography algorithm, called PopTract, that allows effective reconstruction of fiber trajectories in the atlas space by more faithfully preserving fiber crossings. Tractography in the atlas space, as was done in previous works [2], allow the reconstructed trajectories to form a common geometry onto which diffusion properties from the individual images can be projected for tract-based comparison. To this end, a common step involves averaging the tensors [2], the Fiber Orientation Distributions (FODs) [3], or the diffusion signals [4] after spatial normalization, resulting in an atlas on which tractography can be performed after determining the local fiber orientations. We will show, however, that this approach causes significant deviation of the estimated local orientations from the 'true' orientations. The major reason for this is the smearing of the orientation profile, caused by averaging misaligned ODFs.

2 Approach

The main idea of our approach is to model the orientation distribution at each voxel location by simultaneously considering all images in the population. Similar to the average atlas approach, we want to use this orientation information to reconstruct a set of fiber trajectories that will form a common datum, based on which inter-subject tract-based morphometry can be performed. Unlike the average atlas approach, however, orientations are not estimated after, but prior to, averaging in the common space to avoid estimation inaccuracy caused by ODF smearing. The estimated orientation fields are transformed to a common atlas space where the maximum likelihood parameter estimates of the Watson distribution at each voxel location are determined. A probabilistic tractography algorithm is then be used to reconstruct the fiber trajectories.

2.1 Modeling Local Fiber Orientations

Assuming for a moment that the orientation fields computed from N diffusion-weighted images are spatially normalized to a common space and that each voxel location of an orientation field contains only one fiber orientation, our goal now is to accurately relate these orientations to an underlying model. At each voxel location of orientation field $i \in \{1, \ldots, N\}$, the local fiber orientation is denoted by a unit vector $\mathbf{v}_i$. We model the distribution of orientations across subjects using the bipolar Watson distribution, probability density function (PDF) of which is given by [5] $f(\mathbf{v}|\boldsymbol{\mu}, \kappa) = C(\kappa)\mathrm{e}^{\kappa(\boldsymbol{\mu}^{\mathrm{T}}\mathbf{v})^2}$. The parameter $\boldsymbol{\mu}$ is a unit vector called the *mean orientation* and κ is a positive constant called the *concentration parameter*. The density has maxima at $\pm\boldsymbol{\mu}$ and becomes more concentrated around $\pm\boldsymbol{\mu}$ as κ increases. The density is also rotationally invariant around $\pm\boldsymbol{\mu}$. $C(\kappa)$ is a normalizing constant to ensure that the density function integrates to unity over the unit sphere.

Parameter Estimation. (*Single Fiber Orientation*) Assuming that the orientation vectors obtained from the images, $\mathbf{v}_1, \ldots, \mathbf{v}_N$, are random samples from the Watson distribution, the maximum likelihood estimate (MLE) of $\boldsymbol{\mu}$ is the eigenvector corresponding to the largest eigenvalue λ_1 ($\lambda_1 \geq \lambda_2 \geq \lambda_3$) of the symmetric positive definite matrix [5] $\mathbf{A} = \frac{1}{N}\sum_{i=1}^{N} \mathbf{v}_i\mathbf{v}_i^{\mathrm{T}}$. Matrix $\mathbf{A}$ is called the *dyadic tensor*. The MLE of κ is $(1 - \lambda_1)^{-1}$, asymptotically when $\kappa \to \infty$ [5]. When $\kappa = 0$, the distribution is uniform. As κ increases, the PDF becomes more concentrated about $\pm\boldsymbol{\mu}$. Therefore, orientations that depart from $\pm\boldsymbol{\mu}$ are penalized more heavily when there is a strong alignment of the local fiber orientations from all subjects. (*Extension to Multiple Fiber Orientations*) Denoting the maximum number of possible orientations as Ω, we now denote the set of direction vectors at each voxel as $\{\mathbf{v}_i^{[D]}\}$, where $D = 1, \ldots, \Omega$. In the case where there is less than D orientations, the surplus orientation vectors are simply set to nil. For DTI, Ω is limited to 1. For HARDI, however, the representation models used often allows more than one fiber orientations, hence $\Omega \geq 1$. Assuming the presence of multiple compartments in each voxel, each representing the local structure of a fiber bundle, and that there is no exchange between the different compartments, we estimate the parameters of the PDF for each compartment independently. That is, for compartment D, we estimate $\boldsymbol{\mu}^{[D]}$ and $\kappa^{[D]}$ using $\mathbf{v}_1^{[D]}, \ldots, \mathbf{v}_N^{[D]}$.

2.2 Reconstructing Fiber Trajectories

A white matter fiber can be modeled as a finite-length path parameterized by a train of unit length vectors. We use the following notation for such a path: $\mathbf{v}_{(1:T)} = \{\mathbf{v}_{(1)}, \ldots, \mathbf{v}_{(T)}\}$. We further assume that a fiber path can be traced by tracking the trajectory of a particle traveling in an orientation field. Each particle is endowed with an initial speed in an appropriate direction. It then moves with constant speed to position $\mathbf{p}_{(t)}$ according to $\mathbf{p}_{(t+1)} = \mathbf{p}_{(t)} + s\mathbf{v}_{(t)}$, where s is the step length. At each point in space, vector $\mathbf{v}_{(t)}$ is drawn from distribution

$f(\mathbf{v}_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}_{(t)}, D_{(t)})$, where the set of parameters for the Watson distributions are collectively denoted as $\boldsymbol{\theta}_{(t)} = \{\boldsymbol{\mu}_{(t)}^{[1]}, \ldots, \boldsymbol{\mu}_{(t)}^{[\Omega]}, \kappa_{(t)}^{[1]}, \ldots, \kappa_{(t)}^{[\Omega]}\}$. The above distribution is in fact the Watson distribution discussed in Section 2.1. It is, however, now dependent on the prior orientation vector $\mathbf{v}_{(t-1)}$ in determining which orientation compartment $D_{(t)} = 1, \ldots, \Omega$ to follow in the case where a voxel contains multiple orientations. Assuming that the Watson distribution is not directly dependent on $\mathbf{v}_{(t-1)}$ but only on $D_{(t)}$, we can simplify $f(\mathbf{v}_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}_{(t)}, D_{(t)})$ to become $f(\mathbf{v}_{(t)}|\boldsymbol{\theta}_{(t)}, D_{(t)})$. Specifically,

$$f(\mathbf{v}_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}_{(t)}, D_{(t)}) = f(\mathbf{v}_{(t)}|\boldsymbol{\theta}_{(t)}, D_{(t)}) = C\left(\kappa_{(t)}^{[D_{(t)}]}\right) \mathrm{e}^{\kappa_{(t)}^{[D_{(t)}]}((\boldsymbol{\mu}_{(t)}^{[D_{(t)}]})^{\mathrm{T}}\mathbf{v}_{(t)})^2} \tag{1}$$

and

$$f(D_{(t)}|\mathbf{v}_{(t-1)}, \boldsymbol{\theta}_{(t)}) \propto \begin{cases} \bar{\rho}_{(t)}^{[D_{(t)}]} \left[\mathbf{v}_{(t-1)}^{\mathrm{T}} \boldsymbol{\mu}_{(t)}^{[D_{(t)}]}\right]^2, & \left|\mathbf{v}_{(t-1)}^{\mathrm{T}} \boldsymbol{\mu}_{(t)}^{[D_{(t)}]}\right| \geq \cos(\phi) \\ 0, & \text{otherwise} \end{cases} \tag{2}$$

where ϕ is the maximum allowed turning angle and $\bar{\rho}_{(t)}^{[D]}$ is defined as $\bar{\rho}_{(t)}^{[D]} = \frac{1}{N}\sum_{i=1}^{N} \rho_{i,(t)}^{[D]}$, with $\rho_{i,(t)}^{[D]}$ being the value of the ODF at the voxel location determined by $\mathbf{p}_{(t)}$, sampled at orientation $\mathbf{v}^{[D]}$. The tracing is stopped if the trajectory reaches a voxel with *orientation coherence* $\beta^{[D]} = 1 - \sqrt{\frac{\lambda_2^{[D]} + \lambda_3^{[D]}}{2\lambda_1^{[D]}}}$, $\beta^{[D]} \in [0, 1]$ falling below a predefined threshold β_0, or simply when the brain boundary is encountered. The λ's are the eigenvalues of the dyadic tensor at each voxel location. Perfect alignment of the orientations in compartment D results in $\beta^{[D]} = 1$ and an uniform distribution of orientations results in $\beta^{[D]} = 0$.

3 Results

3.1 Synthesized Dataset

To evaluate the effectiveness of the proposed method in preserving fiber crossings and in correctly estimating the orientations, we synthesized an 8×8 image (Camino [6]; two-tensor model) with each voxel containing a crossing with two orientations - one vertical and one horizontal. To simulate inter-subject variability, we perturbed the synthesized diffusion-weighted signals by applying random rotation matrices (angles: 15°, 30°, 45°, 60°) and adding isotropic complex Gaussian noise (SNR = 8, 16) to each voxel. The ODF at each voxel is then computed using Camino [6]. We applied these perturbations 10 times to the no-noise image and then used the resultant images to attempt to recover the 'true' population-based orientations. We show, for qualitative evaluation, the results for the case of 45° and SNR = 16 in Fig. 1. It can be observed that the average atlas method (Fig. 1(b)), generated by averaging the respective ODFs, resulted in loss of fiber crossings and significant deviation in the estimated orientations. The proposed method (Fig. 1(c)) gave a more consistent result.

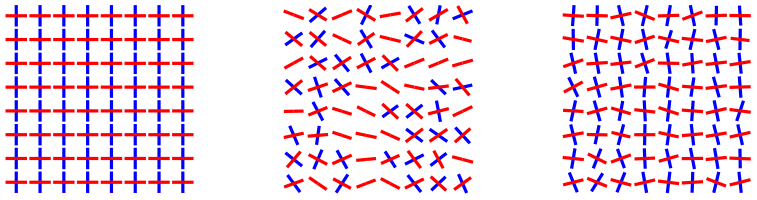

(a) Ground Truth (b) Average Atlas (c) PopTract

Fig. 1. Estimation of orientations using different schemes

For quantitative evaluation, we measured the *orientational discrepency* (OD) of the estimated orientations with respect to the ground truth orientations. Assuming that $\mathbf{U}(\mathbf{x})$ is the set of ground truth directions at voxel location $\mathbf{x}$ and $\mathbf{V}(\mathbf{x})$ is the corresponding set of estimated directions, OD is defined as

$$\mathrm{OD}(\mathbf{x}) = \frac{1}{2}\Big[\max_{\mathbf{u}\in\mathbf{U}(\mathbf{x})}\min_{\mathbf{v}\in\mathbf{V}(\mathbf{x})} d_\theta(\mathbf{u},\mathbf{v}) + \max_{\mathbf{v}\in\mathbf{V}(\mathbf{x})}\min_{\mathbf{u}\in\mathbf{U}(\mathbf{x})} d_\theta(\mathbf{v},\mathbf{u})\Big] \tag{3}$$

where $d_\theta(\mathbf{u},\mathbf{v})$ gives the angle difference between $\mathbf{u}$ and $\mathbf{v}$, i.e., $d_\theta(\mathbf{u},\mathbf{v}) = \cos^{-1}(|\mathbf{u}\cdot\mathbf{v}|)$. The absolute value is taken since diffusion is assumed to be antipodal symmetric. In cases of multiple local maxima, the term $\min_{\mathbf{v}\in\mathbf{V}} d_\theta(\mathbf{u},\mathbf{v})$ returns the angle difference between $\mathbf{u}$ with an orientation $\mathbf{v}$ in $\mathbf{V}$ that is most closely aligned with itself. Evaluating the OD between the estimated orientations with the ground truth under different rotation angles and SNRs, the results, shown in Fig. 2, indicate that the proposed method is capable of estimating orientations which are closer to the ground truth, with improvement especially prominent when the angles of rotation is large.

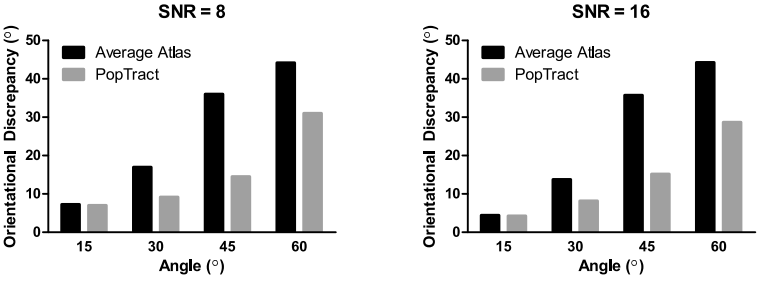

Fig. 2. Orientational dicrepancy between the estimated orientations and the ground truth orientations under different rotation angles and signal-to-noise ratios.

3.2 In Vivo Dataset

Materials. Diffusion-weighted images were acquired for 14 adult subjects using a Siemens 3T TIM Trio MR Scanner with an EPI sequence. Diffusion

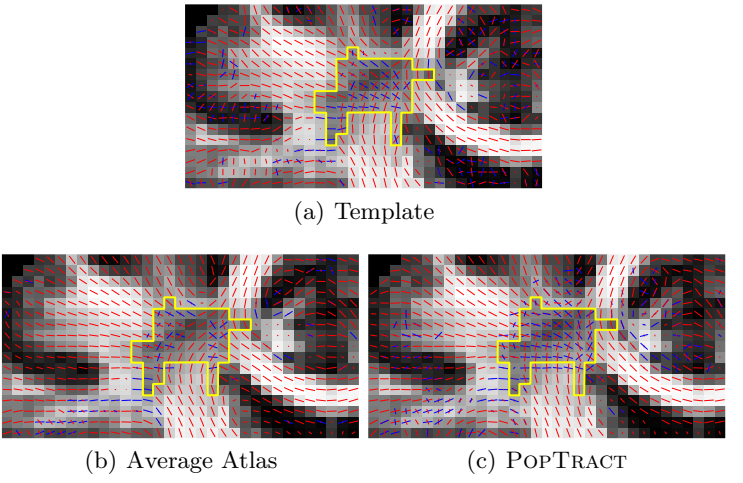

(a) Template

(b) Average Atlas (c) PopTract

Fig. 3. Fiber crossings. A significant loss of crossings can be observed for the average atlas method.

gradients were applied in 120 non-collinear directions with diffusion weighting $b = 2000$ s/mm^2, flip angle = 90°, repetition time (TR) = 12,400 ms and echo time (TE) = 116 ms. The imaging matrix was 128×128 with a rectangular FOV of 256×256 mm^2. 80 contiguous slices with a slice thickness of 2 mm covered the whole brain. Data post-processing includes brain skull removal, motion correction and eddy current correction using algorithms developed and distributed as part of the FMRIB Software Library (FSL) package.

Spatial Normalization and Orientation Estimation. The local maxima of an ODF reflect local fiber orientations. In our case, the orientations were computed with Camino [6] by locating the peaks of the ODFs using Powell's algorithm on a set of sample points generated by randomly rotating a unit icosahedron 1000 times. We limited the maximum number of orientations for each voxel to two. This choice is guided by several previous studies [7,8], where the estimation of two orientations is generally deemed as stable. One image was selected, out of the 14 images, as the template onto which 13 other images were registered using a deformable spherical-harmonics-based HARDI registration algorithm [9] to align the ODFs. The estimated transformations were used to map the respective orientation fields to a common space. When transforming the orientations, they were reoriented using $\mathbf{v}' = \mathbf{F}\mathbf{v}/||\mathbf{F}\mathbf{v}||$, where $\mathbf{F}$ is a local affine transform matrix computed from the image transformation map. For generating the average atlas, transformation of the ODF was performed based on a method similar to that proposed by Rafflet et al. [3]; we approximated the ODF using a number of Point Spread Functions (PSFs), reoriented these PSFs individually, and then recomposed the reoriented PSFs to obtain the transformed ODF. The

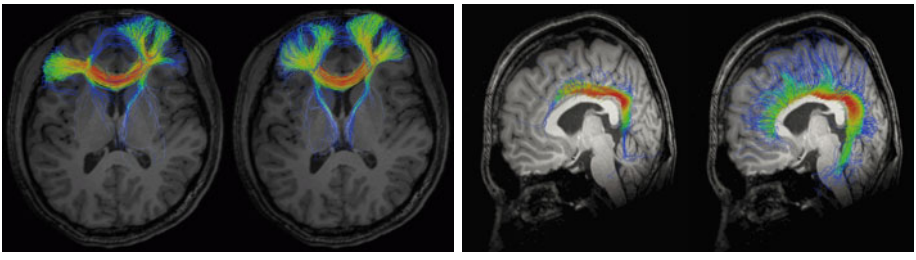

(a) Forceps Minor (b) Cingulum Bundle

Fig. 4. Fiber trajectory reconstruction of the (a) forceps minor and (b) cingulum bundle. The results given by the average atlas method and PopTract are shown on the left and right, respectively. The coloring indicates the probability of finding a fiber at a specific spatial location. Dark red indicates a high probability of a particular location being traversed by fibers, and dark blue indicates otherwise.

same exact set of deformation fields, as applied to the orientation fields, were used to align the ODF fields. The average ODF at each voxel was then used to estimate the local fiber orientations.

Crossings. Similar to the results shown for the synthesized data (Fig. 1), significant loss of fiber crossings can be observed in the case of the *in vivo* data when the average atlas method is used for estimating the orientations (Fig. 3).

Tractography. For evaluation, we considered the following commonly studied fiber bundles: 1) The forceps minor, also known as the anterior forceps, is a fiber bundle which connects the lateral and medial surfaces of the frontal lobes and crosses the midline via the genu of the corpus callosum; 2) The cingulum is a medial associative bundle that runs within the cingulate gyrus all around the corpus callosum. To test whether these fiber bundles were preserved by the proposed method, we placed single voxel seeds at the points where the midline crosses the genu of the corpus callosum, and also a seed along the cingulate gyrus, and performed tractography based on these seeds. The tractography parameters used were $\phi = 70°$, $\beta_0 = 0.1$ and $s = 1$. 3000 fibers were initiated from each seed. An identical set of parameters was used for the average atlas method. The results, shown in Fig. 4, indicate that the average atlas method results in premature termination of the tracking of the fiber bundles. Specifically, in Fig. 4(a), we find that PopTract, as opposed to the average atlas method, not only reconstructed the full forceps minor, with both frontal extensions of the fiber bundles intact, but also part of the anterior thalamic radiations. Similar conclusion can be made from Fig. 4(b), where PopTract shows a complete reconstruction of the cingulum bundle and the U-shaped fibers connecting the medial frontal, parietal, occipital and temporal lobes with different portions of the cingulate cortex. PopTract, therefore, in contrast to the conventional average atlas method, gives a more complete reconstruction of the fiber trajectories, which are more consistent with our anatomical knowledge of the bundles.

4 Conclusion

We have presented a tractography algorithm, called PopTract, that is more effective in preserving fiber crossings and is more accurate in estimating local fiber orientations. PopTract results in more reasonable reconstruction of the fiber trajectories that are in closer agreement with known fiber bundles.

Acknowledgment. This work was supported in part by NIH grants: EB006733, EB008374, EB009634, MH088520, HD05300, MH064065, and NS055754.

References

1. Tuch, D.S., Weisskoff, R.M., Belliveau, J.W., Wedeen, V.J.: High angular resolution diffusion imaging of the human brain. In: ISMRM 1999 (1999)
2. Goodlett, C.B., Fletcher, P.T., Gilmore, J.H., Gerig, G.: Group analysis of DTI fiber tract statistics with application to neurodevelopment. NeuroImage 45(1), S133–S142 (2009)
3. Raffelt, D., Tournier, J., Fripp, J., Crozier, S., Connelly, A., Salvado, O.: Symmetric diffeomorphic registration of fibre orientation distributions. NeuroImage (2011)
4. Bouix, S., Rathi, Y., Sabancu, M.: Building an average population atlas. In: MICCAI 2010 Workshop on Computational Diffusion MRI (2010)
5. Schwartzman, A., Dougherty, R.F., Taylor, J.E.: False discovery rate analysis of brain diffusion direction maps. Annals of Applied Statistics 2(1), 153–175 (2008)
6. Cook, P.A., Bai, Y., Nedjati-Gilani, S., Seunarine, K.K., Hall, M.G., Parker, G.J., Alexander, D.C.: Camino: Open-source diffusion-MRI reconstruction and processing. In: 14th Scientific Meeting of the International Society for Magnetic Resonance in Medicine, p. 2759 (2006)
7. Tuch, D., Reese, T., Wiegell, M., Makris, N., Belliveau, J., Wedeen, V.: High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. Magnetic Resonance in Medicine 48, 577–582 (2002)
8. Peled, S., Friman, O., Jolesz, F., Westin, C.F.: Geometrically constrained twotensor model for crossing tracts in DWI. Magnetic Resonance Imaging 24(9), 1263–1270 (2006)
9. Yap, P.T., Chen, Y., An, H., Yang, Y., Gilmore, J.H., Lin, W., Shen, D.: SPHERE: SPherical Harmonic Elastic REgistration of HARDI data. NeuroImage 55(2), 545–556 (2011)

Segmenting Thalamic Nuclei: What Can We Gain from HARDI?*

Thomas Schultz

Computation Institute, University of Chicago, Chicago IL, USA

Abstract. The contrast provided by diffusion MRI has been exploited repeatedly for in vivo segmentations of thalamic nuclei. This paper systematically investigates the benefits of high-angular resolution (HARDI) data for this purpose. An empirical analysis of clustering stability reveals a clear advantage of acquiring HARDI data at $b = 1000\,\mathrm{s/mm^2}$. However, based on stability arguments, as well as further visual and statistical evidence and theoretical insights about the impact of parameters, HARDI models such as the q-ball do not exhibit clear benefits over the standard diffusion tensor for thalamus segmentation at this b value.

1 Introduction

Since Wiegell et al. [1] demonstrated that the striations within the thalamic nuclei provide sufficient contrast in diffusion MRI to allow for an automated segmentation, several groups have proposed algorithmic methods to segment these nuclei based on the diffusion tensor (DTI) model [2–5]. More recently, segmentations of the thalamic nuclei have been achieved using high angular resolution (HARDI) data and the q-ball model [6, 7]. However, we are not aware of any studies that investigate the relative benefit from using HARDI models for this particular task, or provide guidelines on which model is most suitable for which measurement setup. It is the goal of our investigation to find out what can be gained from using HARDI data and specific HARDI models for this task.

Studying this question is essential to further refine methods for thalamus segmentation, but it is complicated by the fact that final segmentations do not only depend on the chosen model, but also on the algorithmic method and spatial regularization. Moreover, it is certainly possible to verify the overall plausibility of a segmentation by comparing it to known anatomy [1–7], but a lack of exact and reliable ground truth for individual subjects makes it difficult to argue about the relative validity in case of more subtle differences.

Therefore, we only turn to actual segmentations after a visual (Section 2.1) and statistical (Section 2.3) investigation of distance measures that result from different models, and a theoretical analysis of their parameters (Section 2.2). As

* I would like to thank Alfred Anwander (MPI CBS, Leipzig, Germany) for providing data and important feedback on this work. Gordon Kindlmann (University of Chicago) contributed through productive and inspiring discussions. This work has been supported by a fellowship within the DAAD Postdoc Program.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 141–148, 2011.

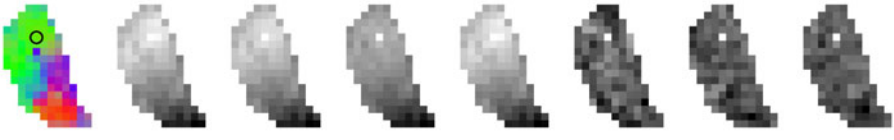

(a) (b) (c) (d) (e) (f) (g) (h)

Fig. 1. Distances from a reference voxel (a) show a large agreement between the diffusion tensor deviatoric (b) and the q-ball model (c). The Sobolev norm (d) weights the higher orders (f)–(h) stronger than the second order (e).

an objective way of judging segmentations, we then perform an analysis of their stability under small perturbations of the input (Sections 2.4 and 3).

2 Evaluating Benefits of the Q-Ball Model

2.1 Distance Maps

Measures of dissimilarity between the values in different voxels are at the center of all segmentation approaches, so it is a natural first step to study their behavior.

Figure 1 shows an axial slice through the right thalamus in one of our subjects. Our data has 60 gradient directions at $b = 1000\,\mathrm{s/mm^2}$, with 1.72 mm isotropic voxel size. Distances to a reference voxel, marked by a circle in (a), are linearly mapped to grayscale, from zero to the largest value in each map. To reduce the impact of partial voluming with the adjacent ventricles, we consider the diffusion tensor deviatoric $\tilde{\mathbf{D}} = \mathbf{D} - \mathrm{tr}(\mathbf{D})/3\mathbf{I}$ instead of the full tensors $\mathbf{D}$ (tr is matrix trace and $\mathbf{I}$ is the identity).

The Frobenius norm $\|\mathbf{D}\| = \sqrt{\mathrm{tr}(\mathbf{D}^{\mathrm{T}}\mathbf{D})}$ of differences in $\tilde{\mathbf{D}}$ (Fig. 1 (b), range $[0, 7.91 \times 10^{-3}\,\mathrm{mm^2/s}]$) produces a similar map as the ℓ_2 norm of q-balls (c, $[0, 0.48]$), modeled by order-8 spherical harmonics (SHs). The Sobolev norm [7] (d, $[0, 0.81]$) is designed to be more sensitive to the alignment of peaks in the q-ball than the ℓ_2 norm. Its map reveals similar overall structure, with highest distances at the posterior end of the thalamus, but exhibits reduced contrast.

2.2 Influence of Model Parameters

We have used Laplace-Beltrami regularization for the q-balls [8], which introduces a regularization parameter ($\lambda = 0.008$). The Sobolev norm involves three additional parameters ($\alpha = 1$, $t = 0$, $\gamma = 0.21$). Even though we carefully calibrated these values according to the synthetic model proposed in [7], modified to match our experimental setup, it is informative to consider how our result depends on their exact choice.

The Sobolev norm amounts to a weighted ℓ_2 norm, in which higher orders are weighted more strongly [7]. This motivates mapping the ℓ_2 distances of the SH coefficients of orders 2 (Fig. 1 (e), $[0, 0.47]$), 4 (f, $[0, 0.062]$), 6 (g, $[0, 0.026]$) and 8 (h, $[0.011]$) separately. Since q-balls integrate to unity by definition [9], there is

no variability in order zero. All parameters of the Sobolev norm lead to squared distances that are linear combinations of the squared values from Fig. 1 (e)–(h).

We observed that the effect of different regularization parameters λ is also well-approximated by a re-weighting of different SH orders. This is explained by the fact that with Laplace-Beltrami regularization, the SH coefficient vector $\mathbf{c} \in \mathbb{R}^C$ is obtained from the diffusion-weighted signal vector $\mathbf{s} \in \mathbb{R}^N$ via

$$\mathbf{c} = (\mathbf{B}^{\mathrm{T}}\mathbf{B} + \lambda\mathbf{L})^{-1}\mathbf{B}^T\mathbf{s}, \tag{1}$$

where the (i,j)th element of $\mathbf{B} \in \mathbb{R}^{N\times C}$ is computed by evaluating the SH basis function Y_j at the position (θ_i, ϕ_i) of gradient vector i. The diagonal matrix $\mathbf{L} \in \mathbb{R}^{C\times C}$ has entries $l_j^2(l_j+1)^2$, l_j being the order of SH coefficient c_j. In the interest of space, the reader is referred to [8] for details on this notation.

Typical HARDI acquisition schemes achieve a near-uniform distribution of gradient directions over the (hemi-)sphere. In this case, up to a factor of $4\pi/N$, dot products between the columns of $\mathbf{B}$ numerically approximate the integral

$$\langle S^{(1)}, S^{(2)} \rangle = \int_{\theta=0}^{\pi} \int_{\phi=0}^{2\pi} S^{(1)}(\theta,\phi)\, S^{(2)}(\theta,\phi) \sin\theta \, d\phi \, d\theta, \tag{2}$$

which defines a scalar product for real-valued functions on the sphere. Given the orthonormality of the SH basis functions from [8] with respect to (2), this leads to the approximation $\mathbf{B}^{\mathrm{T}}\mathbf{B} \approx N/(4\pi)\mathbf{I}$.

Consequently, $(\mathbf{B}^{\mathrm{T}}\mathbf{B} + \lambda\mathbf{L})^{-1}$ in Eq. (1) is approximated by a diagonal matrix with entries $(N/(4\pi) + \lambda l_j^2(l_j+1)^2)^{-1}$. Thus, the regularized $\tilde{c}_{l,m}$ can be approximated by scaling the unregularized $c_{l,m}$,

$$\tilde{c}_{l,m} \approx c_{l,m} \times \left(1 + \frac{4\pi\lambda l^2(l+1)^2}{N}\right)^{-1}. \tag{3}$$

In the ℓ_2 norm $\|\tilde{\mathbf{c}}\|$, this is reflected by a re-weighting of the terms $c_{l,m}$ by a function of SH order l and the regularization parameter λ.

Even replacing the q-balls by a different HARDI model, spherical deconvolution [10], would only amount to a re-weighting of different SH orders in the ℓ_2 norm. The reason is that linear deconvolution is performed by scaling the SH coefficients of the diffusion weighted signal by a function of SH order, which depends on the fiber response function [10]. Analytically evaluating the Funk-Radon transform, which relates the diffusion-weighted signal to the q-ball, reveals that it can be expressed in the same way [8].

These insights illustrate the high relevance of the maps in Fig. 1 (e)–(h) to our problem: They serve as linear building blocks of the distance maps that would result from q-ball or spherical deconvolution models, combined with either the ℓ_2 or the Sobolev norm. We cannot expect any structures that are not present in these maps to be brought forward by changing the parameters of these models.

Visually, the spatial structure of the second order map (e) largely agrees with the structure from the diffusion tensor deviatoric in (b). It is more difficult to discern structure in the higher orders. In this respect, the thalamus appears to differ from white matter, where higher orders carry meaningful structure in regions of fiber crossings, even at moderate b-values [11].

Table 1. A statistical investigation of distance measures (separately in the left/right part of the thalamus) supports the visual impression that distances at spherical harmonics order 2 expose a much clearer spatial structure than at higher orders.

	order 2	order 4	order 6	order 8
$\bar{d}_A^2/\bar{d}_N^2$	8.17/8.19	1.83/1.75	1.63/1.58	1.60/1.58
skewness of d_N^2	2.82/2.65	1.87/1.50	1.30/1.45	1.19/1.26

2.3 Distance Statistics

To quantify the visual impression that the distances at higher orders reveal much less spatial structure than distances at SH order two, we have computed average squared distances $\bar{d}_A^2$ between all pairs of voxels (in a full 3D thalamus mask, for the left and right side separately), and compared them to average squared distances $\bar{d}_N^2$ between face neighbors. If the voxels contained noise without any spatial structure, the expected ratio were $\bar{d}_A^2/\bar{d}_N^2 = 1$. Piecewise smooth or piecewise constant data with clear spatial structure should lead to $\bar{d}_A^2/\bar{d}_N^2 \gg 1$. Table 1 shows that $\bar{d}_A^2/\bar{d}_N^2$ is much larger for order 2 than it is for higher orders.

We have also considered the distribution of d_N^2, the squared distances between face neighbors. In case of homogeneous regions that are separated by clear boundaries, we expect most distances to be small (within regions), with a heavy tail of much larger distances from boundaries. This should lead to a positive skew in the distribution. In fact, Table 1 reveals such a skew, most pronounced in the second order distances.

2.4 Stability Analysis

Even though the distance maps from DTI and q-balls look similar, we found that the resulting segmentations differ to some degree. In order to assess whether these differences are significant, we compute the adjusted RAND index [12] between the segmentations to quantify their overlap, and compare it to the overlap of segmentations from the same model, but slightly perturbed data.

It is widely accepted that successful segmentation of thalamic nuclei requires some sort of regularization that favors spatially connected structures. All methods we are aware of include such regularization, either by combining the data-based distance with a spatial distance [1, 3, 6, 7], through coupling forces in a level set framework [4], by Markovian relaxation [2], or by learning an atlas from the joint segmentation of multiple subjects [5].

Unfortunately, spatial regularization might lead to paradoxical results in a stability analysis: If a model does not reveal clear structure in the data, the resulting segmentation might be dominated by spatial information and consequently appear more stable than an alternative model that does reveal (somewhat fragile) structure. Such misleading results can be difficult to detect, even when checking the plausibility of the obtained segmentations.

To illustrate this, Fig. 2 (a) reproduces the segmentation presented by Grassi et al. [6] based on our own data. We use k-means with a linear combination of

Fig. 2. Given an individual segmentation result (a), it can be difficult to judge the relative influence of data-based (b) and spatial distances (c). The same slice is shown as in Fig. 1, rotated clockwise by 90°.

Euclidean spatial distance and the ℓ_2 distance between q-balls, giving the spatial distance the minimum weight required to obtain connected components. As in [6], an axial slice of the thalamus is segmented into a pulvinar, a lateral, a medial, and an anterior region. It is reassuring that a very similar result is obtained without spatial regularization (Fig. 2 (b)). However, even if we do not make any use of the data (Fig. 2 (c)), spatial distances alone lead to a segmentation that looks similar enough that it might pass a superficial test of plausibility.

Therefore, we run k-means based on the diffusion information alone, on the right thalamus (in 3D) and for different numbers of regions (2–7). Since the result of k-means may depend on the initialization, we perform each clustering ten times, from randomly selected seeds, and keep the result with the lowest sum of squared distances of data points to their cluster centers. In order to create slightly perturbed model fits without changing the noise characteristics of the data, we obtained 150 perturbed models via weighted least squares, where the weights of the 60 available directions were sampled uniformly at random from $[0, 1]$. The average relative distance $\|\mathbf{D}_p - \mathbf{D}\|/\|\mathbf{D}\|$ of a diffusion tensor $\mathbf{D}_p$ that has been perturbed in this way to its unperturbed version $\mathbf{D}$ is around 3%.

The solid lines in Fig. 3 (a) plot the mean adjusted RAND index when comparing segmentations of perturbed data to their unperturbed counterparts. Results from the diffusion tensor are marked with circles, q-ball with the ℓ_2 norm by squares, and q-ball with Sobolev norm by triangles. None of the models consistently produced more stable results than the others, but we noticed a steep drop in stability for more than five regions, especially when using q-balls.

The dashed lines in the same Figure plot the mean adjusted RAND when comparing segmentations from the same weighting of gradients, but different models. The average overlap of DTI vs. q-ball with ℓ_2 norm (squares) and DTI vs. q-ball with Sobolev norm (triangles) on the same data is consistently larger than the overlap when using the same model on perturbed vs. unperturbed data.

This effect is statistically significant. For each set of random weights, we compared the overlap within the same model (perturbed vs. unperturbed) to the overlap across models (DTI vs. q-ball). For all numbers of regions, we found the latter to be greater in so many cases that a one-sided binomial test rejects, with $p < 0.01$ (the largest p was $p = 5.5 \times 10^{-5}$), the null hypothesis that the within-model overlap and the across-model overlap are equally likely to be greater than each other. We conclude that segmentations change significantly more under slight perturbations of the input than when switching between diffusion tensor and q-ball models.

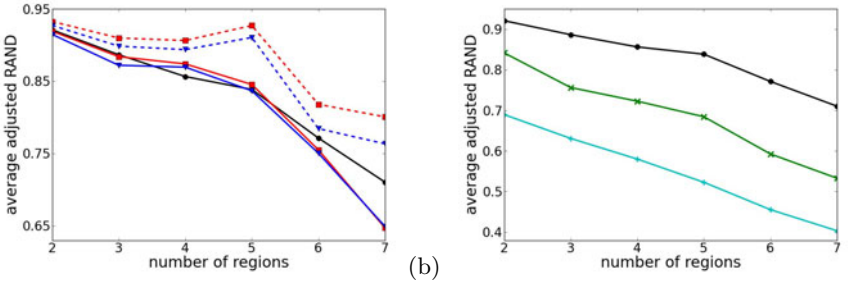

Fig. 3. Solid lines in (a) show segmentation stability under slight perturbations of the input (DTI: black circles, q-ball with ℓ_2 norm: red squares, Sobolev norm: blue triangles). It is consistently lower than the average similarity between results from DTI and q-ball on the same data (dashed lines). However, (b) shows that 60 directions (black circles) improve stability considerably when compared to 40 (green x-es) or 20 directions (cyan + signs), even with the DTI model.

2.5 Segmentation Based on Principal Directions

In their spectral method for thalamus segmentation, Ziyan et al. [2] use angular differences between principal eigenvectors as an alternative to the full diffusion tensor information. We have explored a similar strategy by extracting q-ball maxima. An initial estimate of the maximum has been obtained from a dense sampling on the sphere, and the result has been refined via gradient ascent.

Based on visual inspection of RGB maps (Fig. 4 (a) and (b)), the resulting overall structure appears similar. Within the thalamus mask, the average angular difference (AAD) between both estimates is around 10°. However, principal directions from q-balls are less stable under the perturbation from the previous Section (AAD 6.4°) than those estimated from the diffusion tensor (AAD 4.9°).

In order to average directions in the k-means algorithm, we added their outer products and took the principal eigenvector of the result. In general, we found segmentations that use principal directions to be less stable than those based on the deviatoric of the diffusion tensor, in particular when the estimate was derived from the q-ball model. Interpretation of Fig. 4 (c) is analogous to Fig. 3; stars indicate use of the principal direction from DTI, squares from q-ball. The stability of the deviatoric (circles) is repeated from Fig. 3 for reference.

3 Evaluating Benefits of Acquiring HARDI Data

Since our results in the previous Section did not demonstrate a clear advantage of using the q-ball model, it is natural to ask if it is even worthwhile to acquire 60 gradient directions at $b = 1000\,\mathrm{s/mm^2}$ when the goal is not tractography, but segmentation of the thalamic nuclei. To investigate this question, we have repeated the stability analysis based on the diffusion tensor model and subsampled sets of 40 and 20 gradient directions, respectively. In order to retain a reasonable

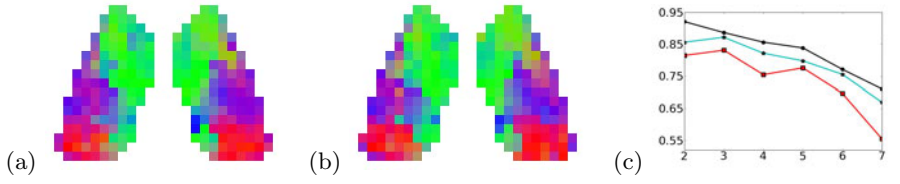

Fig. 4. Principal directions estimated from diffusion tensors (a) and q-balls (b) largely agree. However, (c) shows that segmentations based on principal directions from DTI (cyan stars) and q-ball (red squares) are less stable than ones from the diffusion tensor deviatoric (black circles).

gradient distribution, we generated sets of 40 and 20 directions based on electrostatic repulsion [13], and kept the original measurements that were closest to these directions.

Fig. 3 (b) compares the average overlap of segmentations from perturbed data (random weights of the gradients) to a segmentation based on the same set of gradients with unit weights. The original dataset is marked with circles (and is identical to Fig. 3 (a)), the subsampled data is marked with x (40 directions) and + signs (20 directions). It is obvious that stability of the segmentation suffers considerably when using fewer directions.

4 Conclusion

While it is widely accepted that HARDI models provide useful additional information within the white matter [11, 14], it has not been studied systematically what they can contribute to the segmentation of gray matter structures such as the thalamus. We have shed light on this interesting question from different perspectives, using a mixture of visual and statistical methods, theoretical reasoning about the impact of parameters, and empirical stability analysis.

For the data available to us ($b = 1000\,\mathrm{s/mm^2}$), we conclude that segmentations of thalamic nuclei clearly benefit from acquiring data at high angular resolution, but that this is mainly due to an improved estimate of the information present in the standard diffusion tensor model. We found no clear benefit from switching to more complex models, such as q-ball, where most of the useful information seems to reside in the second-order spherical harmonics coefficients. This finding might not carry over to much higher b-values, as they have been used in [6, 7]. The methods proposed in our work could be used to test for potential benefits of the q-ball model when combined with such measurement setups.

We have only considered segmentations based on local diffusion properties. A different strategy for segmenting thalamic nuclei has involved a probabilistic tracking of the corticothalamic/thalamocortical connections [15, 16]. Since significant advantages of using HARDI models have been reported when tracking non-dominant fibers, even at $b = 1000\,\mathrm{s/mm^2}$ [14], the relative benefit of tractography-based over purely local methods merits further investigation.

References

1. Wiegell, M.R., Tuch, D., Larsson, H.B., Wedeen, V.J.: Automatic segmentation of thalamic nuclei from diffusion tensor magnetic resonance imaging. NeuroImage 19, 391–401 (2003)
2. Ziyan, U., Tuch, D., Westin, C.-F.: Segmentation of thalamic nuclei from DTI using spectral clustering. In: Larsen, R., Nielsen, M., Sporring, J. (eds.) MICCAI 2006. LNCS, vol. 4191, pp. 807–814. Springer, Heidelberg (2006)
3. Duan, Y., Li, X., Xi, Y.: Thalamus segmentation from diffusion tensor magnetic resonance imaging. Int. J. Biomed. Imaging 2007 (2007)
4. Jonasson, L., Hagmann, P., Pollo, C., Bresson, X., Wilson, C.R., Meuli, R., Thiran, J.P.: A level set method for segmentation of the thalamus and its nuclei in DT-MRI. Signal Processing 87(2), 309–321 (2007)
5. Ziyan, U., Westin, C.-F.: Joint segmentation of thalamic nuclei from a population of diffusion tensor MR images. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part I. LNCS, vol. 5241, pp. 279–286. Springer, Heidelberg (2008)
6. Grassi, A., Cammoun, L., Pollo, C., Hagmann, P., Meuli, R., Thiran, J.P.: Thalamic nuclei clustering on high angular resolution diffusion images. In: Proc. Int. Soc. Magn. Reson. Med., p. 1777 (2008)
7. Brunenberg, E., Duits, R., ter Haar Romeny, B., Platel, B.: A sobolev norm based distance measure for HARDI clustering. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. 6361, pp. 175–182. Springer, Heidelberg (2010)
8. Descoteaux, M., Angelino, E., Fitzgibbons, S., Deriche, R.: Regularized, fast, and robust analytical Q-Ball imaging. Magn. Reson. Med. 58, 497–510 (2007)
9. Tuch, D.S.: Q-Ball imaging. Magn. Reson. Med. 52, 1358–1372 (2004)
10. Tournier, J.D., Calamante, F., Gadian, D.G., Connelly, A.: Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. NeuroImage 23, 1176–1185 (2004)
11. Alexander, D.C., Barker, G.J., Arridge, S.R.: Detection and modeling of non-gaussian apparent diffusion coefficient profiles in human brain data. Magn. Reson. Med. 48, 331–340 (2002)
12. Hubert, L., Arabie, P.: Comparing partitions. J. Classif. 2, 193–218 (1985)
13. Jones, D.K., Horsfield, M.A., Simmons, A.: Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. Magn. Reson. Med. 42, 515–525 (1999)
14. Behrens, T.E.J., Johansen-Berg, H., Jbabdi, S., Rushworth, M.F.S., Woolrich, M.W.: Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? NeuroImage 34, 144–155 (2007)
15. Behrens, T.E.J., Johansen-Berg, H., Woolrich, M.W., Smith, S.M., Wheeler-Kingshott, C.A.M., Boulby, P.A., Barker, G.J., Sillery, E.L., Sheehan, K., Ciccarelli, O., Thompson, A.J., Brady, J.M., Matthews, P.M.: Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging. Nat. Neurosci. 6(7), 750–757 (2003)
16. Johansen-Berg, H., Behrens, T.E.J., Sillery, E., Ciccarelli, O., Thompson, A.J., Smith, S.M., Matthews, P.M.: Functional-anatomical validation and individual variation of diffusion tractography-based segmentation of the human thalamus. Cereb. Cortex 15(1), 31–39 (2005)

Resting State fMRI-Guided Fiber Clustering

Bao Ge[1], Lei Guo[1], Jinglei Lv[1], Xintao Hu[1], Junwei Han[1], Tuo Zhang[1,2], and Tianming Liu[2]

[1] School of Automation, Northwestern Polytechnical University, Xi'an, China
{oct.bob,guolei.npu,lvjinglei,xintao.hu,
junweihan2010,zhangtuo.npu}@gmail.com
[2] Department of Computer Science and Bioimaging Research Center,
University of Georgia, Athens, GA
tliu@cs.uga.edu

Abstract. Fiber clustering is a prerequisite step towards tract-based analysis of white mater integrity via diffusion tensor imaging (DTI) in various clinical neuroscience applications. Many methods reported in the literature used geometric or anatomic information for fiber clustering. This paper proposes a novel method that uses functional coherence as the criterion to guide the clustering of fibers derived from DTI tractography. Specifically, we represent the functional identity of a white matter fiber by two resting state fMRI (rsfMRI) time series extracted from the two gray matter voxels to which the fiber connects. Then, the functional coherence or similarity between two white matter fibers is defined as their rsfMRI time series' correlations, and the data-driven affinity propagation (AP) algorithm is used to cluster fibers into bundles. At current stage, we use the corpus callosum (CC) fibers that are the largest fiber bundle in the brain as an example. Experimental results show that the proposed fiber clustering method can achieve meaningful bundles that are reasonably consistent across different brains, and part of the clustered bundles was validated via the benchmark data provided by task-based fMRI data.

Keywords: Resting state fMRI, DTI, fiber clustering.

1 Introduction

Diffusion tensor imaging (DTI), as a powerful tool to image the axonal fibers in vivo, provides rich structural connectivity information that is believed to be closely related to brain function. In order to infer meaningful and comparable information from DTI data of different brains, the large number of fiber trajectories produced by DTI tractography need to be grouped into appropriate fiber bundles for tract-based analysis [4]. Many approaches reported in the literature used geometric, anatomical or structural features, e.g., fiber's Euclidean distances [3, 4], fiber shape information [11], or fiber's end point positions [10], to cluster fiber bundles. Though these methods have their own advantages in clustering meaningful bundles, the functional interpretation of the clustering results remains to be elucidated.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 149–156, 2011.

Recently, resting state fMRI (rsfMRI) has been demonstrated to be an effective modality by which to explore the functional networks in the human brain, because similar low-frequency oscillations in rsfMRI time series between spatially distinct brain regions are indicative of correlated functional activity patterns in the brain [2]. In addition, a variety of recent studies demonstrated that structural connectivity derived from DTI data is closely correlated with the functional connectivity derived from rsfMRI data [8]. Inspired by these studies, we are motivated to apply the criterion of functional coherence to cluster white matter fibers. Our premise is that the clustered fibers within a bundle should have functional homogeneity or coherence. To achieve this goal, we represent a white matter fiber by two rsfMRI time series extracted from the two gray matter (GM) voxels that the fiber's two end points connect, and the functional coherence between white matter fibers is measured by the similarities of their rsfMRI time series. Then, the data-driven affinity propagation (AP) algorithm [7] is applied to cluster fibers into bundle tracts. We currently use the corpus callosum (CC) fibers, which are the largest fiber bundle in the brain, as an example for algorithm development and validation. Our experimental results in seven brains with multimodal rsfMRI and DTI datasets show that the proposed rsfMRI-guided fiber clustering method can achieve meaningful fiber bundles that are reasonably consistent across different brains, and part of the clustered bundles is validated by the benchmark data provided by task-based fMRI data.

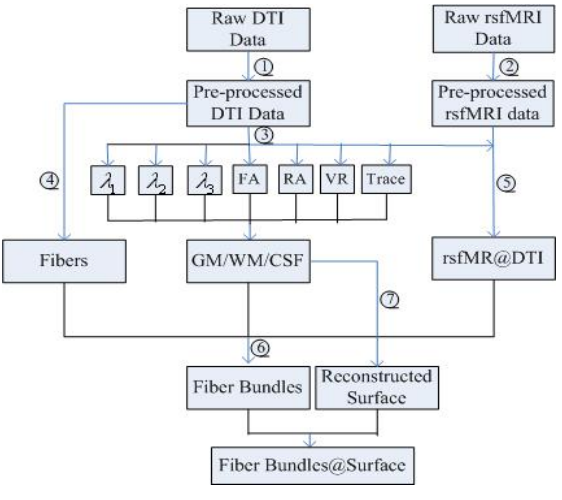

Fig. 1. The flowchart of our framework. (1) and (2): pre-processing steps; (3): segmentation of brain tissue using DTI data; (4): DTI tractography; (5): registration of rsfMRI to DTI images; (6): affinity propagation clustering guided by rsfMRI data; (7): WM (white matter)/GM (gray matter) cortical surface reconstruction from DTI data.

2 Materials and Methods

2.1 Overview

As summarized in Fig.1, our algorithmic pipeline includes the following steps. First, we pre-processed the raw DTI data, and then performed brain tissue segmentation and

fiber tracking based on DTI data. The tracked fiber trajectories were projected to the cortical surface via a similar method in [8] to facilitate the extraction of rsfMRI signals on the gray matter volume. Also, we registered the rsfMRI signals to the DTI space using FSL FLIRT. Then, we clustered fibers into bundles based on fibers' function coherences via the affinity propagation algorithm [7]. Finally, we identified consistent fiber bundles from seven subjects for evaluation and validation.

2.2 Multimodal Data Acquisition and Pre-processing

Seven volunteers were scanned using a 3T GE Signa MRI system. We acquired the rsfMRI data with dimensionality 128*128*60*100, space resolution 2mm*2mm*2mm, TR 5s, TE 25ms, and flip angle 90 degrees. DTI data was acquired using the same spatial resolution as the rsfMRI data; parameters were TR 15.5s and TE 89.5ms, with 30 DWI gradient directions and 3 B0 volumes acquired. For two out of the seven subjects, the working memory OSPAN tasks [12] was used for fMRI data acquisition with the parameters of 64×64 matrix, 4mm slice thickness, 220mm2 FOV, 30 slices, TR=1.5s, TE=25ms, ASSET=2. Pre-processing of the rsfMRI data included skull removal, motion correction, spatial smoothing, temporal pre-whitening, slice time correction, global drift removal, and band pass filtering (0.01Hz~0.1Hz). For the DTI data, pre-processing included skull removal, motion correction, and eddy current correction. After the pre-processing, fiber tracking was performed using MEDINRIA (FA threshold: 0.2; minimum fiber length: 20). Based on pre-processed DTI data, brain tissue segmentation was performed using the multi-channel fusion method akin to that in [5]. DTI space was used as the standard space from which to generate the tissue segmentation and exhibit the functional coherent fiber bundles. Since rsfMRI and DTI sequences are both EPI sequences, their distortions tend to be similar, and thus the misalignment between their images is much less than that between T1 and fMRI images [8]. DTI and fMRI images were registered via FSL FLIRT.

2.3 Fiber Clustering Based on Functional Coherence

Extraction of rsfMRI Signals for a Fiber's Two Ends

It should be noted that the blood supply to the white matter is significantly lower than that of the cortex (less than one fourth) [9], and the blood-oxygen-level dependence (BOLD) contribution of the white matter is relatively low. Hence, the investigation of gray matter rsfMRI signals is more reasonable. Therefore, before extracting rsfMRI signals from GM voxels for a fiber's two ends, we need to project some fibers onto the gray matter cortex in that the DTI-derived fiber trajectories are not necessarily located on the cortex due to two reasons. 1) DTI fiber tractography using the streamline approach has difficulty in tracking inside GM since the FA (fractional anisotropy) values around the boundaries of gray matter and white matter are relatively low. As a result, there are some fibers that cannot touch the GM. 2) There is discrepancy in brain tissue segmentation based on DTI data and the DTI tractography [5]. In this case, the fiber could be either outside the cortex if the gray matter (GM) is over-segmented, or inside the cortex if the GM is under-segmented.

In order to make use of the fiber connection information on the cortex, we projected the fibers onto the cortical surface guided by the tissue segmentaion map.

There are four types of fiber projections here. 1) If the end point of a fiber already lies on a GM voxel in the brain tissue map, no search is conducted, e.g., fiber #1 shown in Fig. 2(a); 2) If the end point of a fiber lies inside the cortex, e.g., the fiber #2 shown in Fig. 2(a), we search forward along the tangent direction until reaching the gray matter. 3) Otherwise, e.g., the fiber #3 shown in Fig. 2(a), we search backward along the tangent direction until reaching the gray matter. The search process stops either when the fiber arrives at a GM voxel or it exceeds a search threshold. 4) In very rare cases when a fiber cannot reach the surface, e.g., the fiber #4 shown in Fig. 2(a), we treat this fiber as an outlier and remove it from the data. Fig. 2(b) shows the positions that the fibers arrive at after the projection. The search was conducted iteratively until at least one GM voxel can be found in the 1-ring surface vertex neighborhood of the current seed point, or the number of iteration exceeds a given threshold. When multiple GM voxels exist, the closest one is used as the projected point. Finally, for each projected fiber, we extract the rsfMRI signals for two ends of the fiber.

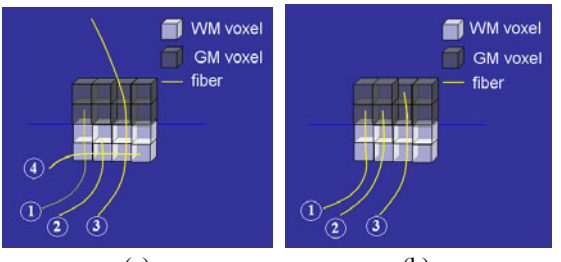

(a) (b)

Fig. 2. Illustration of fiber projection. Gray matter and white matter voxels are represented by gray and white color boxes respectively. Fibers are represented by yellow curves. (a) The four situations before fiber projection; (b) The results of fiber projections for three situations.

Measurement of Functional Coherence Among fibers

As illustrated in Fig. 3, given any pair of fibers with four end points located in the gray matter, the functional coherence between these two fibers is defined as follows:

$$C=0.5*(max(C_{13},C_{14})+max(C_{23},C_{24})) \quad (1)$$

$$C_{13}=PsCor(v1,v3),\ C_{14}=PsCor(v1,v4),\ C_{23}=PsCor(v2,v3),\ C_{24}=PsCor(v2,v4)$$

where *vi* indexes the end points of two fibers, the function *PsCor* is the Pearson correlation coefficient of two end points' rsfMRI signals. Our premise here is that the fibers belonging to the same tract should have higher functional coherence, and those belonging to different tracts should have lower coherence.

It should be noted that the crierion of functional coherence derived from rsfMRI data offers unique capability to cluster functionally coherent fibers into the same bundle and differentiate non-coherent fibers into different bundles. As an example, Fig. 4a shows three fibers that are functionally coherent, and thus they should be clustered into one bundle. However, if we use geometric or shape criteria [3, 4, 11], e.g., the Euclidean distances betweeen neighboring fibers, the blue fiber in Fig. 4a

(highlighted by a red arrow) is very likely to be separated from the bundle composed of the red and green ones. Another example is shown in Fig. 4b, in which the blue fiber (highlighted by a red arrow) has low functional coherence with the green and red ones and thus the blue one can be differentiated from other two fibers via rsfMRI data. However, geometry or shape based fiber clustering methods are likely to have difficulties in differentiating the blue fiber from other two fibers. Hence, the criterion of functional coherence is a powerful approach for fiber clustering.

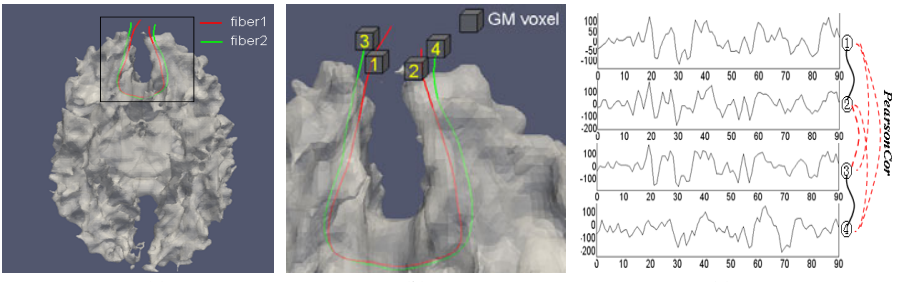

(a) (b) (c)

Fig. 3. The calculation of fibers' functional coherence. (a) Two fibers overlaid on the reconstructed cortical surface (gray mesh); (b) The zoomed-in view of the black rectangle in (a); (c) The rsfMRI signals of the four end points. Their correlations are measured by Eq. (1).

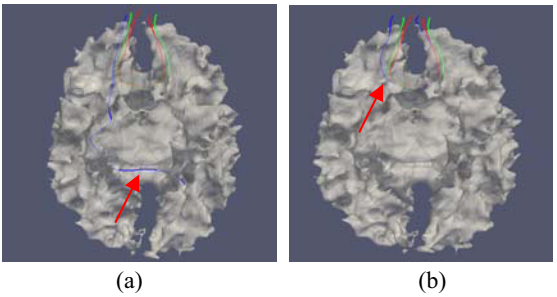

(a) (b)

Fig. 4. (a). An example showing functional coherence can cluster fibers of different shapes or geometries into the same bundle. (b) An example showing functional difference can differentiate neighboring fibers into different bundles.

Fiber Clustering via the Affinity Propagation Algorithm

The affinity propagation (AP) algorithm [7] has been widely used to identify data clusters automatically. In the AP clustering method, each cluster is represented by a data point called a cluster center, or an exemplar, and the method searches for cluster so as to maximize a goal function called net similarity [7]. In this paper, we applied the AP clustering method on the functional similarity matrix of all fibers in the corpus callosum, and achieved the clustered fiber bundles. In particular, each fiber cluster is represented by the fiber exemplar discovered during the AP clustering procedure.

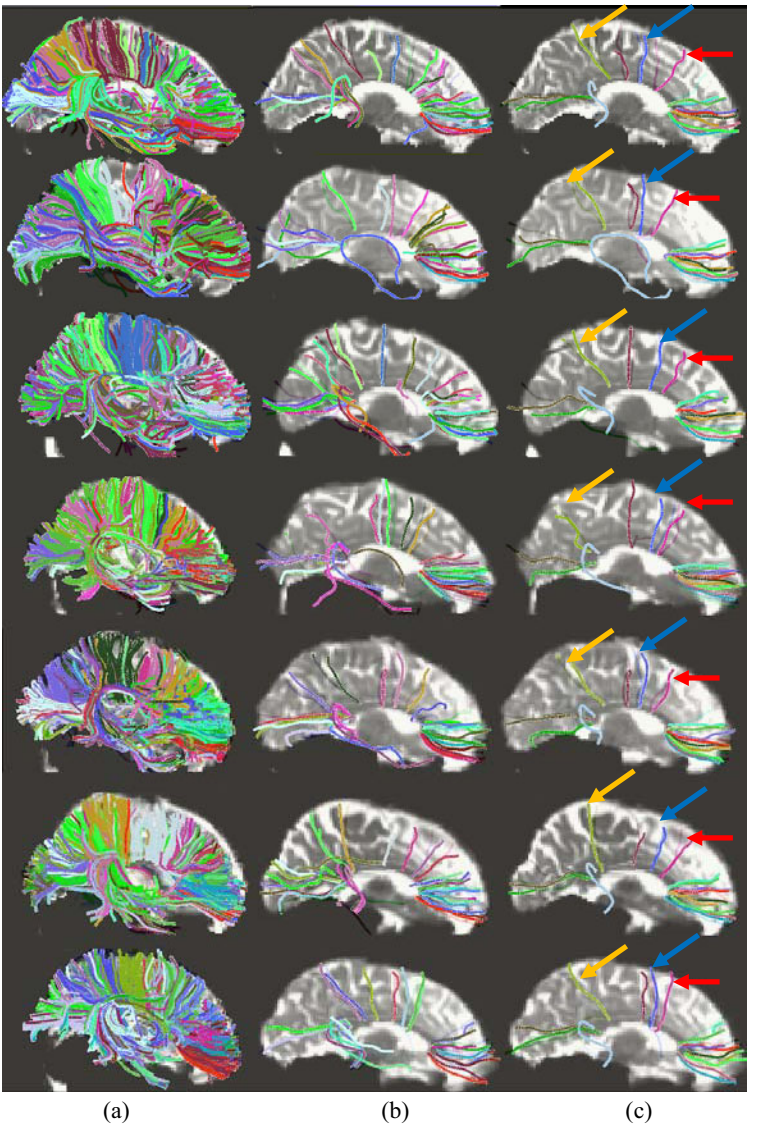

Fig. 5. The clustered fiber bundles for 7 subjects. (a) The fiber clusters with randomly set colors overlaid on the DTI B0 images; (b) The fiber exemplars of all clusters (c) The 16 most consistent fiber exemplars.

3 Experimental Results

3.1 Identification of Functionally Coherent Fiber Bundles

The fibers passing corpus callosum (CC) were clustered into aournd 30 bundles for 7 subjects separately, as shown in the 7 rows in Fig. 5(a). For the sake of visual differentiation, each fiber bundle was represented by the fiber exemplar obtained during the affinity propogation clustering procedure [7], as shown in Fig. 5(b). In order to identify the corresponding fiber bundles in different subjects, we computed the Hausdorff distances between the representative exemplars across subjects and picked out those exemplars that are closest to the representative exemplars in other subjects. We visually confirmed 16 most consistent and representative exemplar fibers in all fiber exemplars from 7 subjects and showed all of them in Fig. 5(c). Each corresponding fiber exemplar in Fig. 5(c) has the same color in different brains, and three of corresponding ones are highlighted by arrows of the same color. It is evident that the distributions of these 16 fiber exemplars are quite reasonable and consistent. As another exmple, Fig. 6(a) visualizes four corresponding bundles from 3 randomly selected subjects, showing that the clustered bundles are quite reasonable.

3.2 Validation by Task-Based fMRI Data

In addition to the qualitative visual evaluation of the clustered CC fiber bundles in Section 3.1, we used working memory task-based fMRI data [12] to examine the functional correspondence of the clustered fiber bundles in Section 3.1. Specifically, the working memory task-based fMRI data provided 16 consistently activated brain regions, as shown by green boxes in Fig. 6(b). These ROIs provide the benchmark data for comparison of functional correspondences of fiber bundles. It is striking that one fiber bundle (blue ones in Figs. 5(c) and Fig. 6(b)) clustered in Section 3.1 coincidently falls into the neighborhoods of two corresponding working memory ROIs of left and right paracingulate gyri (highlighted by yellow arrows) consistently in the testing subjects. These close vicinities indicate that the paracingulate gyri are consistently connected by the blue fiber bundle across individuals, suggesting that the rsfMRI-guided fiber clustering method grouped functionally coherent fibers into the same bundle. This result is considered as a validation of the rsfMRI-guided fiber clustering approach.

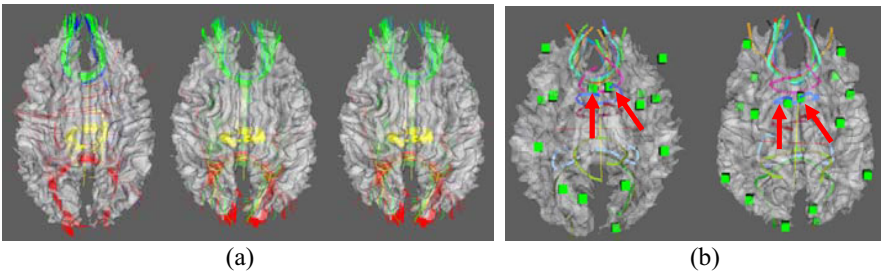

(a) (b)

Fig. 6. (a) Visualization of four corresponding fiber bundles from 3 randomly chosen subjects. They are labeled by the four different colors of blue, green, yellow, and red, respectively. (b) Joint visualization of 16 activated working memory ROIs (represented by green boxes) and clustered fiber bundles (represented by exemplars) for two subjects.

4 Conclusion

This paper presents a novel methodology of using rsfMRI data to guide fiber clustering. The underlying neuroscience basis is that axonal fibers within a bundle should have functional coherence, and our results have shown that functional coherence is a meaningful criterion for fiber clustering. In particular, part of the clustered bundles was validated via task-based fMRI data. Currently, only CC fibers were used for algorithm development and evaluation. In the future, we plan to apply the proposed method to other major fiber bundles such as cortico-cortical and cortical-subcortical pathways, and apply the methods for tract-based analysis of DTI datasets of brain diseases such as schizophrenia and Alzheimer's disease.

References

1. Friston, K.: Modalities, modes, and models in functional neuroimaging. Science 326(5951), 399–403 (2009)
2. Fox, M.D., Raichle, M.E.: Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. Nat. Rev. Neurosci. 8, 700–711 (2007)
3. Maddah, M., Grimson, W., Warfield, S.: Statistical Modeling and EM Clustering of White Matter Fiber Tracts. In: ISBI, vol. 1, pp. 53–56 (2006)
4. Gerig, G., Gouttard, S., Corouge, I.: Analysis of Brain White Matter via Fiber Tract Modeling. IEEE EMBS 2, 4421–4424 (2004)
5. Liu, T., Li, H., Wong, K., Tarokh, A., Guo, L., Wong, S.: Brain Tissue Segmentation Based on DTI Data. NeuroImage 38(1), 114–123 (2007)
6. Liu, T., Nie, J., Tarokh, A., Guo, L., Wong, S.: Reconstruction of Central Cortical Surface from MRI Brain Images: Method and Application. NeuroImage 40(3), 991–1002 (2008)
7. Frey, B.J., Dueck, D.: Clustering by Passing Messages between Data Points. Science 315, 972–976 (2007)
8. Li, K., Guo, L., Li, G., Nie, J., Faraco, C., Zhao, Q., Miller, S., Liu, T.: Cortical surface based identification of brain networks using high spatial resolution resting state FMRI data. In: ISBI, pp. 657–659 (2010)
9. Mezer, A., Yovel, Y., Pasternak, O., Gorfine, T., Assaf, Y.: Cluster analysis of resting-state fMRI time series. NeuroImage 45(4), 1117–1125 (2009)
10. Ge, B., Guo, L., Li, K., Li, H., Faraco, C., Zhao, Q., Miller, S., Liu, T.: Automatic Clustering of White Matter Fibers via Symbolic Sequence Analysis. In: SPIE Medical Image, vol. 7623, pp. 762327.1–762327.8 (2010)
11. Zhang, T., Guo, L., Hu, X., Li, G., Nie, J., Jiang, X., Zhang, D., Liu, T.: Joint analysis of fiber shape and cortical folding patterns. In: ISBI, pp. 1165–1168 (2010)
12. Zhang, T., Guo, L., Hu, X., Li, K., Liu, T.: Predicting Functional Cortical ROIs based on Fiber Shape Models. In: Fichtinger, G., Martel, A., Peters, T. (eds.) MICCAI 2011, Part II. LNCS, vol. 6892, pp. 42–49. Springer, Heidelberg (2011)

Apparent Intravoxel Fibre Population Dispersion (FPD) Using Spherical Harmonics

Haz-Edine Assemlal[1], Jennifer Campbell[2], Bruce Pike[2], and Kaleem Siddiqi[1]

[1] Centre for Intelligent Machines, McGill University,
3480 University Street, Montréal, QC, Canada H3A 2A7
[2] McConnell Brain Imaging Centre, Montreal Neurological Institute,
3801 University Street, Montréal, QC Canada H3A 2B4
assemlal@cim.mcgill.ca

Abstract. The vast majority of High Angular Resolution Diffusion Imaging (HARDI) modeling methods recover networks of neuronal fibres, using a heuristic extraction of their local orientation. In this paper, we present a method for computing the apparent intravoxel Fibre Population Dispersion (FPD), which conveys the manner in which distinct fibre populations are partitioned within the same voxel. We provide a statistical analysis, without any prior assumptions on the number or size of these fibre populations, using an analytical formulation of the diffusion signal autocorrelation function in the spherical harmonics basis. We also propose to extract features of the FPD obtained in the group of rotations, using several metrics based on unit quaternions. We show results on simulated data and on physical phantoms, that demonstrate the effectiveness of the FPD to reveal regions with crossing tracts, in contrast to the standard anisotropy measures.

1 Introduction

Diffusion magnetic resonance imaging (dMRI) allows one to examine the microscopic diffusion of water molecules in biological tissue *in-vivo*. In practice, this imaging modality requires the collection of successive images with magnetic field gradients applied in different directions [1]. A reconstruction step is then used to estimate the 3D diffusion probability density function (PDF) from the acquired images [2]. Recently a Spherical Polar Fourier (SPF) expansion method has been introduced in [3], which takes full advantage of the acquisition protocol. Here the MR signal attenuation E at the diffusion wave-vector $\mathbf{q}$ of the q-space (a subset of Euclidean $\mathbf{R}^3$) is expressed as the following series in the SPF basis:

$$E(\mathbf{q}) = \sum_{n,l,m \in \mathbf{D}} a_{nlm} \sqrt{\frac{2}{\zeta^{3/2}} \frac{n!}{\Gamma(n+3/2)}} \exp\left(-\frac{q^2}{2\zeta}\right) L_n^{1/2}\left(\frac{q^2}{\zeta}\right) y_l^m(\hat{\mathbf{q}}), \qquad (1)$$

where the triplet $\{n, l, m\}$ stands for the index defined in the set $\mathbf{D} \in \mathbf{N}^2 \times \mathbf{Z}$ so that $n \in \mathbf{N}$ is the radial index, and $l \in \mathbf{N}$, $m \in \mathbf{Z}$, $-l \leq m \leq l$ are the angular indexes. The symbols a_{nlm} are the series coefficients, y_l^m are the real

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 157–165, 2011.

spherical harmonics (SH), and R_n is an orthonormal radial basis function made of Gaussian-Laguerre (GL) functions.

Let $L^2(\mathbf{S}^2)$ denote the space of square integrable functions on $\mathbf{S}^2$. It has been shown in [3] that any feature $f \in L^2(\mathbf{S}^2)$ of the diffusion PDF can be directly extracted from the modeled diffusion signal in the SPF expansion [3] as

$$f(\hat{\mathbf{r}}) = \sum_{n,l,m \in \mathbf{D}} a_{nlm} b_{nlm} y_l^m(\hat{\mathbf{r}}) = \sum_{n,l,m \in \mathbf{D}} f_{nlm} y_l^m(\hat{\mathbf{r}}), \tag{2}$$

where the symbol $\hat{\mathbf{r}}$ is a unit vector in the 2-sphere $\mathbf{S}^2$, b_{nlm} stands for the coefficients of the feature projection function expressed in the SPF expansion and $f_{nlm} = a_{nlm} b_{nlm}$ are the spherical harmonics coefficients of f. The feature related to the second-order Orientation Density Function (ODF_2) displacement of water molecules is expressed in the SPF basis as [4]:

$$b_{nlm} = \frac{\delta_{l0}\delta_{m0}}{\sqrt{4\pi}} - \frac{1}{8\pi} \sum_{j=0}^{n} \frac{(-1)^j}{j 2^j} \sqrt{\frac{2}{\zeta^{3/2}} \frac{n!}{\Gamma(n+3/2)}} \binom{n+1/2}{n-j} P_l(0)(-l^2+l). \tag{3}$$

Other spherical features of the diffusion PDF are given in [3,4]. Given a spherical feature $f \in L^2(\mathbf{S}^2)$ such as (3), which represents the orientation features of the diffusion PDF, this paper presents a method to compute the statistical dispersion of fibre populations within a voxel, *without any prior* assumptions on their number or shape (Section 2). We illustrate the validity of the proposed method on synthetic and physical phantom datasets (Section 3).

Among the previous related works, Seunarine *et al*. proposed in [5] to compute the anisotropy of each fibre population within the same voxel. Other relevant methods include the labelling of crossing and fanning fibre populations [6] and the computation of a torsion index [7], both using inter-voxel computations. To the extent of our knowledge, the present paper describes the first statistical method for assessing the partitioning of fibre populations within the same voxel. As such, it may have a significant impact in applications involving the study of neurological diseases from dMRI.

2 Theory

2.1 Fibre Population Dispersion (FPD)

Let $\mathbf{SO(3)}$ denote the rotation group. According to Euler's rotation theorem, any rotation $\Lambda \in \mathbf{SO(3)}$ can be described by three successive rotations by a set of Euler angles $\boldsymbol{\Theta} = (\alpha, \beta, \gamma)$ about three axes, where $0 \leq \alpha, \gamma \leq 2\pi$ and $0 \leq \beta \leq \pi$. There are 24 standard Euler angle conventions depending upon which axes are used and the order in which the rotations are applied. Throughout this paper, we use the zyz-configuration to compose the intrinsic rotations. In this setting, we can parametrize the general rotation Λ as a function of the set of Euler angles $\boldsymbol{\Theta}$.

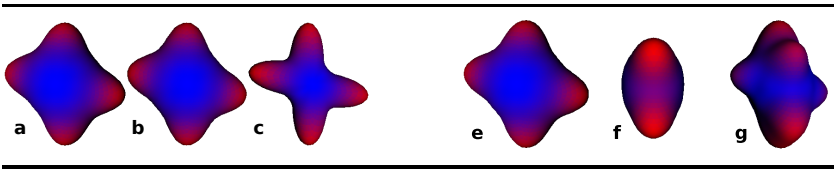

Fig. 1. The Fibre Population Dispersion (FPD) is defined as the autocorrelation of any spherical feature function defined on $\mathbf{S}^2$ of the diffusion PDF. *Left:* (a) Two-dimensional projection of the ODF_2 for two fibre populations crossing at 1.37 rad in the xy-plane. (b) $\mathbf{SO(3)}$ rotation of (a) by the set of Euler angles $\mathbf{\Theta}_1$. (c) The point-wise $\mathbf{S}^2$ multiplication of (a) and (b), from which the autocorrelation is computed as the integral volume. *Right:* The same computations, but with a different set of Euler angles $\mathbf{\Theta}_2$. The FPD of (c) is greater than the FPD of (g), with a ratio equal to 1.14, so that the rotation relates to the distance between intravoxel fibre populations as a function of Euler angles $\mathbf{\Theta}$. Note that in this figure, the functions are scaled for visualization purposes.

Let $\Lambda(\mathbf{\Theta})f(\hat{\mathbf{r}})$ denote a rotation Λ of the function $f \in L^2(\mathbf{S}^2)$ such that it is equivalent to $f(\Lambda^{-1}(\mathbf{\Theta})\hat{\mathbf{r}})$. We define the Fibre Population Dispersion (FPD) as the normalized autocorrelation of the spherical feature function $f \in L^2(\mathbf{S}^2)$, so that $\text{FPD} : \mathbf{SO(3)} \rightarrow \mathbf{R}$ and

$$\text{FPD}(\Lambda) = \int_{\hat{\mathbf{r}} \in \mathbf{S}^2} \frac{(f(\hat{\mathbf{r}}) - \bar{f})}{\sigma_f} \frac{\Lambda(\mathbf{\Theta})(f(\hat{\mathbf{r}}) - \bar{f})}{\sigma_r} \mathrm{d}\hat{\mathbf{r}}, \tag{4}$$

where σ_f and σ_r denote the standard variation of f and its rotated version $\Lambda(\mathbf{\Theta})f$. Fig. 1 clarifies the intuition behind (4), *i.e.*, how the correlation function between a spherical feature and a rotated version of it can recover fibre population dispersion. The left and right parts of Fig. 1 illustrate high and low values of the FPD, respectively.

Since the function f and its rotated version $\Lambda(\mathbf{\Theta})(f)$ have the same standard variations, the normalizing factor is therefore the variance of f, *i.e.*, σ_f^2. Let $\mathbf{D}^* = \mathbf{D} \backslash \{l = 0\}$ be the index domain minus all the terms which nullify the order l. Fortunately, the variance and the mean are directly expressed in the spherical harmonics space, which yields the following simplification of (4):

$$\text{FPD}(\Lambda) = N \int_{\hat{\mathbf{r}} \in \mathbf{S}^2} \sum_{n,l,m \in \mathbf{D}^*} f_{nlm} y_l^m(\hat{\mathbf{r}})\, \Lambda(\mathbf{\Theta}) \sum_{i,j,k \in \mathbf{D}^*} f_{ijk} y_j^k(\hat{\mathbf{r}}) \mathrm{d}\hat{\mathbf{r}}, \tag{5}$$

where the normalization factor N is obtained by inverting the sum of square coefficients f_{nlm}. When subject to a rotation, real spherical harmonics y_l^m can be expressed as a linear combination of real spherical harmonics of the same order l [8]. As a consequence, (5) can be expressed as

$$\text{FPD}(\Lambda) = N \int_{\hat{\mathbf{r}} \in \mathbf{S}^2} \sum_{n,l,m \in \mathbf{D}^*} f_{nlm} y_l^m(\hat{\mathbf{r}}) \sum_{i,j,k \in \mathbf{D}^*} f_{ijk} \sum_{k'=-j}^{j} R_{k'k}^{(j)}(\mathbf{\Theta}) y_j^{k'}(\hat{\mathbf{r}}) \mathrm{d}\hat{\mathbf{r}}, \tag{6}$$

where $R^{(j)}_{k'k}(\mathbf{\Theta})$ are the real Wigner-D functions. On integrating over the unit sphere, all terms vanish except $n = i$, $l = j$ and $m = k'$ to give an expression for the FPD at rotation Λ:

$$\mathrm{FPD}(\Lambda) = \left(\sum_{n,l,m \in \mathbf{D}^*} f^2_{nlm} \right)^{-1} \sum_{n,l,m \in \mathbf{D}^*} f_{nlm} \sum_{k=-l}^{l} f_{nlk} R^{(l)}_{km}(\mathbf{\Theta}), \tag{7}$$

in which the normalization factor N has been replaced by its expression. It is remarkable that (7) can be interpreted as a result of the Wiener-Khinchin theorem which states that the autocorrelation of f is the inverse $\mathbf{SO(3)}$–Fourier transform of the corresponding power spectral density function.

The expression in (7) is a closed-form solution of (4), which directly relates the FPD to the SPF expansion coefficients of f for a single voxel. This avoids the need for cumbersome numerical discretization schemes of f on the 2-sphere for each voxel, by turning the FPD computation into a fast dot product between vectors of SPF coefficients. Note that the result obtained in (7) is *not only* valid for the Spherical Polar Fourier (SPF) expansion method [3] but also for any local reconstruction method of the diffusion signal based on spherical harmonics functions (*e.g.*, Q-Ball Imaging (QBI) [9], the Diffusion Orientation Transform (DOT) [10] and Diffusion Propagator Imaging (DPI) [11]).

2.2 A Distance Metric

Having meaningful statistics on the FPD function $\mathbf{SO(3)} \rightarrow \mathbf{R}$ given in (7) requires the definition of a distance metric on the rotation group $\mathbf{SO(3)}$. Although it is possible to define a metric which uses the Euler angles, such a parametrization will degenerate at some points on the hypersphere, leading to the problem of *gimbal lock*, the loss of one degree of rotational freedom. We avoid this by using the quaternion representation $\mathbf{h} \in \mathbf{H} \subset \mathbf{R}^4$, *i.e.*, a set of four Euclidean coordinates $\mathbf{h} = \{(a, b, c, d) \,|\, a, b, c, d \in \mathbf{R}\}$ using the basis elements $1, i, j, k$ which satisfy the relationship $i^2 = j^2 = k^2 = ijk = -1$. Quaternions can parametrize rotations in $\mathbf{SO(3)}$ by constraining their norm to unity, *i.e.*, $\|\mathbf{h}\| = a^2 + b^2 + c^2 + d^2 = 1$. In this setting, the four scalar numbers have only three degrees of freedom and the unit quaternions form a subset $\mathbf{S}^3 \subset \mathbf{R}^4$. If we regard the quaternions as elements of $\mathbf{R}^4$, any usual L_p norm ρ_s can be used to define a metric on $\mathbf{SO(3)}$ between $\mathbf{h}_1$ and $\mathbf{h}_2$:

$$\rho(\mathbf{h}_1, \mathbf{h}_2) = \min \left\{ \rho_s(\mathbf{h}_1, \mathbf{h}_2), \rho_s(\mathbf{h}_1, -\mathbf{h}_2) \right\}, \tag{8}$$

Throughout this paper, ρ_s stands for the L_2 norm. It follows that $(\mathbf{H}, \rho)$ is a metric space. We define the p-th moment of the FPD about a given quaternion $\mathbf{h}_0 \in \mathbf{H}$ as

$$\mathcal{M}_p(\mathbf{h}_o) = \int_{\mathbf{h} \in \mathbf{H}} \rho(\mathbf{h}, \mathbf{h}_0)^p \, d\,\mathrm{FPD}(\mathbf{h}), \tag{9}$$

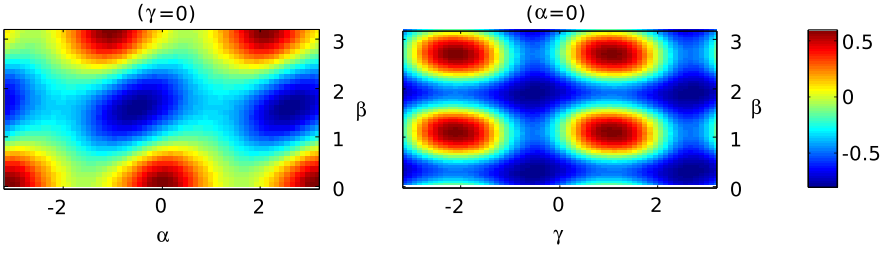

Fig. 2. The Fibre Population Dispersion (FPD) function reveals the autocorrelation of a spherical feature of the diffusion PDF. The FPD is computed on the ODF_2 composed of two fibre populations crossing at 0.5 rad in the xy-plane. Slices of the three-dimensional FPD are are plotted in this figure on a Cartesian grid by selecting a uniform sampling in Euler angles sets (α, β, γ), with $d\alpha = d\beta = d\gamma = 0.1$ rad (see text for a discussion).

with $p \in \mathbf{N}$. The p-th moments of the FPD defined in (9) give a quantitative measure of the shape of the FPD, and yield different values depending on the separation, size and numbers of fibre populations (as discussed further in Section 3). The moment $\mathcal{M}_p$ can be computed about any unit quaternion $\mathbf{h}_0 \in \mathbf{H}$, but the identity quaternion $\mathbf{h}_I = k$ represents a zero rotation and is consequently the most "natural" choice in the sense that it leads to the central moments $\mathcal{M}_p(\mathbf{h}_I)$.

3 Results

3.1 Synthetic Data

Fig. 2 illustrates a simulated diffusion PDF for two fibre populations crossing in the same voxel (left of Fig. 2) and its related Fibre Population Dispersion (FPD) function computed using the closed-form we provide in (7). We chose the Euler zyz-convention, therefore the symbol α denotes the Euler angle around the z-axis, β the angle around the rotated y-axis and γ around the twice-rotated z-axis (right of Fig. 2). Note that although apparent proximity or distance may appear in the FPD image of Fig. 2, this may be misleading because the Euler angles do not form an Euclidean metric. This results in a distorted mapping of the neighbourhood compared to a proper metric such as the quaternion distance $\rho(\mathbf{h}_1, \mathbf{h}_2)$ (8). Nonetheless, it is relatively easy to identify the periodicity of the FPD values of fibre populations which are *conveniently* oriented according to the Euler axis of rotations. In Fig. 2, the variations of the FPD function along β values on the line defined by $\gamma = 1.1$ rad and $\alpha = 0$ rad reveals the presence of two fibre populations at $\beta = \{1.1, 2.6\}$ rad.

The computation of the closed-form FPD expression given in (7) requires one to convert unit quaternions to Euler angles. A very efficient and generic conversion is given in [12].

Table 1. MICCAI 2009 Fibre Cup Phantom: parameters of the diffusion acquisition sequence [13]. The diffusion sensitization was applied along a set of 64 orientations.

Scanner Siemens TrioTim (12 channel)	Sequence=SSTR	Magnetic field=3 T
Voxel size= $3 \times 3 \times 3\text{mm}^3$	Image size=$64 \times 64 \times 3$	TR=5000ms
b-value= $\{650, 1500, 2000\}\text{s mm}^{-2}$	$\|\mathbf{g}\|_{\max}$=40 mT m^{-1}	TE=$\{77, 94, 102\}$ms

(a) Ground Truth (left) (b) GFA [14]. (c) FPD (our result).

Fig. 3. The Fibre Cup phantom consists of fibre bundles crossing at specific locations across the slice. We show a comparison of the generalized fractional anisotropy (GFA) measure and our fibre population dispersion (FPD) measure. The GFA image does not explicitly reveal crossings, while in the FPD image the brightness of a voxel is proportional to the crossing angle.

3.2 MICCAI 2009: Fibre Cup Phantom

For evaluating our proposed analytical computation of the fibre population dispersion (FPD), we present results on the MICCAI 2009 Fibre Cup phantom. It is composed of large bundles made of hydrophobic acrylic fibres whose diameter is of the same order of that of myelinated axons, with a density close to 1900 fibres/mm^2. The parameters of the acquisition are described in Table 1 and the phantom is shown in Fig. 3a. The sampling uses three spheres in the q-space, each having a q radius proportional to the b value, so that $b = (2\pi)^2(\Delta - \delta/3)q^2$. The SPF reconstruction technique [3] is able to naturally take advantage of the full set of samples provided, unlike HARDI reconstruction techniques which are restricted to a subset of samples (*i.e.*, a single sphere in the q-space). It was also recently demonstrated to lead to more accurate and robust reconstruction of the diffusion signal [4].

Fig. 3 illustrates the computation of the FPD on this phantom. As can be seen in the ground truth image (*c.f.*, Fig. 3a), there are a limited number of *crossings* located at specific regions of interest (ROIs) named (a-f). This controlled environment is ideal to assess the potential of the FPD. Furthermore, the ROIs can be sorted by the angle of the crossing, from nearly π rad at ROI (a) to

Table 2. Rat spinal cord phantom: parameters of the diffusion acquisition sequence [15]. The diffusion sensitization was applied along a set of 90 orientations.

Scanner Siemens Sonata	Sequence=SSTR	Magnetic field=1.5 T
Voxel size= $2.5 \times 2.5 \times 2.5\text{mm}^3$	Image size=$128 \times 96 \times 40$	TR=8000ms
b-value= $\{3000\}\text{s}\,\text{mm}^{-2}$	Acquisition time=15 min	TE=$\{110\}$ms

fanning fibres at ROIs (e,f). Fig. 3b shows the widely used generalized fractional anisotropy (GFA) measure [14], which reflects the degree of angular coherency but does not immediately reveal ROIs corresponding to crossings. In contrast, our FPD technique, with the moment $\mathcal{M}_8$ measure proposed in (9), brings to light these ROIs with brightness levels that depend on the angle of crossing, as illustrated in Fig. 3c. In our experiments, the moment p-th order is based on heuristics with a balance between good contrast for high moment order p and robustness to low signal-to-noise ratio for low order p. Our results indicates that $p = 2$ offers a good balance.

3.3 Rat Spinal Cord Phantom

For further evaluation of our FPD technique with HARDI data, Fig. 4 shows results computed on a biological phantom constructed from excised rat spinal cords with known connectivity [15]. Fig. 4a demonstrates that this biological phantom possesses a single region of crossing tracts. As before, we compare the GFA measure and our FPD technique using the moment $\mathcal{M}_6$.

The central voxel of the phantom interestingly exhibits a three dimensional crossing, as a result of one fibre curving over the other. We hypothesize that this diffusion profile results from the averaging of both tracts in a single voxel. Conse-

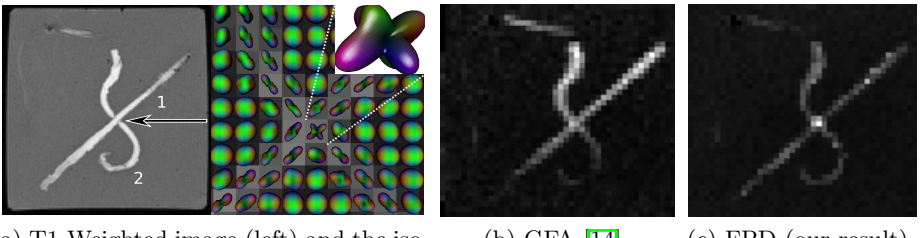

(a) T1-Weighted image (left) and the iso-probability feature [3] (right).

(b) GFA [14].

(c) FPD (our result).

Fig. 4. Rat spinal cord phantom: it is constructed from excised rat spinal cords embedded in 2% agar [15]. (a): an image of the phantom, and a zoom-in on the crossing. (b-c): comparison of the generalized fractional anisotropy (GFA) measure and our fibre population dispersion (FPD) measure. The latter clearly displays a single crossing region, in correspondance with the ground truth.

quently, this yields the highest FPD value in the slice, as shown in Fig. 4, whereas the GFA does not identify this region. The remaining parts of the rat spinal cords have a relatively constant brightness, which is consistent with the fact that their diffusion profiles are very similar, except for the preferred direction of diffusivity.

4 Conclusion

In this paper we have presented a novel method which uses images acquired with dMRI to statistically assess the intravoxel angular dispersion of fibre populations (FPD). We have demonstrated a closed-form expression for any spherical function based on the diffusion Probability Function (PDF), and expressed in the spherical harmonics basis. We have provided a proper metric on the rotation group for analyzing the FPD to obtain the relative distance between intravoxel fibre populations. This naturally leads to the introduction of scalar indices such as the moments of the FPD, which summarizes the intra-voxel dispersion of fibres. The proposed approach is based on a local per-voxel statistical analysis, and is thus fundamentally different from non-local techniques such as tractography. We have illustrated the potential of the FPD technique on both synthetic data and physical phantoms. Our experiments reveal features of the underlying microstructure, and in particular, reveal crossing regions that are not discernible in the widely used GFA images. As such, the FPD may yield interesting new possibilities for the detection and monitoring of neurological disease from dMRI.

References

1. Stejskal, E., Tanner, J.: Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient. J. Chem. Phys. 42, 288–292 (1965)
2. Stejskal, E.: Use of spin echoes in a pulsed magnetic-field gradient to study anisotropic, restricted diffusion and flow. J. Chem. Phys. 43(10), 3597–3603 (1965)
3. Assemlal, H.E., Tschumperlé, D., Brun, L.: Efficient and robust computation of PDF features from diffusion MR signal. Med. Image Anal. 13(5), 715–729 (2009)
4. Cheng, J., Ghosh, A., Deriche, R., Jiang, T.: Model-free, regularized, fast, and robust analytical orientation distribution function estimation. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. 6361, pp. 648–656. Springer, Heidelberg (2010)
5. Seunarine, K.K., Cook, P.A., Hall, M.G., Embleton, K.V., Parker, G.J.M., Alexander, D.C.: Exploiting peak anisotropy for tracking through complex structures. In: ICCV, pp. 1–8. IEEE, Los Alamitos (2007)
6. Savadjiev, P., Campbell, J., Descoteaux, M., Deriche, R., Pike, G., Siddiqi, K.: Labeling of ambiguous subvoxel fibre bundle configurations in high angular resolution diffusion MRI. NeuroImage 41(1), 58–68 (2008)
7. Savadjiev, P., Kindlmann, G., Bouix, S., Shenton, M., Westin, C.: Local white matter geometry from diffusion tensor gradients. NeuroImage 49(4), 3175–3186 (2010)
8. Su, Z., Coppens, P.: Rotation of real spherical harmonics. Found. Cryst. 50(5), 7673 (1994)

9. Anderson, A.: Measurement of fiber orientation distributions using high angular resolution diffusion imaging. Magn. Reson. Med. 54(5) (2005)
10. Özarslan, E., Sherperd, T.M., Vemuri, B.C., Blackband, S.J., Mareci, T.H.: Resolution of complex tissue microarchitecture using the diffusion orientation transform (DOT). NeuroImage 31, 1086–1103 (2006)
11. Descoteaux, M., Deriche, R., Bihan, D., Mangin, J., Poupon, C.: Multiple q-Shell Diffusion Propagator Imaging. Med. Image Anal. (2010)
12. Shoemake, K.: Graphics gems IV. In: Euler Angle Conversion, pp. 222–229. Academic Press Professional, Inc., London (1994)
13. Poupon, C., Rieul, B., Kezele, I., Perrin, M., Poupon, F., Mangin, J.F.: New diffusion phantoms dedicated to the study and validation of high-angular-resolution diffusion imaging (HARDI) models. Magn. Reson. Med. 60(6), 1276–1283 (2008)
14. Tuch, D.: Q-ball imaging. Magnetic Resonance in Medicine 52, 1358–1372 (2004)
15. Campbell, J., Siddiqi, K., Rymar, V., Sadikot, A., Pike, G.: Flow-based fiber tracking with diffusion tensor and q-ball data: validation and comparison to principal diffusion direction techniques. NeuroImage 27(4), 725–736 (2005)

Feasibility and Advantages of Diffusion Weighted Imaging Atlas Construction in Q-Space

Thijs Dhollander[1,2], Jelle Veraart[3], Wim Van Hecke[1,4,5], Frederik Maes[1,2], Stefan Sunaert[1,4], Jan Sijbers[3], and Paul Suetens[1,2]

[1] Medical Imaging Research Center (MIRC), K.U. Leuven, Leuven, Belgium
[2] Center for Processing Speech and Images (PSI), Department of Electrical Engineering (ESAT), Faculty of Engineering, K.U. Leuven, Leuven, Belgium
[3] Vision Lab, Department of Physics, University of Antwerp, Antwerp, Belgium
[4] Department of Radiology, University Hospitals of the K.U. Leuven, Leuven, Belgium
[5] Department of Radiology, University Hospital Antwerp, Antwerp, Belgium

Abstract. In the field of diffusion weighted imaging (DWI), it is common to fit one of many available models to the acquired data. A hybrid diffusion imaging (HYDI) approach even allows to reconstruct different models and measures from a single dataset. Methods for DWI atlas construction (and registration) are as plenty as the number of available models. Therefore, it would be nice if we were able to perform atlas building *before* model reconstruction.

In this work, we present a method for atlas construction of DWI data in q-space: we developed a new multi-subject multi-channel diffeomorphic matching algorithm, which is combined with a recently proposed DWI retransformation method in q-space.

We applied our method to HYDI data of 10 healthy subjects. From the resulting atlas, we also reconstructed some advanced models. We hereby demonstrate the feasibility of q-space atlas building, as well as the quality, advantages and possibilities of such an atlas.

1 Introduction

Diffusion weighted imaging (DWI), a magnetic resonance imaging (MRI) technique, is able to provide us with valuable information about tissue microstructure, for instance in the white matter of the brain, by measuring and imaging the self-diffusion of water within this tissue in vivo. Over the years, many models have been proposed to get a grip on the complex information acquired using a DWI protocol. Different models usually have different requirements on the quality and specifications of the acquired data. The recent strategy of hybrid diffusion imaging (HYDI) [1] accounts for this by proposing to acquire data from multiple shells of constant diffusion weighting in q-space. The resulting data (or the appropriate parts of it) can then easily be used for the reconstruction of different models and the calculation of different measures.

Methods for image registration and atlas building are plenty in number, and many have been devised for different (but not all) possible DWI models. The

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 166–173, 2011.

unique challenge involved in handling the information in DWI datasets is its orientational dependency. This particular fact renders spatial transformation of these data a difficult challenge of its own: apart from a standard interpolation, an additional reorientation or retransformation step has to be performed in every voxel. This is necessary to keep the information in all voxels in correspondence with the underlying orientational structure of the tissue. Different reorientation and retransformation strategies have already been proposed for some, but again not all, possible models. With the use of all these different models and their own unique registration and atlas building methods, the question also raises to what extent these choices might cause a bias on the results.

All of these problems could potentially be evaded and solved at the same time by being able to perform registration or atlas building *before* the reconstruction of any particular model. The result of such a strategy would have the advantage of still allowing any model to be reconstructed from the same transformed data in q-space. Recently, a method has been proposed to correctly transform raw data (sampled from different shells) in q-space [4]. This effectively opens up the path to registration and atlas building in the raw q-space. In this work, we combine this method with a newly developed multi-subject multi-channel diffeomorphic matching algorithm. The result is a true q-space atlas building method. We then apply this method to HYDI data of 10 healthy subjects. From the resulting atlas, we also reconstruct a selection of some more advanced models.

This work should illustrate the feasibility of q-space atlas construction (or registration) and at the same time highlight the advantages of a resulting q-space atlas by presenting some of its possibilities.

2 Methods

2.1 Acquisition

Data were acquired from 10 healthy subjects. This was done using a Siemens 3T scanner, with a $2.5\,mm$ isotropic voxel size. A HYDI approach was taken. For each subject, in addition to 10 non-DWI (B0) volumes (which were averaged), 3 different shells of q-space were sampled: 25 gradient directions at $b = 700\,s/mm^2$ ($q = 24.98\,mm^{-1}$), 40 gradient directions at $b = 1000\,s/mm^2$ ($q = 29.85\,mm^{-1}$) and 75 gradient directions at $b = 2800\,s/mm^2$ ($q = 49.96\,mm^{-1}$).

2.2 Matching

A new multi-subject multi-channel diffeomorphic matching algorithm was developed in order to match all subjects simultaneously in the average subject (atlas) space. The matching process was guided by 7 "channels", which were calculated straight from the raw data. The 1^{st} channel is simply the average B0 of each subject. The other channels are based on the apparent diffusion coefficient (ADC). The ADC of each acquired DWI volume is defined as $ADC = -\ln(S)/b$, where S is its signal intensity normalized by the average B0. The 2^{nd}, 3^{rd} and 4^{th} channel consist of the average ADC over each of the 3 shells. The 5^{th}, 6^{th} and

7^{th} channel contain a generalized fractional anisotropy (GFA) measure for each shell, which is calculated as $GFA = \text{std}(ADC)/\,\text{rms}(ADC)$.

The algorithm itself is a combination of diffeomorphic demons [2] and the idea of SyN [3] (no target image, both images deform), extended towards any number of subjects (N) and channels (M). At each iteration, given the currently deformed subjects I, an unconstrained correspondence update field $\vec{U}$ is calculated for every subject I_i:

$$\vec{U}(I_i) = \sum_{\substack{1 \leq k \leq M \\ 1 \leq j \leq N \\ j \neq i}} \frac{\vec{V}(I_{i,k}, I_{j,k})}{(N-1)M} \tag{1}$$

This is an average, over all $N-1$ other subjects and M channels, of a force $\vec{V}$ acting between this subject's channel $I_{i,k}$ and the corresponding channel $I_{j,k}$ of another subject. In our case, this results in $9 \times 7 = 63$ force fields acting on each subject at each iteration! The force $\vec{V}$ between 2 images I and J is defined as:

$$\vec{V}(I,J) = -\frac{(I-J)\left(\vec{\nabla}I + \vec{\nabla}J\right)/2}{\left\|\left(\vec{\nabla}I + \vec{\nabla}J\right)/2\right\|^2 + (I-J)^2/(2\epsilon)^2} \tag{2}$$

This symmetric ($\vec{V}(I,J) = -\vec{V}(J,I)$) force incorporates the gradients of both images and has theoretical as well as practical advantages [2]. The parameter ϵ can be set as the maximum step size.

The algorithm proceeds as diffeomorphic demons, but with every subject treated as a moving entity. In short: a fluid regularization (Gaussian filter) is applied to each $\vec{U}(I_i)$; the diffeomorphic update step (composition of the deformation fields with the fast vector field exponentials of the correspondence update fields) is performed; an elastic/diffusion regularization (Gaussian filter) is applied to each resulting new deformation field. All subject's channels are finally deformed according to the new total deformation fields. The whole process iterates, starting again from calculating (1).

The only parameters are ϵ and the standard deviations of the regularization kernels. We have them initially all set at 2 (voxels). When convergence is reached, they are all switched to 1, and finally to 0.5.

2.3 Deformation and Averaging

Using the final 10 deformation fields obtained from the matching algorithm, the raw acquired volumes are resampled. However, for DWI data, a retransformation step is also needed in each sampled voxel. In [4], it is shown that just reorienting the gradient directions using the forward affine matrix in each voxel produces wrong results, and a correct new method that also preserves isotropic and anisotropic volume fractions is proposed. We apply this method for each of the 3 shells, based on the Jacobian of the deformation fields in each voxel,

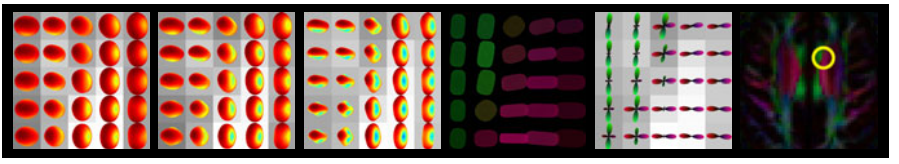

Fig. 1. 5x5 voxel detail of the actual atlas (3 analytic shells representing normalized signals in q-space: b700, b1000, b2800). Reconstructed tensors from b1000 as well as fODF's from b2800 are also shown. The region is indicated on an axial slice.

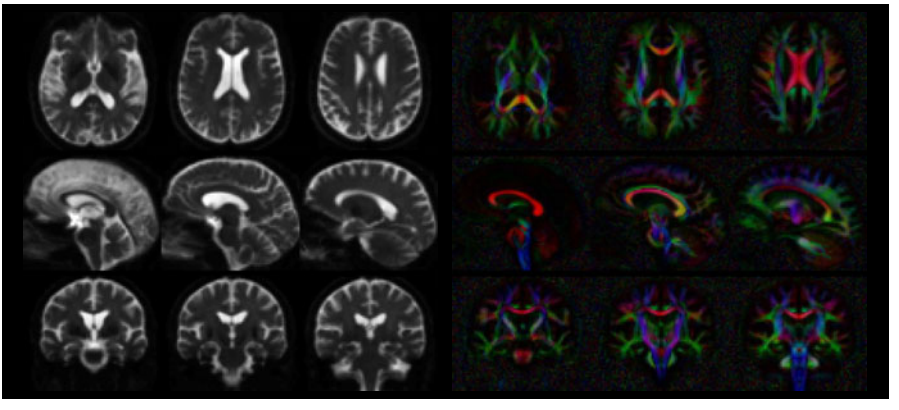

Fig. 2. Non-DWI (B0) (*left*) and CFA (*right*) of some axial, sagittal and coronal slices

using an order 6 spherical harmonics (SH) representation of the normalized signal, 1 isotropic and 300 anisotropic volume fractions, shaped by eigenvalues $0.0018\,mm^2/s$ and $0.0003\,mm^2/s$. The resulting SH coefficients of the 3 shells are averaged over all subjects, producing the final atlas: every voxel contains 3×28 SH coefficients, analytically representing 3 normalized signal shells in q-space. We will further refer to these as b700, b1000 and b2800.

2.4 Reconstructions

Apart from a simple tensor fit from diffusion tensor imaging (DTI) and a fiber orientation distribution function (fODF) reconstruction using spherical deconvolution (SD) [5], we will also demonstrate some more advanced reconstructions on the atlas to explore its possibilities to a larger extent. First we will perform a full brain fiber tractography on the fODF's resulting from constrained spherical deconvolution (CSD) [6]. This puts the angular quality of b2800 to the test. Next, we present different parameter maps resulting from a diffusion kurtosis imaging (DKI) model fit on the full atlas [7,8]. This is more of a test for the quality of the radial information, i.e. the relation between the different shells. Finally, we reconstruct the ensemble average propagator (EAP) using the recently proposed multiple q-shell diffusion propagator imaging (mq-DPI) strategy [9]. This technique exploits the full angular and radial information of the atlas.

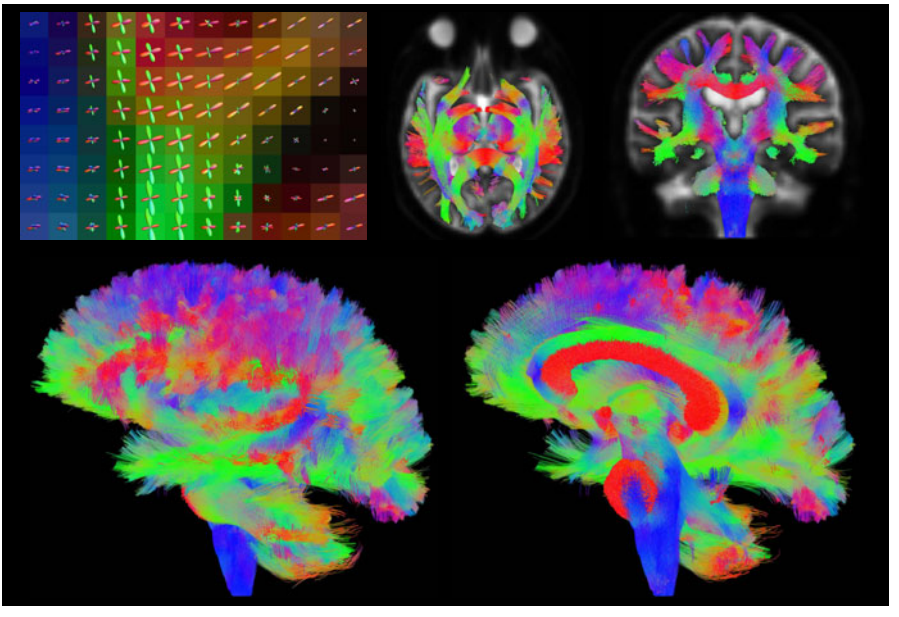

Fig. 3. Fiber tracking: fODF's from CSD (*top left*); 10 *mm* thick axial slab (*top middle*) and 3 *mm* thick coronal slab (*top right*) through the tract volume; sagittal view of the full tract volume (*bottom left*) and cut in half by the mid-sagittal plane (*bottom right*)

3 Results

In this section, we mainly focus on presenting visualisations of the atlas and its reconstructions, allowing for easy assessment and interpretation of the results.

3.1 The Atlas

The atlas itself consists of 3 analytic shells of normalized signals in q-space. For all purposes, these can be sampled along any set of gradient directions, but the SH representations themselves can also prove to be useful (in addition to being a compact representation of the atlas). A small detail (an area where the cingulum and the corpus callosum pass close to each other) of the actual atlas contents is shown in Fig. 1. Already easily exploiting the hybrid nature of the atlas, we also show reconstructed tensors from b1000 and fODF's from b2800 (using SD [5]), as well as some color-encoded fraction anisotropy (CFA) maps in Fig. 2. For the atlas to be complete, we should add the non-DWI volume to it (Fig. 2). If required, it can for instance be used to "de-normalize" samples from the shells, resulting in DWI images as they would be produced by a scanner.

3.2 Full Brain Fiber Tractography on fODF's Resulting from CSD

The MRtrix package (http://www.brain.org.au/software/) was used to perform CSD [6] and fiber tracking. To this end, we sampled b2800 in the original 75

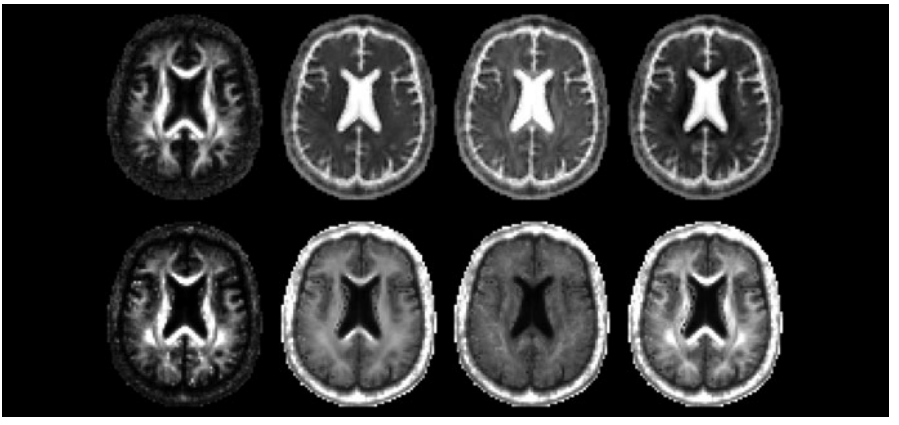

Fig. 4. DKI maps: FA, MD, AD, RD (*top row*) and KA, MK, AK, RK (*bottom row*)

directions, and provided this (just as if it were a normal dataset) to MRtrix. We performed CSD using an order 8 SH basis. Even though the atlas itself has only information "up to order 6", this is still useful because CSD is a nonlinear operation: the higher order can be used to fulfill the non-negativity constraint while preserving the quality of the resulting fODF. Fiber tracking was seeded in a white matter mask and constrained to a full brain mask. The fODF amplitude cutoff was 0.25 to initiate tracks and 0.15 to terminate tracks. A stepsize of $0.2\,mm$ and minimum radius of curvature of $1\,mm$ were used. Minimum and maximum track length were set to $15\,mm$ and $300\,mm$. Using these settings, 50000 deterministic and 50000 probabilistic tracks were generated, resulting in a dense 100000 tract volume. Results are shown in Fig. 3.

3.3 Parametermaps from DKI

The full atlas was sampled for the original 140 directions and the DKI [7] model was fit to the resulting dataset. DKI accounts for non-Gaussian diffusion (whereas DTI assumes Gaussian diffusion). This allows for an improved and b-value-independent estimation of the classical tensor parameters and at the same time produces some new parameters that quantify the non-Gaussianity [8]. We calculated the classical fractional anisotropy (FA), mean diffusivity (MD), axial diffusivity (AD) and radial diffusivity (RD), as well as the kurtosis anisotropy (KA), mean kurtosis (MK), axial kurtosis (AK) and radial kurtosis (RK). Results are shown in Fig. 4.

3.4 Full 3D Signal Fit and the EAP from mq-DPI

Recently, mq-DPI [9] has been proposed as a means of reconstructing the 3D EAP. The idea is to model data of several shells in q-space with solid harmonics (solutions of the Laplace equation in spherical polar coordinates). Using this representation, an analytical solution exists for the Fourier integral that relates the measured signals to the EAP. This time we will use the SH representation

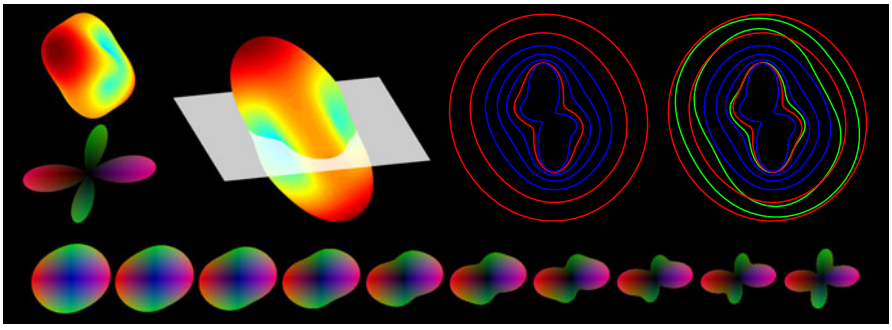

Fig. 5. mq-DPI in a single voxel: b2800 shell, fODF, axial plane through the shell (*top left*); solid harmonic fit in the axial plane for $q = 25, 30, 35, 40, 45, 50, 55\, mm^{-1}$ (*from outside to inside*) and with addition of the original b700, b1000, b2800 (*green*) (*top right*); EAP at radii ranging from $12\,\mu m$ to $21\,\mu m$ (*bottom row, left to right*)

of the 3 shells to fit the new model (solid harmonics) to. The implementation described in [9] can easily be adapted for this purpose, as it also implicitly performs a SH fit (but combined with the radial solid harmonic matrices). The final fit describes the signal in function of angle and radius (expressed by the q-value instead of b-value). We used an order 6 fit, modelling our previous 3×28 with 2×28 new coefficients. In Fig. 5, we try to give an impression of what such a fit looks like. We picked out a voxel containing a very general setting: 2 fiber populations with a different volume fraction and crossing at a certain angle. A view on the fit inside an axial plane is presented. Using the analytical solution of [9], we also reconstructed the EAP at different radii.

4 Discussion and Conclusion

In this work, we presented a method for atlas building of DWI data in the raw q-space. To this end, we combined a newly developed multi-subject multi-channel diffeomorphic matching algorithm with a recently proposed retransformation method [4]. We pursued methods that stay as close as possible to the original data: the channels used in the matching process were all calculated straight from the raw data and the retransformation also operates on the shells in q-space. We should note that the resulting method can just as well be used as a groupwise registration algorithm in different group studies. For the case of 2 subjects, it simply reduces to a (bias/template-free) registration algorithm.

We applied our method to HYDI data of 10 healthy subjects. The result is a DWI atlas, containing analytic representations of 3 shells of q-space data in each voxel. These shells can be sampled for any set of gradient directions, but the analytic SH representations themselves can also be used. We inspected the atlas for its quality and possibilities by performing a simple DTI fit and fODF (from SD [5]) reconstruction, as well as a full brain fiber tractography using fODF's from CSD [6], parametermaps from a DKI model fit [7,8] and a 3D

signal fit and EAP reconstruction from mq-DPI [9]. The outcomes of all these reconstructions are realistic and of good quality. This indicates that the full atlas building method must have brought all subjects in good correspondence, as well spatially as in orientational and radial structure in each voxel. Especially the latter parts are not evident. Because the retransformation is performed as a separate step after the matching and it is very sensitive to the deformation field, a sufficiently smooth deformation field is an absolute requirement. On the other hand, a good matching is of course also necessary. A good balance between both is the key to success. We believe that the fluid and elastic/diffusion regularization *both* play an invaluable part in achieving this.

To conclude, we have shown that DWI atlas construction in q-space is feasible, has advantages and can produce a qualitatively good result.

References

1. Wu, Y.C., Alexander, A.L.: Hybrid Diffusion Imaging. NeuroImage 36(3), 617–629 (2007)
2. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Diffeomorphic Demons: Efficient Non-parametric Image Registration. NeuroImage 45(1), S61–S72 (2009)
3. Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C.: Symmetric Diffeomorphic Image Registration with Cross-Correlation: Evaluating Automated Labeling of Elderly and Neurodegenerative Brain. Medical Image Analysis 12(1), 26–41 (2008)
4. Dhollander, T., Van Hecke, W., Maes, F., Sunaert, S., Suetens, P.: Spatial Transformations of High Angular Resolution Diffusion Imaging Data in Q-space. In: MICCAI 13, CDMRI Workshop, pp. 73–83 (2010)
5. Tournier, J.D., Calamante, F., Gadian, D.G., Connelly, A.: Direct Estimation of the Fiber Orientation Density Function from Diffusion-Weighted MRI Data using Spherical Deconvolution. NeuroImage 23(3), 1176–1185 (2004)
6. Tournier, J.D., Calamante, F., Connelly, A.: Robust Determination of the Fibre Orientation Distribution in Diffusion MRI: Non-negativity Constrained Super-resolved Spherical Deconvolution. NeuroImage 35(4), 1459–1472 (2007)
7. Jensen, J.H., Helpern, J.A., Ramani, A., Lu, H., Kaczynski, K.: Diffusional Kurtosis Imaging: The Quantification of Non-Gaussian Water Diffusion by Means of Magnetic Resonance Imaging. MRM 53(6), 1432–1440 (2005)
8. Veraart, J., Poot, D.H.J., Van Hecke, W., Blockx, I., Van der Linden, A., Verhoye, M., Sijbers, J.: More Accurate Estimation of Diffusion Tensor Parameters Using Diffusion Kurtosis Imaging. MRM 65(1), 138–145 (2011)
9. Descoteaux, M., Deriche, R., Le Bihan, D., Mangin, J.F., Poupon, C.: Multiple Q-shell Diffusion Propagator Imaging. Medical Image Analysis (2010) (in press)

Susceptibility Distortion Correction for Echo Planar Images with Non-uniform B-Spline Grid Sampling: A Diffusion Tensor Image Study

M.O. Irfanoglu[1], L. Walker[3], S. Sammet[2],
C. Pierpaoli[3], and R. Machiraju[1]

[1] Computer Sciences and Engineering Department,
The Ohio State University, USA
irfanoglu.1@osu.edu
[2] Department of Radiology,
The Ohio State University, USA
[3] National Institute of Child Health and
Human Development, NIH, USA

Abstract. In this paper, we propose a novel method for correcting the geometric distortions in diffusion weighted images (DWI) obtained with echo planar imaging (EPI) protocol. Our EPI distortion correction approach employs a deformable registration framework with the B-splines transformation, where the control point distributions are non-uniform and functions of the expected norm of the spatial distortions. In our framework, the amount of distortions are first computed by estimating the B_0 fieldmap from an initial segmentation of a distortion–free structural image and tissue susceptibility models. Fieldmap estimates are propagated to obtain expected spatial distortion maps, which are used in the sampling of active B-spline control points. This transformation is flexible in locations with large distortion expectations, yet with relatively few degrees-of-freedom and does not suffer from local optima convergence and hence does not distort anatomically salient locations. Results indicate that with the proposed correction scheme, tensor derived scalar maps and fiber tracts of the same subject computed from data acquired with different phase encoding directions provide better coherency and consistency compared traditional registration based approaches.

1 Introduction

Diffusion tensor images (DTI) are typically acquired using echo planar imaging (EPI) [1] acquisitions, which are vulnerable to static magnetic field inhomogeneities subsequently leading to image distortions. The amount of distortion in an EPI image is proportional to the static field inhomogeneity. Additionally, images obtained at higher field strengths suffer more from EPI distortion artifacts. The geometric distortions in areas with large magnetic susceptibility gradients, such as the sphenoid sinus, temporal lobe and brain stem can lead to incorrect tensor derived scalar maps and can also result in incorrect fiber tracking.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 174–181, 2011.

EPI distortions and their effects on scalar images is a well investigated problem and several schools of correction algorithms have been proposed. In their pioneering work, Jezzard *et al.* [2] employed a B_0 inhomogeneity map (fieldmap), which was computed from two gradient–echo scans with differing echo times. They showed that distortions are significant along the phase-encoding direction and are directly related to the inhomogeneity maps. Several other works followed, which employ the fieldmapping strategy [3]. Another approach to the problem involves a non-linear registration of the distorted B_0 EPI image to a structural anatomically correct MR image [4,5]

Fieldmapping based methods, in spite of their physical intuitiveness, suffer from the difficulty of calculating the phase maps near edges or in regions of high-field inhomogeneity. They also require additional scans and precise measurements, such as dwell time. Additionally [6,7] show that elastic registration based correction schemes usually outperform fieldmapping approaches. However, elastic registration based approaches usually suffer from the "curse of dimensionality" due to their large parametric space and result in distortion of salient anatomical locations. If the parametric space is reduced to deal with this problem, large distortions at high and ultra–high fields can not be corrected.

In this work, we propose a novel fast and robust algorithm for EPI distortion correction, which combines the strengths of fieldmap approaches and elastic registration. A non-uniform B-spline control grid is constructed, where the image is densely sampled with grid knots at locations with large expected distortions and sparsely sampled at locations where the distortions are homogenous. The expected distortions are obtained by synthesizing an artificial fieldmap based on tissue segmentation maps and tissue susceptibility models. Our methods are physically based due to their relationship to fieldmaps; however, they overcome the shortcomings of fieldmap based techniques given the increased adaptivity. Section 2 describes the main steps of the proposed pipeline, where our fieldmap estimation process is briefly reviewed in Section 2.1 and B-Spline knot sampling is described in Section 2.2. Our experimental setup and validation procedures are described in Section 3 and Section 3.1. The results are presented in Section 4 and the paper is concluded with future directions of Section 5.

2 Methodology

The proposed EPI distortion correction framework is demonstrated in Figure 1. After motion and distortion correction of the DWI data [8], the distortion-free structural (T_2) image is rigidly registered to the specific $b = 0\ s/mm^2$ image. It is then segmented into four classes, white matter, gray matter, cerebro–spinal fluid (CSF) and air in the $b = 0\ s/mm^2$ image's native space. The tissue label image is subsequently fed into a fieldmap estimation routine first proposed in [9], later employed in [10]. The estimated fieldmaps are transformed into deformation fields and these deformation fields are used to determine the non-uniform sampling regimen of B-spline knot points. The resulting transformation is used during elastic registration and the final deformation field is applied to each diffusion weighted volume.

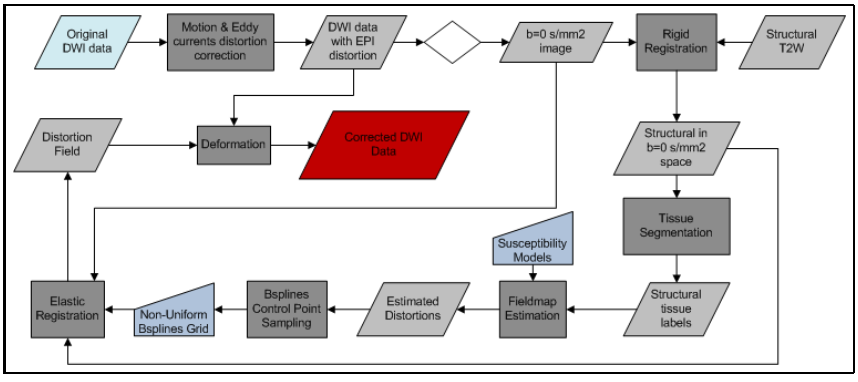

Fig. 1. The flow of the proposed EPI distortion correction framework

2.1 Fieldmap Estimation

Our fieldmap estimation process is based on the work of Jenkinson et al. [9], which models the first order perturbations in the main magnetic field B_z^0. Let the susceptibility, χ, be expanded as : $\chi = \chi_0 + \delta\chi_1$, where χ_0 is the susceptibility of air (4×10^{-7}) , δ is the susceptibility difference between tissue and air and χ_1 is a binary variable describing the tissue type. The first order perturbations in the magnetic field $B_z^{(1)}$ can be written in terms of the main magnetic field $B_z^{(0)}$:

$$B_z^{(1)} = \frac{\chi_1}{3+\chi_0} B_z^{(0)} - \frac{1}{1+\chi_0}((\frac{\partial^2 G}{\partial z^2}) * (\chi_1 B_z^{(0)}))$$

where G is the Green's function with $G(x) = (4\pi r)^{-1}$ with $(r = \sqrt{x^2+y^2+z^2})$. The solution to the convolution operation for a single voxel of resolution (a, b, c):

$$H(x) = (\frac{\partial^2 G}{\partial z^2}) * (\chi_1 B_z^{(0)}) = \sum_{i,j,k \in \{-1,1\}} (ijk) F(x + ia/2, y + jb/2, z + kc/2)$$

where $F(x) = \frac{1}{4\pi} atan(\frac{xy}{zr})$. For a set of voxels, due to the linearity of convolution operation, the perturbation field becomes: $B_z^{(1)}(x) = \sum_{x'} \chi_1(x') H(x - x')$

Assuming the phase encoding direction along the y axis, the deformation due to this fieldmap can be obtained with:

$$\Delta y = \frac{\gamma B_z^{(1)}(x) N_y t_{dwell}}{2\pi}$$

where γ is the gyromagnetic ratio, Δy is the pixel shift, N_y is the number of voxels along the phase encoding direction and t_{dwell} is the dwell time. As described in [10], this model does not account for shimming applied around the edges and might not be adequate for direct use in distortion correction. Figure 2 (a) displays a slice from a structural image and (b) the estimated fieldmap.

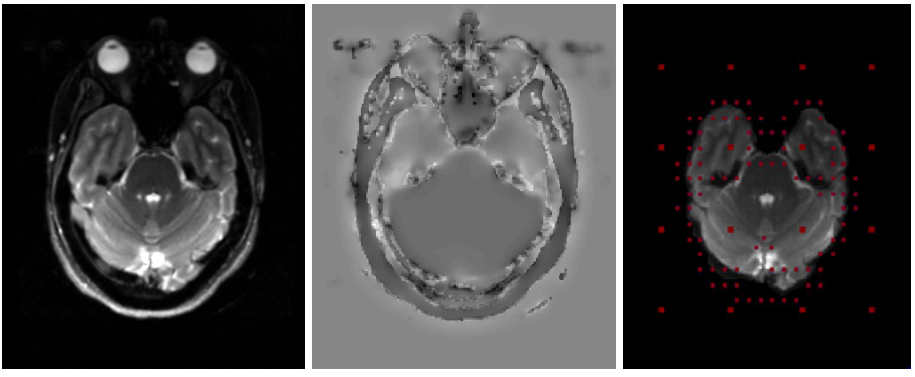

a) Structural image b)Unmasked fieldmap c) Knot point distribution.

Fig. 2. Examples slices for fieldmap estimation from a structural image and the sampling of a B-spline grid points

2.2 Adaptive B-spline Sampling Using Fieldmaps

The power of the proposed approach is its ability to model the complex deformations estimated with fieldmaps in a robust manner through a sophisticated physically-based transformation model. This is achieved by nonuniform sampling of B-spline grid locations as a function of the estimated displacements. The sampling of the knot points is performed algorithmically in a multi–resolution fashion as follows: Let $\mathcal{D}_f$ be the deformation field estimated using the fieldmaps.

1. Place a B-spline grid of $m \times m \times m$ onto $\mathcal{D}_f$. This partitions the image onto $(m-1)^3$ cubes. Typically, we set $m = 7$.
2. Generate $\mathcal{D}_e$ by only using values of $\mathcal{D}_f$ on grid locations and interpolating the other voxel displacements using cubic B-Spline kernels.
3. For each cube Ω:
   - Temporarily place a control point p at the center of Ω.
   - Recompute $\mathcal{D}_e(\Omega)$ for the given cube within the cube Ω.
   - If $||\mathcal{D}_f(\Omega) - \mathcal{D}_e(\Omega)|| > \varepsilon$ (a user–defined threshold),
     - Set p as a new control point.
     - Replace Ω with eight new cubes within itself.
     - Activate the new cubes.
   - Else deactivate Ω.
4. Repeat until no active cubes are present.

This strategy generates a multi-resolution pyramid of B-spline grids, with the initial level containing fully active $7 \times 7 \times 7$ control points (if $m = 7$), the second level containing partially active $15 \times 15 \times 15$ control points, and so on. Figure 2 (c) displays the knot sampling obtained with this procedure with the displacements computed from the fieldmap in (b). This image is brain masked for clarity purposes.

3 Experimental Setup

Five healthy volunteers were scanned on a 3T scanner using an eight channel coil. DWI data were acquired with $FOV = 24 \times 24\ cm$, slice thickness=2.5 mm, matrix size=128×128, 66 axial slices, parallel imaging factor of 2. The DWI data set consisted of ten images with $b = 0\ s/mm^2$, ten images with $b = 300\ s/mm^2$ and 60 images with $b = 1100\ s/mm^2$. For all subjects, two DTI scans were acquired with different phase encode direction (AP and RL) for validation purposes. Structural T_2 weighted anatomical images (T_2W) were also acquired.

During the elastic registration, the deformation field obtained in the previous level of the pyramid is used to initialize the coefficients of the B-spline at the current level. The optimizer used in the registration is a variant of limited memory Broyden-Fletcher-Goldfarb-Shanno (BFGS) method, which lets the user to set lower and upper bounds for the coefficients. At current level of the pyramid, the active grid points are allowed to move freely, whereas the inactive nodes were constrained to move only within 10% of the value obtained from the previous level. Additionally, control points were only allowed to move along the phase encoding direction. Up to four levels of the B-spline coefficient pyramid was employed in registration in practice.

3.1 Validation

For comparison purposes, another elastic registration algorithm with B-spline transformation of uniform grid sampling (e.g., $10 \times 10 \times 10$) was implemented as reference. The performance of such a registration scheme has been previously shown to outperform that of a fieldmap based approach in most cases [6].

The validation of the proposed correction framework was two-folds: First, the overlap of the structural image, the distorted image, image corrected with the reference algorithm and the image corrected with the proposed algorithm were visually assessed. Second, to quantify the overall improvements in diffusion tensor image quality, probabilistic tractography was carried out on cingulum bundle. The tensor images for every subject, for each phase encode direction and correction scheme, were then fed into a tensor field-based elastic registration routine to compute a population average tensor image and the transformations that mapped each data onto this average brain space. The transformations were subsequently applied to the corresponding tract images in such a way that every single tract image resided in the same coordinate framework. These population connectivity images were first visually analyzed. For these tract images, voxel values were subsequently statistically compared with Wilcoxon Rank test to check for the equality of the medians of number of visitations and to deduce consistency differences between AP and RL; and reference v.s proposed method.

4 Results

Figure 3 displays the results for the scalar images. The first column of the figure depicts the original undistorted T_2W structural image, the second column

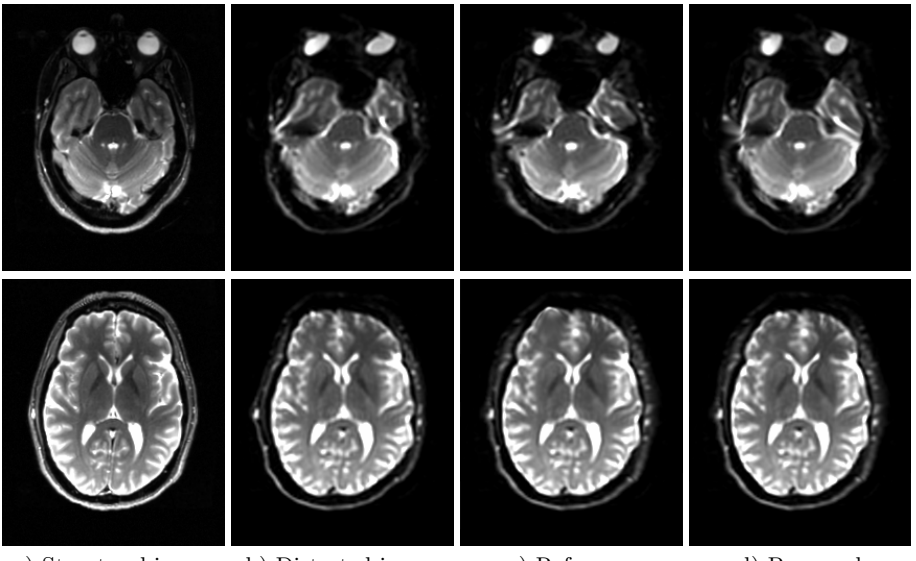

a) Structural image b) Distorted image c) Reference d) Proposed

Fig. 3. Results of the proposed correction algorithm and the reference elastic registration algorithm on the $b = 0\ s/mm^2$ images. Phase encoding direction is RL.

displays slices from the distorted $b = 0\ s/mm^2$ image, third column shows the output of the proposed correction scheme and the last column is for the reference algorithm. The distortion is along the RL direction. In the lower brain, where the distortion is significant, the proposed algorithm performs significantly better than the reference method in the right temporal lobe and left limbic lobe. Both approaches can correct for global displacements but the proposed approach can cope with large local displacements due to its transformation model. For mid-level slices (bottom row), both approaches perform well, with the proposed method showing a slightly better performance around middle frontal gyrus (top left portion of images).

4.1 Tractography Results

The average population tracts for the cingulum bundle over 5 subjects, for data acquired with RL and AP distortion are displayed in Figure 4. In this figure, brightness of tracts indicate the probability of reaching the voxel from the seed ROI. Following conclusions can be drawn from the figure:

- Both correction schemes improve the continuity of the tracts indicated by brighter shades along the tracts. Therefore, one can conclude that EPI distortion correction is an often neglected but crucial step in DTI processing.
- Continuities are also improved with the proposed method.

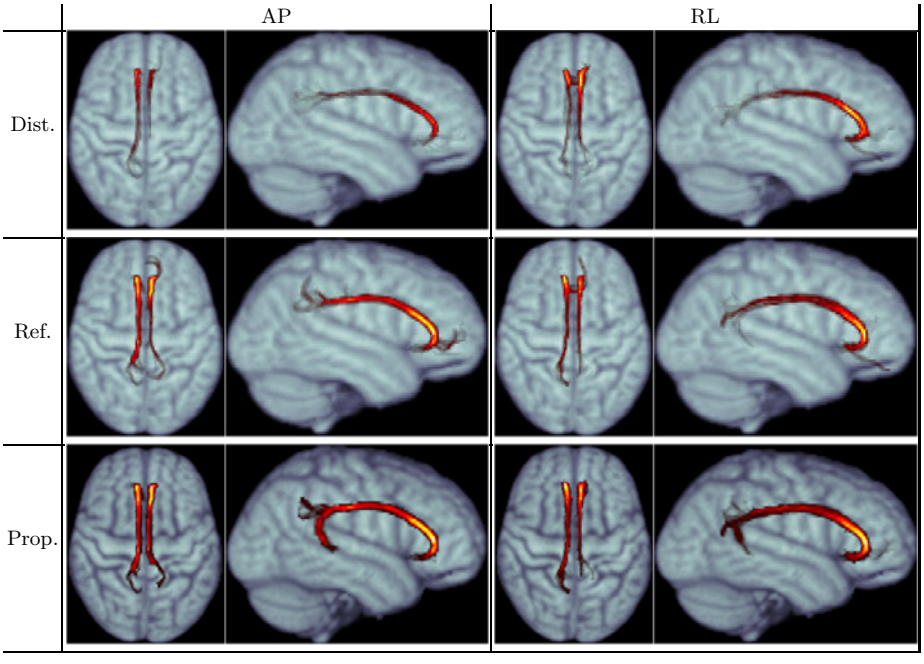

Fig. 4. Population tract averages. Left column displays the results for data acquired with AP distortion and the right column RL distortion. Top row tracts are computed with distorted data, and the bottom row with the proposed correction scheme.

– The proposed method also increases the consistency between data acquired with different distortion directions as the similarity of tracts increases with the reference method compared to the distorted data and with the proposed method compared to the reference approach.

Table 1 displays the results of the statistical tests. These statistics indicate that the proposed correction algorithm results in improved continuity along the tracts. Even though tracts obtained from RL distorted data and AP distorted data are still not statistically significantly similar, a considerable improvement can be observed in their consistency behavior.

Table 1. Statistics on equality of the median of cingulum visitation distributions. The mean values indicate the average number of visitations per voxels along the tracts. 10,000 tracts were initially casted per seed voxel.

H_0	p–value	μ	
$\mu_{RL_{ref}} = \mu_{AP_{ref}}$	$1.19\ 10^{-6}$	AP_{ref}	4275
$\mu_{RL_{prop}} = \mu_{AP_{prop}}$	$8.37\ 10^{-3}$	RL_{ref}	3415
$\mu_{RL_{prop}} = \mu_{RL_{ref}}$	$1.81\ 10^{-7}$	AP_{prop}	4776
$\mu_{AP_{prop}} = \mu_{AP_{ref}}$	0	RL_{prop}	3854

5 Conclusions and Future Work

In this work, we proposed a novel EPI distortion correction approach that combines the strengths of fieldmap based and elastic registration based correction approaches and minimizes their pitfalls. The algorithm has been shown to perform better than a typical elastic registration based approach both in terms of overlaps of single images and tracts computed from a set of diffusion weighted images of a population. As a future direction, the deformation maps generated by this approach will be used in a point spread function scheme and the effects of EPI signal washout at ultra-high fields will be investigated.

References

1. Mansfield, P.: Multi-planar image formation using NMR spin echoes. J. Phys. C. 10, 55 (1977)
2. Jezzard, P., Balban, R.: Correction for geometric distortion in echo planar images from b0 field variations. Magnetic Resonance in Medicine 34, 65–73 (1995)
3. Reber, P.J., Wong, E., Buxton, R., Frank, L.R.: Correction of off resonance related distortion in echo-planar imaging using epi-based field maps. Magn. Res. in Med. 39, 328–330 (1998)
4. Kybic, J., Thevenaz, P., Nirkko, A., Unser, M.: Unwarping of unidirectionally distorted EPI images. IEEE Trans. on Medical Imaging 19(2), 80–93 (2000)
5. Ardekani, S., Sinha, U.: Geometric distortion correction of high-resolution 3T diffusion tensor brain images. Magn. Res. in Med. 54, 1163–1171 (2005)
6. Wu, M., Chang, L.C., Walker, L., Lemaitre, H., Barnett, A.S., Marenco, S., Pierpaoli, C.: Comparison of EPI distortion correction methods in diffusion tensor MRI using a novel framework. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part II. LNCS, vol. 5242, pp. 321–329. Springer, Heidelberg (2008)
7. Tao, R., Fletcher, P.T., Gerber, S., Whitaker, R.T.: A variational image-based approach to the correction of susceptibility artifacts in the alignment of diffusion weighted and structural MRI. Inf. Process Med. Imaging 21, 664–675 (2009)
8. Rohde, G., Barnet, A., Basser, P., Marenco, S., Pierpaoli, C.: Comprehensive approach for motion and distortion correction in diffusion weighted MRI. Magnetic Resonance in Medicine 51, 103–114 (2004)
9. Jenkinson, M., Wilson, J.L., Jezzard, P.: Perturbation method for magnetic field calculations of nonconductive objects. Magn. Res. in Med. 52, 471–477 (2004)
10. Poynton, C., Jenkinson, M., Whalen, S., Golby, A.J., Wells, W.: Fieldmap-free retrospective registration and distortion correction for epi-based functional imaging. Med. Image Comput. Comput. Assist Interv. 11, 271–279 (2008)

Probabilistic ODF Estimation from Reduced HARDI Data with Sparse Regularization

Antonio Tristán-Vega and Carl-Fredrik Westin

Laboratory of Mathematics in Imaging, Brigham and Women's Hospital, Boston
{atriveg,westin}@bwh.harvard.edu

Abstract. High Angular Resolution Diffusion Imaging (HARDI) demands a higher amount of data measurements compared to Diffusion Tensor Imaging (DTI), restricting its use in practice. We propose to represent the probabilistic Orientation Distribution Function (ODF) in the frame of Spherical Wavelets (SW), where it is highly sparse. From a reduced subset of measurements (nearly four times less than the standard for HARDI), we pose the estimation as an inverse problem with sparsity regularization. This allows the fast computation of a positive, unit-mass, probabilistic ODF from 14-16 samples, as we show with both synthetic diffusion signals and real HARDI data with typical parameters.

1 Introduction

Diffusion Magnetic Resonance Imaging (MRI) provides an unparalleled tool to probe the connectivity of the nerve fibers within the white matter of the brain *in vivo*. At the core of this technique is the relationship between the signal $E(\mathbf{q})$, acquired by the MRI scanner when a pulsed gradient encoded by the wave-vector $\mathbf{q}$ is applied, and the random process $P(\mathbf{R})$ driving the restricted diffusion of water molecules. Such relationship is modeled as the following Fourier transform [1]:

$$P(\mathbf{R}) = \iiint_{\mathbb{R}^3} |E(\mathbf{q})| \exp\left(-j2\pi\mathbf{q}\cdot\mathbf{R}\right) d\mathbf{q}. \tag{1}$$

The obvious way to recover $P(\mathbf{R})$ is to sample $E(\mathbf{q})$ and approximate eq. (1) as a discrete Fourier transform, what is known as Diffusion Spectrum Imaging [2]. The main drawback of this technique is the need to sample the entire space of wave-vectors $\mathbf{q}$, since each of them implies the acquisition of a whole MRI volume. Though it has been recently shown [3] how this amount of data can be drastically reduced using Compressed Sensing (CS, [4]), this approach still requires several hundreds of samples to attain a reliable reconstruction.

High Angular Resolution Diffusion Imaging (HARDI, [5]) permits reduction of the measurements to a sampling of the space of orientations with constant $\|\mathbf{q}\|$. This is achieved at the expense of losing the radial information in $P(\mathbf{R})$, computing some sort of projection or Orientation Distribution Function (ODF) [2]:

$$\Phi(\mathbf{r}) = \int_{-\infty}^{\infty} P(R\mathbf{r})R^2 dR = \frac{-1}{8\pi^2} \iint_{\perp\mathbf{r}} \Delta E(\mathbf{q}) d\mathbf{q}. \tag{2}$$

There are some other alternative definitions for the orientation function (see [6]

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 182–190, 2011.

for an excellent review on this topic), but eq. (2) has the advantage of being a true probability law [7]. HARDI typically requires 60-100 wave-vectors $\mathbf{q}$, but considerations similar to those in [3] allow relaxation of this requirement to a few tens (16-30) [8,9], whenever a sparse representation of $E(\mathbf{q})$ may be achieved. This condition is met by representing $E(\mathbf{q})$ in a suitable vector frame, namely Spherical Ridgelets (SR), whose components closely resemble the shape of $E(\mathbf{q})$ at different scales. This frame is adequate in [8] since the object of interest is the diffusion signal. Unfortunately, these plate-like functions do not fit the cigar-shaped appearance of the probabilistic ODF, $\Phi(\mathbf{r})$, and cannot be expected to provide a sparse representation of it.

In **section 2**, we postulate Spherical Wavelets [10] as a suitable frame to sparsely represent the $\Phi(\mathbf{r})$ in eq. (2), as depicted in Fig. 1. In **section 3** we formulate the estimation of the ODF as an inverse problem, by relating the SW functions to their counterparts in the $\mathbf{q}$-space. We propose two alternative solutions based on either ℓ_1 or ℓ_2 regularization, justifying that non-negative ℓ_2 provides sparse representations analogous to ℓ_1. Finally, the results in **section 4** illustrate how positive, unit-mass, probabilistic ODFs can be recovered from as few as 14 gradient directions. With ℓ_2 regularization, the minimization problem can be solved with Newton-like methods, becoming computationally efficient.

2 Spherical Wavelets to Sparsely Represent $\Phi(\mathbf{r})$

The shape of the ODF in Fig. 1 implies that a combination of *a few* Spherical Wavelets (SW, [10]) with different widths and orientations may accurately represent $\Phi(\mathbf{r})$ for an arbitrary ensemble of diffusion compartments, i.e. the coefficients of $\Phi(\mathbf{r})$ in its SW expansion are sparse. The domain of definition of SW is the set $\Omega \equiv \{\mathbf{r} \in \mathbb{R}^3 : \|\mathbf{r}\| = 1\}$ of unit directions (i.e. diffusion gradients). Their definition is founded on the basis of Spherical Harmonics (SH), for which the following summation formula holds:

$$\sum_{m=-l}^{l} Y_l^m(\mathbf{u}) Y_l^m(\mathbf{v}) = \frac{2l+1}{4\pi} P_l(\mathbf{u}\cdot\mathbf{v}), \tag{3}$$

where Y_l^m is the real SH basis function of degree $l \geq 0$ and order $-l \leq m \leq l$, and P_l is the Legendre polynomial of degree l. Note $P_l(\mathbf{u}\cdot\mathbf{v})$ is a rotation of $Y_l^0(\mathbf{u})$ to an auxiliary system for which the 'z' axis is aligned with $\mathbf{v}$. Any symmetric (i.e. $f(\mathbf{u}) = f(-\mathbf{u})$), square-integrable function $f \in \mathbb{L}^2_{\text{sym}}(\Omega)$, can be written as [10]:

$$f(\mathbf{u}) = \int\!\!\int_{\Omega} f(\mathbf{v}) \sum_{l=0,2,\ldots}^{\infty} \frac{2l+1}{4\pi} P_l(\mathbf{u}\cdot\mathbf{v}) d\mathbf{v} = \int\!\!\int_{\Omega} f(\mathbf{v}) \Theta_{\mathbf{v}}(\mathbf{u}) d\mathbf{v}, \tag{4}$$

where $\Theta_{\mathbf{v}}(\mathbf{u})$ is a spherical convolution kernel. To achieve a scale-space representation of $f(\mathbf{u})$, the following scaling kernel $K_{\mathbf{v},j}(\mathbf{u})$ and **semi**-discrete frame $\bigcup_{j=-1}^{\infty} \bigcup_{\mathbf{v}\in\Omega} \Psi_{\mathbf{v},j}(\mathbf{u})$, that spans $\mathbb{L}^2_{\text{sym}}(\Omega)$, are introduced in [10]:

$$K_{\mathbf{v},j}(\mathbf{u}) = \sum_{l=0,2,\ldots}^{\infty} \frac{2l+1}{4\pi} \kappa_\rho\left(2^{-j} l\right) P_l(\mathbf{u}\cdot\mathbf{v}), \tag{5}$$

$$\Psi_{\mathbf{v},j}(\mathbf{u}) = \begin{cases} K_{\mathbf{v},0}(\mathbf{u}), & j = -1; \\ K_{\mathbf{v},j+1}(\mathbf{u}) - K_{\mathbf{v},j}(\mathbf{u}), & j \geq 0, \end{cases} \tag{6}$$

where we choose $\kappa_\rho(x) = e^{-\rho x(x+1)}$. When the resolution $j \to \infty$, $\kappa_\rho\left(2^{-j}l\right) \to 1$ for each l, and we have the exact recovery in eq. (4). When $j \to -\infty$, $\kappa_\rho\left(2^{-j}l\right) \to 0$ for $l \neq 0$, and we have the DC component of $f(\mathbf{u})$. This reasoning clarifies the multiscale nature of $K_{\mathbf{v},j}$, from coarse ($j \to -\infty$) to fine details ($j \to \infty$).

3 ODF Reconstruction from Sparse Representations

From the representation provided by SW as described in the previous section, we aim to estimate the ODF, $\Phi(\mathbf{r})$, as the sparsest combination of SW that can explain the data acquired in the $\mathbf{q}$-space, $E(\mathbf{q})$. Therefore, we need to compute the $\mathbf{q}$-space response of each SW defined in $\mathbf{r} \in \Omega$ (section 3.1), and then discretize the problem to account for the discrete nature of the measurements (section 3.2).

3.1 Response of SW $\Psi_{\mathbf{v},j}(\mathbf{r})$ in the Space of Wave-Vectors q

We are interested in probabilistic ODF estimators [7], so that the response pursued obeys the following form:

$$\Phi(\mathbf{r}) = \frac{1}{4\pi} + \frac{1}{16\pi^2}\mathcal{G}\left\{\Delta_b \zeta\left(E(\mathbf{q})\right)\right\}(\mathbf{r}), \tag{7}$$

Among the chances provided in [7], we choose $\zeta(\cdot) = -E_i(-\log(\cdot))$ as it seems to be a trade-off between angular resolution and noise sensitivity. Also, the unit-mass constraint comes for free with this model, since the Beltrami operator Δ_b eliminates all DC components so that we keep only the $1/4\pi$ factor. We denote by $\mathcal{G}$ the Funk-Radon Transform (FRT, [5]).

Since SH are eigenfunctions for both $\mathcal{G}$ and Δ_b, it is straightforward to prove this property holds also for the rotated versions of $Y_l^0(\mathbf{u})$, i.e. $P_l(\mathbf{u}\cdot\mathbf{v})$. If we insert eqs. (5) and (6) into (7), we get the following duality between the SW, $\Psi_{\mathbf{v},j}(\mathbf{r})$, and their counterparts in the $\mathbf{q}$-space, $\Xi_{\mathbf{v},j}(\mathbf{q})$ (see Fig. 1):

$$\Psi_{\mathbf{v},j}(\mathbf{r}) = \sum_{l=2,4,\ldots}^{\infty} \frac{2l+1}{4\pi}\nu_{\rho,j}(l)P_l(\mathbf{r}\cdot\mathbf{v}) \Leftrightarrow \Xi_{\mathbf{v},j}(\mathbf{q}) = \sum_{l=2,4,\ldots}^{\infty} \frac{4\pi(2l+1)}{-l(l+1)}\frac{\nu_{\rho,j}(l)}{\lambda_l}P_l(\mathbf{q}\cdot\mathbf{v}),$$
$$\nu_{\rho,j}(l) = \begin{cases} \kappa_\rho(l), & j = -1; \\ \kappa_\rho\left(2^{-j-1}l\right) - \kappa_\rho\left(2^{-j}l\right), & j \geq 0, \end{cases} \tag{8}$$

where λ_l and $-l(l+1)$ are respectively the eigenvalues for $\mathcal{G}$ and Δ_b. Note we drop the DC term $l = 0$ from the sums, according to our previous statement.

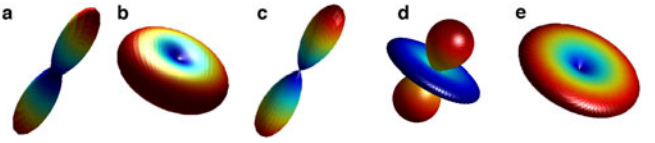

Fig. 1. Examples of: **a)** the true $\Phi(\mathbf{r})$ when $P(\mathbf{R})$ is Gaussian; **b)** its corresponding diffusion signal $E(\mathbf{q})$; **c)** a SW, $\Psi_{\mathbf{v},j}(\mathbf{r})$. These wavelets aim to sparely represent signals of the type a), which seems quite reasonable; **d)** the response $\Xi(\mathbf{q})$ to $\Psi_{\mathbf{v},j}(\mathbf{r})$ in the space of measurements; **e)** a SR, which in [8] aims to sparsely represent signals of the type b). Note this shape prevents the use of SR to represent ODFs like a).

3.2 Optimization Problem with ℓ_p Regularization

Like in [8], we need a fully discrete version of the frame we use for sparse representation: let $\{\mathbf{q}_n\}_{n=1}^{N}$ be the set of wave-vectors in the HARDI data-set; let $\{\mathbf{r}_m\}_{m=1}^{M}$ be the set of directions for which the ODF is computed; let $\{\mathbf{v}_{j(k),k}\}_{k=1}^{K}$ be a collection of discretizations of Ω for each resolution level $-1 \leq j \leq J$. The estimation of $\mathbf{x} \equiv \Phi(\mathbf{r})$ is arranged as the ℓ_p inverse problem:

$$\begin{array}{l} \min_{\mathbf{a}} \|\mathbf{a}\|_p \,, \text{ s.t. } \|A\mathbf{a}-\mathbf{y}\|_2^2 < \eta \\ \text{with: } \mathbf{x} = B\mathbf{a} + \frac{1}{4\pi} \end{array} \;\Bigg|\; \begin{array}{l} A_{n,k} = \Xi_{\mathbf{v}_{j(k),k},j(k)}(\mathbf{q}_n);\ \mathbf{y}_n = E(\mathbf{q}_n); \\ B_{m,k} = \Psi_{\mathbf{v}_{j(k),k},j(k)}(\mathbf{r}_m);\ \mathbf{x}_m = \Phi(\mathbf{r}_m). \end{array} \tag{9}$$

In brief, we discretize the responses $\Xi_{\mathbf{v},j}(\mathbf{q})$ into a $N \times K$ *sampling matrix* A to find that $\mathbf{a}$ best describing the $N \times 1$ vector $\mathbf{y} \equiv E(\mathbf{q})$. These coefficients are used to recover $\mathbf{x}$ from the discretized $M \times K$ *sparsifying matrix* B. We deliberately choose the same names used in [4], given the formal similarity of eq. (9) with CS for $p = 1$. Typically, $N \ll K$, so that $A\mathbf{a} = \mathbf{y}$ is highly underdetermined. This is the reason why $\mathbf{a}$ has to be assumed sparse *a priori*. Though sparsity is only enforced in case $p \leq 1$ [4], we also consider here the case $p = 2$ for the reasons detailed below. Note the ℓ_2 problem has a closed-form solution for eq. (9):

$$\mathbf{a} = A^T \left(AA^T + \tau(\eta) \cdot I_N\right)^{-1} \mathbf{y}, \tag{10}$$

3.3 Positivity Constraints on the Probabilistic ODF

Since $\Phi(\mathbf{r})$ is a probability law, it has to be positive for all $\mathbf{r}$. Given a set of directions $\{\mathbf{r}'_{m'}\}_{m'=1}^{M'}$, we add M' additional linear constraints to eq. (9):

$$B'\mathbf{a} + \tfrac{1}{4\pi} \geq 0, \quad B_{m',k} = \Psi_{\mathbf{v}_{j(k),k},j(k)}(\mathbf{r}'_{m'}). \tag{11}$$

Two important remarks arise here regarding the ℓ_2 problem: 1) the solution in eq. (10) no longer holds, but eq. (9) becomes a quadratic program which is still very efficient to solve compared to ℓ_1. 2) Non-negative least squares in fact provide sparse solutions [11], so that we may reasonably hypothesize (see [9]) that the solution of the constrained ℓ_2 will approach that of the constrained ℓ_1.

4 Results

Synthetic Experiments. We generate a two-tensor model simulating two fibers crossing in angles 45^o, 60^o, or 90^o. The results are the average for 300 trials with random rotations of the tensor ensemble, with the diffusion signal generated as:

$$S(\mathbf{u}) = \sqrt{\left(\frac{1}{2}\exp\left(-b\mathbf{u}^T\mathbf{D}_1\mathbf{u}\right) + \frac{1}{2}\exp\left(-b\mathbf{u}^T\mathbf{D}_2\mathbf{u}\right) + \sigma\cdot n_c\right)^2 + (\sigma\cdot n_s)^2}, \tag{12}$$

for eigenvalues of $\mathbf{D}_n$: $[1.7, 0.3, 0.3]\cdot 10^{-3}$ mm^2/s, $b = 2{,}000$ s/mm^2, $n_{c,s} \sim \mathcal{N}(0,1)$, and Signal to Noise Ratios (SNR): $1/\sigma$ =100, 40.

All the methods are coded in MatLab®, using the ℓ_1-magic library (`http://www.acm.caltech.edu/l1magic/`) to solve ℓ_1 and `quadprog` to solve ℓ_2.

The parameters are fixed following the guidelines in [8]: a width $\rho = 0.75$ of the kernel in eq. (5) is chosen by trial and error for best performance, and so it is the error threshold η in eq. (9) (its value depends on the scenario). For the directions $\mathbf{r}$, $\mathbf{r}'$, and $\mathbf{q}$, we choose regular samplings of Ω in all cases, with respective set sizes $M = 214$, $M' = 50$, and N (the number of HARDI measurements) ranging from 12 to 28. For the discretization of the SW frame, we uniformly sample Ω at each level $-1 \leq j \leq J$, taking $(2^{j+1}m_0 + 1)^2$ evenly spaced samples, with $m_0 = 4$ and $J = 1$.

For the sake of comparison, we have adapted the approach in [8] (namely, SR) to interpolate $\zeta(E)$, and used eq. (7) to evaluate the probabilistic ODF. We have also adapted the method described in section 3 by alternatively representing the ODF in the basis of SH with $l \leq 4$, where $\Phi(\mathbf{r})$ **is not** sparse[1]. The comparison is based on the relative Mean Squared Error (rMSE) between the recovered ODF and the ground-truth, computed using eq. (7) over a SH fitting ($l \leq 8$) of 214 noise-free $\mathbf{q}$-space samples, see [7, Table 2] for details.

Discussion. Numerical results are shown in Fig. 2, using the optimal parameters in each case. It is worth to stress the following remarks:

1) With ℓ_1 regularization, our approach (SW) provides the lowest reconstruction errors (see Fig. 2.a, top). Though this numeric difference may seem subtle, the ODFs reconstructed in Fig. 2.c suggest it might be critical to correctly resolve the fiber crossings in poor SNR scenarios.

2) As the number of gradients acquired increases, the rMSE with SW and SR rapidly decreases, but not with SH. Since the ODFs are not actually sparse in the SH basis (see Fig. 2.b), the *a priori* sparse behavior does not hold.

3) From Fig. 2.a, bottom, our SW approach provides very similar rMSE (and visually identical ODFs, see Fig. 2.c) with either ℓ_1 or ℓ_2 regularization. As stated before, non-negative ℓ_2 yield sparse solutions [11], close to those with ℓ_1 in the space of coefficients. We show in Appendix A that SW, as used here,

[1] We cannot reproduce the respective derivations due to space limits, but they are straightforward by updating eq. (8) using that SH are eigenfunctions of $\mathcal{G}$, Δ_b.

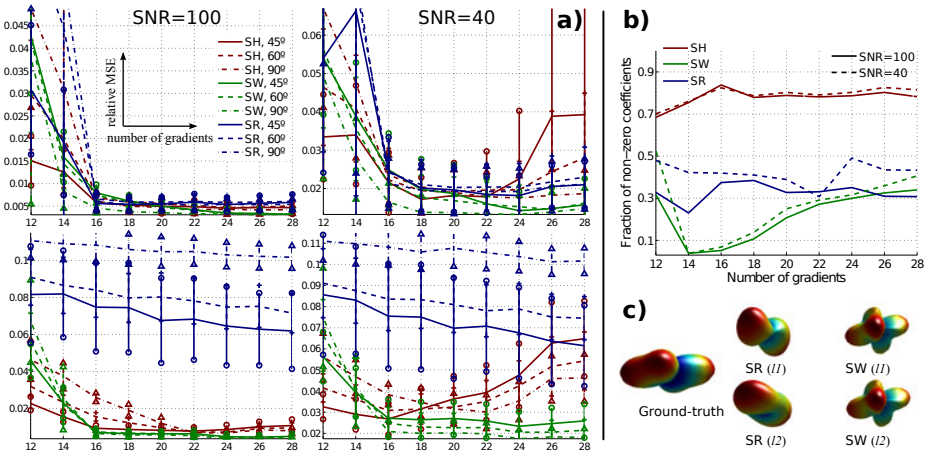

Fig. 2. Results obtained with our approach (SW), our adaptation of [8] to estimate ODFs (SR), and an implementation of our method based on SH. **a)** Mean and variance of the rMSE for ℓ_1 (top row) and ℓ_2 (bottom row). **b)** Sparsity of the solutions for a 60^o crossing. **c)** Examples of the ODFs reconstructed for SNR=40 and a 60^o crossing.

nearly meet the Parseval property: if the ℓ_1 and ℓ_2 solutions are alike, the Euclidean distance between the corresponding ODFs will also be small. The SR in [8] do not exhibit this feature. Moreover, according to eq (7), a differential operator Δ_b is applied to compute Φ from $\zeta(E)$, so the difference in the reconstruction of $\zeta(E)$ (which **is not** bounded by the distance between the ℓ_1 and ℓ_2 solutions) is amplified when computing the ODF. Accordingly, the corresponding errors in Fig. 2.a, bottom (ℓ_2) for [8] are far larger compared to the ℓ_1 case.

4) The minimum number of gradients to attain a reliable ODF reconstruction seems to be 14-16 (Fig. 2.a, top and bottom), and below this threshold the rMSE blows. With SW, this range coincides with the minimum number of non-zero coefficients of the solution (see remark 5), so we may hypothesize that a minimum of 3% of the SW coefficients are required to properly describe ODFs.

5) Our method, SW, provides also the sparsest solution (Fig. 2.b). As the number of measurements grows, we have more information on the ODF, so the solution may be refined with more SW coefficients and it becomes less sparse.

***In vivo* Experiments.** A HARDI data set with 8 unweighted T2 baseline images and 16 evenly spaced gradients ($b = 3,000$ s/mm^2) is used in this section.

Fig. 3 shows an axial slice intersecting the corpus callosum (red lobes aligned with the 'x' axis) and the uppermost part of the cingulum (green lobes, 'y' axis).

We are able to correctly resolve the crossings between these structures in the top row, or even the complex three-compartment crossings in the bottom row (blue lobes along the 'z' axis correspond to the corona radiata).

While for ℓ_1 the improvement with our method over [8] is not dramatic, the reconstructed ODFs are still sharper (see blue square), according to remark 1.

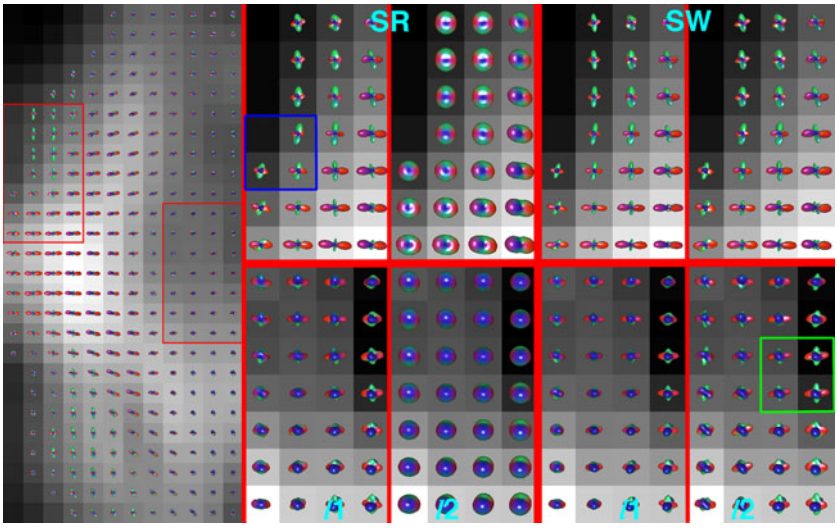

Fig. 3. Axial slice of a real data set with 16 gradients. The ODFs field has been recovered with our method. The highlighted regions have been further analyzed with both our method (SW) and the method in [8] (SR), with either ℓ_1 or ℓ_2 regularization.

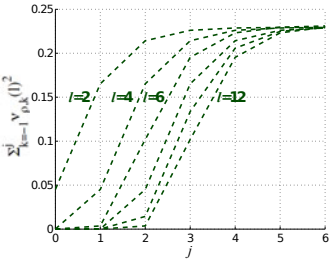

Fig. 4. Partial sums $\sum_{k=-l}^{j} \nu_{\rho,k}(l)^2$ for different degrees of l (for $l = 0$, the curve is a line with constant value 1). We use $\rho = 0.75$ as in all the experiments throughout the results section (this behavior does not hold for all ρ).

With ℓ_2 regularization, the method in [8] completely fails, while our method outputs accurate ODFs, and even the three-compartment crossings are better resolved than with ℓ_1 (see green square).

This is in agreement with remark 3. Our final remark is that:

6) ℓ_1-based reconstruction takes ∼200 s, while ℓ_2-based takes ∼1.5 s, for Fig. 3 (the speedup is ∼130), which is an additional argument supporting ℓ_2.

5 Conclusion

We have shown that as few as 14-16 diffusion directions suffice to obtain accurate probabilistic ODFs estimates. Positivity constraints, which are not considered in [7], can be easily incorporated into our model. Although sparse regularization has been previously used in [8] to fit the diffusion signal, our approach is, to our knowledge, the very first to use this technique to reconstruct the ODF

from HARDI data with reduced measurements. Though we have discussed in the results section a work-around to compute the ODF from the diffusion signal provided by [8], the results with our own approach are clearly better with ℓ_1, and the only usable at all with ℓ_2. This is specially important, since ℓ_2 reconstruction is two orders of magnitude faster than ℓ_1: for this latter approach, the overwhelming duration of the estimation for an entire volume (in the order of 100 days with our Matlab implementation) makes the overall acquisition/reconstruction time unattractive. Our ℓ_2 approach provides virtually identical results as ℓ_1 with reasonable reconstruction times, while allowing an acquisition speedup in the order of 4.

Acknowledgments. Work funded by grant numbers: R01 MH074794, R01 R01MH092862, P41 RR013218 (NIH); TEC2010- 17982, FMECD-2010/ 71131616E (Ministry of Education, Spain/Fulbright Committee).

A Parseval Property in the Semi-Discrete SW Frame

A frame is said to be tight or Parseval if the projections of a vector f in each frame element conserve the energy of f. If C is the frame bound, c_l^m the coefficient of $f(\mathbf{u})$ for the SH basis function $Y_l^m(\mathbf{u})$, and $f_l(\mathbf{u})$ the orthogonal projection of $f(\mathbf{u})$ in the subspace of SH of degree l, this condition reads:

$$\sum_{j=-1}^{\infty} \iint_{\Omega} \left|\langle f(\mathbf{u}), \Psi_{\mathbf{v},j}(\mathbf{u})\rangle\right|^2 d\mathbf{v} \stackrel{?}{=} C \left\|f(\mathbf{u})\right\|_2^2 \tag{13}$$

for the semi-discrete frame of SW. From its definition in eq. (6), eq. (13) reads:

$$\begin{aligned}
&\sum_{j=-1}^{\infty} \iint_{\Omega} d\mathbf{v} \left| \left\langle f(\mathbf{u}), \sum_{l=0,2,\ldots}^{\infty} \nu_{\rho,j}(l) \sum_{m=-l}^{l} Y_l^m(\mathbf{u}) Y_l^m(\mathbf{v}) \right\rangle \right|^2 \\
&= \sum_{j=-1}^{\infty} \iint_{\Omega} d\mathbf{v} \left| \sum_{l=0,2,\ldots}^{\infty} \nu_{\rho,j}(l) \sum_{m=-l}^{l} c_l^m Y_l^m(\mathbf{v}) \right|^2 = \sum_{j=-1}^{\infty} \sum_{l,l'=0,2,\ldots}^{\infty} \nu_{\rho,j}(l) \nu_{\rho,j}(l') \\
&\sum_{m,m'=-l}^{l} c_l^m c_{l'}^{m'} \left\langle Y_l^m, Y_{l'}^{m'} \right\rangle = \sum_{j=-1}^{\infty} \sum_{l=0,2,\ldots}^{\infty} \nu_{\rho,j}(l)^2 \sum_{m=-l}^{l} (c_l^m)^2 = \sum_{l=0,2,\ldots}^{\infty} s_l \|f_l\|_2^2 .
\end{aligned} \tag{14}$$

The condition (13) is equivalent to $s_l = \sum_{j=-1}^{\infty} \nu_{\rho,j}(l)^2$ (see eq. (8)) being constant for all l. Fig. 4 shows the partial sums for the first few l. SW are *near* Parseval (i.e. the curves approach the same value when $j \to \infty$) if $l = 0$ is ignored, but this is the DC component, unnecessary for our derivations.

References

1. Callaghan, P.: Principles of Nuclear Magnetic Resonance Microscopy. Clarendon Press, Oxford (1991)

2. Wedeen, V.J., Hagmann, P., Tseng, W.-Y.I., Reese, T.G., Weisskoff, R.M.: Mapping complex tissue architecture with DSI. Mag. Res. Med. 54, 1377–1386 (2005)
3. Merlet, S., Deriche, R.: Compressed sensing for accelerated EAP recovery in diffusion MRI. In: MICCAI Workshop Comp. Diffusion MRI, Beijing (China) (September 2010)
4. Donoho, D.L.: Compressed sensing. IEEE Trans. Info. Th. 52(4), 1289–1306 (2006)
5. Tuch, D.S.: Q–Ball imaging. Mag. Res. Med. 52, 1358–1372 (2004)
6. Assemlal, H.-E., Tschumperlé, D., Brun, L., Siddiqui, K.: Recent advances in diffusion MRI modeling: Angular and radial reconstruction. Med. Im. Anal. (2011), doi:10.1016/j.media.2011.02.002
7. Tristán-Vega, A., Westin, C.-F., Aja-Fernández, S.: A new methodology for the estimation of fiber populations in the white matter of the brain with the Funk-Radon transform. NeuroIm. 49, 1301–1315 (2010)
8. Michailovich, O., Rathi, Y.: Fast and accurate reconstruction of HARDI data using compressed sensing. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. 6361, pp. 607–614. Springer, Heidelberg (2010)
9. Jian, B., Vemuri, B.C.: Multi-fiber reconstruction from diffusion MRI using mixture of wisharts and sparse deconvolution. In: Karssemeijer, N., Lelieveldt, B. (eds.) IPMI 2007. LNCS, vol. 4584, pp. 384–395. Springer, Heidelberg (2007)
10. Freeden, W., Schreiner, M.: Orthogonal and non-orthogonal multirresolution analysis, scale discrete and exact fully discrete wavelet transform on the sphere. Const. Approx. 14, 493–515 (1998)
11. Lawson, C.L., Hanson, R.J.: Solving Least Squares Problems. Prentice-Hall, Englewood Cliffs (1974)

Sheet-Like White Matter Fiber Tracts: Representation, Clustering, and Quantitative Analysis

Mahnaz Maddah[1], James V. Miller[2],
Edith V. Sullivan[1,3], Adolf Pfefferbaum[1,3], and Torsten Rohlfing[1]

[1] Neuroscience Program, SRI International, Menlo Park, CA 94025, USA
[2] Interventional and Therapy, GE Global Research, Niskayuna, NY 12309, USA
[3] Dept. of Psychiatry and Behavioral Sciences, Stanford University, Stanford, CA 94305, USA

Abstract. We introduce an automated and probabilistic method for subject-specific segmentation of sheet-like fiber tracts. In addition to clustering of trajectories into anatomically meaningful bundles, the method provides statistics of diffusion measures by establishing point correspondences on the estimated medial representation of each bundle. We also introduce a new approach for medial surface generation of sheet-like fiber bundles in order too initialize the proposed clustering algorithm. Applying the new method to a population study of brain aging on 24 subjects demonstrates the capabilities and strengths of the algorithm in identifying and visualizing spatial patterns of group differences.

1 Introduction

Analysis of diffusion-weighted MR imaging enables *in vivo* study of human brain integrity to assess the neurodegeneration and de-myelination of white matter fiber tracts. Developing computational tools to extract quantitative information from diffusion MRI is, therefore, of great interest to the clinical community. Early methods for quantitative DTI analysis were based on the analysis of scalar diffusion measures, either within a region of interest (ROI) or at each acquired voxel. More recent variants incorporate tractography models to increase specificity. In [4], ROIs are constructed based on identified tracts, whereas *TBSS* [11] aligns subjects based on tract skeletons onto which scalar measures are then projected. However, these methods still suffer from some limitations. The former does not preserve the local variations and the latter ignores tract orientation and thus cannot always distinguish between adjacent tracts.

To overcome these limitations, recent tract-based quantitative methods analyze diffusion measures for a group of trajectories that belong to the same fiber tract and report the statistics along a descriptive model (e.g., tract skeleton) [7, 9, 15]. These methods have two main components: clustering and model construction. Within a single subject, clustering of *fiber trajectories* into groups that correspond to macroscopic *fiber tracts* (*bundles*), greatly improves the quality of tract-based analysis as it eliminates outlier trajectories [7, 13]. Across subjects, clustering ensures that the measurements are performed on the same tract in all subjects. By building a representative geometric model for each bundle (e.g., a medial curve), these methods provide a reference system for quantitative analysis of diffusion measures along the clustered trajectories. Statistics of these measures projected onto the medial model typically have lower variance

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 191–199, 2011.

across subjects and result in higher statistical power compared with voxel-wise analysis. On the other hand, unlike ROI-based methods, tract-oriented analysis allows studying of local variation along the bundle, as dimensionality reduction is constrained to the directions perpendicular to the model.

Most tract-based methods proposed so far are limited to tubular-shaped tracts, such as the cingulum bundle (CB), which can be represented by medial curves [1, 7, 9]. Although this one-dimensional model has been successfully used for population studies [8, 9], important tracts such as the body of the corpus callosum (CC) and the rostral part of corticospinal tract (CST) are not tubular and two-dimensional models are more appropriate for these sheet-like bundles. A medial representation for sheet-like fiber bundles has previously been proposed in [15] but that method used manual clustering of fiber trajectories in the group-averaged tensor data. Our method, by contrast, incorporates automatic clustering of the trajectories extracted in the subject space.

In this paper, we address problems of probabilistic clustering, geometric modeling, and quantitative analysis of sheet-like fiber bundles. Clustering is achieved by solving a mixture model on the distances between fiber trajectories and bundle medial surfaces. A novel composite distance measure incorporates spatial distance and orientation difference between trajectories. To our knowledge, this paper is the first to address the probabilistic clustering of sheet-like bundles. We also present a novel method for generating the medial surface models of such bundles. Statistics of diffusion measures for a bundle of interest are achieved by establishing point correspondences between trajectories and medial surface model. We demonstrate the proposed method by identifying spatial patterns of group differences between 12 young and 12 elderly subjects in a brain aging study for the right CST (rCST) and CC as examples of sheet-like tracts.

2 Method

Our method has three main steps: initial medial surface generation, clustering with surface evolution, and quantitative analysis, as detailed in the next subsections.

2.1 Orientational Medial Surface Generation

For each sheet-like fiber tract, a medial surface model is constructed as a triangular mesh. For each vertex in the mesh, we store the orientation of fiber trajectories at that location. Each vertex is thus described by a tuple $(\boldsymbol{\mu}, \boldsymbol{\epsilon})$, where $\boldsymbol{\mu} \in \mathbb{R}^3$ is its location and $\boldsymbol{\epsilon} \in \mathbb{R}^3$ is the local fiber orientation. This orientational medial surface is generated from a binary tract segmentation in an atlas or in a reference subject and serves as the initial tract model, which is then evolved by the clustering algorithm.

Yushkevich *et al.* [14] used a Voronoi skeleton as an initial parametric medial model, which was then fitted to the binary segmentation of the structure using a deformable model. Here, we leverage the fact that sheet-like fiber tracts are thin-walled shapes and propose an alternative approach. The method is based on the chordal axis transform (CAT) [10], wherein the chordal (i.e., medial) surface of a thin-wall-ed shell structure is estimated by connecting the mid-planes of its tetrahedral mesh elements [5].

Figure 1 illustrates the steps to estimate the medial surface of CC. Starting with the binary segmentation (Fig. 1a), which can be supplied by an atlas of fiber tracts [12],

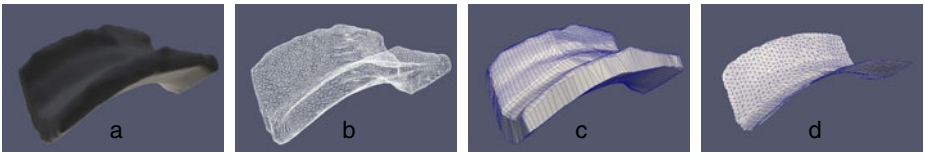

Fig. 1. Steps to estimate the medial surface of a fiber tract. *(a)* Binary segmentation of the tract from an atlas [12], *(b)* surface mesh generation using marching cubes algorithm, *(c)* single-layer volumetric hexahedral mesh generation, and *(d)* the CAT-based medial surface estimation.

we first extract the surface mesh (Fig. 1b) using the marching cubes algorithm [6]. We then construct a single-layer volumetric hexahedral mesh (Fig. 1c), using the publicly available `IA-FEMesh` [3]. From the single-layer hexahedral mesh, we generate the chordal surface mesh by connecting the patches of each hexahedron mid-plane. This requires the cut direction to be determined. For each hexahedron, we first identify the two facing facets that are not shared with other hexahedrons. The cut is the mid plane between these two facets. This leads to a quadrilateral surface mesh, represented as a `QuadEdgeMesh` structure [2], which is then triangulated (Fig. 1d) to facilitate the use of InsightToolkit (ITK) libraries. Optionally, mesh smoothing and decimation are performed to increase the quality of the chordal mesh. Finally, we add orientation information at each point on the surface by storing the principal eigenvector of the diffusion tensor at that point, computed from the atlas or reference subject tensor volume.

2.2 Trajectory-Surface Distance Measure

Each trajectory (i.e., every streamline produced by tractography) is treated as a uniformly-sampled 3-D curve, which is mapped from the subject space into the template space, for example via an affine transformation computed by image registration. Each sample on a trajectory is represented by its coordinate, $\mathbf{r} \in \mathbb{R}^3$, and its unit-length local orientation $\mathbf{e} \in \mathbb{R}^3$. A trajectory is a collection of these samples, i.e., $(\mathbf{r}_i, \mathbf{e}_i) = [(\mathbf{r}_{ij}, \mathbf{e}_{ij})]_{j=1,\ldots,J(i)}$ where $J(i)$ is the number of samples on the ith trajectory. We calculate the trajectory's local orientation at each point from its 3-D representation. Alternatively, the principal diffusion eigenvector can be used as $\mathbf{e}_{ij}$, but explicit vector re-orientation is then needed when the trajectories are mapped into the reference space.

Each cluster (i.e., each fiber bundle) is represented by a medial surface, initially generated as described in the previous subsection and evolved throughout the algorithm. For each trajectory, distances to all cluster medial surfaces are calculated. We define the distance measure between the ith trajectory and the kth cluster medial surface representation as a combination of Euclidean distance and orientation dissimilarity. The Euclidean distance can be calculated efficiently by constructing a distance map from each cluster medial surface. For each trajectory-surface pair, the spatial distance measure is $d_E(\mathbf{r}_i, \boldsymbol{\mu}_k) = \frac{1}{J(i)} \sum_j \mathcal{D}_k(\mathbf{r}_{ij})$, where $\mathcal{D}_k(\mathbf{x})$ is the value of the Euclidean distance map for cluster k at point $\mathbf{x}$.

Orientation dissimilarity is calculated from the angle between the local orientation of the trajectory points and the orientation stored for their corresponding points on the medial surface. The point correspondences are obtained by generating a Voronoi diagram,

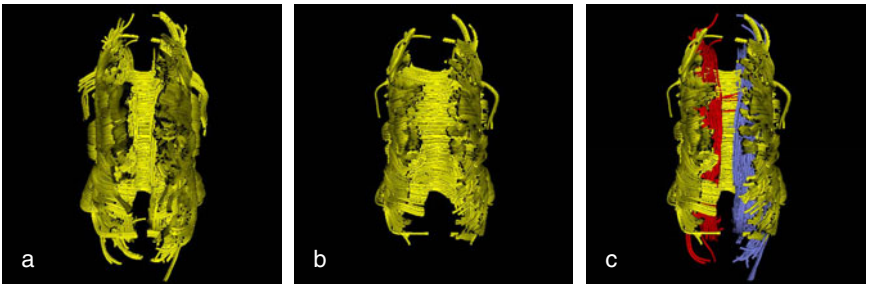

Fig. 2. The importance of orientation information for successful clustering. Clustering of the trajectories seeded from the ROI in Fig. 1a: *(a)* When orientation information is not used, some CB trajectories are falsely assigned to CC. *(b)* These trajectories are removed as outliers when orientation information is used. *(c)* Orientation information enables successful clustering of CB trajectories that are spatially close to CC but differ in shape.

$\mathcal{L}_k(\mathbf{x})$, for each cluster medial representation, which provides the index of the closest sample on the kth cluster medial surface for a given point $\mathbf{x}$ in the space. In an ideal point correspondence, the local orientation of a trajectory point and its closest point on the medial surface should match. The orientation dissimilarity between trajectory i and cluster medial representation k is defined as

$$d_O(\mathbf{e}_i, \boldsymbol{\epsilon}_k) = \frac{1}{J(i)} \sum_j (1 - u(|\langle \mathbf{e}_{ij}, \boldsymbol{\epsilon}_{k,\mathcal{L}_k(\mathbf{r}_{ij})} \rangle| - \tau)), \tag{1}$$

where $\boldsymbol{\epsilon}_{k,\mathcal{L}_k(\mathbf{r}_{ij})}$ is the orientation at the corresponding sample on the kth cluster medial representation to the point $\mathbf{r}_{ij}$, $u(\cdot)$ is the unit step function, and τ is the threshold of acceptable misorientation, i.e., a user-defined minimum bound on the cosine of the angle between each closest-point pair. We define the combined dissimilarity measure, d_{ik}, between trajectory i and cluster k as

$$d_{ik} = d_E(\mathbf{r}_i, \boldsymbol{\mu}_k) + \lambda_{ik} d_O(\mathbf{e}_i, \boldsymbol{\epsilon}_k), \tag{2}$$

where λ_{ik} is a weight factor to correct for the scale difference between the orientation dissimilarity and the distance. Here, we simply set $\lambda_{ik} = d_E(\mathbf{r}_i, \boldsymbol{\mu}_k)$.

2.3 Probabilistic Clustering

The trajectories in a subject can be clustered based on their calculated distances to each cluster medial representation. To this end, we follow the gamma mixture model approach proposed in [7][1]. Assuming that each d_{ik} is drawn from a gamma distribution, we estimate the unknown parametersand the expected value of the hidden data, which indicates the membership probabilities. Specifically, p_{ik} denotes the membership probability of trajectory i to cluster k.

[1] http://www.nitrc.org/projects/quantitativedti/

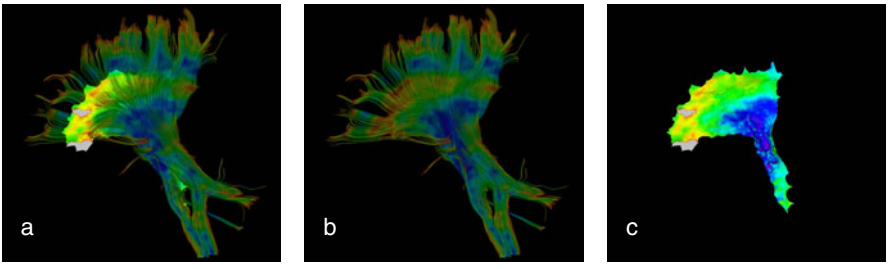

Fig. 3. Projection of scalar measures onto the cluster medial surface enables reduction in the dimensionality without loss of spatial information. Shown here is FA distribution in the rCST of a single subject. The clustered trajectories and the medial surface are shown together in *(a)* to aid the visual comparison, and separately in *(b)* and *(c)* to reveal more details. Note that we used the corticospinal ROI defined in [12] which specifies a region substantially smaller than what is spanned by the trajectories. Portions of the medial surface that do not have any contribution from the trajectories are shown in gray.

An Expectation Maximization formulation has been described in detail in [7] for both maximum-likelihood and maximum-posterior estimation of the mixture-model parameters given the observed data (i.e., d_{ik}). The clustering algorithm is an iterative process that alternates between parameter estimation and medial surface evolution. At each iteration, once the EM part has converged, the cluster medial surfaces are updated based on the new membership probabilities assigned to the trajectories by the EM formulation. Each point on the medial surface is updated as the weighted average of the corresponding points on the trajectories that belong to that cluster. Note that due to partial overlap and fiber tractography errors, it is possible to have trajectories that belong to a given cluster but contain portions that do not resemble the shape of the tract. To exclude these portions, we use a threshold on the mis-orientation between the trajectory points and the corresponding points on the medial surface, i.e.,

$$\boldsymbol{\mu}_{kj} = \frac{\sum_i p_{ik} w_{ikj} \mathbf{r}_{i,n_{ik}(j)}}{\sum_i p_{ik} w_{ikj}}, \tag{3}$$

where $w_{ikj} = 1 - u(\langle \mathbf{e}_{i,n_{ik}(j)}, \boldsymbol{\epsilon}_{kj} \rangle - \tau)$. Here, $n_{ik}(\cdot)$ is the reverse lookup function on $\mathcal{L}_k(\mathbf{x})$, which returns the closest corresponding sample on the ith trajectory to a given index on the medial representation of kth cluster, and τ is the same mis-orientation threshold as in Eq. (1). After each vertex on the mesh has been updated, the mesh is regularized by a Laplacian smoothing filter, available in ITK.

The output of the clustering algorithm is the probabilistic assignment of the trajectories to each cluster and medial representations of all clusters. Outliers are identified as those trajectories that receive membership likelihood lower than a user-specified threshold from all clusters, and these are removed from further processing.

To demonstrate the importance of using orientation information in the clustering, Fig. 2 shows the clustering of CC trajectories seeded from the ROI shown in Fig. 1a and the initial medial surface shown in Fig. 1d. When orientation information is not used (Fig. 2a) some of the CB trajectories are falsely assigned to the CC cluster. These

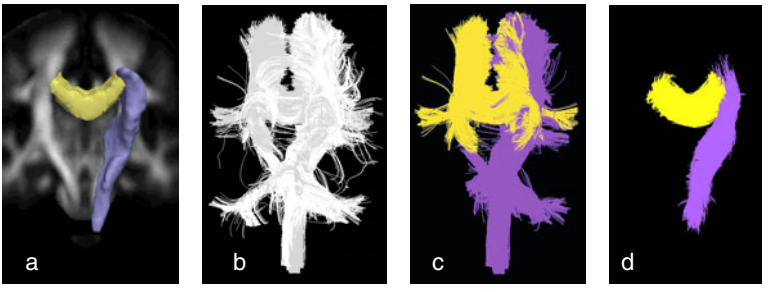

Fig. 4. Group clustering of sheet-like bundles *(a)* Body of CC and rCST as defined in the ICBM atlas, *(b)* trajectories of 24 subjects transformed to the ICBM space, *(c)* successful clustering of CC and rCST; trajectories that belong to other fiber tracts such as CB and pons have been correctly excluded. *(d)* Portions of the clustered trajectories which meet the maximum mis-orientation threshold of 30°, and therefore contribute to the final statistics on the estimated medial surfaces.

trajectories are removed as outliers when orientation information is used (Fig. 2b). In Fig. 2c clustering of the same input trajectories is performed with three initial centers, adding initial medial curves for the left and the right CBs. Note that the same EM formulation works for clustering of tubular bundles given 3D curves as the initial medial representations. Figure 3 shows the clustered trajectories of rCST in a single subject and compares visually the match between the variation of FA along the clustered trajectories with the computed average value over the estimated medial surface.

3 Experiments and Results

To demonstrate the capabilities of the proposed approach, we performed a preliminary population study to quantify how the integrity of white matter fiber tracts is affected in normal aging. Data were acquired from 12 young (age=25.5±4.34) and 12 elderly (age=77.67±4.94) healthy subjects. Echo-planar DWI data was acquired with slice thickness of 2.5 mm, fifteen unique diffusion directions with b=860 s/mm^2, along with 5 baseline scans with b=0. Images were corrected for eddy-current and B0 distortions, and tensors were estimated using the Teem library.

The labeled tracts in the ICBM atlas[2] were mapped to each subject space to seed the tractography by applying the transformation computed from pairwise affine registration on the FA volumes of the subject and the atlas using CMTK's registration tool.[3] Streamline tractography was performed using 3D Slicer, seeded from the mapped labeled regions of CC and rCST (Fig. 4a), and terminated when an FA value less than 0.15, or maximum curvature of 0.8 was reached. The quantitative parameters of interest, such as FA, were computed at each point on the trajectories and stored for subsequent analysis. Subject-specific trajectories were then back-transformed into the atlas space as shown in Fig. 4b. Given the orientational medial surfaces of CC and rCST as initial surfaces, trajectories were successfully clustered by the proposed EM algorithm (Fig. 4c) with

[2] http://www.loni.ucla.edu/ICBM/
[3] http://nitrc.org/projects/cmtk/

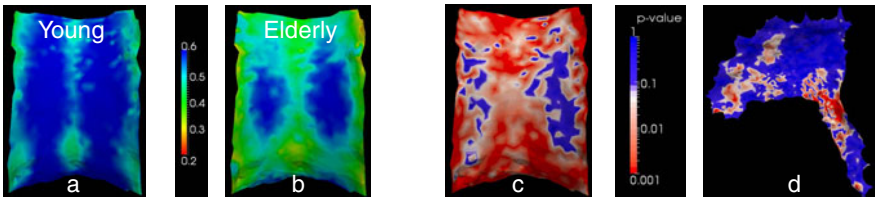

Fig. 5. Example of group-difference statistics. Each point on the medial representation of CC is colored with the weighted average FA over all corresponding points on the clustered trajectories for *(a)* young and *(b)* elderly group. The results clearly show the spatial variation of FA along the cluster and lower FA in the elderly population. The corresponding p-value for the group difference in FA is shown in *(c)* for CC and *(d)* right CST. Red corresponds to regions with high statistical significance ($p \leq 0.001$).

minimum-likelihood threshold of 0.02 for outlier rejection and the orientation threshold of $\tau = 0.85$ (allowing approximately 30° mis-orientation). Once the trajectories of all 24 subjects were clustered, for each subject diffusion measure means were calculated for each vertex on the estimated medial surface of the bundle. For subject s, let I_s be the set of trajectory indices i that originate from that subject. The weighted mean of the feature of interest for subject s at the jth point on the kth cluster medial surface was calculated as $\bar{f}^s_{kj} = \sum_{i \in I_s} p_{ik} w_{ikj} f_{i,n_{ik}(j)} \big/ \sum_{i \in I_s} p_{ik} w_{ikj}$, where $f_{i,n_{ik}(j)}$ is the feature sample at the closest corresponding point on trajectory i. At each point j on every cluster k we performed group comparisons using the per-subject feature values $\bar{f}^s_{kj}$. For each vertex on the medial surface, we performed a two-sample Welch's t-test, assuming unequal variances, to calculate the statistical significance of the group differences. Since we perform tractography in subject image space, as opposed to group-average tractography (e.g., [15]), our method can identify the regions in which a given subject does not contribute to the statistics as shown in Fig. 3 (regions in gray), adding to the reliability of the quantitative analysis. The proposed framework also enables the user to control the extent of coherence in the bundle of interest through probabilistic label assignments from the clustering. Moreover, in the quantitative analysis step, the user has control over inclusion of contributing points in the final statistics by adjusting a threshold of acceptable local orientation difference between the corresponding points. Threshold adjustment capability is important for reliable statistical analysis, because definition of tracts can be indefinite and subjective. Through visual inspection as shown in Fig. 4d, the user can be confident of the region on which the derived statistics are based. Fig. 5 illustrates the final medial representations of body of CC, colored by the average FA of 12 young subjects (Fig. 5a) and 12 elderly subjects (Fig. 5b). In these visualizations, one easily observes the local variation of the diffusion measure over the medial surface and differences between the group means. Figs. 5c and 5d demonstrate the results of our statistical analysis of FA for CC and CST. In general, we observed lower FA and higher diffusivity in the elderly group than the young group, which is consistent with findings in earlier diffusion MRI studies on aging.

4 Conclusion

This work is the extension and generalization of [7] and enables probabilistic clustering and quantitative analysis of sheet-like tracts, either in a single subject or in a population, while still supporting tubular bundles. We also propose a novel method for estimating the orientational medial surface of bundles, which then serve as initial tract models in the proposed clustering scheme. Here, we start from segmentation provided in the ICBM atlas but the method is general and could instead have been started from a manually segmented region in a reference subject. Our experiments demonstrate the strengths of the presented method in computing the spatial summary statistics of diffusion measures on the medial representations of white matter fiber tracts.

Acknowledgments. The authors would like to thank Dr. C.-F. Westin for his constructive feedback. This work was supported by AA005965, AA012388, AA017347, AA017168, EB008381, and NAC P41-RR13218.

References

1. Corouge, I., Gouttard, S., Gerig, G.: A statistical shape model of individual fiber tracts extracted from diffusion tensor MRI. In: Barillot, C., Haynor, D.R., Hellier, P. (eds.) MICCAI 2004. LNCS, vol. 3217, pp. 671–679. Springer, Heidelberg (2004)
2. Gouaillard, A., Florez-Valencia, L., Boix, E.: ItkQuadEdgeMesh: A discrete orientable 2-manifold data structure for image processing. The Insight Journal (2006)
3. Grosland, N.M., Shivanna, K.H., Magnotta, V.A., Kallemeyn, N.A., DeVries, N.A., Tadepalli, S.C., Lisle, C.: IA-FEMesh: An open-source, interactive, multiblock approach to anatomic finite element model development. Computer Methods and Programs in Biomedicine 94(1), 96 (2009)
4. Kanaan, R.A., Shergill, S.S., Barker, G.J., Catani, M., Ng, V.W., Howard, R., McGuire, P.K., Jones, D.K.: Tract-specific anisotropy measurements in diffusion tensor imaging. Psychiatry Research 146(1), 73–82 (2006)
5. Kwon, K.Y., Lee, B.C., Chae, S.W.: Medial surface generation using chordal axis transformation in shell structures. Computers and Structures 84(26-27), 1673 (2006)
6. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3D surface construction algorithm. In: SIGGRAPH, pp. 163–169 (1987)
7. Maddah, M., Grimson, W.E.L., Warfield, S.K., Wells, W.M.: A unified framework for clustering and quantitative analysis of white matter fiber tracts. Med. Image Anal. 12, 191–202 (2008)
8. Maddah, M., Kubicki, M., Wells, W.M., Westin, C.-F., Shenton, M.E., Grimson, W.E.L.: Findings in schizophrenia by tract-oriented DT-MRI analysis. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part I. LNCS, vol. 5241, pp. 917–924. Springer, Heidelberg (2008)
9. O'Donnell, L.J., Westin, C.-F.: Automatic tractography segmentation using a high-dimensional white matter atlas. IEEE Trans. Med. Imag. 26, 1562–1575 (2007)
10. Prasad, L.: Morphological analysis of shapes. CNLS Newsletter 139, 1–18 (1997)
11. Smith, S.M., Jenkinson, M., Johansen-Berg, H., Rueckert, D., Nichols, T.E., Mackay, C.E., Watkins, K.E., Ciccarelli, O., Cader, M.Z., Matthews, P.M., Behrens, T.E.J.: Tract-based spatial statistics: voxelwise analysis of multi-subject diffusion data. NeuroImage 31, 1487–1505 (2006)

12. Wakana, S., Jiang, H., Nagae-Poetscher, L.M., van Zijl, P.C.M., Mori, S.: Fiber tract-based atlas of human white matter anatomy. Radiology 230, 77–87 (2004)
13. Wassermann, D., Bloy, L., Kanterakis, E., Verma, R., Deriche, R.: Unsupervised white matter fiber clustering and tract probability map generation: applications of a gaussian process framework for white matter fibers. NeuroImage 51(1), 228–241 (2010)
14. Yushkevich, P., Zhang, H., Gee, J.: Continuous medial representation for anatomical structures. IEEE Trans. Med. Imag. 25(12), 1547 (2006)
15. Zhang, H., Awate, S.P., Das, S.R., Woo, J.H., Melhem, E.R., Gee, J.C., Yushkevich, P.A.: A tract-specific framework for white matter morphometry combining macroscopic and microscopic tract features. Med. Image Anal. 14(5), 666 (2010)

Diffusion Tensor Image Registration with Combined Tract and Tensor Features

Qian Wang[1,2], Pew-Thian Yap[2], Guorong Wu[2], and Dinggang Shen[2]

[1] Department of Computer Science, University of North Carolina at Chapel Hill
qianwang@cs.unc.edu
[2] Department of Radiology and BRIC, University of North Carolina at Chapel Hill
{ptyap,grwu,dgshen}@med.unc.edu

Abstract. Registration of diffusion tensor (DT) images is indispensible, especially in white-matter studies involving a significant amount of data. This task is however faced with challenging issues such as the generally low SNR of diffusion-weighted images and the relatively high complexity of tensor representation. To improve the accuracy of DT image registration, we design an attribute vector that encapsulates both tract and tensor information to serve as a voxel morphological signature for effective correspondence matching. The attribute vector captures complementary information from both the global connectivity structure given by the fiber tracts and the local anatomical architecture given by the tensor regional descriptors. We incorporate this attribute vector into a multi-scale registration framework where the moving image is warped to the space of the fixed image under the guidance of tract information at a more global level (coarse scales), followed by alignment refinement using regional tensor distribution features at a more local level (fine scales). Experimental results indicate that this framework yields marked improvement over DT image registration using volumetric information alone.

1 Introduction

Diffusion tensor imaging (DTI) plays a vital role in the understanding of water diffusion patterns of brain tissues, providing an effective means for mapping brain structural circuitry *in vivo*. Water molecules are more likely to diffuse along directions tangential to the axons than directions orthogonal to the myelin sheaths. Harnessing this unique diffusion pattern allows tracing of the neuronal trajectories, and hence mapping of brain connectivity.

To make possible the comparison between individual DT images, proper spatial normalization of these images to a common space is often required. The task of DTI registration, however, is more challenging than scalar image registration. Tensors, unlike scalars, live in a space with higher dimensions that requires more complicated metrics for their effective quantification. Moreover, normalizing the DT images to a common space involves an additional step of tensor reorientation for preserving the consistency of local diffusivity pattern. Several DTI registration methods have been reported in the literature. Since scalar image registration (e.g., T1 weighted MR image) has been relatively well

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 200–208, 2011.

studied, a straightforward approach is to directly port state-of-the-art scalar registration algorithms for use of DT images. For example, Park *et. al.* [1] propose a multi-channel strategy for DTI registration by leveraging complementary information from several scalar images. In [2], tensor reorientation is explicitly incorporated in the registration optimization function. Yeo *et. al.* [3] utilize the Finite Strain (FS) differential for constructing a fast diffeomorphic DTI registration algorithm. For the purpose of DTI registration, various tensor similarity/dissimilarity measures have been explored [4]. Examples include mutual information [5], symmetrized Kullback-Leibler (sKL) divergence [6], and Geodesic-Loxodromes [7].

An image voxel can be represented by its *attribute vector*, which is essentially a vector grouping together a set of features giving descriptive and possibly distinctive information about the voxel. The effectiveness of attribute vector has been demonstrated in the registration of both scalar images and DT images. Yang *et. al.* [8], for instance, use prolateness, oblateness, and sphericity as the tensor attributes. In [9], a more sophisticated form of attribute vector, composing of statistical features (local mean and variance), geometrical features (edge information, and tensor principal diffusivity), is used.

Tractography, another important topic in DTI studies, streamlines fibers in the white matter (WM) and reveals neural pathways communicating individual brain regions [10]. These fibers can be grouped into bundles or tracts, each of which contains fibers with similar structural and functional characteristics. Bundles are more reproducible and consistent across subjects [11], hence suggesting their usages to guide registration of DT images. An early approach can be found in [12], where a carefully designed linear operator is iterated over the FA maps to detect bundles occurring as tubular or sheet-like structures, generating feature maps that can be used for registration using a scalar image registration method. In [13], correspondences between images are identified via tractography clustering using a polyaffine transformation to describe the deformation field.

In this paper, we design a multi-scale DTI registration framework that leverages attribute vectors for capturing both tract and tensor information. Registration is formulated as minimizing the overall difference between the attribute vectors of a pair of images. We further decompose the registration framework into two consecutive stages, each of which utilizes a different scale of features. In the first (coarse) stage, the moving image will be deformed to the fixed image under the guidance of tract derived features. Then the deformation field is transferred to the next (fine) stage and refined by using tensor related features. Details of the algorithm are provided in Section 2. Experimental results in Section 3 show that the proposed method yields improved performance in terms of registration accuracy.

2 Method

In the following, we will firstly brief the complete registration framework for DT images proposed in this paper (Section 2.1). Then, we will give details on the two registration stages in Section 2.2 and Section 2.3, respectively.

2.1 Algorithm Framework

In tractography, fibers are streamlined by sequentially tracing the principal diffusion directions of neighboring tensors. The generated fibers can be grouped into bundles according to their topologies and anatomical properties. To register a pair of DT images, two criteria are enforced: (1) tensors with similar local features and (2) fiber tracts with identical anatomical properties should both be matched across the images. We denote the attribute vector of voxel x as $\vec{a}(x)=[\vec{a}_B(x),\vec{a}_T(x)]$, where $\vec{a}_B(x)$ and $\vec{a}_T(x)$ incorporate features derived from the tracts and tensors, respectively. The elements in $\vec{a}_B(x)$ describe the relationship between a specific bundle and voxel x. In particular, each element of $\vec{a}_B(x)$ corresponds to a bundle and can be set if the voxel is incorporated by the bundle; otherwise it is left unset. More details about $\vec{a}_B(x)$ are provided in Section 2.2. Elements of $\vec{a}_T(x)$ represent the tensor features computed from x and its local neighborhood [9]. The registration of a fixed image F and a moving image M is cast as a problem of minimizing the overall distances, $D_B(F, M \circ s)$ for tract derived features and $D_T(F, M \circ s)$ for the tensor derived features, with a certain smooth constraint imposed on the deformation field s.

Multi-scale strategy is commonly used in image registration, owing to its effectiveness in avoiding local minima that often result in suboptimal solutions in high-dimensional optimization problems. Different anatomical structures are best represented at their appropriate scales. For instance, fiber bundles reflect brain connectivity and give structural information at a more global scale than an individual voxel or its adjacent neighborhood. Based on this observation, we minimize the two distances $D_B(F, M \circ s)$ and $D_T(F, M \circ s)$ separately at different scales, decomposing the registration problem into two consecutive stages – the tract-guided stage and the tensor-guided stage. The tract-guided stage provides a coarse estimation of deformation field for further refinement in the tensor-guided stage. A flow chart of the proposed algorithm is shown in Fig. 1. More details on the tract-guided stage and tensor-guided stage are given in Section 2.2 and Section 2.3, respectively.

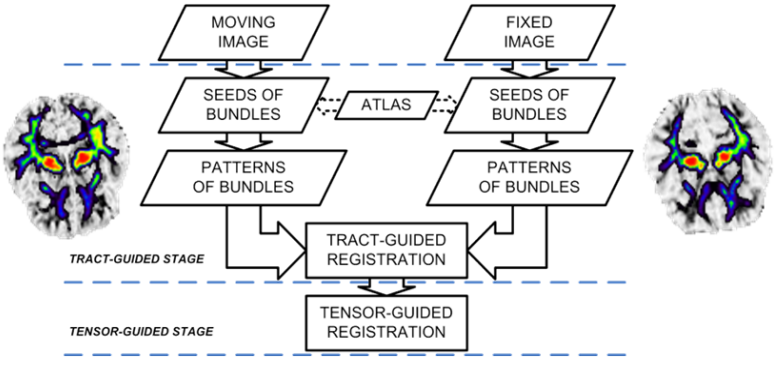

Fig. 1. The proposed DTI registration method consists of two consecutive stages: (1) tract-guided registration, and (2) tensor-guided registration

2.2 Tract-Guided Registration

For tract-guided registration of a pair of DT images, tractography is first applied to the images independently to yield two separate sets of fibers. After bundles are detected, the goal of registration is then to align the corresponding bundles of different images. As in Ziyan *et. al.* [13], bundle correspondence resulting from tractography clustering is utilized to guide DTI registration. However, spatial variation of fibers can be large, due to (1) the intrinsic complexity of brain anatomy; (2) the limited SNR in DTI scans; and (3) additional errors introduced by tractography algorithms. Especially in the case of DTI where fiber crossings cannot be modeled correctly, many fibers might be prematurely terminated. Correspondence between bundles in these situations is hence unreliable.

We take a different approach and use representations of a subset of easily identifiable tracts to help guide registration. Specifically, we have selected the corpus callosum (CC), the left/right thalamus radiata (TR), and the left/right corticospinal tracts (CST) as the driving bundles, since they traverse a large portion of the brain and represent the major WM pathways. The high reproducibility of these bundles [11] allows us to estimate their correspondences between images more reliably.

To identify the bundles, we warp the JHU WMPM [14] atlas (referred to as the *"atlas"* for simplicity) to the spaces of the fixed and the moving DT images by registration of the FA maps. ROIs corresponding to the selected bundles are used to determine the fiber bundles. Specifically, within the ROI for each bundle, we identify the core of the bundle by locating the voxels with fiber traversing density greater than a specified threshold. The qualified voxels form the set of *bundle seeds*, which we will use to generate the *bundle pattern* by allowing the seeds to *diffuse* along the underlying fibers. The *bundle diffusion* mechanism is explained with more details in the following. The elements of the attribute vector $\vec{a}_B(x)$ derived from the diffused bundle patterns are corresponding to individual bundles and used for registration.

Bundle Diffusion

Bundle diffusion is initiated at the seed voxels and proceeds along the underlying fibers. For this purpose, we have adopted the fast marching (FM) method [15]. In Fig. 2, for instance, the fibers traversing this region (in the horizontal direction) are represented by the light blue strips. The seed region is represented by the dark red area, which is circled by the green dashed curve. In the diffusion process, the interface of the seed region propagates outward to generate the pattern P after a certain amount of time, as shown in Fig. 2. For each voxel location x, $P(x)$ records the time when the voxel is traversed by the evolving interface. At each point on the interface, the diffusion always proceeds along the surface normal. The diffusion velocity $V(x)$ at location x is calculated by integrating the total contributions of fibers passing through x. The contribution of each fiber is equal to the cosine of the angle between the tangential direction of the fiber and the normal direction of the evolving interface. By iteratively solving the Eikonal equation $V\|\nabla_x P\|=1$ in FM, the propagation of the region interface can be simulated numerically and the bundle pattern P can thus be acquired. For corresponding tracts in both the fixed and the moving images, we allow the diffusion to last for an equal amount of time. We then reverse the sign of the

traversing times and scale all values to the range of [0, 1]. Hence, seed voxels have a value of 1, which gradually decreases to zero at the stopping interface of P. Two real examples of the diffused bundles patterns are also shown in Fig. 1.

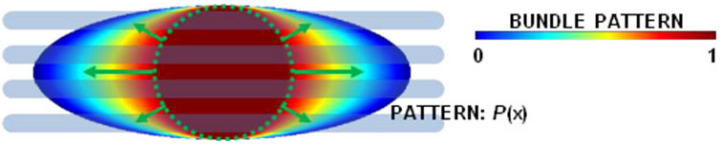

Fig. 2. Starting from the seed points within the green dashed circle, the interface propagates outward to generate the bundle pattern . Each point in the space is colored according to the interface traversing time (sign-inverted and linearly scaled to [0, 1]).

Bundle-Derived Attribute Vector and Registration

To register the fixed image F and the moving image M, we attempt to match the attribute vector $\vec{a}_B(x)$ of both images. We define that the i-th entry of the attribute vector $\vec{a}_B(x)$ records the value from the diffused pattern of the i-th bundle. Thus the length of the attribute vector $\vec{a}_B(x)$ is equal to the number of bundles used to guide registration. The term $D_B(F, M \circ s)$ captures the overall Euclidean distance between tract-derived attribute vectors of both images, and can be minimized iteratively. In each iteration, the increment u of the deformation s is estimated as [16]:

$$u = ({}^t\mathcal{J}_x \cdot \mathcal{J}_x + \sigma I)^{-1} \left(\vec{a}_B^F(x) - \vec{a}_B^M\big(s(x)\big)\right) {}^t\mathcal{J}_x \tag{1}$$

where σ enforces the constraint that u should be infinitesimal with respect to s, and the term $\mathcal{J}_x$ (with its transpose ${}^t\mathcal{J}_x$) is:

$$\mathcal{J}_x = \frac{1}{2}\left(\nabla_x \vec{a}_B^F(x) + \nabla_x \vec{a}_B^M\big(s(x)\big)\right). \tag{2}$$

Based on (1) and (2), the deformation increment u is estimated to refine the deformation s iteratively. The deformation s is updated based on the composition rule $s \leftarrow s \circ \exp(u)$ in order to keep the generated deformation diffeomorphic. Upon estimating the deformation field, we reorient the tensors using the method described in [17]. The deformation s output by the tract-guided stage here will be transferred to the following tensor-guided stage for more refinement.

2.3 Tensor-Guided Registration

Using tract-guided registration as initialization, further refinement to the estimated deformation s is performed using tensor features. To this end, we have adopted the approach described in [9]. Specifically, tensor features, including regional features (means and variances), edge features (tensor edges and FA map edges), and geometrical features (FA values and principal diffusivities), are extracted for each voxel and incorporated into the attribute vector component $\vec{a}_T(x)$. A subset of voxels with distinctive attribute vectors is detected to drive the correspondence matching.

Thin-Plate Splines (TPS) is then used to generate a dense deformation field based on the estimated displacements of the driving voxels. Unlike [9], however, only fine-level (or the highest level) registration is employed in the tensor-guided stage, since the coarse deformation is already estimated from the tract-guided stage.

3 Experimental Results

To demonstrate the performance of the proposed method, we performed the proposed method on both real and simulated data, and compared with DT-ITK [2] as well as F-TIMER [9] in the following.

In-Vivo Dataset

Totally, 15 DT images were acquired for our validation (Siemens Allegra scanner, b=1000s/mm^2, flip angle 90°, TR/TE=13,640/82ms, matrix 128×128, FoV 256×256mm^2, slice thickness 2mm, 80 contiguous slices). One image is randomly selected to be fixed, to which other 14 moving images are normalized by affine registration. We then apply all three registration methods – DT-ITK [2], F-TIMER [9], and our method – to deform the 14 images to the fixed space. For quantitative evaluation, we use the Frobenius norm to measure the distance between a pair of tensors. Specifically, we compute the tensor distance for each corresponding voxel after the moving image has been deformed to the fixed space, where tissue segmentation is available. In Table 1, the average tensor distances, as well as the respective standard deviations, of both white and grey matter voxels are listed. For white matter, our method reduces the average tensor distance by 40% over DT-ITK and 19% over F-TIMER, as $p<0.01$ in t-tests for both. The significantly lower tensor distance values suggest that our method is capable of DTI registration at a higher accuracy.

Table 1. The average voxelwise tensor distances

	DT-ITK	F-TIMER	Our Method
White Matter	0.375±0.021	0.247±0.025	0.199±0.030
Grey Matter	0.267±0.010	0.219±0.031	0.178±0.029

Fig. 3 shows a specific moving image in the different stages of registration to the fixed space, via both the proposed method (top) and F-TIMER (bottom). All slices are extracted from the same position and show part of cingulums. After the tract-guided registration stage in our method, the left and the right structures are apart and thus more similar to the fixed image. The output of the tract-guided stage is further refined using tensor features in our method. On the contrary, the middle level registration in F-TIMER fails to provide a good initialization as indicated by the still-connected left and right cingulums.

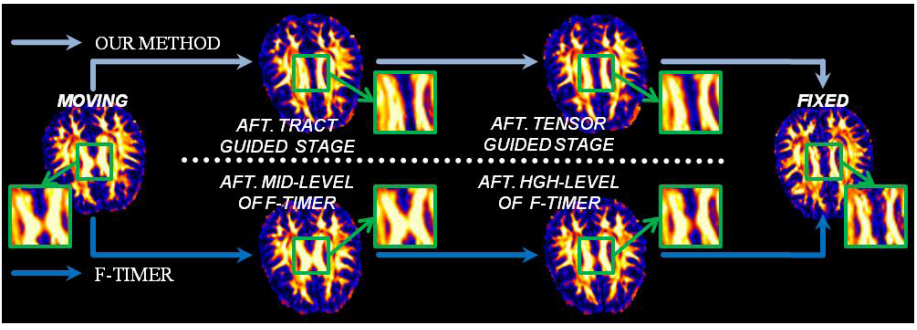

Fig. 3. The moving image is registered by both our method and F-TIMER to the fixed image. The tract-guided stage of our method yields better initialization than the middle level of F-TIMER. This can be observed from the fact that the left and the right cingulums are apart using our method.

Simulated Dataset

We further evaluate the performance of the proposed method by gauging the accuracy of the estimated deformation fields. To generate a set of realistic deformation fields, we register the FA map of a randomly selected fixed image to FA maps of the other 10 images using the Demons algorithm. The estimated deformation fields are then applied to the fixed DT image to generate 10 simulated moving images. The inverse of the deformation field can hence be used as ground truth for evaluation. Across 10 simulated images, our method achieves an average deformation error of 0.637mm, much lower than 1.929mm of DT-ITK and 0.905mm of F-TIMER. The average deformation errors for the individual bundles are shown in Fig. 4, as the improvements on all bundles achieved by our method are significant statistically compared with the other two methods. It is worth noting that the deformation model in simulating data is different with any of the three methods.

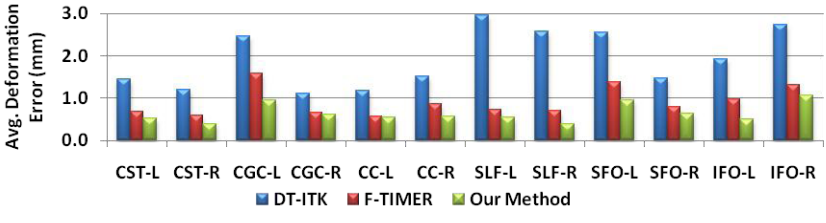

Fig. 4. The proposed method produces the lowest deformation errors in average among all three registration methods under comparison.

4 Conclusion

We have proposed a novel DTI registration framework by combining complementary information from both tracts and tensors. Since tract and tensor features represent

information of different structural scales, we decompose the registration task into two consecutive stages: (1) tract-guided registration and (2) tensor-guided registration. In the tract-guided stage, attributes extracted from bundle patterns are used to give a coarse estimation of the deformation field. This aligns the major bundles and provides a good initialization for the following tensor-guided stage, where tensor features are used to refine the registration. Evaluations with both *in-vivo* and simulated datasets indicate that the proposed method gives superior performance when compared with DTI registration using volumetric information alone.

References

1. Park, H.-J., Kubicki, M., Shenton, M.E., Guimond, A., McCarley, R.W., Maier, S.E., Kikinis, R., Jolesz, F.A., Westin, C.-F.: Spatial normalization of diffusion tensor MRI using multiple channels. NeuroImage 20, 1995–2009 (2003)
2. Zhang, H., Yushkevich, P.A., Alexander, D.C., Gee, J.C.: Deformable registration of diffusion tensor MR images with explicit orientation optimization. Medical Image Analysis 10, 764–785 (2006)
3. Yeo, B.T.T., Vercauteren, T., Fillard, P., Pennec, X., Gotland, P., Ayache, N., Clatz, O.: DTI registration with exact finite-strain differential. In: ISBI (2008)
4. Ruiz-Alzola, J., Westin, C.F., Warfield, S.K., Alberola, C., Maier, S., Kikinis, R.: Nonrigid registration of 3D tensor medical data. Medical Image Analysis 6, 143–161 (2002)
5. Van Hecke, W., Leemans, A., D'Agostino, E., De Backer, S., Vandervliet, E., Parizel, P.M., Sijbers, J.: Nonrigid Coregistration of Diffusion Tensor Images Using a Viscous Fluid Model and Mutual Information. IEEE Trans. Medical Imaging 26, 1598–1612 (2007)
6. Chiang, M.-C., Leow, A.D., Klunder, A.D., Dutton, R.A., Barysheva, M., Rose, S.E., McMahon, K.L., de Zubicaray, G.I., Toga, A.W., Thompson, P.M.: Fluid Registration of Diffusion Tensor Images Using Information Theory. IEEE Trans. Medical Imaging 27, 442–456 (2008)
7. Irfanoglu, M.O., Machiraju, R., Sammet, S., Pierpaoli, C., Knopp, M.: Automatic Deformable Diffusion Tensor Registration for Fiber Population Analysis. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part II. LNCS, vol. 5242, pp. 1014–1022. Springer, Heidelberg (2008)
8. Yang, J., Shen, D., Davatzikos, C., Verma, R.: Diffusion Tensor Image Registration Using Tensor Geometry and Orientation Features. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part II. LNCS, vol. 5242, pp. 905–913. Springer, Heidelberg (2008)
9. Yap, P.T., Wu, G., Zhu, H., Lin, W., Shen, D.: F-TIMER: Fast Tensor Image Morphing for Elastic Registration. IEEE Trans. Medical Imaging 29, 1192–1203 (2010)
10. Mori, S., Zijl, P.C.M.v.: Fiber Tracking: Principles and Strategies - A Technical Review. NMR in Biomedicine 15, 468–480 (2002)
11. Wakana, S., Caprihan, A., Panzenboeck, M.M., Fallon, J.H., Perry, M., Gollub, R.L., Hua, K., Zhang, J., Jiang, H., Dubey, P., Blitz, A., van Zijl, P., Mori, S.: Reproducibility of quantitative tractography methods applied to cerebral white matter. NeuroImage 36, 630–644 (2007)

12. Goodlett, C., Davis, B., Jean, R., Gilmore, J., Gerig, G.: Improved Correspondence for DTI Population Studies Via Unbiased Atlas Building. In: Larsen, R., Nielsen, M., Sporring, J. (eds.) MICCAI 2006. LNCS, vol. 4191, pp. 260–267. Springer, Heidelberg (2006)
13. Ziyan, U., Sabuncu, M., O'Donnell, L., Westin, C.-F.: Nonlinear Registration of Diffusion MR Images Based on Fiber Bundles. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) MICCAI 2007, Part I. LNCS, vol. 4791, pp. 351–358. Springer, Heidelberg (2007)
14. Mori, S., Oishi, K., Jiang, H., Jiang, L., Li, X., Akhter, K., Hua, K., Faria, A.V., Mahmood, A., Woods, R., Toga, A.W., Pike, G.B., Neto, P.R., Evans, A., Zhang, J., Huang, H., Miller, M.I., van Zijl, P., Mazziotta, J.: Stereotaxic white matter atlas based on diffusion tensor imaging in an ICBM template. NeuroImage 40, 570–582 (2008)
15. Sethian, J.A.: A Fast Marching Level Set Method for Monotonically Advancing Fronts. PNAS 93, 1591–1595 (1996)
16. Vercauteren, T., Pennec, X., Malis, E., Perchant, A., Ayache, N.: Insight into Efficient Image Registration Techniques and the Demons Algorithm. In: Karssemeijer, N., Lelieveldt, B. (eds.) IPMI 2007. LNCS, vol. 4584, pp. 495–506. Springer, Heidelberg (2007)
17. Xu, D., Mori, S., Shen, D., van Zijl, P.C.M., Davatzikos, C.: Spatial normalization of diffusion tensor fields. Magnetic Resonance in Medicine 50, 175–182 (2003)

Probabilistic Tractography Using Q-Ball Modeling and Particle Filtering

Julien Pontabry and François Rousseau

LSIIT, UMR 7005 CNRS-Université de Strasbourg

Abstract. By assuming that orientation information of brain white matter fibers can be inferred from Diffusion Weighted Magnetic Resonance Imaging (DWMRI) measurements, tractography algorithms provide an estimation of the brain connectivity in-vivo. The two key ingredients of tractography are the diffusion model (tensor, high-order tensor, Q-ball, etc.) and the way to deal with uncertainty during the tracking process (deterministic vs probabilistic). In this paper, we investigate the use of an analytical Q-ball model for the diffusion data within a well-formalized particle filtering framework. The proposed method is validated and compared to other tracking algorithms on the MICCAI'09 contest Fiber Cup phantom and on in-vivo brain DWMRI data.

1 Introduction

In the last decade, Magnetic Resonance Imaging (MRI) has become a popular and powerful tool for medical imaging and brain understanding. In particular, Diffusion Weighted MRI (DWMRI) is a non-invasive imaging system, which gives information on water diffusion in human brain. These indirect observations of the white matter geometry in-vivo [3] allow nerve fiber reconstruction by using tractography algorithms. The two keypoints of a tractography algorithm are the diffusion model and the mathematical framework describing the tracking process.

Water diffusion is historically modeled by tensors [3], which has been quite successful in brain region with homogeneous intra-voxel fiber orientation. However, Berhens et al. have detected complex fiber architecture in approximately a third of voxels of the brain [4]. The standard tensor model does not handle such fiber bundle geometry well. A way to deal with this issue is to use High Angular Resolution Diffusion Imaging (HARDI) with appropriate diffusion models such as Q-ball or high-order tensors (HOT).

The second keypoint of tractography concerns the tracking process modeling. Tractography methods can be mainly categorized into two classes: deterministic and probabilistic. Deterministic solutions are subdivided into two groups: streamline and optimization based. In the first one, at each voxel, next direction(s) of fiber paths is(are) locally determined by mean direction(s) indicated by the diffusion model [14,2]. Such algorithms consist in a step-by-step construction of the solution. As only local mean directions are followed, there is a possible accumulation of errors during the tracking process. Optimization-based methods treat the tractography as an energy minimization problem [16,18]. The purpose of these algorithms is to find the global solution, i.e. the solution that minimizes errors according to the diffusion model.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 209–216, 2011.

Nevertheless, noise in MRI [11], partial volume effect [1] and possible ambiguity induced by the diffusion model can lead to uncertainty in the tracking process. Since this uncertainty may be represented in the form of probability density functions (PDF), probabilistic framework represents an attractive option for brain tractography [15,13,9,4]. A way to produce probabilistic maps of fibers consists in applying a streamline propagation process with directions randomly sampled from the PDF. The fraction of these sampled paths passing through each voxel provides a stochastic measure of the strength of connectivity between that voxel and the seed points. In such framework, a keypoint is the effectiveness of the sampling stage [15,20]. In particular, as fiber paths can be modeled as Markov chains, sequential Monte Carlo and Markov Chain (MCMC) methods can be efficiently used for the sampling stage [20].

Extending previous work of Zhang et al. [20], we investigate the use of Q-ball based modeling within a particle filtering framework. The reliability of Q-Ball models allows us to extract more information from data, such as the orientation distribution function (ODF) which provides a precise idea of the underlying fiber architecture at every voxel. Then, a non-linear state-model is used for the tracking modeling and the probabilistic maps of fibers are estimated using a particle filtering algorithm.

2 Proposed Method

2.1 Diffusion Model

Let consider an image domain $\Omega \subset \mathbb{R}^3$ and diffusion weighted measurement in N directions. The diffusion signal is decomposed into spherical harmonics by linear regression [6]. Then, applying Funk-Radon transform to the signal decomposition, the diffusion ODF can be analytically computed at point $x \in \Omega$ and in direction $u = (\theta, \phi)$ [6]. In order to get a sharper ODF, the fiber orientation distribution function (fODF) $\psi_x(\theta, \phi)$ is computed from diffusion ODF by a spherical deconvolution [7]. Let the estimated diffusion signal S_x and the fODF ψ_x be the diffusion observations at $x \in \Omega$.

2.2 Fiber Tracking Model

In space state model, a fiber path $v_{0:k}$ can be modeled as a sequence of vectors in a volume Ω and can be built iteratively. Fiber paths are assumed to have Markovian nature. From a given starting point $x_0 \in \Omega$, points of path $v_{0:k}$ are defined as $x_{k+1} = x_k + \lambda v_k$, where $\lambda \in \mathbb{R}$ is the step size which is assumed to be constant and v_k is a unit vector.

The principle of particle filtering is to sample sequentially a set of M paths from the starting point $x_0 \in \Omega$, given observations $z_k = \{S_{x_k}, \psi_{x_k}\}$. This means that M particles are placed at point x_0 at step $k = 0$ and are propagated as time progresses. Given the set of particles $\{v_{0:k}^{(m)}, w_k^{(m)}\}_{m=1}^{M}$ at step k, the propagation to the next step $k+1$ is performed following three stages: prediction, weighting and selection.

In prediction stage, since the posterior $p(v_{0:k}|z_{1:k})$ cannot be evaluated, the importance density $\pi(v_{0:k}|z_{1:k})$ (which is an approximation of the posterior density) is be used to simulate each of paths' vectors. At step k and considering Markovian nature of fiber paths, the vector v_k of a path $v_{0:k-1}$ is sampled according to the recursive form of importance density [8]: $\pi(v_k|v_{0:k-1}, z_{1:k})$.

After prediction stage, a weighting stage giving a measurement of reliability of the approximation of the posterior density is performed. A particle's weight is written recursively as [8]

$$\tilde{w}_k^{(m)} \propto \tilde{w}_{k-1}^{(m)} \frac{p(z_k|v_k^{(m)})p(v_k^{(m)}|v_{k-1}^{(m)})}{\pi(v_k^{(m)}|v_{0:k-1}^{(m)}, z_{1:k})} \quad , \tag{1}$$

where $\tilde{w}_t^{(m)}$ are normalized weights, k is the current state, $p(z_k|v_k^{(m)})$ and $p(v_k^{(m)}|v_{k-1}^{(m)})$ are respectively likelihood and prior densities.

Choosing recursive importance distribution leads to an increasing variance of the weights as the time progresses [12]. The purpose of the final stage selection is to avoid this degeneracy. We first measure degeneracy of the cloud of paths using the effective sample size (ESS): $N_{\text{ESS}} = 1/\sum_{m=1}^{M} \tilde{w}_{(m)}^2$. When N_{ESS} decreases below a fixed threshold ε_{ESS}, a resampling procedure is used in order to eliminate particles with low weight [8].

2.3 Densities

The Prior Distribution. As in [20], the von Mises-Fisher (vMF) distribution has been selected as prior because it is a parametric distribution for directional data. This distribution is parametrized by μ and κ, respectively called mean direction and concentration. The greater the value of κ, the stronger is the concentration of the distribution around mean direction μ. Considering the tractography problem and the tracking model, the value of concentration parameter κ is a smoothness constraint of fiber path.

The Importance Density. The iterative solution estimation relies on the approximation of the posterior, i.e. on the importance density. At each step, knowing a path's vector sequence and its observation information, next direction added to path is sampled according to the importance density. The optimal importance density is $p(v_k|v_{k-1}, z_{1:k})$, because conditionally upon v_k and $z_{1:k}$, the variance of the importance weights w_k is then minimal [8]. Since it is difficult to sample from $p(v_k|v_{k-1}, z_{1:k})$, a usual choice is to use for importance the same distribution as prior. Nevertheless, since no observation information is used, generated particles using the prior are often outliers of the true posterior distribution. As importance function, we choose a local linearization using observations. The ODF functions model water diffusion in the brain and give an estimation of the underlying fiber architecture. Each ODF's maxima should indicate fibers' directions. Let Λ_k be a set of directions where ψ_{x_k} is locally maximum. Then, importance density is defined in dimension 3 as vMF mixture:

$$\pi(v_k|v_{0:k-1}, z_{1:k}) = \begin{cases} \sum_{\mu \in \Lambda_k} \omega_\mu \, \mathrm{f}_3(u_k|\, \mu, \kappa_\mu) & \text{if } \Lambda_k \neq \emptyset \\ \mathrm{f}_3(u_k|\, u_{k-1}, \kappa) & \text{else} \end{cases} \quad , \tag{2}$$

where f_3 is the vMF on the 2-sphere in $\mathbb{R}^3$, u_k is the unit vector in spherical coordinates corresponding to the unit vector v_k in Cartesian coordinates, κ_μ is the concentration depending on observations, κ is the concentration parameter and ω_μ are mixture proportions such that $0 \leq \omega_\mu \leq 1$ and $\sum_{\mu \in \Lambda_k} \omega_\mu = 1$. Each ω_μ is proportional to the ODF value in direction μ.

In [20], the concentration parameter of importance is computed as an exponential function of fractional anisotropy (FA). For ODF functions, FA in one direction can be measured as the mean curvature of the function in that direction [5]. Thus, the concentration parameter κ_μ is proportional to ODF curvature in direction μ: $\kappa_\mu \propto H(\mu)$ where $H(\mu)$ is the mean curvature of ψ_{x_k} in direction μ. The use of a vMF mixture allows us an efficient sampling procedure of the importance distribution [19].

This definition of importance distribution approximate likelihood distribution [20]. The paths cloud is then weighted by importance weights proportionally to prior distribution. This means that each particle is projected backward in order to know its probability to have a predecessor among paths cloud at previous step.

The Likelihood. Partially due to noise, there is uncertainty in MRI diffusion information. The observation density is supposed to be a measure of this uncertainty. In term of probabilities, likelihood defines a measure of how observations match the current model. Therefore, we model likelihood density as a distance error between measured observations and observations matching perfectly the current state model. At step k, u_k is the direction sampled according to the importance and μ is assumed to be the mean vector of the sampled vMF distribution (in order to get a direction u_k). According to section 2.1, the diffusion signal at position x_k is described by S_{x_k}. Let $\bar{S}_{x_k}$ be the diffusion signal at position x_k if the sampled direction u_k was exactly the mean direction μ, i.e. $\mu = u_k$. Then, $\bar{S}_{x_k}$ is determined from S_{x_k} by a rotation of angle $\widehat{\mu u_k}$ in three dimensional space: $\bar{S}_{x_k} = \mathrm{rot}(\widehat{\mu u_k}, S_{x_k})$, where $\widehat{\mu u_k}$ denotes the angle between μ and u_k.

Similary to [20], let consider $\bar{S}_{x_k}$ as the diffusion signal and the MRI intensities $\mathcal{U}_{x_k}$, a noisy measure of it, i.e. $\mathcal{U}_{x_k} = \bar{S}_{x_k} + \varepsilon$. Noise in MRI images can be described closely approximated for moderate-large SNR by a normal distribution [11] with a zero-mean and standard deviation of noise in MRI diffusion weighted image, i.e. $\varepsilon = \mathcal{U}_{x_k} - \bar{S}_{x_k} \sim \mathcal{N}(0, \Sigma^2)$. The standard deviation of diffusion MRI image in each gradient directions σ_i is estimated by least square estimation and pseudo-residuals [10]. Thus, likelihood is written as a multiplication of error distribution in each gradient direction g_i:

$$p(z_k|v_k) = \prod_{i=1}^{N} \mathcal{N}(\,\mathcal{U}_{x_k}(g_i) - \bar{S}_{x_k}(g_i)|\;0, \sigma_i^2) \quad . \tag{3}$$

3 Results

The tractography method has been evaluated on the MICCAI'09 Fiber Cup phantom and on in-vivo brain DWMRI data. Since the sampled particles evolve in a continuous domain and data is acquired in a discrete manner, a third-order B-Spline interpolation is applied to dataset when required. Each tractography is performed by sampling 1000 particles. The result of the algorithm is an estimation of the posterior density which corresponds to a connectivity map. Fiber bundles can be obtained by estimating the maximum a posteriori (MAP) of the posterior density.

3.1 Fiber Cup

The Fiber Cup is a tractography contest proposed at the MICCAI conference help in London in 2009[1]. By containing several crossing, kissing, splitting and bending fiber configurations, this phantom is a very convenient way to compare the proposed method with other existing approaches. The Fiber Cup phantom size is $64 \times 64 \times 3$ voxels with a resolution of $3 \times 3 \times 3$ mm. It contains 2 repetitions of 65 gradient directions each, including 1 baseline direction for each repetition. The b-value is 1500 for each direction. The seeds' locations of the contest are exposed in Fig. 1.

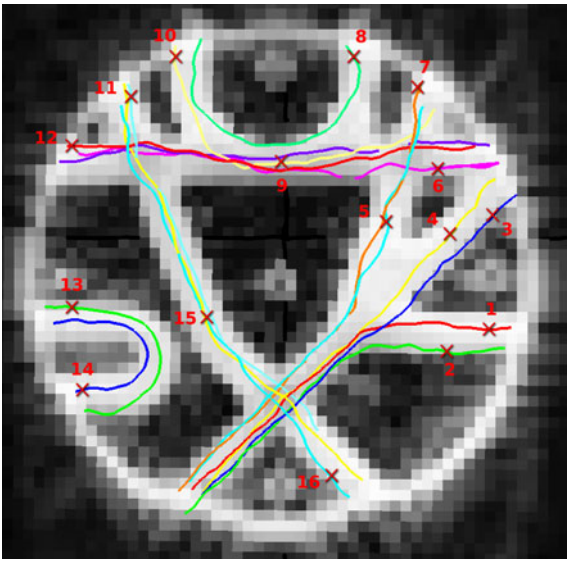

Fig. 1. Fibers estimations of our algorithm on the Fiber Cup phantom. Seeds are marked with red cross.

Table 1. Quantitative results of the proposed algorithm on the Fiber Cup phantom. When the proposed algorithm reaches the first three places of the contest, the corresponding results are marked in bold.

	L2	tan	curv
F1	**1.68952**	14.8774	0.155911
F2	**2.93385**	25.9358	2.81907
F3	**1.64951**	13.8554	0.165466
F4	**2.12385**	16.8663	0.157994
F5	**2.72598**	**16.3259**	0.160928
F6	**3.51537**	23.713	1.02897
F7	**3.58313**	**21.245**	2.54886
F8	**2.34624**	21.7216	3.82264
F9	**3.03358**	17.6071	0.150982
F10	8.96364	31.2289	0.147886
F11	3.53476	17.6234	0.136861
F12	**3.87773**	**20.7672**	0.183536
F13	**2.60168**	16.2984	0.17025
F14	**2.47026**	**16.2079**	0.133395
F15	**2.24925**	18.6442	0.157172
F16	**4.26736**	17.4609	0.136369

The quantitative results of our algorithm are summarized in the Tab. 1 and estimated fibers from the particles cloud are shown in Fig 1. The algorithm gets 35 points according to the Fiber Cup contest and would be at the 3^{rd} position in the ranking. Considering only the spatial metric, our algorithm places itself to the first position with 31 points, before [17] which gets 22 points. As shown in Table 1, although the method has pleasing results with the spatial metric (L2 norm), it suffers from a lack of precision in both tangent and curvature metric. This might be due to the quality of the estimation of the MAP of the posterior density and not to the particle filtering algorithm itself. Nevertheless, satisfactory results are generated in a short computation time. For instance, the

[1] The website http://www.lnao.fr/spip.php?rubrique79 provides details and results about this competition. The phantom used to compare tractographies and the comparison program are also available on this website.

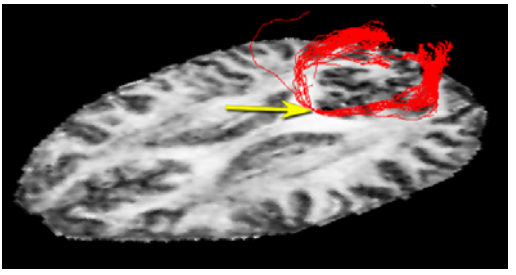

(a) Particles cloud

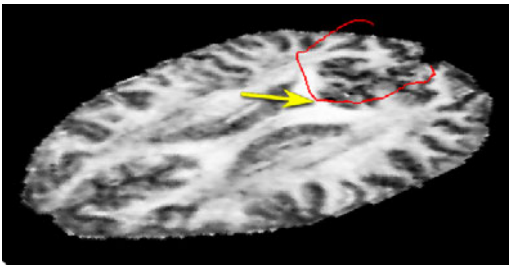

(b) MAP estimate of the posterior density

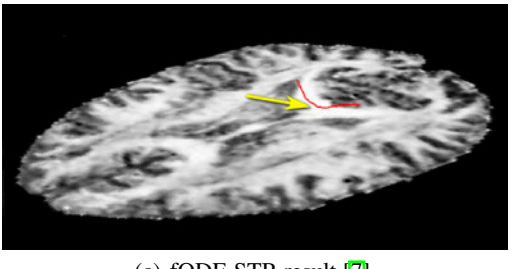

(c) fODF-STR result [7]

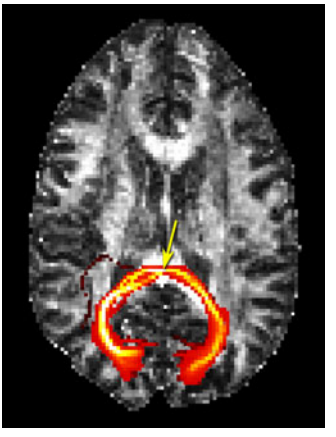

(d) Connectivity map

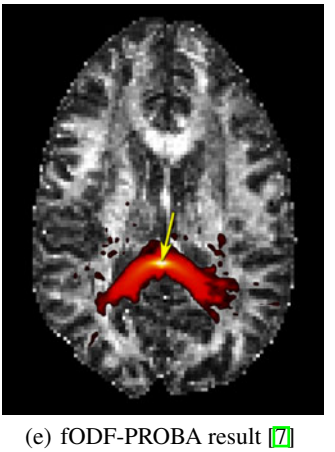

(e) fODF-PROBA result [7]

Fig. 2. Results of tractographies of adult human brain starting from a seed point in the corpus callosum indicated by a yellow arrow. Each result is displayed on a FA image slice. Tractographies were performed using fODF-STR (Fig. 2(c)), fODF-PROBA (Fig. 2(e)) and the proposed algorithm (Fig. 2(a), 2(b) and 2(d)).

winner of this contest [17] uses a global algorithm which performed the tractography in one day of computation time whereas ours implementation did run in 8 minutes on a computer with 8 processor Intel Xeon 2.4GHz (including the preprocessing, i.e. the model estimation, the noise estimation, the maxima extraction, etc.).

3.2 In-vivo Data

In-vivo brain dataset comes from the community MIDAS / National Alliance for Medical Image Community (NAMIC)[2] and was acquired from a healthy adult volunteer

[2] http://insight-journal.org/midas/collection/view/190

using a 3 Tesla GE system. It contains a $144 \times 144 \times 85$ volume image with $1.7 \times 1.7 \times 1.7$ mm voxel resolution. The diffusion signal was measured in 51 directions with $b = 900$ s.mm^{-2} and there are 8 baseline images ($b = 0$ s.mm^{-2}). During the model estimation, the 8 baseline images have been used by averaging them. The model order for the spherical harmonics is $l = 4$. The proposed algorithm was applied to this volume image with a step length of 0.85 mm, a resampling threshold $\varepsilon_{ESS} = 60\%$ and 1000 particles. The tractography has been performed on a brain mask defined by voxels with a fractional anisotropy (FA) value greater than 0.2. In order to show the contribution of the particle filtering framework into the tractography results, the proposed algorithm is compared with two other tractography frameworks (using the same diffusion model): fODF streamline (fODF-STR) and fODF "probabilistic" (or random walk, named later fODF-PROBA) [7].

Figure 2 shows the results of a tractography starting from a seed point in the corpus callosum. An axial slice of FA map is shown as a reference background image. The particle cloud is shown in Fig. 2(a). From this cloud, both fiber path and posterior density map (projected on the axial slice) have been generated (respectively Fig. 2(b) and 2(d)). Using the same experimental parameters, our implementations of fODF based algorithms, were applied to the dataset (results are displayed respectively in Fig. 2(c) and 2(e)). The propagation of fODF-STR is early stopped because it reaches the boundary of the brain mask, resulting in a partial tractography (Fig. 2(c)). The fODF-PROBA algorithm leads to a very smooth posterior density approximation (Fig. 2(e)) due to the diffusion-based random walk strategy used in this algorithm.

4 Discussion

In this article, a probabilistic tractography method has been presented based on the well-formalized particle filtering framework and using an analytical Q-ball model for diffusion data. The output of the algorithm is an estimate of the posterior density of the white matter fibers. The keypoints of the method are the three densities: prior, likelihood, importance. First, the prior density constrains the solution. Then, the likelihood density ensures the reasonableness of the estimation with a noise model of DWMRI data. Finally, fast and efficient sampling is realized by a vMF mixture.

As shown by the experiments performed on the Fiber Cup phantom, the use of a Q-ball data modeling within the particle filtering framework leads to accurate estimations of complex fiber configurations. Experiments on in-vivo data have shown the contribution of the particle filtering framework compared to streamline approach or random walk like techniques. Further work could study the whole diffusion information available, without any diffusion model constraint. Indeed, techniques for estimating models induce approximations that can exaggerate uncertainty in the data.

Acknowledgment. The research leading to these results has received funding from the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013 Grant Agreement no. 207667).

References

1. Alexander, A.L., Hasan, K.M., Lazar, M., Tsuruda, J.S., Parker, D.L.: Analysis of partial volume effects in diffusion-tensor MRI. Magnetic Resonance in Medicine 45(5), 770–780 (2001)
2. Basser, P.J., Pajevic, S., Pierpaoli, C., Duda, J., Aldroubi, A.: In-vivo fiber tractography using DT-MRI data. Magnetic Resonance in Medicine 44(4), 625–632 (2000)
3. Basser, P., Mattiello, J., LeBihan, D.: MR diffusion tensor spectroscopy and imaging. Biophysical Journal 66(1), 259–267 (1994)
4. Behrens, T., Berg, H.J., Jbabdi, S., Rushworth, M., Woolrich, M.: Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? NeuroImage 34(1), 144–155 (2007)
5. Bloy, L., Verma, R.: On computing the underlying fiber directions from the diffusion orientation distribution function. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part I. LNCS, vol. 5241, pp. 1–8. Springer, Heidelberg (2008)
6. Descoteaux, M., Angelino, E., Fitzgibbons, S., Deriche, R.: Regularized, fast, and robust analytical q-ball imaging. Magnetic Resonance in Medicine 58(3), 497510 (2007)
7. Descoteaux, M., Deriche, R., Knosche, T., Anwander, A.: Deterministic and probabilistic tractography based on complex fibre orientation distributions. IEEE Transactions on Medical Imaging 28(2), 269–286 (2009)
8. Doucet, A., Godsill, S., Andrieu, C.: On sequential monte carlo sampling methods for bayesian filtering. Statistics and Computing 10(3), 197–208 (2000)
9. Friman, O., Farneback, G., Westin, C.: A bayesian approach for stochastic white matter tractography. IEEE Transactions on Medical Imaging 25(8), 965–978 (2006)
10. Gasser, T., Sroka, L., Jennen-Steinmetz, C.: Residual variance and residual pattern in nonlinear regression. Biometrika 73(3), 625–633 (1986)
11. Gudbjartsson, H., Patz, S.: The rician distribution of noisy mri data. Magnetic Resonance in Medicine 34(6), 910–914 (1995)
12. Kong, A., Liu, J.S., Wong, W.H.: Sequential imputations and bayesian missing data problems. Journal of the American Statistical Association 89(425)
13. Lazar, M., Alexander, A.L.: Bootstrap white matter tractography (BOOT-TRAC). NeuroImage 24(2), 524–532 (2005)
14. Mori, S., Crain, B.J., Chacko, V.P., Zijl, P.C.M.V.: Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. Annals of Neurology 45(2), 265–269 (1999)
15. Parker, G., Alexander, D.: Probabilistic monte carlo based mapping of cerebral connections utilising whole-brain crossing fibre information. In: Information Processing in Medical Imaging, p. 684695 (2003)
16. Parker, G., Wheeler-Kingshott, C., Barker, G.: Estimating distributed anatomical connectivity using fast marching methods and diffusion tensor imaging. IEEE Transactions on Medical Imaging 21(5), 505–512 (2002)
17. Reisert, M., Mader, I., Anastasopoulos, C., Weigel, M., Schnell, S., Kiselev, V.: Global fiber reconstruction becomes practical. NeuroImage 54(2), 955–962 (2011)
18. Staempfli, P., Jaermann, T., Crelier, G., Kollias, S., Valavanis, A., Boesiger, P.: Resolving fiber crossing using advanced fast marching tractography based on diffusion tensor imaging. NeuroImage 30(1), 110–120 (2006)
19. Ulrich, G.: Computer generation of distributions on the m-Sphere. Journal of the Royal Statistical Society. Series C (Applied Statistics) 33(2), 158–163 (1984)
20. Zhang, F., Hancock, E.R., Goodlett, C., Gerig, G.: Probabilistic white matter fiber tracking using particle filtering and von Mises-Fisher sampling. Medical Image Analysis 13(1), 5–18 (2009)

Bessel Fourier Orientation Reconstruction: An Analytical EAP Reconstruction Using Multiple Shell Acquisitions in Diffusion MRI

Ameer Pasha Hosseinbor*, Moo K. Chung, Yu-Chien Wu, and Andrew L. Alexander

University of Wisconsin-Madison, Madison, WI, USA
hosseinbor@wisc.edu

Abstract. The estimation of the ensemble average propagator (EAP) directly from **q**-space DWI signals is an open problem in diffusion MRI. Diffusion spectrum imaging (DSI) is one common technique to compute the EAP directly from the diffusion signal, but it is burdened by the large sampling required. Recently, several analytical EAP reconstruction schemes for multiple **q**-shell acquisitions have been proposed. One, in particular, is Diffusion Propagator Imaging (DPI) which is based on the Laplace's equation estimation of diffusion signal for each shell acquisition. Viewed intuitively in terms of the heat equation, the DPI solution is obtained when the heat distribution between temperatuere measurements at each shell is at steady state.

We propose a generalized extension of DPI, Bessel Fourier Orientation Reconstruction (BFOR), whose solution is based on heat equation estimation of the diffusion signal for each shell acquisition. That is, the heat distribution between shell measurements is no longer at steady state. In addition to being analytical, the BFOR solution also includes an intrinsic exponential smootheing term. We illustrate the effectiveness of the proposed method by showing results on both synthetic and real MR datasets.

1 Introduction

The main aim of diffusion-weighted imaging (DWI) is to non-invasively recover information about the diffusion of water molecules in biological tissues, in particular white matter (WM). The EAP contains the full information about the diffusion process, which reflects the complex tissue micro-structure, and its estimation lies at the heart of diffusion MRI. When the narrow pulse condition is met, the EAP is related to the **q**-space diffusion signal E by the Fourier transform:

$$P(\mathbf{p}) = \int E(\mathbf{q}) e^{-2\pi \mathbf{q}\cdot\mathbf{p}} d^3\mathbf{q}, \tag{1}$$

* Corresponding author.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 217–225, 2011.

where $\mathbf{p}$ is the displacement vector in propagator-space and $\mathbf{q}$ is the diffusion wave-vector in signal-space. We denote $\mathbf{q} = q\mathbf{u}$ and $\mathbf{p} = p\mathbf{r}$, where $\mathbf{u}$ and $\mathbf{r}$ are 3D unit vectors.

In DTI [2], the diffusion signal is modeled with a Gaussian function. Consequently, DTI can only map a single fiber orientation within a voxel, and fails in voxels with orientational heterogeneity [11]. Diffusion Spectrum Imaging (DSI) [12] is a well known model-free method that uses the Fourier transform to numerically estimate the EAP. It is burdened, however, by the dense sampling required on a Cartesian grid in $\mathbf{q}$-space, which results in long acquisition times.

As a result of the constraints of DSI, several analytical technqiues have arisen that derive an expression for the EAP based on a particular assumption of the diffusion signal. Spherical Polar Fourier Imaging (SPFI) [1,3] models the signal in terms of an orthonormal basis comprising spherical harmonics (SH) and Gaussian-Laguerre polynomials, and was proposed to sparsely represent $E(\mathbf{q})$. The authors in [10] expanded the diffusion signal in terms of another orthonormal basis that appears in the 3D quantum mechanical harmonic oscillator problem. Their basis is closely related to the SPFI basis. A technique that does not model signal in terms of an orthonormal basis is Diffusion Propagator Imaging (DPI) [7], where the diffusion signal is assumed to be a solution to the 3D Laplace's equation $\bigtriangledown^2 E = 0$. It seems to work well with only a small number samples. However, according to DPI, $E(0)$ does not exist, which makes the assumption of Laplacian equation modeling unrealistic for $E(\mathbf{q})$. Thus, DPI may work well within a range of $\mathbf{q}$ values, but not for the entire $\mathbf{q}$-space.

In this paper, we develop Bessel Fourier Orientation Reconstruction (BFOR), which is a generalized extension of DPI. Rather than assuming the signal satisfies Laplace's equation, we reformulate the problem into a Cauchy problem and assume the signal satisfies the heat equation. The heat equation is a generalization of Laplace's equation, which the latter approaches in the steady state (i.e. $t \to \infty$). It provides an analytical reconstruction of the EAP profile from diffusion signal and models the diffusion signal in terms of an orthonormal basis. In addition, BFOR contains an intrinsic exponential smoothening term that allows one to control the amount of smoothening in the EAP estimation. The last point is significant because, although the Laplacian modeling intrinsically smoothens the diffusion signal, the amount of smoothening can not be controlled, and hence it may oversmooth the signal. We test our method on both synthetic and in vivo datasets.

2 Bessel Fourier Orientation Reconstruction

2.1 Estimation of EAP

Consider the eigenvalue/boundary condition problem

$$\Im_x \psi_i(x) = -\lambda_i \psi_i(x), \quad \psi_i(x = \tau) = 0 \tag{2}$$

which we use to solve the Cauchy problem

$$\frac{\partial}{\partial t}g(x,t) - \Im_x g(x,t) = 0, \quad g(x,0) = f(x), \tag{3}$$

where $f(x)$ is simply the acquired signal and $\Im$ is some self-adjoint linear operator. Chung et al. in [4] derived a unique solution for (3):

$$g(x,t) = \sum_{i=0}^{\infty} a_i e^{-\lambda_i t}\psi_i(x), \tag{4}$$

where $e^{-\lambda_i t}$ is a smoothening term controlled by parameter $t \geq 0$ and the coefficients are given by $a_i = \langle f, \psi_i \rangle$. The implication of (4) is that the solution decreases exponentially as t increases and smoothes out high spatial frequency noise much faster than low-frequency noise. In DPI, however, the steady state assumption permanently removes any temporal term, which governs the extent of smoothening, so there is no smoothening control mechanism. Note that $t = 0$ corresponds to unsmoothened solution.

Within the context of our problem, $g(x,t)$ is **q**-space diffusion signal. The assumption of a 3D Laplacian operator in spherical coordinates allows us to solve (2) via separation of variables, and hence obtain the orthonormal basis $\psi_{nj}(\mathbf{q}) = j_{l(j)}(\frac{\alpha_{nl(j)}q}{\tau})Y_j(\mathbf{u})$,[1] where $\alpha_{nl(j)}$ is n^{th} root of l^{th} order spherical Bessel function of first kind j_l and τ is the radial distance in q-space at which the Bessel function (and hence signal) goes to zero. Y_j are a modified real and symmetric SH basis proposed in [7] to reflect the symmetry and realness of the diffusion signal. The eigenvalues are $\lambda_{nl(j)} = \frac{\alpha^2_{nl(j)}}{\tau^2}$.

The assumption of a Laplacian operator results in (3) becoming the heat equation: $\bigtriangledown^2 E = \frac{\partial E}{\partial t}$. From (4) then, the **q**-space signal can be expanded in terms of the spherical orthonormal basis ψ_{nlm}:

$$E(\mathbf{q},t) = \sum_{n=1}^{N}\sum_{j=1}^{R} C_{nj} e^{\frac{-\alpha^2_{nl(j)}t}{\tau^2}} j_{l(j)}(\frac{\alpha_{nl(j)}q}{\tau})Y_j(\mathbf{u}), \tag{5}$$

where C_{nj} are the expansion coefficients, $R = \frac{(L+1)(L+2)}{2}$ is the number of terms in the modified SH basis of truncation order L, and N is the truncation order of radial basis. Thus, the total number of coefficients in the expansion is $W = \frac{N(L+1)(L+2)}{2}$. Note that the actual acquired signal from scanner is given at $t = 0$.

To derive the EAP, we express the Fourier kernel in (1) as a plane wave expansion in spherical coordinates

$$e^{-2\pi i \mathbf{q}\cdot\mathbf{p}} = 4\pi \sum_{j=1}^{\infty} (-i)^{l(j)} j_{l(j)}(2\pi qp) Y_j(\mathbf{u}) Y_j(\mathbf{r}) \tag{6}$$

Then substituting (5) and (6) into (1), we obtain

[1] For detailed derivations, visit http://brainimaging.waisman.wisc.edu/∼ameer/Derivations.pdf

$$P(\mathbf{p},t) = 4\pi \int \sum_{n=1}^{N}\sum_{j=1}^{R} C_{nj} e^{\frac{-\alpha_{nl(j)}^2 t}{\tau^2}} j_{l(j)}(\frac{\alpha_{nl(j)} q}{\tau}) Y_j(\mathbf{u}) \sum_{j'=1}^{\infty} (-i)^{l(j')} j_{l(j')}(2\pi qp) Y_{j'}(\mathbf{u}) Y_{j'}(\mathbf{r})\, d^3\mathbf{q}$$
$$= 4\pi \sum_{n=1}^{N}\sum_{j=1}^{R} (-1)^{l(j)/2} C_{nj} e^{\frac{-\alpha_{nl(j)}^2 t}{\tau^2}} Y_j(\mathbf{r}) I_{nl(j)}(p), \tag{7}$$

where we use the orthonormal property of SH, i.e. $\int Y_j(\mathbf{u}) Y_{j'}(\mathbf{u}) d^2\mathbf{u} = \delta_{jj'}$ and define

$$I_{nl(j)}(p) = \int_0^\infty q^2 j_{l(j)}(\frac{\alpha_{nl(j)} q}{\tau}) j_{l(j)}(2\pi qp) dq \approx \int_0^\tau q^2 j_{l(j)}(\frac{\alpha_{nl(j)} q}{\tau}) j_{l(j)}(2\pi qp) dq \tag{8}$$

After some algebra and exploiting properties of the spherical Bessel function, we can write the EAP, for $p > 0$, as

$$P(\mathbf{p},t) = \sqrt{\frac{2\pi^3\tau}{p}} \sum_{n=1}^{N}\sum_{j=1}^{R} (-1)^{\frac{l(j)}{2}} C_{nj} e^{\frac{-\alpha_{nl(j)}^2 t}{\tau^2}} Y_j(\mathbf{r}) \frac{\sqrt{\alpha_{nl(j)}} J_{l(j)-1/2}(\alpha_{nl(j)}) J_{l(j)+1/2}(2\pi\tau p)}{\left(4\pi^2 p^2 - \frac{\alpha_{nl(j)}^2}{\tau^2}\right)} \tag{9}$$

2.2 Implementation of Methods

The task is to estimate coefficients C_{nj} from the observed signal $E(q, \mathbf{u}, t = 0)$. We achieve this by carrying out a linear least square (LS) fitting with regularization in the radial and angular parts. For the LS estimation, denote signal vector by $E = [E(\mathbf{q}_k, t = 0]_{S\times 1}$, the coefficient vector by $C = [C_{nj}]_{W\times 1}$, and the basis matrix by $M = [j_{l(j)}(\frac{\alpha_{nl(j)} q_k}{\tau}) Y_j(\mathbf{u}_k)]_{S\times W}$, where $k = 1, ..., S$. The angular regularization matrix, denoted by $\tilde{L}$, is the Laplace-Beltrami diagonal matrix with $l^2(l+1)^2$ entries on the diagonal and the radial regularization matrix, denoted by $\tilde{N}$, is a diagonal matrix with entries $n^2(n+1)^2$ along its diagonal. The angular and radial regularization matrices penalize, respectively, high degrees of the angular and radial parts of (5) in the estimation under the assumption that they are likely to capture noise [1]. Then the coefficients are $C = (M^T M + \lambda_l \tilde{L} + \lambda_n \tilde{N})^{-1} M^T E$, where λ_l and λ_n are the regularization terms for angular and radial bases, respectively.

Lastly, once we extract the coefficents, the EAP profile at some radius and some instant of smoothening is computed by interpolating Z points along the

equator of a sphere of the same radius (i.e. the polar angle is $\pi/2$ and the azimuthal angle is varied from 0 to 2π). Denoting the EAP vector at given radius p_o and at a given instant t_o of smoothening by $P = [P(p_o, \mathbf{r}_o, t_o)]_{Z\times 1}$ and the $Z \times W$ EAP basis matrix evaluated at p_o and t_o by F, we have $P = FC = F(M^T M + \lambda_l \tilde{L} + \lambda_n \tilde{N})^{-1} M^T E$. It is important to note that we apply smoothening on the EAP, itself, and not on the diffusion signal.

3 Results on Synthetic and Real Data

Synthetic Data. The monoexponential mixture model is frequently used to generate synthetic data to validate a given EAP reconstruction, such as in [1,3], where the maximum b-value used exceeded 3000 s/mm^2. However, diffusion MR imaging experiments using high b-values (> 2000 s/mm^2) have shown that the diffusion signal decay is no longer monoexponential. Studies in normal human brain, with b-values over an extended range of up to 6000 s/mm^2, have shown that the signal decay is better described with a biexponetial curve [9,5]. Thus, we apply BFOR and DPI to simulations of crossing fiber configurations generated via a biexponential mixture model.

In biexponential mixture, $E(\mathbf{q}) = \sum_{k=1}^{N_b} [f_{kf} e^{-b\mathbf{u}^T \mathbf{D}_{kf} \mathbf{u}} + f_{ks} e^{-b\mathbf{u}^T \mathbf{D}_{ks} \mathbf{u}}]$, assuming no exchange between compartments [5, 13]. We look at two equally weighed fibers and set eigenvalues of each diffusion tensor to be [1.6,0.4,0.4]e-3. Diffusion measurements in the corpus callosum were used to simulate fast and slow Gaussian diffusion functions [8]. A hybrid sampling scheme [13] was used and consisted of one baseling image acquired at b=0 s/mm^2 and 6 shells, with $(Ne, b) = \{(6, 690), (21, 2780), (24, 6250), (24, 1.11e4), (100, 1.74e4)\, (100, 2.5e4)\}$, where Ne denotes number of encoding directions, and $q_{max}/\Delta q$=76/15.2 mm^{-1}. Since EAP reconstruction is sensitive to angular resolution, the number of encoding directions is increased with each shell to increase the angular resolution with the level of diffusion weighting. We then add Rician noise the same way as in [6], with $SNR = 1/\sigma$, which is defined as the ratio of maximum signal intensity $S(0) = 1$ to the standard deviation σ of complex Gaussian noise. At $SNR = 10$, 200 trials were simulated. The BFOR parameters are $\{N = 8, L = 6, \tau = 106.4$ mm^{-1}, $\lambda_l = 10^{-6}, \lambda_n = 10^{-6}\}$ and the DPI parameters $\{L = 6, \lambda = 0\}$.

Fig. 1 shows that the BFOR basis fits the diffusion signal nearly perfectly, while Figs. 2 and 4 demonstrate that BFOR successfully captures the geometry and orientation of the EAP profile. The corresponding DPI estimated EAP profiles, shown in Figs. 3 and 5, are not as accurate as BFOR. Fig. 6 shows the Euclidean squared error, averaged over all noise simulations, at various radii for DPI and BFOR, and indicates both methods have a similar robustness to noise. As the smoothening parameter t is increased, the squared error is gradually reduced.

Real Data. We tested our method using healthy, adult human data, with $q_{max}/\Delta q$=76/15.2 mm^{-1} and TE/TR/matrix=122ms/11700ms/128 $\times$ 128 $\times$ 30. The sampling scheme consisted of two baseline images acquired at $b = 0$ s/mm^2 and 5 shells, with (Ne, b)={(6,375),(21,1500),(24,3375),(24,6000),

(a) b=690 (b) b=2780 (c) b=6250 (d) b=1.11e4 (e) b=1.74e4 (f) b=2.50e4

Fig. 1. The ground truth diffusion signal (green) and BFOR estimated signal (red) when noise was absent are compared using all 6 available shells. Two equally weighted fibers were simulated at $t = 0$ crossing at 60°.

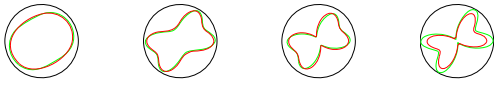

(a) $p = 5$ (b) $p = 15$ (c) $p = 25$ (d) $p = 30$

Fig. 2. The EAP estimated by BFOR (red) and actual EAP (green) in absence of noise for two equally weighed fibers crossing at 60° at $t = 0$ for radii $p = 5$, 15, 25, and 30 μm.

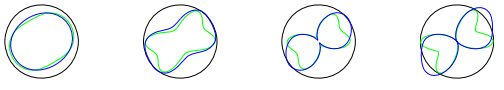

(a) $p = 5$ (b) $p = 15$ (c) $p = 25$ (d) $p = 30$

Fig. 3. The EAP estimated by DPI (blue) and actual EAP (green) in absence of noise for two equally weighed fibers crossing at 60° for radii $p = 5$, 15, 25, and 30 μm

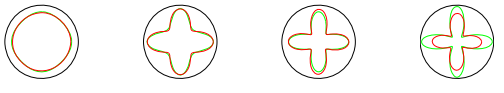

(a) $p = 5$ (b) $p = 15$ (c) $p = 25$ (d) $p = 30$

Fig. 4. The EAP estimated by BFOR (red) and actual EAP (green) in absence of noise for two equally weighed fibers crossing at 90° at $t = 0$ for radii $p = 5$, 15, 25, and 30 μm

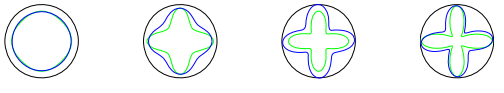

(a) $p = 5$ (b) $p = 15$ (c) $p = 25$ (d) $p = 30$

Fig. 5. The EAP estimated by DPI (blue) and actual EAP (green) in absence of noise for two equally weighed fibers crossing at 90° for radii $p = 5$, 15, 25, and 30 μm

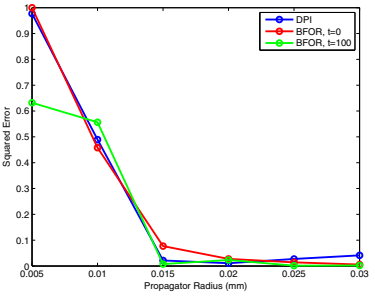

Fig. 6. Plotted is the point-wise mean Euclidean squared error of BFOR and DPI as a function of propagator radius for two fibers crossing at 60° when $SNR = 10$. The squared error is normalized by maximum squared error from BFOR.

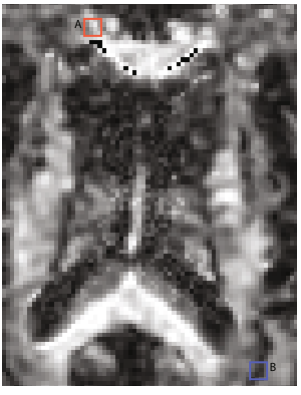

Fig. 7. Axial slice of FA map (b=1500) of adult human brain. `ROI A` on the corpus callosum is a region where we expect single fibers, while `ROI B` is one where we expect crossing fibers (i.e. dark streaks).

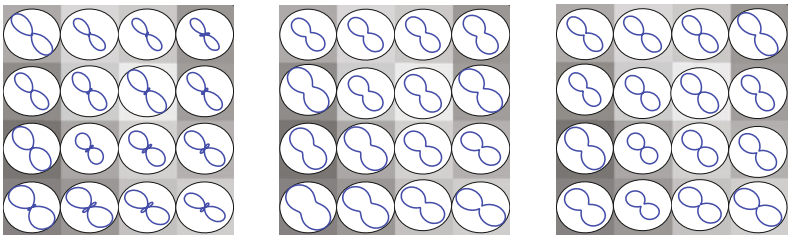

(a) BFOR at $t = 0$ (b) BFOR at $t = 400$ (c) DPI

Fig. 8. Plotted is the EAP profile at $p = 10$ μm overlaid on FA map in `ROI A` using BFOR at (a) $t = 0$, (b) $t = 400$, and (c) DPI

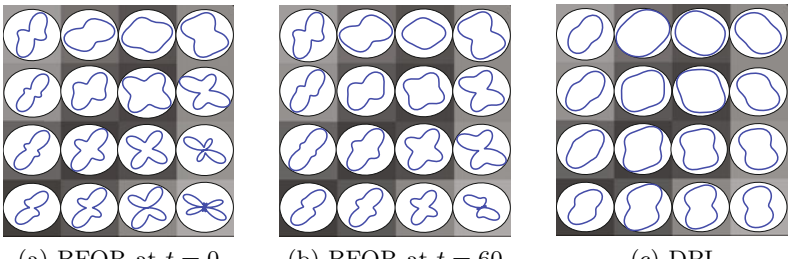

(a) BFOR at $t = 0$ (b) BFOR at $t = 60$ (c) DPI

Fig. 9. Plotted is the EAP profile at $p = 10$ μm overlaid on FA map in `ROI B` using BFOR at (a) $t = 0$, (b) $t = 60$, and (c) DPI. Note how DPI falsely indicates certain WM voxels exhibiting near isotropic diffusion.

(50,9375)} [13]. The number of directions in the outer shells were increased to better characterize complex tissue organization. We set $L = 4$, $N = 4$, $\tau = 91.2$ mm^{-1}, $\lambda_l = 10^{-6}$, $\lambda_n = 10^{-6}$. As shown in Fig. 7, two 4×4 ROIs were drawn on the same slice from a FA map: one on corpus callosum where we expect single fibers, and another on area with black streaks where we expect crossing fibers. The EAP profile was estimated at $p = 10$ μm for each ROI. Based on results shown in Figs. 8 and 9, we see that BFOR performs well in corpus callosum and can successively resolve multiple fiber orientations in voxels with orientational heterogenity. The DPI reconstruction in Fig. 9c, however, erroneously indicates certain WM voxels exhibiting near isotropic diffusion.

4 Discussion

In BFOR, we derive a general solution for the heat equation that is a function of the amount of smoothening t, while DPI solves the heat equation at steady state ($t \to \infty$). This is the first study to include biexponential diffusion model for EAP reconstruction, and the results from the numerical phantom show that BFOR basis models the diffusion signal very well and can successfully reproduce the ground truth EAP. However, the sampling scheme used in the numerical simulations may not be clinically feasible because of the extremely low SNR at b_{max}=2.5e4. The application of BFOR to real data revealed that it can successfully retrieve multiple fiber orientations. Comparison with DPI depicts our model in a favorable light. Future work includes applying compressed sensing to BFOR to find a sparser sampling scheme, and computing quantitative properties of the EAP using BFOR basis.

References

1. Assemlal, H.E., Tschumperlé, D., Brun, L.: Efficient and robust computation of pdf features from diffusion mr signal. Med. Image Anal. 13, 715–729 (2009)
2. Basser, P.J., Mattiello, J., LeBihan, D.: Mr diffusion tensor spectroscopy and imaging. Biophysical Journal 66, 259–267 (1994)
3. Cheng, J., Ghosh, A., Jiang, T., Deriche, R.: Model-free and analytical EAP reconstruction via spherical polar fourier diffusion MRI. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. 6361, pp. 590–597. Springer, Heidelberg (2010)
4. Chung, M.K., Dalton, K.M., Shen, L., Evans, A.C., Davidson, R.J.: Weighted fourier series representation and its application to quantifying the amount of gray matter. IEEE Transac. Med. Imaging 26, 566–581 (2007)
5. Clark, C.A., Le Bihan, D.: Water diffusion compartmentation and anisotropy at high b values in the human brain. Magn. Reson. Med. 44, 852–859 (2000)
6. Descoteaux, M., Angelino, E., Fitzgibbons, S., Deriche, R.: Regularized, fast, and robust analytical q-ball imaging. Magn. Reson. Med. 58, 497–510 (2007)
7. Descoteaux, M., Deriche, R., LeBihan, D., Mangin, J.F., Poupon, C.: Multiple q-shell diffusion propagator imaging. Med. Image Anal (2010)

8. Maier, S.E., Vajapeyam, S., Mamata, H., Westin, C.F., Jolesz, F.A., Mulkern, R.V.: Biexponential diffusion tensor analysis of human brain diffusion data. Magn. Reson. Med. 51, 321–330 (2004)
9. Mulkern, R.V., Gudbjartsson, H., Westin, C.F., Zengingonul, H.P., Gartner, W., Guttmann, C.R., Robertson, R., Kyriakos, W., Schwartz, R., Holtzman, D., Jolesz, F.A., Maier, S.E.: Multi-component apparent diffusion coefficients in human brain. NMR Biomed. 12, 51–62 (1999)
10. Ozarslan, E., Koay, C., Shepherd, T.M., Blackband, S.J., Basser, P.J.: Simple harmonic oscillator based reconstruction and estimation for three-dimensional q-space mri. In: Proc. Intl. Soc. Mag. Reson. Med. (2009)
11. Tuch, D.S., Reese, T.G., Wiegell, M.R., Makris, N., Belliveau, J.W., Weeden, V.J.: High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. Magn. Reson. Med. 48, 577–582 (2002)
12. Weeden, V.J., Hagmann, P., Tseng, W.Y.I., Reese, T.G., Weisskoff, R.M.: Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. Magn. Reson. Med. 54, 1377–1386 (2005)
13. Wu, Y.C., Alexander, A.L.: Hybrid diffusion imaging. NeuroImage 36, 617–629 (2007)

Parallel MRI Noise Correction: An Extension of the LMMSE to Non Central χ Distributions

Véronique Brion[1,2], Cyril Poupon[1,2], Olivier Riff[1,2], Santiago Aja-Fernández[3], Antonio Tristán-Vega[3], Jean-François Mangin[1,2], Denis Le Bihan[1,2], and Fabrice Poupon[1,2]

[1] CEA I2BM NeuroSpin, Gif-sur-Yvette, France
[2] IFR 49, Paris, France
[3] LPI, Universidad de Valladolid, Spain

Abstract. Parallel MRI leads to magnitude data corrupted by noise described in most cases as following a Rician or a non central χ distribution. And yet, very few correction methods perform a non central χ noise removal. However, this correction step, adapted to the correct noise model, is of very much importance, especially when working with Diffusion Weighted MR data yielding a low SNR. We propose an extended Linear Minimum Mean Square Error estimator (LMMSE), which is adapted to deal with non central χ distributions. We demonstrate on simulated and real data that the extended LMMSE outperforms the original LMMSE on images corrupted by a non central χ noise.

1 Introduction

As many other imaging techniques, Magnetic Resonance Imaging (MRI) is very sensitive to noise in the data. Noise decreases image quality in MRI and particularly in diffusion weighted MRI (dMRI) where the diffusion indicator in the tissues is the measured signal loss. Moreover, noise has higher corrupting effects when using high b-values.

The MR signal is complex with two real and imaginary channels each corrupted by a zero-mean uncorrelated Gaussian noise [1]. It is most common to work on the magnitude of the measured signal in order to avoid the artefacts due to the phase. Because of the non linear square root function used to get it, the measured magnitude will not follow a Gaussian, but a more complex distribution. In case of a multiple-channel acquisition, with a Sum-of-squares (SoS) reconstruction, the noisy magnitude follows a non-central χ (nc-χ) distribution [2]. In the case of Generalized Autocalibrating Partially Parallel Acquisition (GRAPPA) reconstruction, [3] reminds that the noise is non-stationnary, so in this case, the nc-χ hypothesis becomes a good approximation. A particular case of the nc-χ distribution, namely the Rician distribution, appears in case of a single-channel acquisition, or when using the Sensitivity Encoding for Fast MRI (SENSE) algorithm [4]. For high signal to noise ratios (SNRs), a nc-χ distribution can be approximated by a Gaussian distribution, and least square estimators can be used efficiently. At high b-values however, the SNR drops and

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 226–233, 2011.

the signal depicts a systematic bias. Least square estimators cannot be used anymore because the mean operator does not converge to the noise-free magnitude. This makes the nc-χ process more difficult to correct than a Gaussian one. High Angular Resolution Diffusion Imaging (HARDI) or HYbrid Diffusion Imaging (HYDI) are usually performed at those high b-values yielding low SNR data on clinical systems. That is why a correction taking into account the nc-χ process is essential to improve the data quality and analysis.

Numerous techniques, adapted to a Rician noise, were proposed to correct MR data. A review of many of them can be found in [5]. But, most of these methods are not adapted to correct a nc-χ noise. To our knowledge, only two techniques deal with this more complex process. The first one, the analytically exact correction scheme developed in [6] relies on a fixed point formula of SNR and a correction factor to extract the noise-free signal. The second one is based on a framework which transforms the magnitude signals so that they follow Gaussian distributions, in order to use the least square approach on them [7].

In this paper, we propose to extend the method proposed by Aja-Fernández et al in [5] to the correction of nc-χ noise. The original technique relies on a Linear Minimum Mean Square Error estimator (LMMSE) that calculates the noise-corrected signals thanks to local estimations of mean and variance. As the LMMSE corrects the magnitude signals directly, any diffusion model can be applied on the corrected diffusion weighted (DW) data, making the method generic. Moreover, the algorithm is easy and fast to compute. Because of all these reasons, the LMMSE was used in many studies ([8], [9], [10], [11]). That is why an adaptation of the LMMSE to the nc-χ noise model might be of great interest for all works dealing with Parallel MRI data corrupted by a nc-χ noise and where the LMMSE is involved. This implies that the noise estimation must also be adapted, as done in [12]. First, we explain the extension of the LMMSE to the nc-χ distribution. Then, we describe the methods to infer the noise standard deviation. Finally, we validate the technique on simulated and real data, applying it to the popular Diffusion Tensor (DTI) and analytical Q-ball models.

2 Methods

2.1 Extended LMMSE Model

The LMMSE of [5] relies on a simple expression of the squared noise-free magnitude:

$$\hat{S^2} = E(S^2) + \mathrm{Cov}(S^2, M^2)\mathrm{Var}(M^2)^{-1} \times \left(M^2 - E(M^2)\right), \tag{1}$$

with S^2 being the squared noise-free magnitude, $\hat{S^2}$ being its estimation and M being the measured magnitude. This expression is true for any voxel and any diffusion orientation, and to simplify we do not add their corresponding values. $\mathrm{Cov}(S^2, M^2)$ is the covariance of S^2 and M^2 and $\mathrm{Var}(M^2)$ is the variance of M^2. (1) allows to estimate the noise-free magnitude and relies on the knowledge of $E(S^2)$, $\mathrm{Cov}(S^2, M^2)$ and $\mathrm{Var}(M^2)$. Their expressions for a Rician noise are

given in [5]. We propose to follow the same process with the assumption of a multiple-coil acquisition yielding a nc-χ noise. The nc-χ distribution is given by:

$$P(M;S,\sigma,n)=\frac{M}{\sigma^2}\left(\frac{M}{S}\right)^{n-1}\exp\left(-\frac{M^2+S^2}{2\sigma^2}\right)\cdot I_{n-1}\left(\frac{SM}{\sigma^2}\right), \tag{2}$$

with σ being the noise standard deviation of the Gaussian noises present on the complex acquisition channels, n being the number of channels used in the acquisition and I_{n-1} being the $(n-1)$th order modified Bessel function of the first kind. In such multiple-channel acquisitions, the measure of the magnitude of the magnetization is generally performed using an n element phased array antenna. First, we consider the commonly used "Sum-of-Squares" (SoS) reconstruction method to get $M=\left(\sum_{k=1}^{n}\left(S_{r_k}+B_{r_k}\right)^2+\left(S_{i_k}+B_{i_k}\right)^2\right)^{1/2}$, where $S=\left(\sum_{k=1}^{n}S_{r_k}^2+S_{i_k}^2\right)^{1/2}=\left(\sum_{k=1}^{n}S_k^2\right)^{1/2}$. S_{r_k} and S_{i_k} are the real and imaginary parts of the noise-free signal S_k received by the coil k; B_{r_k} and B_{i_k} are the real and imaginary parts of the noise corrupting the signal S_k. They are supposed to be zero-mean, uncorrelated and independant Gaussian noises. Thanks to these assumptions, the variance and covariance terms can be written:

$$\begin{cases}\mathrm{Var}(M^2)=E(M^4)-E(M^2)^2,\\ \mathrm{Cov}(S^2,M^2)=E\left(S^4\right)+2n\sigma^2E(S^2)-E(S^2)E(M^2).\end{cases} \tag{3}$$

Injecting both variance and covariance expressions in the equation (1), we obtain:

$$\hat{S^2}=E(S^2)+\frac{\left(E\left(S^4\right)+2n\sigma^2E[S^2]-E(S^2)E(M^2)\right)}{\left(E(M^4)-E(M^2)^2\right)}\times\left(M^2-E(M^2)\right) \tag{4}$$

Now, if we use the 2nd and 4th order moments of the nc-χ distribution, we have:

$$\begin{cases}E(S^2)=E(M^2)-2n\sigma^2\\ E(S^4)=E(M^4)-4(n+1)\sigma^2E(M^2)+4n(n+1)\sigma^4,\end{cases} \tag{5}$$

and the LMMSE equation (4) for a nc-χ noise can finally be expressed as:

$$\hat{S^2}=\langle M^2\rangle-2n\sigma^2+\left(1-\frac{4\sigma^2\left[\langle M^2\rangle-n\sigma^2\right]}{\langle M^4\rangle-\langle M^2\rangle^2}\right)\times\left(M^2-\langle M^2\rangle\right) \tag{6}$$

As done in [5], under the assumption of local ergodicity, we replaced the expectation $E(\cdot)$ by $\langle\cdot\rangle$, a local spatial mean calculated on a neighborhood. For a single channel acquisition (i.e. $n=1$), (6) simplifies into its Rician form, as expected.

Practically, employing equation (6) only requires a good estimate of the noise standard deviation σ. Obtaining such a good estimate is challenging and we propose to address this point in the following subsection.

2.2 Noise Standard Deviation Estimation

In this subsection, 2 methods are introduced to infer σ; a first one relying on the analysis of the signal present in the background of the data; and a second one that does not depend on the presence of a background in the data.

With the Necessity of a Background. σ can easily be estimated using the information provided by the background of the image. The method proposed by Aja-Fernández *et al* [12], which does not necessarily need a background segmentation, produces the estimate $\hat{\sigma}_{AF}$ of σ such that:

$$\hat{\sigma}_{AF} = \left(\sqrt{2}(n)_{\frac{1}{2}}\right)^{-1} \text{mode}\left(\langle M_{bg}(\mathbf{v})\rangle\right), \tag{7}$$

where $M_{bg}(\mathbf{v})$ is the measured magnitude at the voxel $\mathbf{v}$ in the background region; mode $(\langle M_{bg}\rangle)$ is the distribution mode of the mean of $M_{bg}(\mathbf{v})$; $(x)_n$ is the Pochhammer symbol defined as $(x)_n = \frac{\Gamma(x+1)}{\Gamma(x-n+1)}$ with Γ the Gamma function.

Without the Necessity of a Background. An interesting method to estimate σ without background knowledge was developed by Rajan *et al* [13], but for Rician noise only. We have extended it to nc-χ noise. The technique requires to estimate the variance σ_M^2 and the skewness γ, at each voxel of the image, using:

$$\sigma_M^2 = E(M^2) - E(M)^2 \tag{8}$$

$$\gamma = \left(2E(M)^3 - 3E(M)E(M^2) + E(M^3)\right)/\sigma_M^3, \tag{9}$$

where the expectations $E(\cdot)$ can also be replaced by a local spatial mean $\langle\cdot\rangle$. The method relies on the computation of a local correction factor φ which tunes the proximity of the nc-χ distribution toward the central χ distribution for very low SNRs, or toward the Gaussian distribution for very high SNRs. First, a lookup table $\varphi(\gamma)$ is created between the local correction factor φ and the local skewness γ for a nc-χ distribution with a given n, by varying the value of S and keeping σ constant. In order to build this lookup table, the 3 first nc-χ moments have been calculated from the range of values of S and σ:

$$\begin{aligned}
E(M) &= \sqrt{2}\sigma(n)_{\frac{1}{2}} \times {}_1F_1\left(-\tfrac{1}{2}, n, -\tfrac{S^2}{2\sigma^2}\right)\\
E(M^2) &= 2n\sigma^2 + S^2\\
E(M^3) &= \tfrac{2}{3}\sqrt{2}\sigma^3(n)_{\frac{1}{2}} \times {}_1F_1\left(-\tfrac{3}{2}, n, -\tfrac{S^2}{2\sigma^2}\right),
\end{aligned}$$

with ${}_1F_1\,(a, b, c)$ being the confluent hypergeometric function. Then, γ is obtained by injecting the 3 expressions above in (9). In the same manner, the local variance σ_M is computed using (8). Finally, φ is calculated from $\varphi = \sigma^2/\sigma_M^2$. Once the $\varphi(\gamma)$ lookup table is computed and starting from the local estimates σ_M^2 (8) and γ (9), the final noise standard deviation can be estimated from:

$$\hat{\sigma}_R = \sqrt{\text{mode}\left(\varphi \times \sigma_M^2\right)}, \tag{10}$$

where mode $\left(\varphi \times \sigma_M^2\right)$ is the distribution mode of $\varphi \times \sigma_M^2$ calculated on the image.

3 Results and Discussion

We used our extended LMMSE method on both simulated and real DW data perturbed by a nc-χ noise, and we compared it to the original LMMSE method adapted to a Rician noise.

3.1 Standard Deviation Estimations

First, we tested the estimation methods described in section 2.2 on a simulated T1 data taken from the BrainWeb database [14]. The original image is considered as the ground truth. We added a high nc-χ noise using $n = 4$ channels that correspond to the standard case of the 12 channel head coil antenna available on TimTrio systems in which the 12 channels are coupled in 4 associations of 3 channels, and we used a standard deviation $\sigma = 70.0$ leading to a very unfavorable situation. We made the same experiment on a synthetic DW field depicting crossing fiber bundles simulated using a Gaussian mixture model with nc-χ noise addition ($n = 4$, $\sigma = 30.0$). Table 1 shows that very good estimations can be found by the first and second methods, with the knowledge of the background-signal. Without it, the second method still gives good results.

Table 1. Different σ estimations. $\hat{\sigma}_{AF}$ is obtained with the method of [12] using the background (bg) of the image. $\hat{\sigma}_R^{w/bg}$ and $\hat{\sigma}_R^{wo/bg}$ are obtained with the method using equation (10), performed respectively with and without the bg.

Simulated data	Ground truth σ	w/bg necessity	wo/bg necessity
T1 image	70.0	$\hat{\sigma}_{AF} = 70.0$	$\hat{\sigma}_R^{w/bg} = 70.1$ and $\hat{\sigma}_R^{wo/bg} = 73.4$
DW field	30.0	$\hat{\sigma}_{AF} = 29.5$	$\hat{\sigma}_R^{w/bg} = 30.0$ and $\hat{\sigma}_R^{wo/bg} = 30.9$

3.2 Simulated Data

We used both simulated data already mentioned in section 3.1. Using a $5 \times 5 \times 5$ neighborhood, we tested both versions of the LMMSE: the LMMSE correction introduced in [5], adapted for a Rician noise (Rice LMMSE) and the extended LMMSE correction introduced in this work following equation (6), adapted for a nc-χ noise (nc-χ LMMSE). Before the computation of the LMMSEs, σ was estimated using (7), with $n = 1$ for the Rice LMMSE and with $n = 4$ for the nc-χ LMMSE. Concerning the DW field, the Orientation Distribution Functions (ODFs) were processed using the analytical Q-ball model [15], with the maximum spherical harmonic (SH) order $N = 8$ and the Laplace-Beltrami regularization factor $\lambda = 0.006$. In comparison to the Rice LMMSE, Fig. 1 shows that the nc-χ LMMSE presents much better performance: good noise cleaning and low smoothing effects. This reveals the importance of taking the correct noise model into account in the LMMSE tool. To quantitatively validate this result, we measured the quadratic errors:

$$\epsilon = \sum_{\mathbf{v}=1}^{M} \left(\sum_{j=1}^{N} \left(\mathbf{C}_{\mathbf{v}}^{\mathbf{DWI}}(j) - \tilde{\mathbf{C}}_{\mathbf{v}}^{\mathbf{DWI}}(j) \right)^2 \right), \tag{11}$$

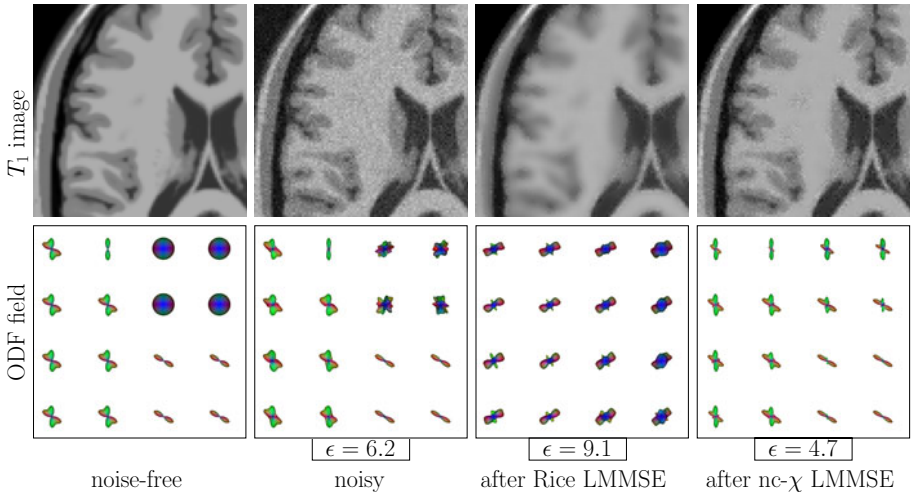

Fig. 1. Comparison between both LMMSEs on simulations of a T1 image and of ODFs

where $\mathbf{C_v^{DWI}}$ is the SH coefficient vector of the noise-free DW signal of voxel $\mathbf{v}$ (M is the number of voxels in the field) and $\tilde{\mathbf{C}}_\mathbf{v}^{\mathbf{DWI}}$ is either the noisy or the corrected coefficient vector. Whereas ϵ increased after the Rice LMMSE, it dramatically decreased after the nc-χ LMMSE.

3.3 Real DW Data

DW data were collected on a Magnetom Tim Trio 3T MRI system (Siemens Medical Solutions, Erlangen, Germany), using a gradient sampling scheme of 60 orientations uniformly distributed over a sphere at $b = 4500 s.mm^{-2}$. A further reference T2 weighted volume was acquired at $b = 0 s.mm^{-2}$. The acquisition parameters were as follows: $T_E/T_R = 116ms/14s$, field of view $FOV = 198mm$, matrix 128×128, resolution $1.7 \times 1.7 \times 1.7 mm^3$, GRAPPA factor of 2, read bandwidth $RBW = 1628Hz/pixel$ and a number of coils of 12 in an architecture of 4×3 coupled coils, so $n = 4$ in this case. We tested our method on the DTI model (color-encoded (RGB) map) and on the analyticall Q-ball model (Generalized Fractional Anisotropy (GFA) and ODFs) with $N = 8$ and $\lambda = 0.006$. Fig. 2 shows the results obtained for the Rice LMMSE and the nc-χ LMMSE. As for the simulated data, the GFA and RGB maps emphasize the importance of using the nc-χ noise model, yielding a preservation of the fine details and a limited smoothing effect. On the ODF maps, the Rice LMMSE seems to provide better results, but this is only due to a visual perception of a better coherence induced by an oversmoothing effect. Concerning the ODFs in the red circle, the Rice LMMSE completely overlooks the raw information in contrary to the nc-χ LMMSE. The latter has a better behavior since it corrects the ODFs from the influence of noise while keeping the right information.

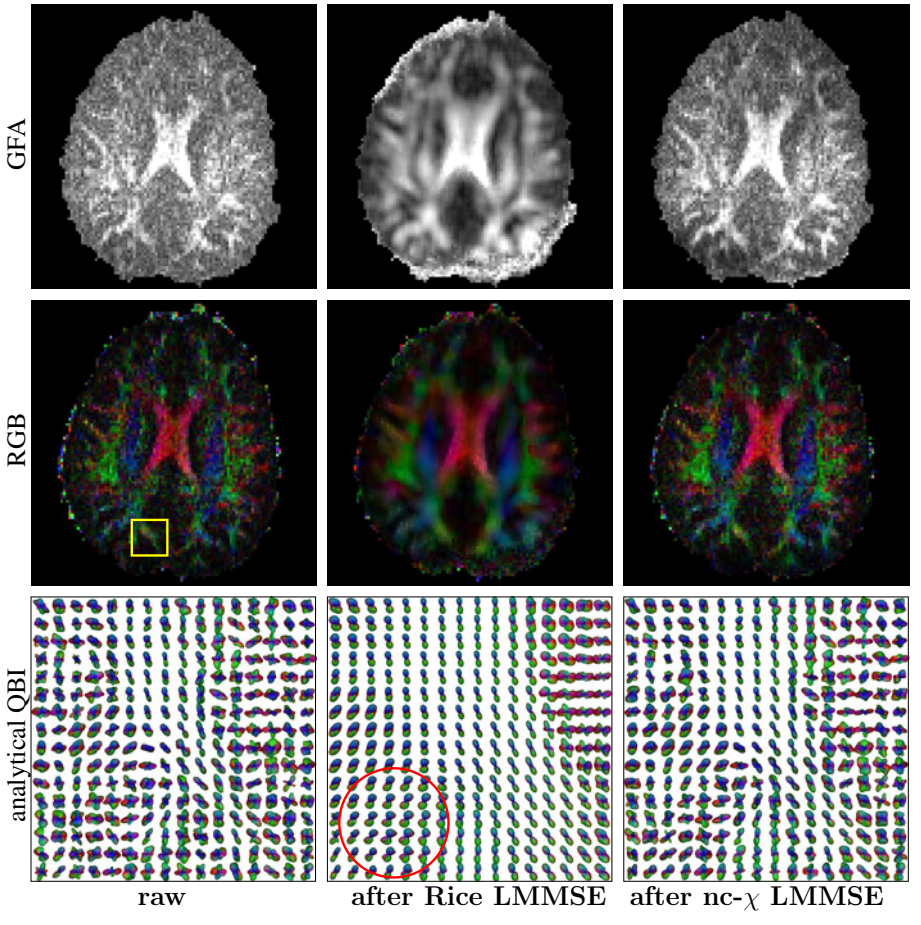

Fig. 2. Comparison between both LMMSEs on GFA, RGB and ODF maps. The ODFs are computed in the yellow ROI shown on the RGB map.

4 Conclusion

We have proposed an extension of the LMMSE to nc-χ noise correction. The comparison between the Rice and the nc-χ LMMSEs on synthetic and real data successfully assesses the choice of the nc-χ LMMSE in case of a multiple-coil acquisition yielding a nc-χ noise with $n > 1$. Whereas the Rice LMMSE over-smoothes the images, the nc-χ LMMSE performs a clearly visible correction with respect to the details, even in very defavorable conditions as at $b = 4500 s.mm^{-2}$. The real data we used were obtained with GRAPPA reconstruction, and yet, it has been recently discussed in [3] that an effective number of channels, as well as an effective variance of noise, must be calculated to get a nc-χ noise model that better fits the data. Both effective calculated parameters are not stationnary and should be evaluated at each voxel of the image. Future work will be done

to apply this effective parameters' calculations to the nc-χLMMSE correction. Last, our method can easily be extended to its recursive form as it was done in [5] for the Rice LMMSE, but this was beyond the scope of this work.

Acknowledgements. The authors thank Mr. Rajan for discussions about the noise estimation method developed in [13] and acknowledge the CONNECT grant and MCIN for grant TEC2010-17982.

References

1. Henkelman, R.M.: Measurement of signal intensities in the presence of noise in MR images. Med. Phys. 12, 232–233 (1985)
2. Constantinides, C.D., Atalar, E., McVeigh, E.R.: Signal-to-Noise Measurements in Magnitude Images from NMR Phased Arrays. Magn. Reson. Med. 38, 852–857 (1997)
3. Aja-Fernández, S., Tristán-Vega, A., Hoge, W.S.: Statistical noise analysis in GRAPPA using a parametrized noncentral chi approximation model. Magn. Reson. Med. 65, 1195–1206 (2011)
4. Dietrich, O., Raya, J.G., Reeder, S.B., Ingrisch, M., Reiser, M.F., Schoenberg, S.O.: Influence of multichannel combination, parallel imaging and other reconstruction techniques on MRI noise characteristics. MRI 26, 754–762 (2008)
5. Aja-Fernández, S., Niethammer, M., Kubicki, M., Shenton, M.E., Westin, C.F.: Restoration of DWI Data Using a Rician LMMSE Estimator. IEEE Trans. Med. Imag. 27, 1389–1403 (2008)
6. Koay, C.G., Basser, P.J.: Analytically exact correction scheme for signal extraction from noisy magnitude MR signals. J. Magn. Reson. 179, 317–322 (2006)
7. Koay, C.G., Özarslan, E., Basser, P.J.: A signal transformational framework for breaking the noise floor and its applications in MRI. J. Magn. Reson. 197, 108–119 (2009)
8. Krissian, K., Aja-Fernández, S.: Noise-driven anisotropic diffusion filtering of MRI. IEEE Trans. Image Process. 18, 2265–2274 (2009)
9. Caan, M.W.A., Khedoe, G., Poot, D., den Dekker, A., Olabarriaga, S., Grimbergen, K., van Vliet, L., Vos, F.: Adaptive noise filtering for accurate and precise diffusion estimation in fiber crossings. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. 6361, pp. 167–174. Springer, Heidelberg (2010)
10. Tristán-Vega, A., Aja-Fernández, S.: DWI filtering using joint information for DTI and HARDI. Med. Image Anal. 14, 205–218 (2010)
11. Brion, V., Kezele, I., Riff, O., Descoteaux, M., Mangin, J.F., Poupon, C., Poupon, F.: Real-time Rician noise correction applied to real-time HARDI and HYDI. In: Workshop CDMRI of MICCAI, pp. 2–13 (2010)
12. Aja-Fernández, S., Tristán-Vega, A., Alberola-López, C.: Noise estimation in single- and multiple-coil magnetic resonance data based on statistical models. MRI 27, 1397–1409 (2009)
13. Rajan, J., Poot, D., Juntu, J., Sijbers, J.: Noise measurement from magnitude MRI using local estimates of variance and skewness. Phys. Med. Biol. 55, N441–N449 (2010)
14. Collins, D., Zijdenbos, A., Kollokian, V., Sled, J., Kabani, N., Holmes, C., Evans, A.: Design and construction of a realistic digital brain phantom. IEEE Trans. Med. Imag. 17, 463–468 (1998)
15. Descoteaux, M., Angelino, E., Fitzgibbons, S., Deriche, R.: Regularized, Fast, and Robust Analytical Q-Ball Imaging. Magn. Reson. Med. 58, 497–510 (2007)

HARDI Based Pattern Classifiers for the Identification of White Matter Pathologies

Luke Bloy[1], Madhura Ingalhalikar[1], Harini Eavani[1,2], Timothy P.L. Roberts[3], Robert T. Schultz[2], and Ragini Verma[1]

[1] Section of Biomedical Image Analysis, Department of Radiology, University of Pennsylvania
[2] Center for Autism Research, Children's Hospital of Philadelphia
[3] Lurie Family Foundation's MEG Imaging Center, Department of Radiology, Children's Hospital of Philadelphia
{Luke.Bloy,Ragini.Verma}@uphs.upenn.edu

Abstract. The paper presents a method for creating abnormality classifiers from high angular resolution diffusion imaging (HARDI) data. We utilized the fiber orientation distribution (FOD) diffusion model to represent the local WM architecture of each subject. The FOD images are then spatially normalized to a common template using a non-linear registration technique. Regions of homogeneous white matter architecture (ROIs) are determined by applying a parcellation algorithm to the population average FOD image. Orientation invariant features of each ROI's mean FOD are determined and concatenated into a feature vector to represent each subject. Principal component analysis (PCA) was used for dimensionality reduction and a linear support vector machine (SVM) classifier is trained on the PCA coefficients. The classifier assigns each test subject a probabilistic score indicating the likelihood of belonging to the patient group. The method was validated using a 5 fold validation scheme on a population containing autism spectrum disorder (ASD) patients and typically developing (TD) controls. A clear distinction between ASD patients and controls was obtained with a 77% accuracy.

Keywords: Diffusion Imaging, HARDI, FOD, Classification, SVM.

1 Introduction

High dimensional pattern classification methods like support vector machines (SVM) identify brain abnormality patterns that enhance group separability while quantifying the degree of pathological abnormality associated with each individual. This paper proposes a HARDI based pattern classification framework that creates abnormality classifiers using information concerning white matter (WM) architecture derived from homogeneous WM regions. This classification framework not only elucidates regions that are affected by pathology but also assigns each individual with a score indicating the degree of abnormality. Such a score may prove useful in conjunction with other clinical measures as a diagnostic tool

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 234–241, 2011.

to predict the extent of disease as well as act as a biomarker to assess disease progression or treatment efficacy.

While the bulk of classification work concerning medical imaging analysis has centered on structural imaging[4,5], the application of classifiers to diffusion imaging is relatively recent [9,7,14,3] , while there is no literature using HARDI data. Caan et al [3] performed principal component analysis (PCA) and linear discriminant analysis (LDA) on linear and fractional anisotropy (FA) images, derived from diffusion tensor images, to classify schizophrenia patients. Wang et al. [14] used a k nearest neighbor (kNN) classifier trained on full brain FA images, while Ingalhalikar et al [7] used non-linear SVM trained on regional FA and diffusivity for classification in autism spectrum disorder (ASD) and schizophrenia. Finally, Lange et al [9] used quadratic discriminant analysis and SVM to perform hypothesis driven classification in an ASD population based on a-priori regions.

While these methods have had success, they are limited by the use of the diffusion tensor (DTI) data model which is known to be ineffective in modeling regions of complex white matter, i.e. multiple fibers with different orientations, different partial volume fractions. High dimensional diffusion data models, such as the the fiber orientation distribution function (FOD) [12], have been developed to make use of new acquisition protocols which acquire HARDI data. These new data models are better able to model complex WM regions and thus should prove useful for studying WM pathology. While Schnell et al [11], proposed a tissue segmentation method based on HARDI classification, we believe this is the first work to utilize HARDI data models to perform subject classification between healthy controls and a disease population

We propose a classification framework that makes use of the FOD HARDI data model to extract orientation invariant features from regions of homogeneous WM architecture. These regions are determined from a WM parcellation algorithm applied to the average FOD image of the control population. Principal component analysis (PCA) is used for dimensionality reduction and a linear support vector machine (SVM) is used to perform the classification. The linearity of the framework allows for the examination of the SVM decision weights in the original feature space to aid in interpretability of the classification results. The framework is applied and cross-validated using a 5-fold cross-validation paradigm using a dataset comprised of children diagnosed with Autism Spectrum Disorder (N=23) and typically developing controls (N=22). The high accuracy, specificity and sensitivity (all $\sim$77%) establish the applicability of classifier scores for aiding diagnosis and prognosis.

2 Methods

The process of training and validating a classification framework like the one proposed here consists of a number of steps. Namely feature extraction and dimensionality reduction/feature selection, followed by training and cross-validation.

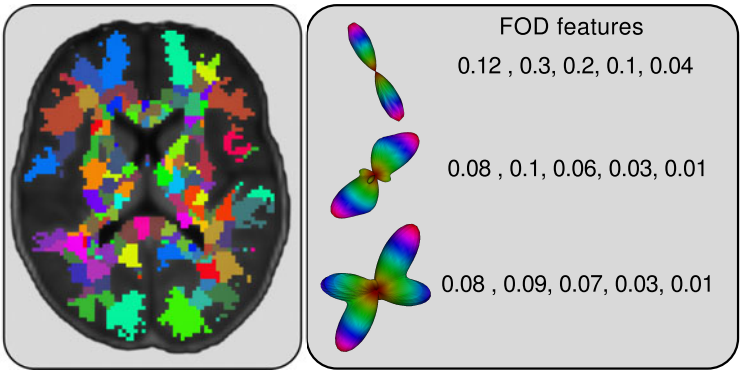

Fig. 1. Feature Extraction: Spatial regions of homogeneous WM are determined by parcelating the population average FOD image into regions of low variance, shown on the left. For each region the mean FOD is computed from which 5 orientation invariant FOD features are computed. The mean FOD and corresponding feature vectors are shown, on the right, for regions in the corpus callosum, as well as in anterior and posterior complex WM.

2.1 Feature Extraction

For each subject, feature extraction must be performed to extract a salient representation of the subject that will serve as a means of comparison. We are interested in identifying pathologies that manifest as abnormalities in the WM. We therefore concentrate on extracting features from spatially localized regions of homogeneous WM. This process entails modeling the diffusion process in each voxel using the fiber orientation distribution function (FOD) followed by spatially normalizing all subjects into a common template reference frame. Once all subjects are spatially normalized an average FOD image of the control population is computed and used to determine homogeneous WM regions of interest (ROIs). Finally orientation invariant features from each ROI are extracted and concatenated to build a feature vector representation of each subject.

Our process begins by using constrained spherical deconvolution[12] to compute an FOD image for each subject. The FOD diffusion model represents the voxel's DW-MRI signal as the spherical convolution of the FOD and the DW signal that would be measured for a single fiber bundle aligned along the z-axis. Information concerning both the orientation and partial volume of any constituent fiber bundles present in a voxel is represented via the FOD [12]. In this work we use the order 8 real spherical harmonic (RSH) representation of the FOD, thus each voxel is represented by the 45 RSH coefficients of its FOD.

Each FOD image is then spatially normalized to a template FOD image, an individual normal subject in this case, using a diffeomorphic demons-based FOD registration algorithm [1]. The process of spatial normalization defines a spatial and anatomical correspondence between subjects that allows a set of ROIs to be determined and to confidently represent corresponding areas of anatomy in each subject.

A population average FOD image, in the template space, is then computed from the registered FOD images of the normal subjects of the population. From this average FOD image we determine a set of WM ROIs with a homogeneous WM architecture using the methods described in [2]. This method utilizes normalized cuts spectral clustering to partition the WM into spatially connected regions which have a small FOD variance. Using this method we determined 883 WM ROIs, with an FOD variance below 0.08 (Shown in Figure 1).

With spatial ROIs determined, we compute the feature vector representation for each subject by first determining the mean FOD in each ROI. From the mean FOD we compute the L2 norm of the RSH coefficients in each order (l level). Because rotations of the FOD will transfer energy within an order but not from one order to another the collection of these L2 norms provides an orientation invariant representation [6] of the mean FOD. Thus for an order 8 FOD model, as the one used here, each ROI is represented by the p_l for l in $0, 2, 4, 6, 8$. Where $p_l = \sqrt{\sum_{m=-l}^{m=l} \tilde{f}^2}$ is defined in terms of the RSH coefficients of the mean FOD ($\tilde{f}$) of the ROI. These 5 features are computed for each of the 883 WM ROIs and concatenating them yields a representation of each subject's WM by a 4415 element feature vector.

2.2 Dimensionality Reduction

When using high dimensional feature representations of subjects within a classification framework such as this, a critical task is that of feature selection or dimensionality reduction. Particularly when using small sample sizes, such as those commonly available to medical imaging studies, the reduction of the feature space dimensionality is essential to avoid over-fitting and for minimizing classification error. In this work, we use principal component analysis (PCA) to obtain a concise basis of the original feature space while still accounting for the majority of the population variance ($> 90\%$). The PCA features are linear combinations of the original FOD features, and the PCA basis describes an invertible linear operator which relates the two feature spaces. This allows for new subjects to be mapped into the PCA feature space and for vectors in the PCA feature space, such as the decision boundaries, to be mapped back into the original FOD feature space.

2.3 Support Vector Machine Training and Cross-Validation

Linear support vector machines (LSVMs) determine a hyperplane in the feature space, defined by the PCA features in this case, which optimally separates the dataset into patients and controls. Once the SVM has been trained, a new test subject (x) can be labeled, assigned to either the control or patient group, based on the distance between the subject and separating hyperplane. This distance is given by the decision function: $f(x) = w^t\, x + b$, where w is the SVM feature weights, describing the contribution of each feature, and b is the bias of the hyperplanes from the origin. The distance is used by the classifier to determine,

5 Fold Validation Results

Accuracy	78%
Specificity	77%
Sensitivity	78%
ROC area	0.81

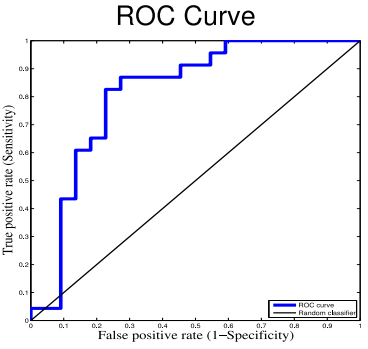

Fig. 2. The classification framework was applied to an ASD dataset consisting of 23 ASD patients and 22 controls. Results from a five-fold cross-validation paradigm are shown on the left, while the receiver operating characteristic (ROC) curve is shown on the right.

via Platt's method [10], the probabilistic score for the subject and the subject is labeled, either patient or control, based on the sign of the score.

The goal of a classification method is to accurately and robustly classify unseen test subjects. In medical imaging studies the small sample size (N=45 in this case) often prohibits the division of the dataset into a single training and testing dataset, each of which accurately represent the entire dataset. For this reason we use a stratified 5-fold cross-validation method to validate our framework. This entails partitioning the original dataset into 5 pieces or folds each of which contain roughly the same proportion of patients and controls. A fold is chosen to serve as the test dataset while the classifier is trained on the remainder of the dataset. This process is repeated until each fold has acted as test data, yielding an abnormality score for each subject as well as a classification accuracy for each fold.

The overall classification method can then be evaluated based on the average accuracy achieved across the 5 folds, as well as by examining the receiver operating characteristic (ROC) curve. The ROC curve is a plot of the sensitivity vs (1-specificity) of the classifier as the discrimination threshold is varied.

3 Application of Classification Framework to a Clinical Population

In this work we apply our framework to the problem of classifying a population of children diagnosed with Autism Spectrum Disorder (ASD). The dataset consisted of 22 typically developing controls (TDC) and 23 ASD patients. Whole brain HARDI was acquired using a Siemens 3T Verio MRI scanner using a spin-echo, echo-planar imaging sequence (TR/TE=14.7s/110ms, 2mm isotropic voxels, $b = 3000s/mm^2$, number of diffusion directions=64, 2 b0 images). Total acquisition time was 18 minutes per subject.

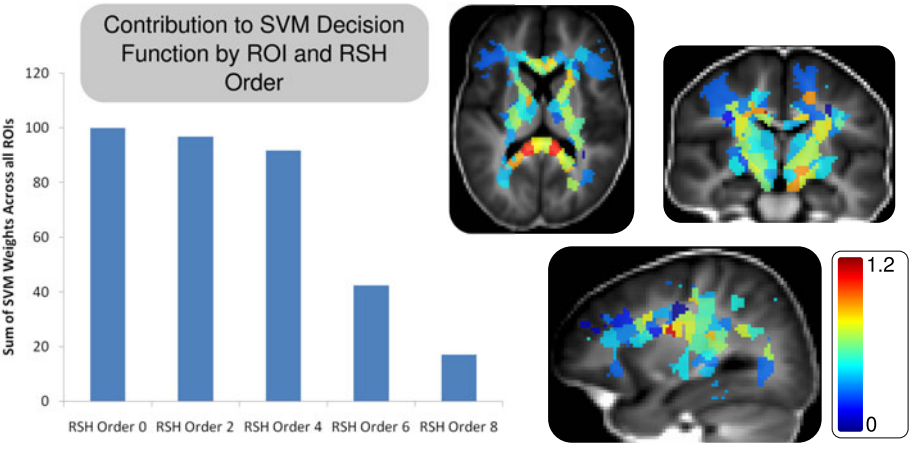

Fig. 3. The SVM weights are mapped into FOD feature space yielding the contribution of each feature to the classification score. The contribution of each RSH order (shown on Left) is obtained by summing the individual contributions across all of the ROIS. Similarly, by averaging the contributions across the rotation invariant features the importance of each spatial region to the classification score (Shown on Right) can be obtained . Higher contributions are indicative a larger group difference within that feature.

The diffusion weighted images (DWI) were first filtered using a joint linear minimum mean squared error filter for removal of Ricean noise [13]. Eddy current correction was then performed using affine registration of each DWI volume to the unweighted b0 image [8]. The feature extraction process, described in section 2.1, was then applied using a 12-year old male TDC subject as the registration template as the age of this subject was closest to the population average, yielding a labeled dataset of 45 subjects each represented by a 4415 element feature vector.

The dataset was then divided into 5 folds and cross-validation was performed using the folds as test datasets and the remainder of the dataset for training. For each validation procedure (i.e. each fold) PCA was applied to the training dataset using a variance threshold of 90% and a linear SVM was trained on the resulting PCA coefficients. The SVM classifier was then applied to each subject in the test dataset yielding a classifier score. Additionally the SVM weights, which describe the contribution of each PCA coefficient to the classifier score, were computed and mapped back into the original feature space. In this way a classifier score and a predicted label for each subject in the dataset was obtained, as were 5 measures of the classifier weights mapped into the original feature space.

Figure 2 shows the classification results. The average classification accuracy was 78% (10 subjects misclassified) with a specificity of 77% (low type I error) and a sensitivity of 78% (low type II error). The receiver operating characteristic (ROC) curve shows a steep balanced curve with an area under the curve of 0.81.

These accuracy numbers are in-line with the existing DTI based classification methods [9,7,14,3], which achieve accuracies in the 70-95% range, where 95% mark was obtained using a hypothesis driven feature selection process designed specific for ASD[9], as opposed to the whole brain approach taken here. However such a comparison, based on published accuracies, is greatly hindered by the fact that each study utilizes a different dataset, with different patient characteristics, and in this case multiple diseases (ASD and schizophrenia).

In addition to computing accurate classifier scores our framework has the capability of determining the degree that each of the originally extracted features contribute to the classification score. By averaging the classifier weights in the original feature space, we obtained the mean contribution of each feature to the classification score. Our original feature space consisted of 5 orientation invariant features, one for each order of the RSH expansion, derived from 883 spatial ROIs. The relative contribution of each orientation invariant feature was determined by summing the contributions of that feature across all of the ROIs. Similarly the contribution of each ROI was determined by the sum across the orientation invariant features. These are shown in figure 3.

By examining the contributions of each orientation invariant feature, figure 3-left, we see that the first 3 RSH orders (0, 2 and 4) are predominant in determining the classification score. This suggests that the higher angular frequency information, contained in the higher RSH orders, is perhaps more variable across the population or inherently less reliable. The regional contributions to the classification score shows large contributions from portions of internal capsule (figure 3-right) as well as from the splenium of the corpus callosum, regions that have been previously implicated in ASD. While these results suggest that regional contributions may be useful in localizing WM areas that are affected in ASD a full investigation of these results is beyond the scope of this paper.

4 Conclusion

We have presented a classification methodology based on regional measures of white matter architecture and fidelity as derived from the FOD diffusion data model. The FOD model when coupled with an atlas-free parcellation algorithm yields a physiologically interpretable feature representation, as it contains the orientation and relative proportion of the underlying anatomical fibers in homogeneous WM regions. We demonstrate that this feature representation, when coupled with PCA and a linear SVM yields robust and accurate classification results in an ASD population. In addition to producing a classification score, which may aid in diagnosis, this framework provides the SVM feature weightings which identify the spatial regions and the FOD features contribute to the score. The feature weights elucidate the regions that are affected in the patient population, providing possible insight into the pathology as well as suggesting future directions for the development of hypothesis based classifiers and studies.

While classification based on DTI features has been attempted this is the first classification work that utilizes features derived from HARDI data models to perform patient classification. The high specificity, sensitivity and accuracy demonstrate the feasibility of HARDI based patient classification.

References

1. Bloy, L., Verma, R.: Demons registration of high angular resolution diffusion images. In: 2010 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, pp. 1013–1016 (2010)
2. Bloy, L., Ingalhalikar, M., Verma, R.: Neuronal white matter parcellation using spatially coherent normalized cuts. In: 2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro (2011)
3. Caan, M.W.A., Vermeer, K.A., van Vliet, L.J., Majoie, C.B.L.M., Peters, B.D., den Heeten, G.J., Vos, F.M.: Shaving diffusion tensor images in discriminant analysis: a study into schizophrenia. Med. Image Anal. 10(6), 841–849 (2006)
4. Ecker, C., Rocha-Rego, V., Johnston, P., Mourao-Miranda, J., Marquand, A., Daly, E.M., Brammer, M.J., Murphy, C., Murphy, D.G., Consortium, M.R.C.A.: Investigating the predictive value of whole-brain structural mr scans in autism: a pattern classification approach. Neuroimage 49(1), 44–56 (2010)
5. Fan, Y., Shen, D., Gur, R.C., Gur, R.E., Davatzikos, C.: Compare: classification of morphological patterns using adaptive regional elements. IEEE Trans. Med. Imaging 26(1), 93–105 (2007)
6. Frank, L.R.: Characterization of anisotropy in high angular resolution diffusion-weighted mri. Magnetic Resonance in Medicine 47(6), 1083–1099 (2002)
7. Ingalhalikar, M., Kanterakis, S., Gur, R., Roberts, T.P.L., Verma, R.: DTI based diagnostic prediction of a disease via pattern classification. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. 6361, pp. 558–565. Springer, Heidelberg (2010)
8. Jezzard, P., Barnett, A.S., Pierpaoli, C.: Characterization of and correction for eddy current artifacts in echo planar diffusion imaging. Magn. Reson. Med. 39(5), 801–812 (1998)
9. Lange, N., Dubray, M.B., Lee, J.E., Froimowitz, M.P., Froehlich, A., Adluru, N., Wright, B., Ravichandran, C., Fletcher, P.T., Bigler, E.D., Alexander, A.L., Lainhart, J.E.: Atypical diffusion tensor hemispheric asymmetry in autism. Autism Res. (December 2010)
10. Platt, J.: Probabilistic outputs for support vector machines and comparison to regularized likelihood methods. In: Advances in Large Margin Classifiers (1999)
11. Schnell, S., Saur, D., Kreher, B.W., Hennig, J., Burkhardt, H., Kiselev, V.G.: Fully automated classification of hardi in vivo data using a support vector machine. Neuroimage 46(3), 642–651 (2009)
12. Tournier, J.D., Calamante, F., Connelly, A.: Robust determination of the fibre orientation distribution in diffusion mri: Non-negativity constrained super-resolved spherical deconvolution. NeuroImage 35(4), 1459–1472 (2007)
13. Tristán-Vega, A., Aja-Fernández, S.: Dwi filtering using joint information for dti and hardi. Med. Image Anal. 14(2), 205–218 (2010)
14. Wang, P., Verma, R.: On classifying disease-induced patterns in the brain using diffusion tensor images. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part I. LNCS, vol. 5241, pp. 908–916. Springer, Heidelberg (2008)

Detrend-Free Hemodynamic Data Assimilation of Two-Stage Kalman Estimator

Hu Zhenghui[1] and Shi Pengcheng[2]

[1] State Key Laboratory of Modern Optical Instrumentation, Zhejiang University, Hangzhou 310027, China
[2] B. Thomas Golisano College of Computing and Information Sciences, Rochester Institute of Technology Rochester, NY 14623, USA
zhenghui@zju.edu.cn

Abstract. Spurious temporal drift is abundant in fMRI data, and its removal is a critical preprocessing step in fMRI data assimilation due to the sparse nature and the complexity of the data. Conventional data-driven approaches rest upon specific assumptions of the drift structure and signal statistics, and may lead to inaccurate results. In this paper we present an approach to the assimilation of nonlinear hemodynamic system, with special attention on drift. By treating the drift variation as a random-walk process, the assimilation problem was translated into the identification of a nonlinear system in the presence of time varying bias. We developed two-stage unscented Kalman filter (UKF) to efficiently handle this problem. In this framework the assimilation can implement with original fMRI data without detrending preprocessing. The fMRI responses and drift were estimated simultaneously in an assimilation cycle. The efficacy of this approach is demonstrated in synthetic and real fMRI experiments. Results show that the joint estimation strategy produces more accurate estimation of physiological states, fMRI response and drift than separate processing due to no assumption of structure of the drift that is not available in fMRI data.

1 Introduction

Over the past decade, the neuroimaging community has witnessed an explosive rise in research that melds observed fMRI data with hemodynamic response models to generate accurate forecasts on underlying physiological states and/or parameters [4,6]. In general, a comprehensive assimilation scheme is suitable for most nonlinear phenomena and distinct spatio-temporal scales. Even though the modeling may differ, the approach chosen is general [9].

However, there is still something unique to fMRI data that needs additional concern. The original fMRI measurements contain abundant slowly varying drift. These undesirable drift may be caused by instrumental instability, spontaneous head movement, as well as aliasing of physiological pulsations. The amount and direction of drift is difficult to predict, whereas estimation and removal of these components have a strong impact on assimilation performance and ensuing statistical analysis due to the inherent low signal to noise ratio (SNR) of fMRI

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 242–250, 2011.

data. And, from a frequency perspective, the nuisance drift is also difficult to be discriminated from *true* fMRI response that often has similar low frequency components. Detrending therefore is an important preprocessing step in fMRI data analysis.

Although the conventional data driven methods have provided good performance, and there have been numerous attempts on drift function design and optimization for their improvement, it remains true that all data-driven detrending methods are based on specific assumptions of the drift curve and signal statistics [1]. In practice, however, the drift has complicated structures due to the interaction of various possible sources. Real drift is noisy, not a smooth curve supposed by data driven methods. In this context, such assumptions may not capture the structure of drift well. As a result, the quality of the assimilation will be subject to degradation imposed by these assumption on drift formulation.

Since the assimilation processing always occurs in activation areas, the actual trended fMRI signal can be explained as a linear combination of the hemodynamic response and drift. Considering a more general case, in which both process and observation are biased,

$$\dot{\mathbf{x}} = f(\mathbf{x}, \boldsymbol{\theta}) + \mathbf{B}\mathbf{b} + \mathbf{v} \qquad \mathbf{v} \sim N(0, \mathbf{Q}) \tag{1}$$

$$\mathbf{y} = h(\mathbf{x}, \boldsymbol{\theta}) + \mathbf{C}\mathbf{b} + \mathbf{w} \qquad \mathbf{w} \sim N(0, \mathbf{R}) \tag{2}$$

where f is nonlinear hemodynamic response function, h is nonlinear measurement function that describes the transformation from physiologic states to the fMRI observation at a given brain region, $\mathbf{x}$ represents physiological state, $\boldsymbol{\theta}$ is physiological parameters, $\mathbf{B}$, $\mathbf{C}$ are time-variant coefficient matrices. $\mathbf{b}$ is biases state and $\mathbf{y}$ is the measurement. The process and measurement noise $\mathbf{v}$,$\mathbf{w}$ are zero-mean white Gaussian noise. It is noted that the drift is low frequency, slowly varying component, and can be treated as stochastic variation in assimilation procedure. In this sense, the assimilation problem naturally fall into the problem of estimating the state variables of a nonlinear system in the presence of unknown random bias.

Eq. (1) and (2) form a state-space-like representation for the fMRI data assimilation problem, with (1) describing the physiological process and (2) expressing the observation. However, the bias terms make the problem in a nonstandard state-space formulation. A convenient approach to this problem is to concatenate the original state vector and the bias terms into a higher-order state vector. The filter then implements on the augmented state space. However, it may suffer from computational burden and numerical problems when state dimensions are large because of solving the inverse of the covariance matrix. This method therefore is reasonably effective only when the number of states is relatively small. In 1969, Friedland introduced a Kalman estimation concept applicable to the case in which state variables contain unknown, constant bias [3]. The resulting two-stage estimator has since received considerable attention because of its inherent computational efficiency and stability. It was also extended to handle nonlinear

systems and more general cases (e.g. time varying bias) [5]. However, most of them cast the standard extended Kalman filter (EKF) into a two-stage structure to deal with nonlinearity, therefore easily lead to the problem of numerical stability due to the linearization of the nonlinear system in practical application and their method is reliable only for near linear system on the time scale of the updates. The unscented Kalman filter (UKF) has been developed to address the deficiencies of EKF [8]. In this paper, we develope a two-stage unscented Kalman estimator for the nonlinear system in the presence of unknown random bias. In this framework, the fMRI response, system states associated with physiological variables and the drift can be forecasted at the same assimilation procedure, thereby realizing detrend-free assimilation approach to original fMRI data without detrending preprocessing.

2 Unscented Kalman Estimator

Two-stage estimation separates the estimation of the bias from that of the dynamic state, thereby reducing the dimensionality of states involved in the computations. Two separate, uncoupled filter run in parallel to generate the optimal estimate of the bias and of the "bias-free" state. A thorough description of the linear two-stage estimator can be found in (not presented here for lack of space) [3]. For a nonlinear dynamic system, the bias-free estimator can be immediately replaced with the standard UKF formulation [8]:

$$\mathcal{X}_{k-1} = [\hat{\mathbf{x}}_{k-1}\ \hat{\mathbf{x}}_{k-1} + \eta\sqrt{\mathbf{P}_{k-1}}\ \hat{\mathbf{x}}_{k-1} - \eta\sqrt{\mathbf{P}_{k-1}}] \tag{3}$$

$$\mathcal{X}_{k|k-1} = F[\mathcal{X}_{k-1}, \mathbf{u}_{k-1}] \tag{4}$$

$$\hat{\mathbf{x}}_k^- = \sum_{i=0}^{2L} W_i^{(m)} \mathcal{X}_{i,k|k-1} \tag{5}$$

$$\mathbf{P}_k^- = \sum_{i=0}^{2L} W_i^{(c)} [\mathcal{X}_{i,k|k-1} - \hat{\mathbf{x}}_k^-][\mathcal{X}_{i,k|k-1} - \hat{\mathbf{x}}_k^-]^T + \mathbf{R}^{\mathbf{v}} \tag{6}$$

$$\mathcal{Y}_{k|k-1} = \mathbf{H}[\mathcal{X}_{k|k-1}] \tag{7}$$

$$\hat{\mathbf{y}}_k^- = \sum_{i=0}^{2L} W_i^{(m)} \mathcal{Y}_{i,k|k-1} \tag{8}$$

$$\mathbf{P}_{\tilde{\mathbf{y}}_k \tilde{\mathbf{y}}_k} = \sum_{i=0}^{2L} W_i^{(c)} [\mathcal{Y}_{i,k|k-1} - \hat{\mathbf{x}}_k^-][\mathcal{Y}_{i,k|k-1} - \hat{\mathbf{y}}_k^-]^T + \mathbf{R}^{\mathbf{n}} \tag{9}$$

$$\mathbf{P}_{\mathbf{x}_k \mathbf{y}_k} = \sum_{i=0}^{2L} W_i^{(c)} [\mathcal{X}_{i,k|k-1} - \hat{\mathbf{x}}_k^-][\mathcal{Y}_{i,k|k-1} - \hat{\mathbf{y}}_k^-]^T \tag{10}$$

$$\mathcal{K}_k = \mathbf{P}_{\mathbf{x}_k\mathbf{y}_k}\mathbf{P}^{-1}_{\tilde{\mathbf{y}}_k\tilde{\mathbf{y}}_k} \tag{11}$$

$$\hat{\mathbf{x}}_k = \hat{\mathbf{x}}^-_k + \mathcal{K}_k(\mathbf{y}_k - \hat{\mathbf{y}}^-_k) \tag{12}$$

$$\mathbf{P}_k = \mathbf{P}^-_k - \mathcal{K}_k\mathbf{P}_{\tilde{\mathbf{y}}_k\tilde{\mathbf{y}}_k}\mathcal{K}^T_k \tag{13}$$

having the following variable definitions: $\mathbf{P}$ is the covariance matrix, $\alpha = 0.01$ determines the size of the sigma-point distribution, $\beta = 2$ for a Gaussian distribution, L is the states dimension, $\lambda = L(\alpha^2 - 1)$ and $\eta = \sqrt{(L + \lambda)}$ is the scaling parameter, $\{W_i\}$ is a set of scalar weights ($W^{(m)}_0 = \lambda/(L + \lambda)$, $W^{(c)}_0 = \lambda/(L + \lambda) + (1 - \alpha^2 + \beta)$, $W^{(m)}_i = W^{(c)}_i = 1/\{2(L + \lambda)\}$, $i = 1,\ldots,2L$). $\mathbf{R_v}$, $\mathbf{R_n}$ is the process and observation noise. The recursive algorithm given by (3-13) defines the first stage of a two-stage estimator.

The bias b is time-varying term added to nonlinear functions. It is clear that the bias estimator have the same forms as in the original linear separate-bias estimation. It is realized in another separate recursive procedure:

$$\hat{b}^-_k = \hat{b}^-_{k-1} \tag{14}$$

$$P^-_b(k) = P_b(k - 1) \tag{15}$$

$$K_b(k) = P^-_b(k)S^T_k[H_k\mathbf{P}^-_k H_k + S_k P^-_b(k)S^T_k + R_n]^{-1} \tag{16}$$

$$P_b(k) = [I - K_b(k)S_k]P^-_b(k) \tag{17}$$

$$\hat{b}_k = \hat{b}^-_k + K_b(k)(r_k - S_k\hat{b}^-_k) \tag{18}$$

where $\hat{b}$ is the bias estimate, K_b is the bias gain equation, P_b is the bias estimation error covariance matrix, $r_k = \mathbf{y}_k - \hat{\mathbf{y}}^-_k$ is the measurement residual of the bias-free estimation, r_k and $\mathbf{P}^-_k$ can be available from the bias-free estimator, and S_k is available from the sensitivity functions.

Note that $H_k\mathbf{P}^-_k H_k$ is the measurement covariance, by the definitions given by Equation (9), Equation (16) becomes

$$K_b(n) = P^-_b(n)S^T_n[\sum_{i=0}^{2L} W^{(c)}_i[\mathcal{Y}_{i,k|k-1} - \hat{\mathbf{x}}^-_k][\mathcal{Y}_{i,k|k-1} - \hat{\mathbf{y}}^-_k]^T + S_n P^-_b(n)S^T_n + R_n]^{-1} \tag{19}$$

The recursive algorithm defined by (14-15) and (17-19) constitute the second stage of the two-stage Kalman estimator.

Suppose that the bias is perfectly known, the optimal estimate of b would be the constant value, and the recursive estimate of x may be expressed as

$$\hat{X}_k^- = f(\hat{x}_{k-1}) + B_k b \tag{20}$$

$$\hat{X}_k = \hat{X}_k^- + \mathcal{K}_k(y_k - h(\hat{X}_k^-) - C_k b) \tag{21}$$

where $\hat{X}_k^-$ and $\hat{X}_k$ are the *priori* and *posteriori* estimates of x when b is perfectly known, and b is the true value of the bias vector. Its gain equation is clearly identical to that of the bias-free estimator. Since the bias term is linear in nature, the component of the general X estimate due to the forcing function y is identical to the bias-free x estimate. It allows the following relationship as given [3]:

$$\hat{X}_k^- = \hat{x}_k^- + U_k b \tag{22}$$

$$\hat{X}_k = \hat{x}_k + V_k b \tag{23}$$

where U and V is *priori* and *posterior* sensitivity functions, respectively. The key step for developing the bias aware nonlinear estimator is to design the sensitivity functions, which connect two parallel filters. From Equation (22), we have

$$\hat{X}_k^- - \hat{x}_k^- = f(\hat{X}_{k-1}) - f(\hat{x}_{k-1}) + B_k b = \left(\left. \frac{\partial f}{\partial x} \right|_{x=\hat{x}_{k-1}} V_{k-1} + B_k \right) b \tag{24}$$

Therefore, a *priori* sensitivity function becomes

$$U_k = \left. \frac{\partial f}{\partial x} \right|_{x=\hat{x}_{k-1}} V_{k-1} + B_k \tag{25}$$

Similarly, the posterior error is written as

$$\hat{X}_k - \hat{x}_k = V_k b = \left[U_k - \mathcal{K}_k \left(\left. \frac{\partial h}{\partial x} \right|_{x=\hat{x}_k^-} U_k + C_k \right) \right] b \tag{26}$$

The a *posterior* sensitivity becomes

$$V_k = U_k - \mathcal{K}_k S_k \tag{27}$$

where S_k is given by

$$S_k = \left. \frac{\partial h}{\partial x} \right|_{x=\hat{x}_k^-} U_k + C_k \tag{28}$$

Equations (25), (27), and (28) develop as the counterpart of the sensitivity functions in the linear bias aware estimation. Taken together they provide a recursive algorithm, which links two separate stages. Thus, it is possible to employ the estimate of b for correcting the bias-free estimate $\hat{x}$ in the nonlinear bias system. The optimal estimate output give as (22) (23).

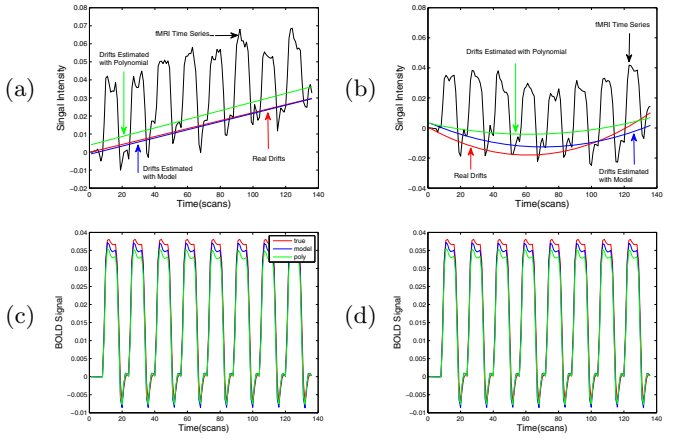

Fig. 1. Estimated drift (a,b) and fMRI signal (c,d) from synthetic data with linear drift and quadratic drift

3 Experimental Validation and Discussion

We consider the actual fMRI measurements in which only the observations are biased, $\mathbf{B} = 0$. The physiological processes underlying the BOLD response were charactered using the hemodynamic Balloon model in this study [2,4]. The resting blood volume fraction $V_0 = 0.02$ [6] and stiffness parameter $\alpha = 0.33$ are assumed known [7] in the assimilation procedure.

Synthetic Data. Since the ground truth is unavailable in real fMRI data, synthetic data are chosen to examine the proposed approach. The synthetic time series contains a known activation response, a known drift term b_j, and Gaussian white noise e_j, with signal-to-noise ratio (SNR) of 3dB. We implement linear drift and quadratic drift in the experiment. The artificial BOLD response is generated by Balloon model, where $\epsilon = 0.59$, $\tau_s = 1.38$, $\tau_f = 2.7$, $\tau_0 = 0.89$, $\alpha = 0.33$, $E_0 = 0.3$, and $V_0 = 0.02$ within their typical range [4]. The experimental condition of synthetic data is consistent with real fMRI experiments below.

Figure 1 (a), (b) show the synthetic time series and the known and estimated drift by the proposed method and the polynomial method [1]. The hemodynamic response of detrend-free assimilation, assimilation after polynomial detrending and true response are shown in Figure 1 (c), (d). Polynomial detrending generates an elevated drift estimate due to the whitening assumption. The recalibration of the baseline can only partly alleviate this effect, shown as a green line in Figure 1. The artificial elevated drift result in underestimated fMRI response (green line), and degrade the performance of data assimilation. In contrast, with the constraint of the hemodynamic model, our two-stage assimilation strategy tracks the known drift better, illustrating the advantage of introducing physiologically meaningful prior into detrending. As a result, the two-stage assimilation

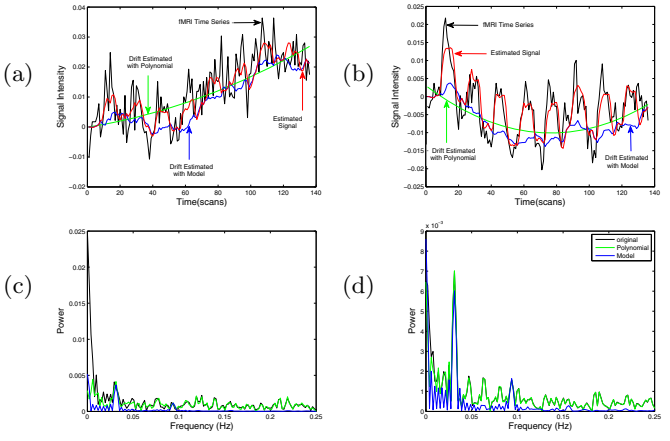

Fig. 2. Real fMRI signal and estimated drifts (a,b), as well as their spectra (c,d)

Subject	F-Statistics ($F_{59,59}$)		
	Original	Polynimal	Model
1	1.4335	1.5458	1.6926
2	1.1046	1.2271	3.2548
3	1.6961	2.5561	3.1550
4	1.6701	1.7248	2.1788
5	1.2970	1.3327	1.5098
6	1.1409	1.3153	1.4276
7	1.0064	1.6961	1.8957
8	1.0083	1.7504	2.3548

Table 1. F-statistics of assimilation with the original signal, assimilation after polynomial detrending, and detrend-free assimilation for the greatest activated area of the group in superior temporal gyrus (GTs). For a significance level of $P < 0.05$, $F_{0.95,59,59} = 1.54$.

generates better estimation of physiological states and reconstructed responses than separate processing (blue line). The statistical analysis effectively shows this case as well. It is noted that the significant level is greater by detrend-free assimilation ($F_{59,59} = 3.22$ for linear drift, $F_{59,59} = 2.51$ for quadratic drift, average from 250 simulations), which shows better detection performance in comparison with assimilation after polynomial detrending ($F_{59,59} = 3.00$ for linear drift, $F_{59,59} = 2.48$ for quadratic drift, average from 250 simulations).

Real Data. The real data was acquired from 8 healthy subjects. The condition for successive blocks alternated between rest and auditory stimulation, starting with additional 8 rest scans. Total 136 acquisitions were made (RT=2s), in blocks of 8, giving 16 16-second blocks. We select the largest activated voxels of the group in superior temporal gyrus (GTs) to implement assimilation.

Figure 2 (a), (b) show the real fMRI time series and the estimated drift by polynomial detrending method and the proposed method from two subjects. Intuitively, the estimated drift by detrend-free assimilation (blue line) can track the complicated drift variation with good approximation while the polynomial

approach (green line) does not work quite well. The estimated drift by detrend-free strategy has more freedom in selecting the structure of the drift component due to the introduction of physiological constraint. As a result, the detrend-free assimilation yields greater estimates of the fMRI response than separate processing (not presented here). Table 1 lists results of statistical analysis from all subjects. The statistic-F is obviously higher for detrend-free method than assimilation after polynomial detrending. In addition, in order to understand the effect of detrending on the spectrum of the fMRI signal, the power spectrum of the original time series is compared with detrended time series by the proposed method and polynomial method shown in Figure 2. It is noted that under the restriction of the hemodynamic model, detrend-free method has impact on all frequency contexts that are responsible to drift, whereas polynomial detrending only removes part of the low frequency component.

In this paper, we have developed a two-stage unscented estimator for nonlinear fMRI data assimilation, which casts the nonlinear unscented transform into the linear separate bias estimator, to account for the presence of time varying bias. It can deal with simultaneously the fMRI responses and drift in an assimilation cycle. It provides more accurate estimation of fMRI signal and physiological states than separate processing, thereby produces higher F value for statistic activation detection. On the other hand, it makes no assumption of the structure of the drift. As proper prior hemodynamics is adopted to guide the detrending procedure, information of the underlying physiological process is included, more reasonable drift estimates can be obtained. It is therefore particularly suited for fMRI imaging where the formulation of real drift remains difficult to acquire. Moreover, other prior knowledge can also be incorporated into the estimation procedure, such as cardiac and respiratory prior from external monitoring, and baseline drift of MR scanner.

Acknowledgments. This work is supported in part by the National Basic Research Program of China (2010CB732500) and in part by the National Natural Science Foundation of China (30800250).

References

1. Bazargani, N., Nasratinia, A., Gopinath, K., Briggs, R.W.: FMRI Baseline Drift Estimation Method by MDL Principle. In: 4th IEEE International Symposium on Biomedical Imaging (ISBI), Washington, D.C., USA, pp. 472–475 (2007)
2. Buxton, R.B., Wong, E.C., Frank, L.R.: Dynamics of Blood Flow and Oxygenation Changes During Brain Activation: The Balloon Model.. Magnetic Resonance in Medicine 39, 855–864 (1998)
3. Friedland, B.: Treatment of Bias in Recursive Filtering. IEEE Transactions on Automatic Control 14(4), 359–367 (1969)
4. Friston, K.J., Mechelli, A., Turner, R., Price, C.J.: Nonlinear Responses in fMRI: The Balloon Model, Volterra Kernels, and Other Hemodynamics.. NeuroImage 12, 466–477 (2000)
5. Hsieh, C.S.: General Two-stage Extended Kalman Filters. IEEE Transactions on Automatic Control 48(2), 289–293 (2003)

6. Hu, Z.H., Shi, P.C.: Nonlinear Analysis of BOLD Signal: Biophysical Modeling, Physiological States, and Functional Activation. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) MICCAI 2007, Part II. LNCS, vol. 4792, pp. 734–741. Springer, Heidelberg (2007)
7. Hu, Z.H., Shi, P.C.: Sensitivity Analysis for Biomedical Models. IEEE Transactions on Medical Imaging 29(11), 1870–1881 (2010)
8. Julier, S.J., Uhlmann, J.K.: Unscented Filtering and Nonlinear Estimation. Proceeding of the IEEE 92(3), 401–422 (2004)
9. Moradkhani, H., Sorooshian, S., Gupta, H.V., Houser, P.R.: Dual State-parameter Estimation of Hydrological Models Using Ensemble Kalman Filter. Advance in Water Resouces 28, 135–147 (2005)

Fiber-Centered Granger Causality Analysis

Xiang Li[1], Kaiming Li[1,2], Lei Guo[2], Chulwoo Lim[1], and Tianming Liu[1]

[1] Department of Computer Science and Bioimaging Research Center,
The University of Georgia, Athens, GA
[2] School of Automation, Northwestern Polytechnic University, Xi'an, China

Abstract. Granger causality analysis (GCA) has been well-established in the brain imaging field. However, the structural underpinnings and functional dynamics of Granger causality remain unclear. In this paper, we present fiber-centered GCA studies on resting state fMRI and natural stimulus fMRI datasets in order to elucidate the structural substrates and functional dynamics of GCA. Specifically, we extract the fMRI BOLD signals from the two ends of a white matter fiber derived from diffusion tensor imaging (DTI) data, and examine their Granger causalities. Our experimental results showed that Granger causalities on white matter fibers are significantly stronger than the causalities between brain regions that are not fiber-connected, demonstrating the structural underpinning of functional causality seen in resting state fMRI data. Cross-session and cross-subject comparisons showed that our observations are reproducible both within and across subjects. Also, the fiber-centered GCA approach was applied on natural stimulus fMRI data and our results suggest that Granger causalities on DTI-derived fibers reveal significant temporal changes, offering novel insights into the functional dynamics of the brain.

Keywords: Granger Causality Analysis, fMRI, DTI, Time Series, Brain State.

1 Introduction

Proposed by Clive Granger in 1969 (Granger 1969), Granger causality analysis (GCA) has been widely applied to analyze the relationships between time series. Briefly, a time series X is said to Granger-cause time series Y if the values of X provide statistically significant information about future values of Y. The GCA is very useful in functional MRI (fMRI) signal analysis, since different brain regions are supposed to connect together and have causal influence upon each other. Thus in recent years, it has been widely used in the brain imaging field [2-4], in order to obtain a hierarchical understanding of the interaction and correlation between different brain regions.

Despite wide application of GCA in fMRI, however, the structural underpinnings of GCA remain unclear, e.g., how structural connectivity is related to Granger causality? In addition, existing approaches of GCA on brain networks [2-4] assume temporal stationarity, where Granger causalities are computed over the entire scan and used to characterize the causality across regions. However, accumulating literature [e.g., 10], have shown that functional brain connectivity is under dynamical changes. Other literature has reported decoding brain states using sliding-window approach [12]. In responses to the above-mentioned issues, this paper presents a fiber-centered GCA

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 251–259, 2011.

approach to examine resting state and natural stimulus fMRI datasets, in order to elucidate the structural substrates and functional dynamics. Specifically, we extract the fMRI BOLD signals from the two ends of a DTI-derived fiber, and measure their Granger causalities. Our premise is that as axonal fibers are the structural substrates of functional connections between computational centers of cortical regions, the fMRI time series along the fibers should reflect the functional causality between brain regions, if any such functional causality exists.

Our experimental results showed that Granger causalities on white matter fibers are significantly stronger than the causalities between brain regions that are not connected by fibers, suggesting the structural underpinning of functional causality observed in resting state fMRI data. This observation is replicated in different imaging sessions of the same subject and across individual subjects. In addition, our experimental results of applying the fiber-centered GCA approach on natural stimulus fMRI data suggest that Granger causalities on fibers reveal significant temporal changes, offering novel insights into the functional dynamics of the brain.

2 Method

2.1 Overview of the Method

Fig. 1 shows the flowchart of this analysis pipeline. The GCA model is akin to the work flow of the program "MATLAB toolbox for Granger causal connectivity"

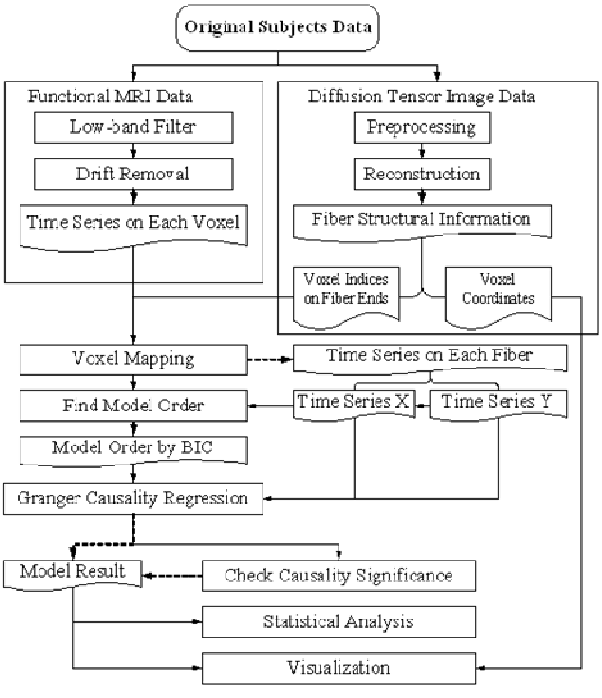

Fig. 1. The flowchart of the fiber-centered GCA framework

developed by Anil Seth (Seth 2005) and in previous studies in applying GCA to EEG data (Andreas, Dean et al. 2009). Details of other methodology components are described in Section 2.2 and 2.3.

2.2 Data Acquisition and Preprocessing

A 3T GE Signa MRI system was used for data acquisition of fMRI data analyzed in this study: resting-state and natural-stimulus fMRI data. Resting state fMRI data were acquired with dimensionality 128*128(matrix)*60 (slices)*100 (volumes), in-plane resolution 2mm*2mm*2mm isotropic, TR 5s, TE 25ms, and flip angle 90 degrees. DTI data were acquired at the same spatial resolution, with a 15.5s TR, TE of 89.5ms, and generated 30 gradient channel DWI volumes. In the natural stimulus fMRI scan, we randomly selected video shots from the TRECVID 2005 database [11], which were presented to the subject during a two-session scan. The acquisition parameters were as follows: dimensionality 128*128*60*240, spatial resolution 2mm*2mm*2mm, TR 5s, TE 25ms, and flip angle 90.

For preprocessing, we registered fMRI data to the DTI space by the FSL FLIRT tool. It should be noted that because DTI and fMRI sequences are both echo planar imaging (EPI) sequences, their distortions tend to be similar [7]. So the misalignment between DTI and fMRI images is much less than that between T1 and fMRI images [7]. DTI pre-processing included skull removal, motion correction and eddy current correction. Then fiber tracking was performed using MEDINRIA. Brain tissue segmentation was conducted on DTI data by a similar method in [8] and the cortical surface was reconstructed using the marching cubes algorithm. FMRI preprocessing steps included motion correction, spatial smoothing, temporal prewhitening, slice time correction, global drift removal, and band pass filtering. After the above preprocessing, we used white matter fibers to guide the fiber-centered GCA, which has been applied in previous studies on brain network and convincing results were obtained, showing the feasibility of the inter-modality data registration [13]. In our work, fMRI BOLD signals of the grey matter voxel pairs connected by a white matter fiber were extracted for the Granger causality analysis, as illustrated in Fig. 2. The number of voxel pairs connected by each fiber varies in different datasets. In average case there are tens of voxel pairs connected by one fiber.

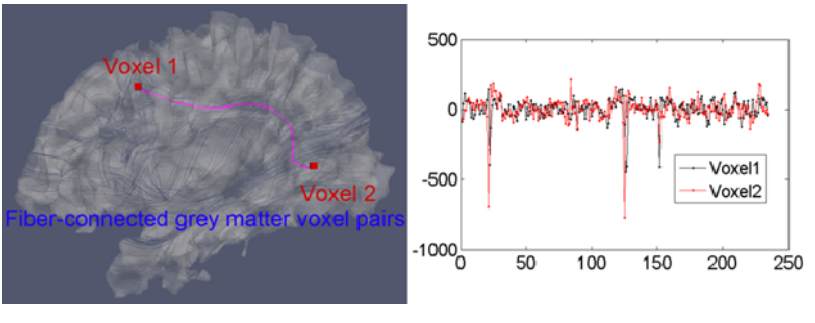

Fig. 2. Two grey matter voxels (marked in red) connected by a fiber (in purple). fMRI time series at these two voxels are shown in the right figure.

2.3 Granger Causality Analysis

Given two stochastic processes X and Y, if they are stationary, each of the process can be expressed as an auto-regression of their lagged values:

$$X_t = \sum_{i=1}^{P} a_i X_{t-i} + e1_t \tag{1}$$

$$Y_t = \sum_{i=1}^{P} d_i Y_{t-i} + e2_t \tag{2}$$

where e1 and e2 are prediction errors and their variances describe the accuracy of the prediction. Assume they have potential causality influences upon each other, there is:

$$X_t = \sum_{i=1}^{P} a_i X_{t-i} + \sum_{i=1}^{P} b_i Y_{t-i} + e3_t \tag{3}$$

$$Y_t = \sum_{i=1}^{P} c_i X_{t-i} + \sum_{i=1}^{P} d_i Y_{t-i} + e4_t \tag{4}$$

where e3 and e4 are prediction errors and a, b, c, d are linear regression coefficients. In order to study the dependency between X and Y, the null hypothesis H0: {b}=0 was made, which means Y will not significantly cause X. According to the null hypothesis, we can construct the F-statistics:

$$F_{Y \to X} = \frac{\text{var}(e1) - \text{var}(e3)}{\text{var}(e3)} \tag{5}$$

When there is no causality caused by Y to X, the value of $F_{Y \to X}$ will approach zero since the additional Y terms will not influence the explanation power in Eq. (3). And if the value is greater than the given threshold, we will reject the null hypothesis, which means there is a significant causality caused by Y to X.

The original GCA model only gives the result of whether there is causality, which is limited for the brain imaging research since there are reciprocal polysynaptic connections between brain areas (Friston 2009). Here we applied the conditional GCA (Seth 2005) which was able to evaluate the Causality Magnitude (CM):

$$CM_{Y \to X} = \ln\left(\frac{\text{var}(e3)}{\text{var}(e1)}\right) \tag{6}$$

This value is used in the following analysis to evaluate the strength of Granger causality; higher CM value indicates greater causal influence.

In this study, we applied the above results obtained by performing GCA on time series extracted from voxel pairs to study the causal characteristic of fibers that connects them. Since more than one voxel pairs may be connected by each fiber, the causality magnitude of that fiber is defined by:

$$CM_{\text{fiber i}} = \left(\sum_{v1,v2\in V_i} CM_{v1\to v2} + CM_{v2\to v1}\right)/\text{sizeof}(V_i) \quad (7)$$

where V_i is the set of voxel pairs connected by the ith fiber, and v1, v2 are pair of voxels in that set. The causality magnitude of fibers is in the range of (0, 2).

The F-statistics was used to determine whether there is a significant (given P=0.01) causal connectivity between pair of voxels, thus a fiber connecting any one pair of voxels with significant causal connectivity at any direction is considered to be significantly casual-connected. Also, we defined the causality phase (CP) based on the F-statistics and the direction of the connectivity:

$$CP_{\text{voxel v1,v2}} = \begin{cases} 1, \text{if causality from v1 to v2 is significant} \\ -1, \text{if causality from v2 to v1 is significant} \\ 0, \text{ if there is no significant causality} \end{cases} \quad (8)$$

And

$$CP_{\text{fiber i}} = \left(\sum_{v1,v2\in V_i} CP_{v1,v2}\right)/\text{sizeof}(V_i) \quad (9)$$

where V_i is the set of voxel pairs connected by the ith fiber, and v1, v2 are a pair of voxels in that set. The causality phase of a fiber is in the range of (-1, 1).

3 Results

3.1 Results in Resting State fMRI Data

Totally, resting state fMRI dataset of 12 imaging sessions from 9 subjects was analyzed. We selected significantly casual-connected fibers and visualized the causality magnitude of them in Fig. 3. In this study, subject #1, #2 and #3 had two scanning sessions, both under the same experimental circumstances. The data collected from two sessions of the same subject enabled us to compare the results and to see whether the inferred causalities were stable within subjects. As shown in Fig.3, the structures of significantly casual-connected fibers were almost identical between two sessions of all the three subjects, suggesting that certain fibers that were significantly casual-connected, might had more importance in causality connection than other fibers. Also, it could be seen that casual-connected fibers formed certain pattern in the brain, especially around frontal lobe and visual cortex in all the subjects, which indicated stronger connection and causal related activities (such as control and motivation) in these regions. Finally, as the statistics analysis of the results listed in Table 1 showed, the number of significantly causal-connected, as well as the causality magnitude, varied very little between two sessions. This provided good evidence that fiber-centered GCA on a whole-brain scale was highly reproducible within-subject.

Table 1. Similarity analyses between two sessions of each subject;

	Fibers Total	Causal-connected fibers in session 1	Causal-connected fibers in session 2	Mean of CM in session 1	Mean of CM in session 2
Subject1	13073	622	589	0.48	0.48
Subject2	17671	1054	1053	0.49	0.49
Subject3	13482	1067	966	0.49	0.5

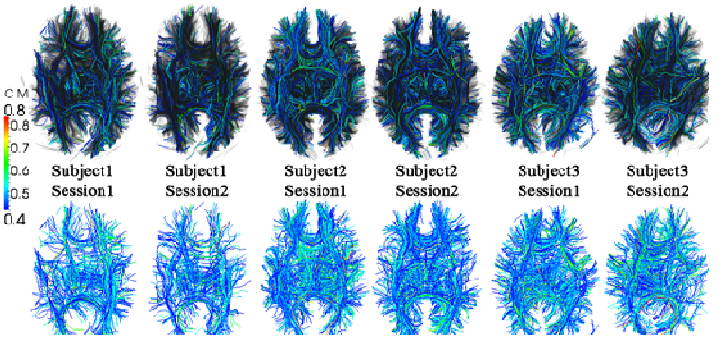

Fig. 3. Visualization of significant casual-connected fibers of three subjects (2 sessions). Images at the first and second row are identical, except that fibers that are not significantly causal-connected were colored in black in first row, and were removed in second row.

3.2 Reproducibility Study: Cross-subjects

In addition to the within-subject reproducibility (stability) study of our method, a cross-subjects reproducibility study was also conducted, to see whether the results were consistent. Fig. 4 showed the significantly casual-connected fibers of subject #4 to #9. It can be seen in the figure that there are consistency of causality pattern across subjects, despite the fact that fiber shapes and global causality levels varied among different subjects. For example, subject 9 had higher causality magnitude than subject 5 (the average CM of subject 9 is 0.54, while subject 5 is 0.45), but the structures of the significantly casual-connected fibers were similar, suggesting the relatively stable and consistent oscillation architecture of the human brain in resting state.

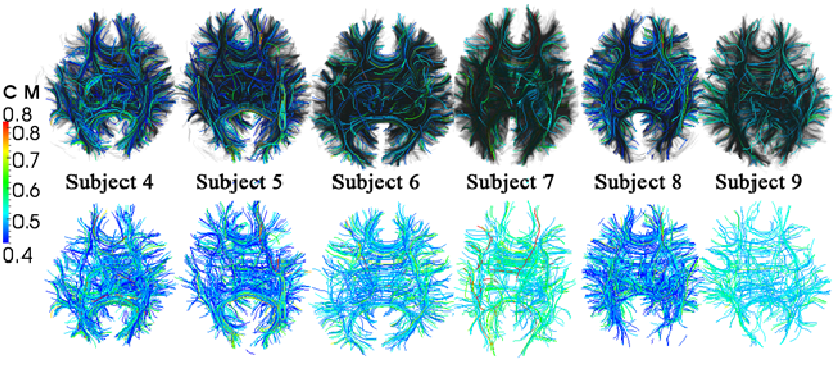

Fig. 4. Visualization of significant casual-connected fibers in 6 subjects. Images at the first and second row are identical, except that fibers that are not significantly causal-connected were colored in black in first row, and were removed in second row for comparison.

3.3 Granger Causality Dynamics in Natural Stimulus fMRI Data

In this study, we also applied sliding window approach on time series scanned from natural stimulus fMRI data, and studied temporal change of causality strength and phase along sliding windows. We constructed 223 consecutive sliding windows with the length of 13 time points and perform GCA on voxel pairs connected by fibers, obtained causality magnitude as well as causality phases. The temporal dynamics of causality magnitude on all the fibers of a subject are shown in Fig. 5. We can see that the causality magnitude varies through the time, and forms patterns by the majority of fibers. Statistically, between each sliding window, the causality magnitude of a fiber will change 14.76% on average, and can be as high as 172%, indicating dramatic temporal dynamical change. This result is reproducible in other 4 subjects analyzed.

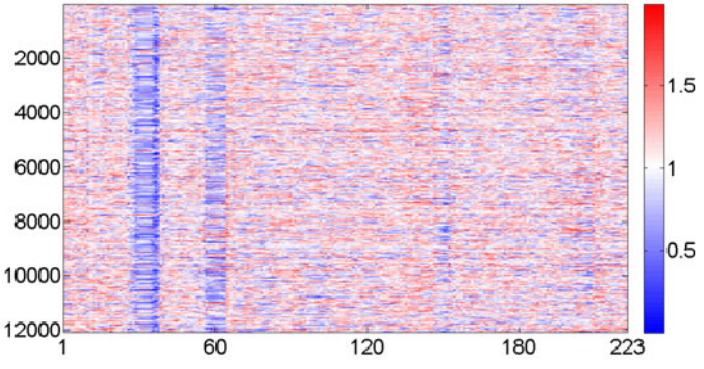

Fig. 5. Visualization of causality magnitude of all fibers in 223 sliding windows. Each row vector is the CM dynamics of one fiber through the whole time period, and each column vector is the causality magnitude state vector in that sliding window.

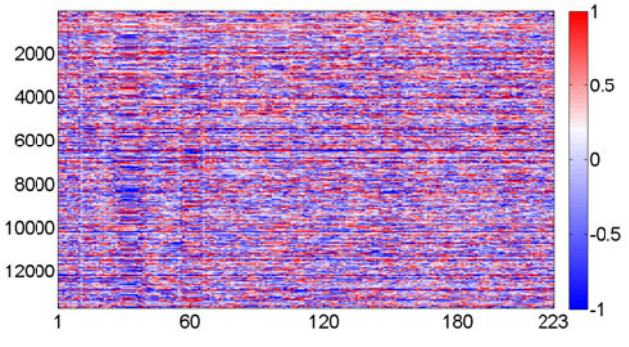

Fig. 6. Visualization of causality phase of all fibers in 223 sliding windows. Each row vector is the causality phase of one fiber through the whole time period, and each column vector is the causality phase state vector in that sliding window.

In addition to the causality magnitude dynamics in Fig. 5, the dynamics of causality phases were also analyzed. As shown in Fig. 6, the causality phases are also changing dramatically along the scan period. In average, one fiber undergoes 50 times of phase change through the whole scan, i.e., during 22.4% of the time a fiber is changing its phase of causality. The frequency can be as high as 93 times, meaning that between every 2 time points the fiber changes its causality phase. The histogram of the frequencies of causality phase changes is shown in Fig. 7. This result reveals the dynamics in natural stimulus fMRI data and is replicated in other 4 subjects studied, imposing challenges to the assumption of temporal stationarity in traditional GCA.

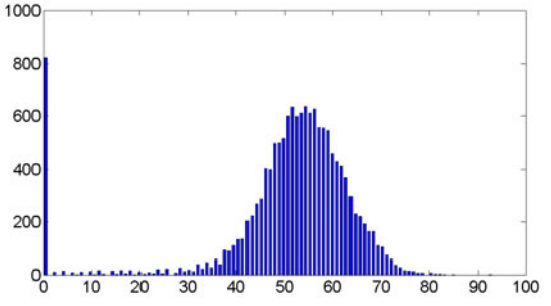

Fig. 7. Histogram of the phase change frequencies of all fibers. The phase change frequency is defined as total number of times a fiber changes its sign of causality phase.

4 Conclusion

This paper presented novel fiber-centered GCA studies on resting state and natural stimulus fMRI datasets for the purpose of elucidating the structural underpinnings and functional dynamics. In this approach, we extracted fMRI BOLD signals from voxel pairs connected by white matter fibers derived from DTI data, and examined their Granger causalities. Results in resting state fMRI data showed that Granger causalities on white matter fibers are significantly stronger than causalities between brain regions not fiber-connected, suggesting the structural connectivity underpinning of functional causality. In addition, the proposed approach was applied on natural stimulus fMRI data and our results suggest that Granger causalities revealed dramatic temporal dynamics, in terms of both causality magnitude and phase, warranting re-examination of the assumption of temporal stationarity of traditional GCA.

References

1. Granger, C.W.J.: Investigating Causal Relations by Econometric Models and Cross-spectral Methods. Econometrica 37(3), 424–438 (1969)
2. Goebel, R., Roebroeck, A., Kim, D.-S., Formisano, E.: Investigating directed cortical interactions in time-resolved fMRI data using vector autoregressive modeling and Granger causality mapping. Magnetic Resonance Imaging 21(10), 1251–1261 (2003)

3. Roebroeck, A., Formisano, E., Goebel, R.: Mapping directed influence over the brain using Granger causality and fMRI. NeuroImage 25(1), 230–242 (2005)
4. Ramsey, J.D., Hanson, S.J., Hanson, C., Halchenko, Y.O., Poldrack, R.A., Glymour, C.: Six problems for causal inference from fMRI. NeuroImage 49(2), 1545–1558 (2010)
5. Seth, A.K.: Causal connectivity of evolved neural networks during behavior. Network: Computation in Neural Systems 16(1), 35–54 (2005)
6. Keil, A., Sabatinelli, D., Ding, M., Lang, P.J., Ihssen, N., Heim, S.: Re-entrant projections modulate visual cortex in affective perception: Evidence from Granger causality analysis. Human Brain Mapping 30(2), 532–540 (2007)
7. Li, K., Guo, L., Li, G., Nie, J., Faraco, C., Zhao, Q., Miller, L.S., Liu, T.: Cortical surface based identification of brain networks using high spatial resolution resting state FMRI data. In: Proceedings of the 2010 IEEE International Conference on Biomedical Imaging, pp. 656–659. IEEE Press, Piscataway (2010)
8. Liu, T., Li, H., Wong, K., Tarokh, A., Guo, L., Wong, S.T.C.: Brain tissue segmentation based on DTI data. NeuroImage 38(1), 114–123 (2007)
9. Friston, K.: Causal Modelling and Brain Connectivity in Functional Magnetic Resonance Imaging. PLoS Biol. 7(2), e1000033 (2009)
10. Chang, C., Glover, G.H.: Time-frequency dynamics of resting-state brain connectivity measured with fMRI. NeuroImage 50(1), 81–98 (2010)
11. TRECVID 2005 Evaluation (2005), http://www-nlpir.nist.gov/projects/tv2005/tv2005.html
12. Richiardi, J., Eryilmaz, H., Schwartz, S., Vuilleumier, P., Van De Ville, D.: Decoding brain states from fMRI connectivity graphs. NeuroImage 56(2), 616–626 (2011)
13. Lv, J., Guo, L., Hu, X., Zhang, T., Li, K., Zhang, D., Yang, J., Liu, T.: Fiber-Centered Analysis of Brain Connectivities Using DTI and Resting State FMRI Data. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. 6362, pp. 143–150. Springer, Heidelberg (2010)

Variational Solution to the Joint Detection Estimation of Brain Activity in fMRI

Lotfi Chaari[1,2], Florence Forbes[1,2], Thomas Vincent[3], Michel Dojat[2,4], and Philippe Ciuciu[3]

[1] INRIA, MISTIS, Grenoble, France
[2] Grenoble University, LJK, Grenoble, France
[3] CEA/DSV/I^{2}BM/Neurospin, LNAO, Gif-Sur-Yvette, France
[4] INSERM, U836, GIN, Grenoble, France

Abstract. We address the issue of jointly detecting brain activity and estimating underlying brain hemodynamics from functional MRI data. We adopt the so-called Joint Detection Estimation (JDE) framework that takes spatial dependencies between voxels into account. We recast the JDE into a missing data framework and derive a Variational Expectation-Maximization (VEM) algorithm for its inference. It follows a new algorithm that has interesting advantages over the previously used intensive simulation methods (Markov Chain Monte Carlo, MCMC): tests on artificial data show that the VEM-JDE is more robust to model mis-specification while additional tests on real data confirm that it achieves similar performance in much less computation time.

1 Introduction

Functional Magnetic Resonance Imaging (fMRI) is a powerful tool to non-invasively study the relation between a cognitive task and an evoked neural activity through neurovascular coupling and the BOLD effect [10]. To localize which parts of the brain are activated by a given stimulus type, most approaches assume a single canonical *a priori* model for the impulse response of the neurovascular coupling also known as the hemodynamic response function (HRF) [5]. However, there has been evidence that this response can vary in space and between subjects [6,1] so that both issues of properly detecting evoked activity and estimating the HRF play a central role in fMRI data analysis. They are usually dealt with independently with no possible feedback although they are strongly connected. To account for these sources of hemodynamic variability, a novel approach referred to as the Joint Detection Estimation (JDE) framework has been introduced in [9] and extended in [13] to account for spatial correlation between neighboring voxels in the brain volume. Since robust and accurate HRF estimation can only be achieved in regions that elicit an evoked response to an experimental condition [7], the JDE approach has been defined at an intermediate spatial resolution corresponding to *parcels* in which a fair compromise between homogeneity of the BOLD signal and reproducibility of the HRF shape is achieved. The JDE approach mainly rests upon: *i.)* a non-parametric or

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 260–268, 2011.

FIR modelling of the HRF at this parcel-level for an unconstrained HRF shape; *ii.)* prior information about the temporal smoothness of the HRF to guarantee a physiologically plausible shape; and *iii.)* the modelling of spatial correlation between neighboring voxels within each parcel using condition-specific discrete hidden Markov fields. In [9,13], posterior inference is carried out in a Bayesian setting using Markov Chain Monte Carlo (MCMC) methods, which requires fine tuning and is time consuming. Several other attempts to segregate neurological and hemodynamic events from fMRI time series have been proposed (see references in [13]). Among them lies an interesting bilinear dynamical system formulation [8] that deals with unknown HRF and uses Variational Bayes (VB) approximation for tractability. However, from a spatial viewpoint, this work remains univariate and ignores spatial correlation between voxels.

In this paper, we reformulate the complete JDE framework [13] as a missing data problem and propose a simplification of its estimation procedure. Akin to [8], we resort to a variational approximation using a Variational Expectation Maximization (VEM) algorithm in order to derive estimates of the HRF and stimulus-related activity. Experiments on artificial and real data demonstrate the good performance of our VEM algorithm. In particular, we provide a comparison with its MCMC counterpart and show the advantages of VEM both in terms of computation time and robustness to noise model deviations. This potentially increases considerably the impact of the JDE framework and makes its application to fMRI studies in neuroscience easier and more valuable.

2 A Joint Detection-Estimation Model

Capital letters indicate random variables and lower case their realizations. Matrices are denoted with bold upper case letters and the transpose with^t.

Observed and Missing Variables. We first recast the parcel-based JDE model of [9,13] in a missing data framework. For a given parcel $\mathcal{P}$, the *observed* data is denoted by $Y = \{Y_i, i \in \mathcal{P}\}$ where $Y_i \in \mathbb{R}^N$ is the fMRI time series measured in voxel $i \in \mathcal{P}$ at times $(t_n)_{n=1:N}$, where $t_n = nTR$, N being the number of scans and TR, the time of repetition. Additional *unobserved* variables are introduced: 1) The Neural Response Levels (NRLs) $A = \{A_m, m = 1 : M\}$ with $A_m = \{A_{mi}, i \in \mathcal{P}\}$ and M the number of experimental conditions involved in the paradigm. We will also make use of $A_i = [A_{mi}, m = 1 : M]^t$. 2) The HRF shape $H = [H_{d\Delta t}, d = 0 : D]^t \in \mathbb{R}^{D+1}$; 3) The activation class assignments $Z = \{Z_m, m = 1 : M\}$ with $Z_m = \{Z_{mi}, i \in \mathcal{P}\}$ represent the *activation classes* for each voxel, in each of the M experimental conditions. $Z_{mi} = k$ means that voxel i lies in activation class k for the m^{th} experimental condition. The number of classes is here $K = 2$ for activating and non activating voxels. An additional deactivation class ($K = 3$) may actually be added depending on the experiment and all provided formulas are general enough to cover this case. The observed and missing variables are then linked by the following relationship involving additional parameters to be estimated:

$$\forall i \in \mathcal{P}, \quad Y_i = \sum_{m=1}^{M} A_{mi} \boldsymbol{X}_m H + \boldsymbol{P}\ell_i + \varepsilon_i, \tag{1}$$

where $\boldsymbol{X}_m = (x^m_{t_n - d\Delta t})_{n=1:N, d=0:D}$ denotes the $N \times (D+1)$ binary matrix that codes the onsets of the m^{th} experimental condition on a Δt-sampled grid, where Δt is the sampling period of the HRF ($\Delta t < TR$); the ε_i's stand for the noise and are independent and normally distributed in space with $\varepsilon_i \sim \mathcal{N}(0, \boldsymbol{\Gamma}_i^{-1})$, and $\boldsymbol{P}$ is the low frequency orthogonal $N \times L$ matrix which accounts for physiological artifacts. Let $\ell = \{\ell_i, i \in \mathcal{P}\}$ be the set of low frequency drifts, where $\ell_i \in \mathbb{R}^L$ have to be estimated and let $\boldsymbol{\Gamma} = \{\boldsymbol{\Gamma}_i, i \in \mathcal{P}\}$ be the set of all unknown precision matrices (see Section 4 for its definition).

Hierarchical Model of the Complete Data Distribution. Using standard additional assumptions and omitting the dependence on the parameters, the distribution of both the observed and unobserved data writes: $p(y, a, h, z) = p(y \,|\, a, h)\; p(a \,|\, z)\; p(h)\; p(z)$. To fully define the model, we now specify each term.
The $p(y \,|\, a, h)$ **term.** From (1), it comes that $p(y \,|\, a, h) = \prod_{i \in \mathcal{P}} p(y_i \,|\, a_i, h)$ with $Y_i \,|\, (A_i = a_i, H = h) \sim \mathcal{N}\left(\sum_{m=1}^M a_{mi} \boldsymbol{X}_m h + \boldsymbol{P}\ell_i, \boldsymbol{\Gamma}_i^{-1}\right)$.
The $p(a \,|\, z)$ **term.** Akin to [9,13], the NRLs A_{mi} are assumed statistically independent across condition types. The allocation variables Z_{mi} are then introduced to segregate activating voxels from non-activating ones in condition-specific mixture models. Also, the A_{mi}'s are supposed independent in space conditionally on Z_m so that putting together all conditions we get: $p(a \,|\, z) = \prod_{m=1}^M \prod_{i \in \mathcal{P}} p(a_{mi} \,|\, z_{mi})$, where we further assume that $p(A_{mi} \,|\, Z_{mi} = k) = \mathcal{N}(\mu_{mk}, \sigma^2_{mk})$. The Gaussian parameters are unknown and denoted by $\mu = \{\mu_m, m = 1 : M\}$ with $\mu_m = [\mu_{m1} \dots \mu_{mK}]^t$ and $\sigma = \{\sigma_m, m = 1 : M\}$ with $\sigma_m = [\sigma_{m1} \dots \sigma_{mK}]^t$. Also, $k = 1$ is assigned to non-activating voxels with $\mu_{m1} = 0$.
The $p(h)$ **term.** Akin to [9,13], we introduce constraints in the prior that favor smooth variations in h: $H \sim \mathcal{N}(0, \sigma_h^2 \boldsymbol{R})$ with $\boldsymbol{R} = (\Delta t)^4\; (\boldsymbol{D}_2^t \boldsymbol{D}_2)^{-1}$ where $\boldsymbol{D}_2$ is the second-order finite difference matrix and σ_h^2 is a parameter to be estimated or fixed. Moreover, $H_0 = H_{D\Delta t} = 0$ as in [9,13].
The $p(z)$ **term.** As in [13], we assume prior independence between experimental conditions regarding the activation class assignments. It follows that $p(z) = \prod_{m=1}^M p(z_m; \beta_m)$ where we assumed in addition that $p(z_m; \beta_m)$ is a spatial Markov prior, namely a K-class Potts model with interaction parameter β_m [13]. The unknown parameters are then $\beta = \{\beta_m, m = 1 : M\}$. For the complete model, the whole set of parameters denoted by $\theta \in \Theta$ is $\theta = \{\boldsymbol{\Gamma}, \ell, \mu, \sigma, \sigma_h, \beta\}$.

3 Estimation by Variational EM

We propose to use an Expectation-Maximization (EM) framework to deal with the missing data namely, $A \in \mathcal{A}$, $H \in \mathcal{H}$, $Z \in \mathcal{Z}$. At iteration (r), denoting current parameter values by $\theta^{(r)}$, the E-step involves the posterior $p(a, h, z \,|\, y; \theta^{(r)})$, which is intractable for our model. Hence, we resort to a variational EM variant in which the intractable posterior is approximated as a product of three pdfs on $\mathcal{A}$, $\mathcal{H}$ and $\mathcal{Z}$ respectively. Previous attempts to use VB inference [2]

in fMRI [14,11] have been successful with this type of approximations usually validated by assessing its fidelity to its MCMC counterpart. In Section 4, we will also provide such a comparison. It follows then that our E-step becomes an approximate E-step, which can be further decomposed into three stages consisting in updating the three pdfs, denoted by q_H, q_A and q_Z, in turn. Let $q_A^{(r-1)}, q_Z^{(r-1)}$ and $\theta^{(r)}$, be the current estimates at the r^{th} iteration and $\mathrm{E}_q[.]$ denotes the expectation with respect to (wrt) some pdf q, the first E-H step reads:

$$\texttt{E-H:} \quad q_H^{(r)}(h) \propto \exp\left(\mathrm{E}_{q_A^{(r-1)} q_Z^{(r-1)}}\left[\log p(h \,|\, y, A, Z; \theta^{(r)})\right]\right) .$$

The following `E-A` and `E-Z` steps have similar expressions obtained by exchanging the role of H and A (resp. of H and Z) and replacing $q_A^{(r-1)}$ by $q_H^{(r)}$ (resp. $q_A^{(r-1)} q_Z^{(r-1)}$ by $q_A^{(r)} q_H^{(r)}$). For the `E-H` and `E-A` steps, it follows from standard algebra that $q_H^{(r)}$ and $q_A^{(r)}$ are both Gaussian pdfs: $q_H^{(r)} \sim \mathcal{N}(m_H^{(r)}, \boldsymbol{\Sigma}_H^{(r)})$ and $q_A^{(r)} = \prod_{i \in \mathcal{P}} \mathcal{N}(m_{A_i}^{(r)}, \boldsymbol{\Sigma}_{A_i}^{(r)})$.

• `E-H step.` The expressions for $m_H^{(r)}$ and $\boldsymbol{\Sigma}_H^{(r)}$ are similar to those derived in the MCMC case [9, Eq. (B.1)] with expressions involving the a_{mi}'s replaced by their expectations wrt $q_{A_i}^{(r-1)}$: $m_H^{(r)} = \boldsymbol{\Sigma}_H^{(r)} \Big(\sum_{i \in \mathcal{P}} \boldsymbol{S}_i^{(r-1)^t} \widetilde{y}_i^{(r)} \Big)$ and

$$\boldsymbol{\Sigma}_H^{(r)-1} = \boldsymbol{R}^{-1}/\sigma_h^{2(r)} + \sum_{i \in \mathcal{P}} \Big(\sum_{m,m'} \sigma_{A_{mi} A_{m'i}}^{(r-1)} \boldsymbol{X}_m^t \boldsymbol{\Gamma}_i^{(r)} \boldsymbol{X}_{m'} + \boldsymbol{S}_i^{(r-1)^t} \boldsymbol{\Gamma}_i^{(r)} \boldsymbol{S}_i^{(r-1)} \Big),$$

with $\boldsymbol{S}_i^{(r-1)} = \sum_{m=1}^{M} m_{A_{mi}}^{(r-1)} \boldsymbol{X}_m$ and $\widetilde{y}_i^{(r)} = \boldsymbol{\Gamma}_i^{(r)} (y_i - \boldsymbol{P} \ell_i^{(r)})$. Here, $m_{A_{mi}}^{(r-1)}$and $\sigma_{A_{mi} A_{m'i}}^{(r-1)}$ denote the m^{th} and $(m,m')^{\text{th}}$ entries of the mean vector and covariance matrix of the current $q_{A_i}^{(r-1)}$, respectively.

• `E-A step.` Here, the relationship with the MCMC update of a is not straightforward. In [9,13], the a_{mi}'s are sampled independently and conditionally on the z_{mi}'s. This is not the case in the VEM framework while some similarity appears if we set the probabilities $q_{Z_{mi}}(k)$'s either to 0 or 1 and consider only the diagonal part of $\boldsymbol{\Sigma}_{A_i}$. The update of q_A reads:

$$m_{A_i}^{(r)} = \boldsymbol{\Sigma}_{A_i}^{(r)} \Big(\sum_{k=1}^{K} \boldsymbol{\Delta}_{ki}^{(r)} \mu_k^{(r)} + \widetilde{\boldsymbol{X}}_i^{(r)^t} m_H^{(r)} \Big) \text{ and } \boldsymbol{\Sigma}_{A_i}^{(r)} = \Big(\sum_{k=1}^{K} \boldsymbol{\Delta}_{ki}^{(r)} + \widetilde{\boldsymbol{H}}_i^{(r)} \Big)^{-1}$$

where $\mu_k^{(r)} = \left[\mu_{1k}^{(r)} \ldots \mu_{Mk}^{(r)}\right]^t$, $\widetilde{\boldsymbol{X}}_i^{(r)}$ is a $D+1 \times M$ matrix whose m^{th} column is $\boldsymbol{X}_m^t \widetilde{y}_i^{(r)}$, $\boldsymbol{\Delta}_{ki}^{(r)}$ is a $M \times M$ diagonal matrix whose $(m,m)^{\text{th}}$ entry is $q_{Z_{mi}}^{(r-1)}(k)/\sigma_{mk}^{2(r)}$ and $\widetilde{\boldsymbol{H}}_i^{(r)}$ is a $M \times M$ matrix whose $(m,m')^{\text{th}}$ entry is $\mathrm{tr}\Big((\boldsymbol{\Sigma}_H^{(r)} + m_H^{(r)} m_H^{(r)^t}) \boldsymbol{X}_m{}^t \boldsymbol{\Gamma}_i^{(r)} \boldsymbol{X}_{m'} \Big)$.

• `E-Z step.` From $p(a|z)$ and $p(z)$ in Section 2, the (A_m, Z_m) couples correspond to independent hidden Potts models with Gaussian class distributions. It comes an approximation that factorizes over conditions: $q_Z^{(r)}(z) = \prod_{m=1}^{M} q_{Z_m}^{(r)}(z_m)$ where $q_{Z_m}^{(r)}(z_m) = p_m(z_m \,|\, A_m = m_{A_m}^{(r)}; \mu_m^{(r)}, \sigma_m^{(r)}, \beta_m^{(r)})$ is the posterior of Z_m in a modified hidden Potts model, p_m, in which the observations a_{mi}'s are replaced by their mean values $m_{A_{mi}}^{(r)}$ and an external field $\{\sigma_{A_{im} A_{im}}^{(r)} \left[1/\sigma_{m1}^{2(r)} \ldots 1/\sigma_{mK}^{2(r)}\right]^t, i \in \mathcal{P}\}$ is added to the prior Potts model $p(z_m; \beta_m^{(r)})$. The Potts expression above is intractable but we use a mean field-like technique (see [4] for details) to approximate $q_{Z_m}^{(r)}(z_m)$ by a factorized pdf $\widetilde{q}_{Z_m}^{(r)}(z_m) = \prod_{i \in \mathcal{P}} \widetilde{q}_{Z_{mi}}^{(r)}(z_{mi})$.

`M-step.` It is also divided into four sub-steps involving separately (μ, σ), σ_h, β and $(\ell, \boldsymbol{\Gamma})$. The first two maximizers admit closed-form expressions:

- `M-`(μ, σ) `step:` Let $\bar{q}_{mk}^{(r)} = \sum_{i \in \mathcal{P}} q_{Z_{mi}}^{(r)}(k)$, then $\mu_{mk}^{(r+1)} = \sum_{i \in \mathcal{P}} q_{Z_{mi}}^{(r)}(k)\, m_{A_{mi}}^{(r)} / \bar{q}_{mk}^{(r)}$ and $\sigma_{mk}^{2(r+1)} = \sum_{i \in \mathcal{P}} q_{Z_{mi}}^{(r)}(k)\, \big((m_{A_{mi}}^{(r)} - \mu_{mk}^{(r+1)})^2 + \sigma_{A_{im}A_{im}}^{(r)}\big) / \bar{q}_{mk}^{(r)}$.
- `M-`σ_h^2 `step:` $\sigma_h^{2(r+1)} = (D-1)^{-1} \mathrm{tr}\,((\boldsymbol{\Sigma}_H^{(r)} + m_H^{(r)} m_H^{(r)^t}) \boldsymbol{R}^{-1})$.

The two other M-steps require iterative maximization procedures. Updating β consists in making further use of a mean field-like approximation [4]. Regarding the pair $(\ell, \boldsymbol{\Gamma})$, it satisfies some fixed point equation, which simplifies in case of white noise. For autoregressive (AR) noise models, we found some similarity with [9, Eq. (B.2)] when replacing h and a by m_H and m_A, respectively.

4 Experiments

Simulation Results. We simulated data according to Eq. (1)[1] and $p(a \mid z)$ with a white Gaussian noise $\boldsymbol{\Gamma}_i^{-1} = 0.5\, \boldsymbol{I}_N$ ($\boldsymbol{I}_N$ is the $N \times N$ identity matrix), $M = 2$ experimental conditions and stimulus-varying contrast-to-noise ratios (CNR): $\mu_{12} = 2.8, \sigma_{12} = 0.5$ and $\mu_{22} = 1.8, \sigma_{22} = 0.6$ so that $\mu_{12}/\sigma_{12} > \mu_{22}/\sigma_{22}$. The *initial* artificial paradigm comprised 15 stimulus events for each condition. The simulation process finally yielded time-series lasting 152 scans. Condition-specific activating and non-activating voxels were defined as 20×20 2D slices shown in Fig. 1 (right) and superimposed to the estimated label probabilities in white solid line. The parameters β_1 and β_2 had been set to fixed values ($\beta_1 = \beta_2 = 0.8$ for the two algorithms). $\boldsymbol{\Gamma}$ and ℓ are estimated as in [9].

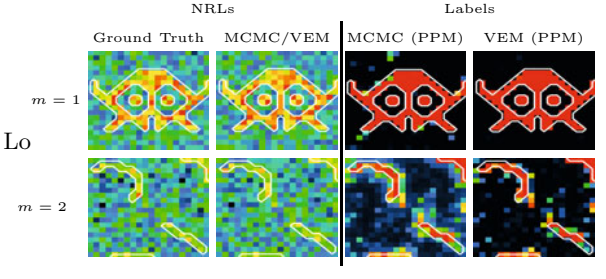

Fig. 1. Left: Ground truth and estimated Neural Response Levels (NRLs) by MCMC and VEM (same results); **Right:** Posterior probability maps (PPM) given by the approximation q_{Z_m} (VEM) and by the MMSE estimator (MCMC)

In Fig. 1, the VEM is compared to the MCMC alternative developed in [13]: both algorithms report similar NRL maps while some difference is exhibited on the posterior activation probability map (PPM) for the low CNR condition ($m = 2$, bottom row). This illustrates the gain in robustness achieved using the variational approximation under the true noise model.

[1] $\boldsymbol{P}$ was defined from a cosine transform basis.

To perform a quantitative comparison, several experiments with different stimuli densities (from 5 to 30), noise variance and autocorrelation ($\boldsymbol{\Gamma}_i^{-1}$) have been conducted. Fig. 2(a) illustrates the evolution of the Mean Square Error (MSE) of NRL estimates wrt the stimulus density in the experimental paradigm when a second order autoregressive noise (AR(2)) is considered. This figure shows that at low stimulus density (i.e. low Signal to Noise Ratio (SNR)[2]), the proposed VEM algorithm is more robust than the MCMC one to model discrepancy. Indeed, here the two inference algorithms were compared for a white and Gaussian noise modelling in Eq. (1). In contrast, at high stimulus density (≥ 20), the two methods perform similarly. Interestingly, Fig. 2(b)-(c) depict the shapes of the ground truth and estimated HRF shapes inferred by the VEM and MCMC schemes wrt the stimulus density: Note that the main HRF features (peak value (PV), time-to-peak (TTP) and time-to-undershoot (TTU)) remain well estimated by both methods. However, at low stimulus density, Fig. 2(b) shows that the VEM algorithm is less accurate than its MCMC counterpart close to the undershoot position. Similar experiments have been conducted while changing the ground truth HRF properties (PV, TTP, TTU), and coherent results have been obtained. Other comparisons performed under the true noise model did not reveal any significant difference between VEM and MCMC.

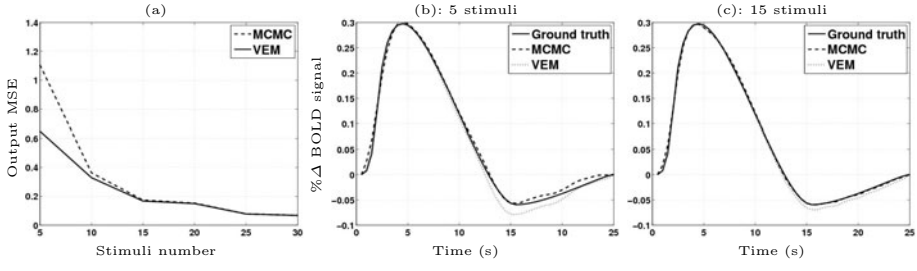

Fig. 2. (a): MSE evolution of estimated NRLs wrt stimuli number. **(b)-(c):** Ground truth and HRF estimates inferred by the VEM and MCMC schemes for two stimulus densities corrupted by AR(2) noise.

In Fig. 3(a)-(b) the output MSE is plotted against the input SNR when varying the noise variance and its amount of autocorrelation, respectively. In the latter case, the two AR parameters are varied while maintaining a stable AR(2) process. As already observed in [3], at fixed input SNR, the impact of large autocorrelation is stronger than that of large noise variance irrespective of the inference scheme. Moreover, the two inference methods perform very similarly on a large scale of input SNR (SNR > 5 dB). In terms of computational time, VEM results have been obtained 30 times faster than with MCMC on an Intel Core 2 - 2.26 GHz - 2 Gb RAM architecture.

[2] The SNR is given by: $\text{SNR} = 10 \log \sum_{i \in \mathcal{P}} \| \sum_{m=1}^{M} A_{mi} \boldsymbol{X}_m h \|^2 / \sum_{i \in \mathcal{P}} \|\varepsilon_i\|^2$.

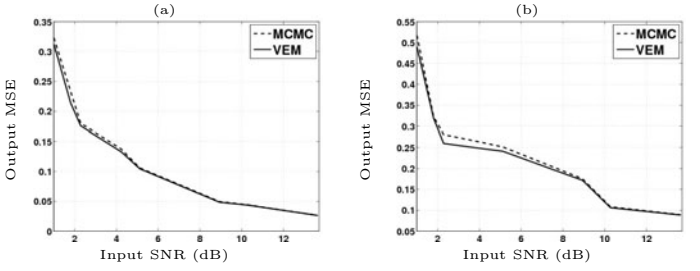

Fig. 3. MSE evolution of NRL estimates wrt input SNR (AR(2) noise) by varying the noise variance **(a)** and the amount of AR(2) noise autocorrelation **(b)**

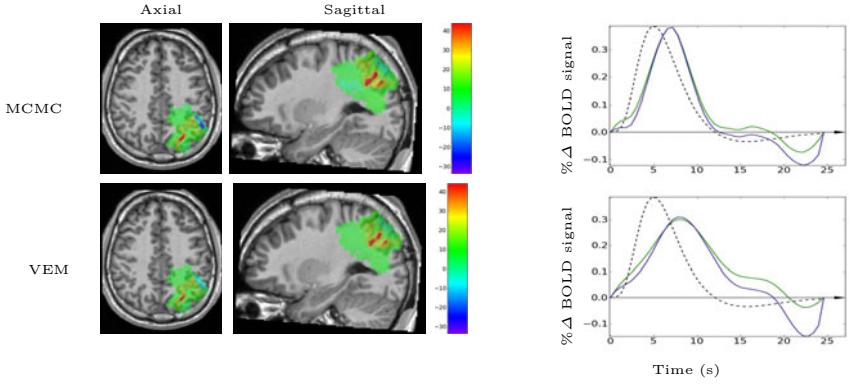

Fig. 4. Left: Estimated contrast **Computation-Sentences** by MCMC and VEM; **Right:** HRF estimates by MCMC (green) and VEM (blue) at maximum intensity peak (top) and in a neighboring parcel (bottom). Canonical HRF with dashed line.

Real Data Processing. fMRI data were recorded at 3 T (Siemens Trio) using a gradient-echo EPI sequence (TE=30 ms/TR=2.4 s/FOV=192 mm^2) during a *Localizer* experiment [12]. The acquisition consisted of a single session of $N = 128$ scans, yielding 3-D volumes with a spatial resolution of $2 \times 2 \times 3$ mm^3. The paradigm was a fast event-related design comprising sixty auditory, visual and motor stimuli, defined in ten experimental conditions (auditory and visual sentences, auditory and visual calculations, left/right auditory and visual clicks, horizontal and vertical checkerboards).

We focused on the `Computation-Sentences` contrast differentiating the activations induced by the calculation and sentence conditions in the left intraparietal sulcus, a region known to elicit hemodynamic response that departs from the canonical HRF. As shown in Fig. 4, the contrasted NRL estimates for the MCMC and VEM inference schemes are very similar and follow the underlying sulco-gyral anatomy. Note that only the most activating slice is considered for visualization purpose. The corresponding HRF estimates in the most acti-

vating parcel of about 200 voxels are also depicted in Fig. 4: they appear very similar and both quite different from the canonical shape regarding the TTP and TTU parameters. More oscillations arise in the VEM inference close to the undershoot, however we may have less confidence in the HRF tail than in its peak since it involves less signal strength. Moreover, the event-related nature of the paradigm is not suited to properly study the undershoot properties. Finally, in terms of computational efficiency, the variational approximation runs also 30 times faster than the MCMC inference in this parcel.

5 Conclusion

We proposed a Variational EM algorithm as an alternative solution to intensive stochastic sampling for inferring upon the JDE parameters. Illustrations on simulated data showed that our approach achieved similar and even better results than the MCMC-based inference scheme at low input SNR or stimuli density and was more robust to noise model mismatch. Also, in contrast to the hybrid MCMC in [13], the VEM algorithm only requires a simple stopping criterion as convergence diagnosis tool. Another advantage of the variational approach lies in its flexibility to adapt to more complex situations such as accounting for higher AR noise order, habituation modelling or including model selection steps using the log-evidence as information criterion. Other future work will focus on analyzing the impact of the VEM-JDE to group-level analysis as done in [1].

References

1. Badillo, S., Vincent, T., Ciuciu, P.: Impact of the joint detection-estimation approach on random effects group studies in fMRI. In: 7th International Symposium on Biomedical Imaging, Chicago, IL, pp. 376–380 (April 2011)
2. Beal, M., Ghahramani, Z.: The variational Bayesian EM algorithm for incomplete data: with application to scoring graphical model structures. Bayesian Statistics, vol. 7, pp. 453–464. University of Oxford Press, Oxford (2003)
3. Casanova, R., Ryali, S., Serences, J., Yang, L., Kraft, R., Laurienti, P., Maldjian, J.: The impact of temporal regularization on estimates of the BOLD hemodynamic response function: a comparative analysis. Neuroimage 40(4), 1606–1618 (2008)
4. Celeux, G., Forbes, F., Peyrard, N.: EM procedures using mean field-like approximations for Markov model-based image segmentation. Pat. Rec. 36, 131–144 (2003)
5. Friston, K., Jezzard, P., Turner, R.: Analysis of functional MRI time-series. Hum. Brain Mapp. 1, 153–171 (1994)
6. Handwerker, D.A., Ollinger, J.M., D'Esposito, M.: Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. Neuroimage 21(4), 1639–1651 (2004)
7. Kershaw, J., Ardekani, B.A., Kanno, I.: Application of Bayesian inference to fMRI data analysis. IEEE Trans. Med. Imag. 18(12), 1138–1153 (1999)
8. Makni, S., Beckmann, C., Smith, S., Woolrich, M.: Bayesian deconvolution of fMRI data using bilinear dynamical systems. Neuroimage 42(4), 1381–1396 (2008)

9. Makni, S., Idier, J., Vincent, T., Thirion, B., Dehaene-Lambertz, G., Ciuciu, P.: A fully Bayesian approach to the parcel-based detection-estimation of brain activity in fMRI. Neuroimage 41(3), 941–969 (2008)
10. Ogawa, S., Lee, T.M., Kay, A.R., Tank, D.W.: Brain magnetic resonance imaging with contrast dependent on blood oxygenation. Nat. Acad. Sci. 87, 9868–9872 (1990)
11. Penny, W.D., Kiebel, S., Friston, K.J.: Variational Bayesian inference for fMRI time series. Neuroimage 19(3), 727–741 (2003)
12. Pinel, P., Thirion, B., Mériaux, S., Jobert, A., Serres, J., Le Bihan, D., Poline, J.B., Dehaene, S.: Fast reproducible identification and large-scale databasing of individual functional cognitive networks. BMC Neurosci. 8(1), 91 (2007)
13. Vincent, T., Risser, L., Ciuciu, P.: Spatially adaptive mixture modeling for analysis of within-subject fMRI time series. IEEE Trans. Med. Imag. 29, 1059–1074 (2010)
14. Woolrich, M., Behrens, T.: Variational Bayes inference of spatial mixture models for segmentation. IEEE Trans. Med. Imag. 25(10), 1380–1391 (2006)

Adaptively and Spatially Estimating the Hemodynamic Response Functions in fMRI

Jiaping Wang[1], Hongtu Zhu[1,4], Jianqing Fan[5],
Kelly Giovanello[3,4], and Weili Lin[2,4]

[1] Department of Biostatistics
[2] Department of Radiology
[3] Department of Psychology
[4] Biomedical Research Imaging Center
University of North Carolina at Chapel Hill, Chapel Hill, NC 27599, USA
[5] Department of Operations Research & Financial Engineering, Princeton University, Princeton, NJ 08540, USA

Abstract. In an event-related functional MRI data analysis, an accurate and robust extraction of the hemodynamic response function (HRF) and its associated statistics (e.g., magnitude, width, and time to peak) is critical to infer quantitative information about the relative timing of the neuronal events in different brain regions. The aim of this paper is to develop a multiscale adaptive smoothing model (MASM) to accurately estimate HRFs pertaining to each stimulus sequence across all voxels. MASM explicitly accounts for both spatial and temporal smoothness information, while incorporating such information to adaptively estimate HRFs in the frequency domain. One simulation study and a real data set are used to demonstrate the methodology and examine its finite sample performance in HRF estimation, which confirms that MASM significantly outperforms the existing methods including the smooth finite impulse response model, the inverse logit model and the canonical HRF.

1 Introduction

The functional MRI (fMRI) study commonly uses blood oxygenation level-dependent (BOLD) contrast to measure the hemodynamic response (HR) related to neural activity in the brain or spinal cord of humans or animals. Most fMRI research correlates the BOLD signal elicited by a specific cognitive process with the underlying unobserved neuronal activation. Therefore, it is critical to accurately model the evoked HR to a neural event in the analysis of fMRI data. See [6] for an overview of different methods to estimate HRF in fMRI. A linear time invariant (LTI) system is commonly implemented to model the relationship between the stimulus sequence and BOLD signal where the signal at time t and voxel $\mathbf{d}$, $Y(t, \mathbf{d})$, is the convolution of a stimulus function $X(t, \mathbf{d})$ and the HR function (HRF) $h(t, \mathbf{d})$ plus an error process $\epsilon(t, \mathbf{d})$. While nonlinearities in the BOLD signal are predominant for stimuli with short separations, it has been shown that LTI is a reasonable assumption in a wide range of situations [6].

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 269–276, 2011.

Almost all HRF models estimate HRF on a voxel-wise basis which ignores the fact that fMRIs are spatially dependent in nature. Particularly, as the case in many fMRI studies, we observe spatially contiguous effect regions with rather sharp edges. There are several attempts to address the issue of spatial dependence in fMRI. A possible approach is to apply a smoothing step before individually estimating HRF in each voxel of fMRI data. Most smoothing methods, however, are independent of the imaging data and apply the same amount of smoothness throughout the whole image [7]. These smoothing methods can blur the information near the edges of the effect regions and thus dramatically increase the number of false positives and false negatives. An alternative approach is to model spatial dependence among spatially connected voxels by using conditional autoregressive (CAR), Markov random field (MRF) or other spatial correlation priors [9,11]. However, calculating the normalizing factor of MRF and estimating spatial correlation for a large number of voxels in the 3D volume are computationally intensive [2]. Moreover, it can be restrictive to assume a specific type of correlation structure for the whole 3D volume (or 2D surface).

The goal of this paper is to develop a multiscale adaptive smoothing model (MASM) to construct an accurate nonparametric estimate of HRF across all voxels pertaining to a specific cognitive process in the frequency domain. Compared with all existing methods, we make several major contributions. (i) To temporally smooth HRF, MASM incorporates an effective method for carrying out locally adaptive bandwidth selection across different frequencies; (ii) To spatially smooth HRFs, MASM builds hierarchically nested spheres by increasing the radius of a spherical neighborhood around each voxel and utilizes information in each of the nested spheres across all voxels to adaptively and simultaneously smooth HRFs; (iii) MASM integrates both spatial and frequency smoothing methods together; (iv) MASM uses a backfitting method [3] to adaptively estimate HRFs for multiple stimulus sequences across all voxels. The group-level inference like testing the significant regions based on estimated HRFs will be further studied in the future.

2 Model Formulation

2.1 Multiscale Adaptive Smoothing Model

Suppose that we acquire a fMRI data set in a three dimensional (3D) volume, denoted by $\mathcal{D} \subset R^3$, from a single subject. In real fMRI studies, it is common that multiple stimuli are present [11]. Under the assumption of the LTI system, the BOLD signal is the individual response to the sum of all stimuli convoluted with their associated HRFs. Let $\mathbf{X}(t) = (X_1(t), \ldots, X_m(t))^T$ be the sequence vector of m different stimuli and its associated HRF vector $\mathbf{H}(t, \mathbf{d}) = (H_1(t, \mathbf{d}), \ldots, H_m(t, \mathbf{d}))^T$. Specifically, in the time domain, most statistical models to estimate HRF under the presence of m different stimuli assume that

$$Y(t, \mathbf{d}) = \mathbf{H}(\cdot, \mathbf{d}) \otimes \mathbf{X}(t) + \epsilon(t, \mathbf{d}) = \int < \mathbf{H}(\mathbf{d}, t-u), \mathbf{X}(u) > du + \epsilon(t, \mathbf{d}), \quad (1)$$

where $\epsilon(t, \mathbf{d})$ is a spatial and frequent error process.

Instead of directly using model (1), our MASM focuses on the discrete Fourier coefficients of $Y(t,\mathbf{d})$, $\mathbf{H}(t,\mathbf{d})$, $\mathbf{X}(t)$, and $\epsilon(t,\mathbf{d})$, which are, respectively, denoted by $\phi_{\mathbf{Y}}(f_k,\mathbf{d})$, $\phi_{\mathbf{H}}(f_k,\mathbf{d})$, $\phi_{\mathbf{X}}(f_k)$, and $\phi_{\epsilon}(f_k,\mathbf{d})$ at the fundamental frequencies $f_k = k/T$ for $k = 0,\cdots,T-1$. Specifically, in the frequency domain, MASM assumes that

$$\phi_Y(f,\mathbf{d}) =< \phi_{\mathbf{H}}(f,\mathbf{d}), \phi_{\mathbf{X}}(f) > +\phi_{\epsilon}(f,\mathbf{d}), \tag{2}$$

where $\phi_{\mathbf{H}}(f,\mathbf{d}) = (\phi_{H_1}(f,\mathbf{d}),\ldots,\phi_{H_m}(f,\mathbf{d}))^T$, $\phi_{\mathbf{X}}(f) = (\phi_{X_1}(f),\ldots,\phi_{X_m}(f))^T$. An advantage of MASM in (2) is that the temporal correlation structure can be reduced since the Fourier coefficients are approximately asymptotically uncorrelated across frequencies under some regularity conditions [8,1]. Moreover, $\phi_{\epsilon}(f,\mathbf{d})$ is assumed to be a complex process with the zero mean function and a finite spatial covariance structure. MASM assumes that for each stimuli j and each voxel d, $\phi_{H_j}(f,\mathbf{d})$ is close to $\phi_{H_j}(f',\mathbf{d}')$ for some neighboring voxels $\mathbf{d}'$ and frequencies f'. This is essentially a spatial and frequent smoothness condition, which allows us to borrow information from neighboring voxels and frequencies. Equation (2) and different shape neighborhoods at different voxels are the two key novelties compared to the existing literatures [6].

2.2 Weighted Least Square Estimate

Our goal is to estimate the unknown functions $\{\phi_{\mathbf{H}}(f,\mathbf{d}) : \mathbf{d}\in\mathcal{D}, f\in[0,1]\}$ based on MASM and the Fourier transformed fMRI data $\mathcal{F}(\mathbf{Y}) = \{\phi_Y(f_k,\mathbf{d}) : k = 0,\cdots,T-1, \mathbf{d}\in\mathcal{D}\}$. To estimate $\phi_{\mathbf{H}}(f,\mathbf{d})$, we may combine all information at fundamental frequencies $f_k \in \eta_f(r) = (f-r, f+r)\cap\{k/T : k = 0,1,\ldots,T-1\}$ with $r > 0$ and voxels $\mathbf{d}'\in B(\mathbf{d},s)$, where $B(\mathbf{d},s)$ is a spherical neighborhood of voxel $\mathbf{d}$ with radius $s \geq 0$ to construct an approximation equation as follows:

$$\phi_Y(f_k,\mathbf{d}') \approx< \phi_{\mathbf{H}}(f,\mathbf{d}), \phi_{\mathbf{X}}(f_k) > +\phi_{\epsilon}(f_k,\mathbf{d}'). \tag{3}$$

Then to estimate $\phi_{H_j}(f,\mathbf{d})$ for $j = 1,\cdots,m$, respectively, we construct m local weighted functions $L_{[-j]}(\phi_{H_j}(f,\mathbf{d});r,s)$ as

$$\sum_{f_k\in\eta_f(r)}\sum_{\mathbf{d}'\in B(\mathbf{d},s)} |\phi_{Y[-j]}(f_k,\mathbf{d}') - \phi_{H_j}(f,\mathbf{d})\phi_{X_j}(f_k)|^2\tilde{\omega}_j(\mathbf{d},\mathbf{d}',f,f_k;r,s) \tag{4}$$

for $j = 1,\cdots,m$, where $\phi_{Y[-j]}(f_k,\mathbf{d}') = \phi_Y(f_k,\mathbf{d}') - \sum_{l\neq j}\phi_{H_l}(f_k,\mathbf{d}')\phi_{X_j}(f_k)$. Moreover, weight $\tilde{\omega}_j(\mathbf{d},\mathbf{d}',f,f_k;r,s)$ characterizes the physical distance between $(f,\mathbf{d})$ and $(f_k,\mathbf{d}')$ and the similarity between $\phi_{H_j}(f,\mathbf{d})$ and $\phi_{H_j}(f_k,\mathbf{d}')$. The procedure for determining all weights $\tilde{\omega}_j(\mathbf{d},\mathbf{d}',f,f_k;r,s)$ will be given later. We derive a recursive formula to update the estimates $\hat{\phi}_{H_j}(f,\mathbf{d})$ and $\mathrm{Var}(\hat{\phi}_{H_j}(f,\mathbf{d}))$ for $j = 1,\ldots,m$ based on any fixed weights $\{\tilde{\omega}_j(\mathbf{d},\mathbf{d}',f,f_k;r,s) : \mathbf{d}'\in B(\mathbf{d},s), f_k\in\eta_f(r)\}$. We obtain $\hat{\phi}_{H_j}(f,\mathbf{d})$ by differentiating $L_{[-j]}(\phi_{H_j}(f,\mathbf{d});r,s)$ as follows:

$$\frac{\sum_{f_k\in\eta_f(r)}\sum_{\mathbf{d}^T\in B(d,s)}\tilde{\omega}_j(\mathbf{d},\mathbf{d}',f,f_k;r,s)\overline{\phi_{X_j}(f_k)}\phi_{Y[-j]}(f_k,\mathbf{d}')}{\sum_{f_k\in\eta_f(r)}\sum_{\mathbf{d}'\in B(\mathbf{d},s)}\tilde{\omega}_j(\mathbf{d},\mathbf{d}',f,f_k;r,s)\phi_{X_j}(f_k)\overline{\phi_{X_j}(f_k)}} \tag{5}$$

where $\overline{\phi_{X_j}(f_k)}$ is the conjugate of $\phi_{X_j}(f_k)$. We approximate $\mathrm{Var}(\hat{\phi}_{H_j}(f,\mathbf{d}))$ as

$$\frac{\{\sum_{f_k\in\eta_f(r)} |\sum_{\mathbf{d}'\in B(\mathbf{d},s)} \tilde{\omega}_j(\mathbf{d},\mathbf{d}',f,f_k;r,s)\overline{\phi_X(f_k)}\hat{\phi}_{j\epsilon}(f_k,\mathbf{d}')|^2\}}{\{\sum_{f_k\in\eta_f(r)}\sum_{\mathbf{d}'\in B(\mathbf{d},s)} \tilde{\omega}_j(\mathbf{d},\mathbf{d}',f,f_k;r,s)\phi_X(f_k)\overline{\phi_X(f_k)}\}^2}, \tag{6}$$

where $\hat{\phi}_{j\epsilon}(f_k,\mathbf{d}') = \phi_Y(f_k,\mathbf{d}') - \sum_{l\neq j}\hat{\phi}_{H_l}(f_k,\mathbf{d}')\phi_{X_j}(f_k)$. Based on $\hat{\phi}_{H_j}(f,\mathbf{d})$ for $j=1,\cdots,m$, we can get

$$\tilde{H}_j(t,\mathbf{d}) = \frac{1}{T}\sum_{k=0}^{T-1}\hat{\phi}_{H_j}(f_k,\mathbf{d})\exp(i2\pi t f_k) \text{ for any } \mathbf{d}\in\mathcal{D} \text{ and } t. \tag{7}$$

2.3 Mutliscale Adaptive Estimation Procedure

We use a multiscale adaptive estimation (MAE) procedure to determine $\{\tilde{\boldsymbol{\omega}}_j(\cdot): j=1,\cdots,m\}$ and then estimate $\{\phi_{\mathbf{H}}(f,\mathbf{d}):\mathbf{d}\in\mathcal{D}, f\in[0,1]\}$. MAE borrows the multiscale adaptive strategy from the well-known Propagation-Separation (PS) approach [10,5]. MAE starts with building two sequences of nested spheres with spatial radii $s_0=0<s_1<\cdots<s_S$ and frequent radii $0<r_0<r_1<\ldots<r_S$. The key idea of MAE for multiple stimuli is to sequentially and recursively compute $\hat{\phi}_{H_j}(f,\mathbf{d})$ from $j=1$ increasing to m. Generally, MAE consists of four key steps: initialization, weight adaption, recursive estimation and stopping check. In the initialization step, we set $s_0=0$, $r_0>0$, say $r_0=5/T$, and the weighting scheme $\tilde{\omega}_j(\mathbf{d},\mathbf{d},f,f_k;r_0,s_0)=K_{loc}(|f-f_k|/r_0)$. We also set up another series $\{r_s=r_{s-1}+b_r: s=1,\cdots,S\}$ as the frequent radii with a constant value b_r, say, $b_r=2/T$. Then we apply the backfitting algorithm to iteratively update $\hat{\phi}^{(0)}_{H_j}(f,\mathbf{d})$ and obtain an estimate of $\mathrm{Var}(\hat{\phi}^{(0)}_{H_j}(f,\mathbf{d}))$ for $j=1,\cdots,m$ until convergence.

In the weight adaptation step, for $s>0$, we set $\tilde{\omega}^{(i)}_j(\mathbf{d},\mathbf{d}',f,f_k;r_s,s_s)$ as

$$K_{loc}(||\mathbf{d}-\mathbf{d}'||_2/s_s)K_{loc}(|f-f_k|/r_s)K_{st}\left(\frac{||\hat{\phi}^{(i-1)}_{H_j}(f,\mathbf{d})-\hat{\phi}^{(i-1)}_{H_j}(f_k,\mathbf{d}')||}{\sqrt{||\mathrm{Var}(\hat{\phi}^{(i-1)}_{H_j}(f,\mathbf{d}))||}}\right), \tag{8}$$

where $||\cdot||$ is the norm operator and $||\cdot||_2$ is the L_2 norm. The functions $K_{loc}(x)$ and $K_{st}(x)$ are two kernel functions such as the Epanechnikov kernel [10,5].

In the recursive estimation step, at the i^{th} iteration, we compute $\hat{\phi}^{(i)}_{j\epsilon}(f,\mathbf{d}) = \phi_Y(f,\mathbf{d}) - \sum_{l\neq j}\hat{\phi}^{(i-1)}_{H_l}(f,\mathbf{d})\phi_{X_l}(f)$. Then based on weights $\tilde{\omega}^{(i)}_j(\mathbf{d},\mathbf{d}',f,f_k;r_s,s_s)$, we sequentially calculate $\hat{\phi}^{(i)}_{H_j}(f,\mathbf{d})$ and approximate $\mathrm{Var}(\hat{\phi}^{(i)}_{H_j}(f,\mathbf{d}))$ according to (5) and (6).

In the stop checking step, after the i_0-th iteration, we calculate the adaptive Neyman test statistic, denoted by $W^{(i)}_j(\mathbf{d};s_s,r_u)$, for the j-th stimulus to test difference between $\hat{\phi}^{(i)}_{H_j}(\mathbf{d}) = \{\hat{\phi}^{(i)}_{H_j}(f_0,\mathbf{d}),\ldots,\hat{\phi}^{(i)}_{H_j}(f_{T-1},\mathbf{d})\}$ and $\hat{\phi}^{(i-1)}_{H_j}(\mathbf{d}) = \{\hat{\phi}^{(i-1)}_{H_j}(f_0,\mathbf{d}),\ldots,\hat{\phi}^{(i-1)}_{H_j}(f_{T-1},\mathbf{d})\}$. If $W^{(i)}_j(\mathbf{d};s_s,r_s)$ is significant, then we set $\hat{\phi}^{(s)}_{H_j}(f,\mathbf{d}) = \hat{\phi}^{(i-1)}_{H_j}(f,\mathbf{d})$ for all $s\geq i$ at voxel d.

Finally, when $s = S$, we report the final $\hat{\phi}_{H_j}^{(S)}(f, \mathbf{d})$ at all fundamental frequencies and substitute them into (7) to calculate $\hat{H}_j^{(S)}(t, \mathbf{d})$ across voxels $\mathbf{d} \in \mathcal{D}$ for all $j = 1, \cdots, m$. After obtaining HRFs for all stimuli, we may calculate their summary statistics including amplitude/height (H), time-to-peak (T), and full-width at half-max (W) and then carry out group-level statistical inference, say to test whether H significantly differs from 0, on the images of these estimated summary statistics [6].

3 Results

Simulation. We conducted a set of Monte Carlo simulations to examine the finite sample performance of MASM and MAE and compared them with several existing HRF models. We simulated the data at 200 time points (i.e., $t = 1, 2, \cdots, 200$) with a 40×40 phanton image containing 9 regions of activation-circles with varying radius at each time point. These 9 regions were also grouped into three different BOLD patterns with each group consisting of three circles, which have the same true signal series. The three true HRFs were set as zeros for $t > 15$ and otherwise for $t > 0$ according to

$$h_j(t) = A_j \cdot (t/d_{j1})^{a_{j1}} \exp\left(-(t - d_{j1})/b_{j1}\right) - c(t/d_{j2})^{a_{j2}} \exp\left(-(t - d_{j2})/b_{j2}\right)$$

with $(A_1, A_2, A_3) = (1, 5, 3)$, $c = 0.35$, $(a_{11}, a_{12}) = (6, 12)$, $(a_{21}, a_{22}) = (4, 8)$, $(a_{31}, a_{32}) = (5, 10)$, $(b_{j1}, b_{j2}) = (0.9, 0.9)$ and $(d_{j1}, d_{j2}) = (a_{j1} * b_{j1}, a_{j2} * b_{j2})$ for $j = 1, 2, 3$. The boxcars consisting of either zeros or ones were independently generated from a Bernoulli random generator with the successful rate=0.15, denoted by $X_j(t)$, $j = 1, 2, 3$. So the true BOLD signals were simulated as $Y(t) = \sum_{j=1}^{3} H_j \otimes X_j(t)$. The signals in each group of the activation-circles were scaled to be $Y_1(t) = Y(t)/6$, $Y_2(t) = Y(t)/4$ and $Y_3(t) = Y(t)/2$, respectively. The noise $\epsilon(t)$ were generated from a Gaussian distribution with mean zero and standard deviation $\sigma = 0.2$. Note that it is straightforward to embed AR noise to simulate the serial autocorrelation. Finally, the simulated BOLD signal was set as $Y_j(t) + \epsilon(t)$ for $j = 1, 2, 3$. In this simulation, the smallest signal-to-noise rate (SNR) is around 0.6.

In order to determine the signal patterns, we implemented some EM-based clustering method with ignoring the details for the sake of space and then computed the average of the estimates in each cluster. The estimates of the clustered HRFs are displayed in Fig. 1 (b.1-3 and c.1-3). It seems that our algorithm can simultaneously recover the correct HRFs in all active regions.

To evaluate our method, we compared MASM with some state-of-art methods in [6], which include (i) SPMs canonical HRF (denoted as GAM); (ii) the semi-parametric smooth finite impulse response (FIR) model (sFIR); and (iii) the inverse logit model (IL). Subsequently, we evaluated the HRF estimates by computing H, T, and W as the potential measure of response magnitude, latency and duration of neuronal activity. It has been reported in [6] that IL is one of the best methods in accurately estimating H, T, and W.

Let $D = \sum_{i=1}^{N}\sum_{j=1}^{M}(|\hat{x}_{ij} - x_0| - |\hat{y}_{ij} - x_0|)/(NM)$, where $\hat{x}_{ij}$ and $\hat{y}_{ij}$, respectively, denote the statistics H, T, or W, calculated from MASM and from the other three methods, and x_0 represents the true value of H (or T, W) in the different active regions corresponding to the different event sequences. Moreover, N and M, respectively, represent the numbers of replications and voxels in active regions with $N = 100$. We computed the average absolute error differences between our method and the other three methods. Generally, the negative value of D indicates that our method outperforms other methods. Table 1 reveals that MASM can provide more accurate estimates of the HRF summary statistics than the other three methods.

Table 1. Comparisons of the differences of the absolute errors between our method with sFIR, IL and GAM, respectively. C1, C2 and C3 denote the 1st, 2nd and 3rd sequences of events, respectively. A1, A2 and A3 denote the 1st, 2nd and 3rd active regions. Values in the parentheses are the standard errors. H=Height, W=Width, T=Time-to-Peak.

Param		sFIR				IR				GAM		
		A1	A2	A3		A1	A2	A3		A1	A2	A3
H	C1	-0.02 (0.045)	-0.04 (0.045)	-0.09 (0.060)	C1	-0.11 (0.0819)	-0.16 (0.115)	-0.32 (0.218)	C1	-0.13 (0.039)	-0.20 (0.044)	-0.42 (0.059)
	C2	-0.02 (0.041)	-0.04 (0.043)	-0.10 (0.057)	C2	-0.09 (0.0699)	-0.14 (0.096)	-0.27 (0.175)	C2	-0.11 (0.041)	-0.17 (0.045)	-0.36 (0.057)
	C3	-0.01 (0.046)	-0.02 (0.047)	-0.07 (0.064)	C3	-0.07 (0.0673)	-0.10 (0.089)	-0.21 (0.161)	C3	-0.07 (0.038)	-0.11 (0.045)	-0.25 (0.066)
		A1	A2	A3		A1	A2	A3		A1	A2	A3
T	C1	-0.08 (0.722)	0.05 (0.326)	0.01 (0.073)	C1	-3.74 (3.313)	-3.49 (3.309)	-3.19 (3.292)	C1	-2.50 (0.425)	-2.66 (0.298)	-2.84 (0.070)
	C2	-0.05 (0.685)	0.07 (0.292)	0.01 (0.069)	C2	-3.50 (3.440)	-3.34 (3.475)	-2.88 (3.419)	C2	-2.57 (0.452)	-2.75 (0.287)	-2.91 (0.069)
	C3	-0.55 (1.159)	-0.10 (0.513)	-0.10 (0.513)	C3	-3.54 (3.404)	-3.26 (3.430)	-3.03 (3.415)	C3	-2.46 (0.577)	-2.66 (0.427)	-2.87 (0.254)
		A1	A2	A3		A1	A2	A3		A1	A2	A3
W	C1	-0.20 (0.671)	-0.28 (0.596)	-0.42 (0.518)	C1	-1.70 (2.122)	-1.73 (2.127)	-1.66 (2.094)	C1	-3.33 (0.623)	-3.41 (0.576)	-3.46 (0.515)
	C2	-0.38 (0.760)	-0.41 (0.597)	-0.49 (0.513)	C2	-1.78 (2.143)	-1.80 (2.099)	-1.79 (2.018)	C2	-3.35 (0.634)	-3.41 (0.575)	-3.49 (0.512)
	C3	-0.32 (0.870)	-0.33 (0.741)	-0.48 (0.658)	C3	-1.79 (2.179)	-1.85 (2.123)	-2.08 (2.221)	C3	-3.30 (0.713)	-3.42 (0.677)	-3.63 (0.550)

From Table 1, amongst the tested HRF estimation alternatives, the sFIR seems to provide the closest results. Thus we pick sFIR to make another comparison. We applied the Gaussian smoothing with FWHM equal 4mm to the original simulated data before running sFIR and compared them to MASM without using the Gaussian smoothing. An evaluation statistics for the j^{th} voxel is given by $D_j = \frac{1}{100}\sum_{i=1}^{100}(|\hat{x}_{ij} - x_0| - |\hat{y}_{ij} - x_0|)$. The comparison results for the parameter H given in Fig. 1 (d.1-3) as a representative reveals that MASM outperforms sFIR, especially on the boundary voxels as Gaussian smoothing blurred them.

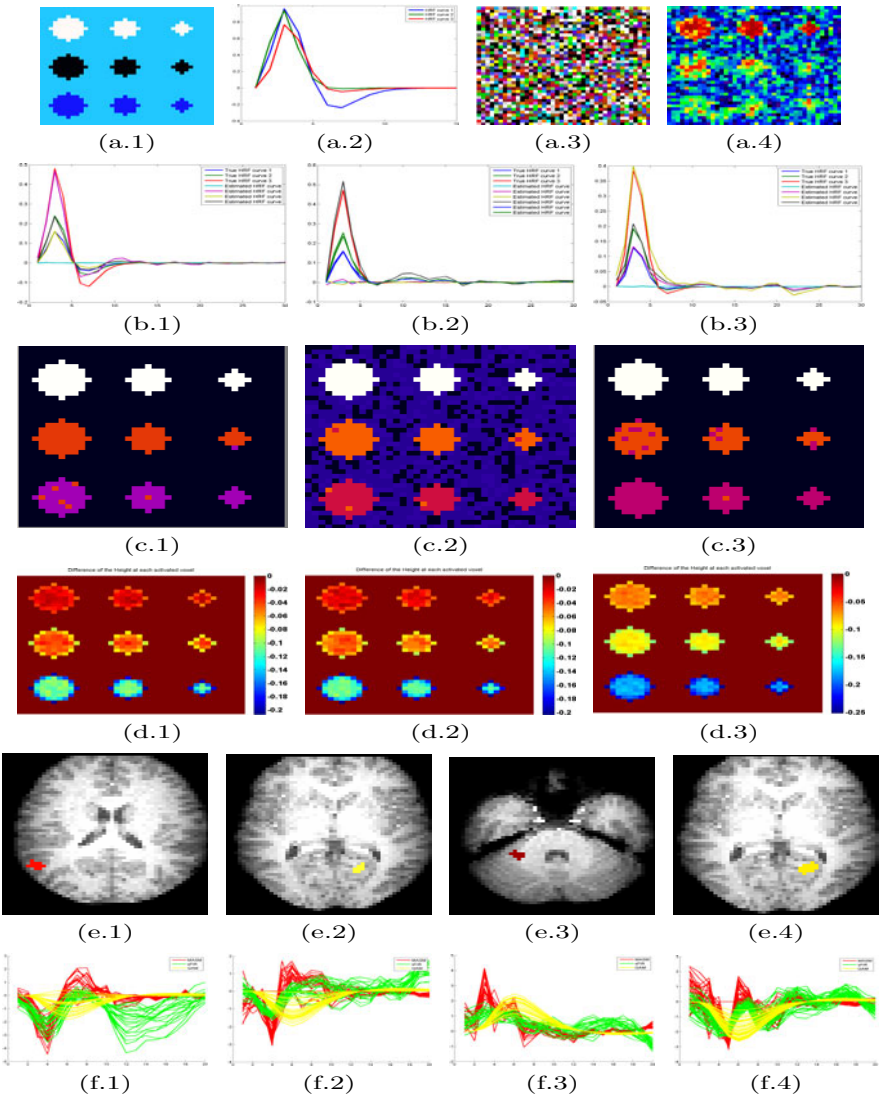

Fig. 1. Set-up of Simulation: (a.1) a temporal cut of the true images; (a.2) the true curves of HRF: $h_1(t)$, $h_2(t)$ and $h_3(t)$; (a.3) a temporal cut of the simulated images; (a.4) the Gaussian smooth result. The estimated results: estimates of HRF for the (b.1 and c.1) 1st; (b.2 and c.2) 2nd; (b.3 and c.3) 3rd sequence of events. The row (b.1-3) is the average estimated HRF in each cluster. The row (c.1-3) is the recovered pattern relative to each sequence of events. The comparison statistics D_j with sFIR: (d.1-3) the difference of estimated Height(H) at each voxel for the three stimulus sequences. The color bar denotes the value of D_j for the j^{th} voxel. Data analysis results: (e.1)-(e.4) the slices containing ROIs (colored ones) of the F maps for the 1st-4th stimulus sequences, respectively; (f.1)-(f.4) estimated HRFs at the significant ROIs corresponding each condition from MASM (red), sFIR (green) and GAM(yellow).

Real Data. We used a subject from a study designed for the investigation of the memory relationship with four different stimulus sequences. We used Statistical Parametric Mapping (SPM) [4] to preprocess the fMRI and MRI images and apply a global signal regression method to detrend the fMRI time series. The F-statistics maps were generated by SPM to test the activation regions triggered by four sequences of stimulus events. For each stimulus, we set a threshold with p value less than 0.01 and the extension $K = 20$ to find significant regions of interest (ROIs). We plotted the estimated HRFs by using MASM, sFIR and GAM in these ROIs and chose one of them in each brain mapping as a demonstration to compare MASM with sFIR and GAM (Fig. 1 (e.1-4)). Based on the SPM findings with GAM, we found the deactive ROIs in Figures (e.1), (e.2) and (e.4) and the active ROIs in Figure (e.3). They are consistent with those ROIs obtained from the other two methods. HRFs calculated from MASM and sFIR have similar H, W, and T, which are different with those statistics obtained from HRFs based on GAM (Fig. 1 (f.1-4)). This result is consistent with our simulation result (see Table 1). Furthermore, we found that HRFs calculated from sFIR had big variation at their tails compared to those calculated from MASM. It may indicate that MASM is an accurate estimation method for reconstructing HRFs in fMRI.

References

1. Bai, P., Truong, Y., Huang, X.: Nonparametric estimation of hemodynamic response function: a frequency domain approach. In: Optimality: The Third Erich L. Lehmann Symposium. IMS Lecture Notes-Monograph Series, vol. 57, pp. 190–215 (2009)
2. Bowman, F.D., Patel, R., Lu, C.: Methods for detecting functional classifications in neuroimaging data. Human Brain Mapping 23, 109–119 (2004)
3. Breiman, L., Friedman, J.H.: Estimating optimal transformations for multiple regression and correlation. Journal of American Statistical Assocication 80, 580–598 (1985)
4. Friston, K.J.: Statistical Parametric Mapping: the Analysis of Functional Brain Images. Academic Press, London (2007)
5. Katkovnik, V., Spokoiny, V.: Spatially adaptive estimation via fitted local likelihood techniques. IEEE Transactions on Signal Processing 56, 873–886 (2008)
6. Lindquist, M.A., Wager, T.D.: Validity and Power in hemodynamic response modeling: A comparison study and a new approach. Human Brain Mapping 28, 764–784 (2007)
7. Makni, S., Idier, J., Vincent, T., Thirion, B., Dehaene-Lambertz, G., Ciuciu, P.: A fully Bayesian approach to the parcel-based detection-estimation of brain activity in fMRI. NeuroImage 41, 941–969 (2008)
8. Marchini, J.L., Ripley, B.D.: A new statistical approach to detecting significant activation in functional MRI. NeuroImage 12, 366–380 (2000)
9. Risser, L., Vincent, T., Ciuciu, P., Idier, J.: Extrapolation scheme for fast ISING field partition functions estimation. In: Yang, G.-Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (eds.) MICCAI 2009. LNCS, vol. 5761, pp. 975–983. Springer, Heidelberg (2009)
10. Tabelow, K., Polzehl, J., Voss, H.U., Spokoiny, V.: Analyzing fMRI experiments with structural adaptive smoothing procedure. NeuroImage 33, 55–62 (2006)
11. Vincent, T., Risser, L., Ciuciu, P.: Spatially adaptive mixture modeling for analysis of fMRI time series. IEEE Transactions on Medical Imaging 29, 1059–1074 (2010)

Identification of Individuals with MCI via Multimodality Connectivity Networks

Chong-Yaw Wee[1], Pew-Thian Yap[1], Daoqiang Zhang[1], Kevin Denny[2], Lihong Wang[2], and Dinggang Shen[1]

[1] Department of Radiology and Biomedical Research Imaging Center (BRIC), University of North Carolina at Chapel Hill, NC, USA
[2] Brain Imaging and Analysis Center (BIAC), Duke University Medical Center, Durham, NC, USA

Abstract. Mild cognitive impairment (MCI), often an early stage of Alzheimer's disease (AD), is difficult to diagnose due to the subtlety of cognitive impairment. Recent emergence of reliable network characterization techniques based on diffusion tensor imaging (DTI) and resting-state functional magnetic resonance imaging (rs-fMRI) has made the understanding of neurological disorders at a whole-brain connectivity level possible, providing new avenues for brain classification. Taking a multi-kernel SVM, we attempt to integrate these two imaging modalities for improving classification performance. Our results indicate that the multimodality classification approach performs better than the single modality approach, with statistically significant improvement in accuracy. It was also found that the prefrontal cortex, orbitofrontal cortex, temporal pole, anterior and posterior cingulate gyrus, precuneus, amygdala, thalamus, parahippocampal gyrus and insula regions provided the most discriminant features for classification, in line with the results reported in previous studies. The multimodality classification approach allows more accurate early detection of brain abnormalities with larger sensitivity, and is important for treatment management of potential AD patients.

1 Introduction

Alzheimer's disease (AD) is one of the most prevalent dementia in older adults worldwide, characterized by cognitive and intellectual deficits that is serious enough to interfere daily life. Accurate diagnosis of AD is crucial for early treatment. Mild cognitive impairment (MCI), often an early stage of AD, is a good target for early diagnosis and therapeutic interventions of AD. Nevertheless, MCI is difficult to diagnose due to the subtlety of cognitive impairment.

Recently, several imaging modalities have been proven to be effective in AD and MCI diagnosis, including diffusion tensor imaging (DTI) [6], magnetic resonance imaging (MRI) [9], resting-state functional MRI (rs-fMRI) [8] and positron emission tomography (PET) [5]. Nevertheless, most existing pattern classification methods use only each of these modalities independently for AD and MCI

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 277–284, 2011.

diagnosis. More effort should be made in integrating of two or more modalities since combining complementary information of different biomarkers can be useful for improving diagnosis.

To the best of our knowledge, DTI and rs-fMRI have not been combined at a network level for identifying individuals with MCI, although they have been employed separately with reasonably good classification performance. The current study is the first attempt to integrate these two modalities to identify individuals with MCI from normal controls. We seek to validate whether complementary structural and functional information can be combined to improve classification performance. We will also report brain regions that contribute most to the classification performance. While confirming findings of previous studies, this paper sheds new light on the effectiveness of applying multimodality information for diagnosis of progressive neurodegenerative disorders.

2 Materials and Methods

The current study involved 27 participants, 10 individuals with MCI and 17 socio-demographically matched healthy controls. Informed consent was obtained from all participants, and the experimental protocols were approved by the institutional ethics board. Confirmation of diagnosis for all subjects was made via expert consensus panels. Demographic information of the participants is shown in Table 1.

Table 1. Demographic information of the participants involved in this study. Group difference was assessed using two-sample t-tests.

Group	**MCI**	**Normal**	p**-value**
No. of subjects	10	17	–
Gender	5M/5F	8M/9F	–
Age (mean ± SD)	74.2 ± 8.6	72.1 ± 8.2	0.5372
Years of education (mean ± SD)	17.7 ± 4.2	16.3 ± 2.4	0.2804
MMSE (mean ± SD)	28.4 ± 1.5	29.4 ± 0.9	0.0405

2.1 Data Acquisition

Data acquisition was performed using a 3.0-Tesla GE Signa EXCITE scanner. Diffusion-weighted images of each participant were acquired with 25-direction diffusion-weighted whole-brain volumes with the following parameters: $b = 0$, 1000 s/mm^2, flip angle = 90°, TR/TE = 17000/78 ms, 72 slices, imaging matrix = 128×128, FOV = 256×256 mm^2, resulting in a voxel dimension of 2×2×2 mm^3 reconstructed resolution. Resting-state functional images were acquired using a SENSE inverse-spiral pulse sequence with the following parameters: flip angle = 77°, TR/TE = 2000/32 ms, 34 slices, imaging matrix = 64 × 64, FOV = 256 × 256 mm^2, resulting in a voxel resolution of 4 × 4 × 4 mm^3. All the subjects were told to keep their eyes open and stare at a fixation cross in the middle of the screen during scanning, which lasted for 5 minutes.

2.2 Data Post-processing and Network Construction

The DTI images were first parcellated into 90 regions by propagating the automated anatomical labeling (AAL) ROIs [11] to each image using a deformable DTI registration algorithm. Whole-brain streamline fiber tractography was then performed on each image using ExploreDTI [7] with the following parameters: starting/stopping FA = 0.45/0.25, minimum/maximum fiber length = 20/400 mm. The number of fibers passing through each pair of regions was counted. Two regions were considered as anatomically connected if fibers passing through their respective masks were present, giving us connection topology of the network. On top of the fiber count based connectivity network, averages of on-fiber fractional anisotropy (FA), mean diffusivity (MD) and principal diffusivity values were computed to form another 5 connectivity networks with the same topology but conveying different biophysical properties. Examples of the constructed connectivity maps are shown in Figure 1.

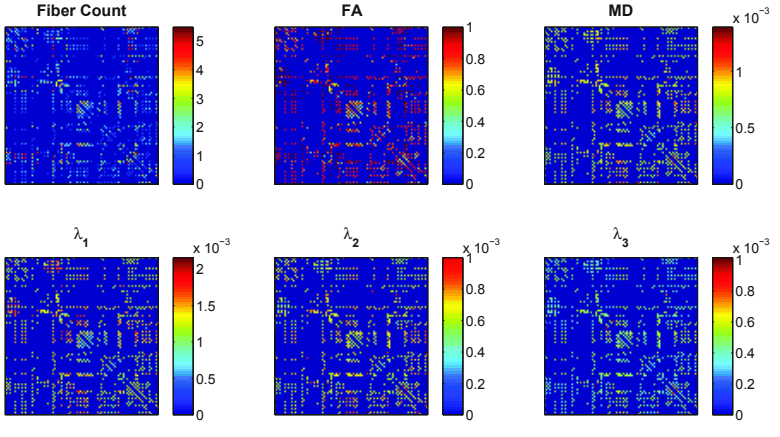

Fig. 1. Connectivity maps constructed from DTI different physiological measures

Post-processing of the fMRI images, such as slice timing correction and head-motion correction were performed using the Statistical Parametric Mapping (SPM8, http://www.fil.ion.ucl.ac.uk.spm) software package. The images were then masked with their respective gray matter (GM) masks, which were created by segmenting the GM regions from their T1-weighted images. This eliminated the white matter and cerebrospinal fluid from contributing to the fMRI time series, which contain a relatively high proportion of noise caused by cardiac and respiratory cycles [12].

Then, we parcellated the brain space into 90 ROIs by warping the fMRI images to the AAL template. For each subject, the mean time series was computed for each GM-masked region. Temporal band-pass filtering of frequency interval ($0.025 \leq f \leq 0.100$Hz) was then performed to minimize the effects of low-frequency drift and high-frequency noise since the fMRI dynamics of neuronal

activities are most salient within this frequency interval. This frequency interval was further decomposed into 5 equally divided, non-overlapping frequency sub-bands, enabling a relatively frequency specific analysis of the regional mean time series via a multi-spectral characterization.

Functional connectivity, which examines interregional correlations in neuronal variability, was measured using pairwise Pearson correlation coefficients between the ROI pairs. Given a set of N random variables, the Pearson correlation matrix is a symmetric matrix in which each off-diagonal element is the correlation coefficient between a pair of variables. We considered the brain regions as a set of nodes and the correlation coefficients as signed weights on the set of edges. Fisher's r-to-z transformation was applied on the elements of the Pearson correlation matrix to improve the normality of Pearson correlation coefficients. Examples of the constructed functional connectivity maps for a normal control (NC) and an MCI patient are shown in the top and bottom rows of Figure 2, respectively.

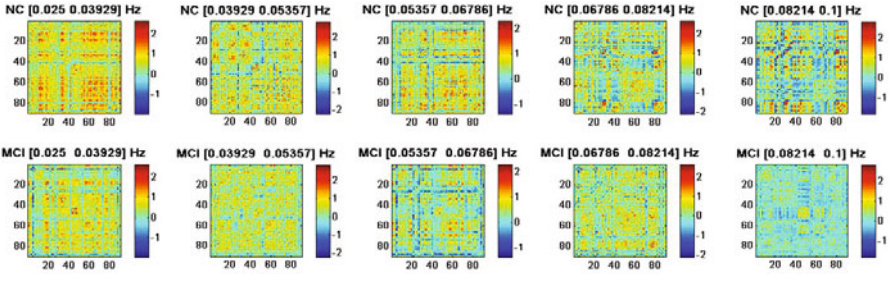

Fig. 2. Multi-spectral functional connectivity maps for NC (top) and MCI (bottom), respectively

2.3 Multi-Kernel SVM Multimodality Classification

The proposed framework for integrating DTI and rs-fMRI is divided into 3 stages: feature extraction, feature selection and multimodality data fusion. In the first stage, the weighted local clustering coefficient, a measure that quantifies the cliquishness of the nodes, is extracted from all connectivity maps as

$$f(p) = \frac{\sum_{q:q\neq p\in\zeta} 2t(p,q)}{k_p(k_p-1)}, \tag{1}$$

where ζ is the subnetwork comprising of k_p nodes directly connected to the p-th node, and $t(p,q)$ is the edge weight between the p-th node and q-th node. Hence, a total of 90 features can be obtained from each connectivity map, producing for each subject a pool of 540 and 450 features for DTI and rs-fMRI, respectively.

In the second stage, statistical t-test was performed to select features for classification. Features with p-values smaller than a predefined threshold will be selected from each individual modality before they were integrated using a

multi-kernel SVM algorithm. In the third stage, for n training samples with each of them is of M modalities, the multi-kernel SVM solves the following primal problem

$$\begin{aligned} &\min_{\mathbf{w}^{(m)},b,\xi_i} \quad \frac{1}{2}\sum_{m=1}^{M}\beta_m\left\|\mathbf{w}^{(m)}\right\|^2 + C\sum_{i=1}^{n}\xi_i, \qquad (2)\\ &\text{s.t. } y_i\left[\sum_{m=1}^{M}\beta_m\Big((\mathbf{w}^{(m)})^T\phi^{(m)}(\mathbf{X}_i^{(m)})+b\Big)\right] \geq 1-\xi_i;\ \text{with } \xi_i \geq 0, i=1,\ldots,n \end{aligned}$$

with $\mathbf{X}_i^{(m)} = \{\mathbf{x}_{i,1}^{(m)},\ldots,\mathbf{x}_{i,D}^{(m)}\}$ denotes a feature vector of the m-th modality of the i-th sample (D = number of maps in the m-th modality, $\mathbf{x}_{i,d}^{(m)} = \{f_{i,d}^{(m)}(1),\ldots,f_{i,d}^{(m)}(90)\}$ and $d=1,\ldots,D$) and $y_i \in \{-1,1\}$ as its corresponding class label. Parameters ξ_i, C, b, $\mathbf{w}^{(m)}$, $\phi^{(m)}$ and $\beta_m \geq 0$ denote the distance of the i-th misclassified observation from its correct side of the margin, the model parameter that controls for the amount of constraint violations introduced by ξ_i, the bias term, the normal vector of hyperplane, the kernel-induced mapping function and the weighting factor of the m-th modality, respectively.

The dual form of multi-kernel SVM is solved as

$$\begin{aligned} &\max_{\alpha} \quad \sum_{i=1}^{n}\alpha_i - \frac{1}{2}\sum_{i,j}\alpha_i\alpha_j y_i y_j \sum_{m=1}^{M}\beta_m k^{(m)}(\mathbf{X}_i^{(m)},\mathbf{X}_j^{(m)}), \qquad (3)\\ &\text{s.t.} \quad \sum_{i=1}^{n}\alpha_i y_i = 0;\ \text{with } 0 \leq \alpha_i \leq C, i=1,\ldots,n \end{aligned}$$

where α_i is the Lagrange multiplier and $k^{(m)}(\mathbf{X}_i^{(m)},\mathbf{X}_j^{(m)}) = \phi^{(m)}(\mathbf{X}_i^{(m)})\,\phi^{(m)}(\mathbf{X}_j^{(m)})$ is the kernel function for a pair of training samples of the m-th modality.

Given a new test sample $\mathbf{X} = \{\mathbf{X}^{(1)},\ldots,\mathbf{X}^{(M)}\}$ and let the kernel between the new test sample and each training sample of the m-th modality be $k^{(m)}(\mathbf{X}_i^{(m)}, \mathbf{X}^{(m)}) = \phi^{(m)}(\mathbf{X}_i^{(m)})\phi^{(m)}(\mathbf{X}^{(m)})$, the decision function for the predicted label can be determined as

$$F(\mathbf{X}) = \text{sign}\left(\sum_{i=1}^{n} y_i\alpha_i \sum_{m=1}^{M}\beta_m k^{(m)}(\mathbf{X}_i^{(m)},\mathbf{X}^{(m)}) + b\right). \qquad (4)$$

The multi-kernel SVM can be naturally embedded into the conventional single kernel SVM by interpreting $k(\mathbf{X}_i,\mathbf{X}) = \sum_m \beta_m k^{(m)}(\mathbf{X}_i^{(m)},\mathbf{X}^{(m)})$ as a mixed kernel between the multimodality training sample $\mathbf{X}_i$ and the test sample $\mathbf{X}$. A linear kernel SVM classifier based on the LIBSVM library [1] was employed to demonstrate that the improvement obtained is due mainly to the complementary information of different modalities.

3 Experimental Results

The proposed multi-kernel SVM based multimodality classification approach was compared with the single modality approach and the direct data fusion method. In the single modality approach, only features selected from a single imaging modality (DTI or rs-fMRI) were applied for SVM classifier training. While in the direct data fusion method, all 990 features, which including the DTI and rs-fMRI features, were first concatenated into a long vector before feature selection using the t-test. In the multi-kernel approach, the optimal weighting factor, β_m, and SVM parameter, C, were determined via grid search over a fixed range. In all compared approaches, the p-value of 0.01 was used for each training set during cross-validation. Classification was performed based on the z-scores of the features.

We employed various measures to evaluate the diagnostic power of the compared methods. The Youden's index, Balanced ACcuracy (BAC) and F-score are defined respectively as [10]

$$\mathbf{Y} = \text{Sensitivity} + \text{Specificity} - 1 = \frac{\text{TP}}{\text{FP}} + \frac{\text{TN}}{\text{FN}} - 1, \tag{5}$$

$$\mathbf{BAC} = \frac{1}{2} \times \left[\frac{\text{TP}}{\text{TP} + \text{FN}} + \frac{\text{TN}}{\text{TN} + \text{FP}}\right], \tag{6}$$

$$\mathbf{F} = 2 \times \left[\frac{\text{precision} \times \text{recall}}{\text{precision} + \text{recall}}\right], \tag{7}$$

where

$$\text{precision} = \frac{\text{TP}}{\text{TP} + \text{FP}};\ \text{recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}$$

with TP, TN, FP and FN denoting the true positive, true negative, false positive and false negative, respectively.

Comparison was performed via leave-one-out cross-validation due to limited number of available samples. Classification performance for individuals with MCI using single and multimodality connectivity networks are summarized in Table 2. The proposed method yields a classification accuracy of 96.59%, which is an increment of at least 7.41% from that of the single modality approach and the direct data fusion method.

Table 2. Classification performance of single and multimodality connectivity networks. (ACC = ACCuracy; SEN = SENsitivity; SPE = SPEcificity; Youden = Youden's index)

Method	ACC (%)	SEN(%)	SPE(%)	AUC	Youden	BAC	F-score
fMRI	55.56	60.00	52.94	0.6059	0.1294	0.5647	0.5000
DTI	88.89	80.00	94.12	0.9353	0.7412	0.8706	0.8421
Direct	88.89	90.00	88.24	0.9118	0.7824	0.8912	0.8571
kernel	96.30	100.00	94.12	0.9529	0.9412	0.9524	0.9706

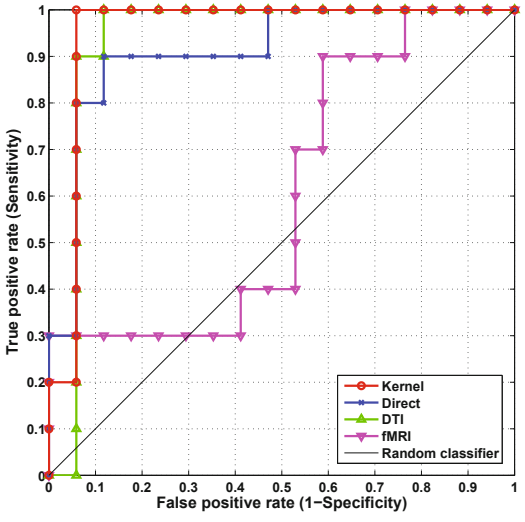

Fig. 3. ROC curves.

A cross-validation estimation of the generalization performance shows an area of 0.9529 under the receiver operating characteristic (ROC) curve (AUC), indicating excellent diagnostic power. ROC curves for all compared methods are shown in Figure 3.

The most discriminant regions that were selected in the course of the classification include the prefrontal cortex and orbitofrontal cortex [5], temporal pole [4], anterior and posterior cingulate gyrus, precuneus and insula [3], amygdala [2], thalamus [13], parahippocampal gyrus [8], in line with results reported in previous studies. The discriminative regions were selected based on whether their clustering coefficients were selected to be employed for classification. The order of the regions is determined by their p-values; a smaller p-value indicates higher discriminative power.

4 Discussions and Conclusion

We investigated the diagnostic power of multimodality classification of the DTI and rs-fMRI connectivity maps in identifying individuals with MCI. In this framework, linear SVM classifiers were trained using a mixed kernel that was constructed from the individual kernels of multiple modalities. This framework shows better classification performance, and justifies the hypothesis that the DTI and rs-fMRI contain complementary information, each of them indispensable particularly for achieving better diagnostic power. Furthermore, the higher sensitivity rate of the proposed approach is an important improvement since the cost of misclassifying individuals with MCI is significantly larger than that of misclassifying normal controls. Correct diagnosis of individuals with MCI enables early treatment management of potential AD patients possible and thus reduces the MCI to AD conversion rate. The promising results indicate that the

proposed framework can provide an alternative and complementary approach for clinical diagnosis of brain degeneration.

References

1. Chang, C.C., Lin, C.J.: LIBSVM: a library for support vector machines (2001), software, available at http://www.csie.ntu.edu.tw/~cjlin/libsvm
2. Dai, W., Lopez, O.L., Carmichael, O.T., Becker, J.T., Kuller, L.H., Gach, H.M.: Mild cognitive impairment and Alzheimer disease: Patterns of altered cerebral blood flow at MR imaging. Radiology 250, 856–866 (2009)
3. Davatzikos, C., Bhatt, P., Shaw, L.M., Batmanghelich, K.N., Trojanowski, J.Q.: Prediction of MCI to AD conversion, via MRI, CSF biomarkers and pattern classification. Neurobiol Aging (2010)
4. Flavio, N., Dario, S., Silvia, M., Nicola, G., Arnoldo, P., Andrea, B., Barbara, D., Stig, A.L., Guido, R., Marco, P.: Principal component analysis of FDG PET in amnestic MCI. Eur. J. Nucl. Med. Mol. Imaging 35(12), 2191–2202 (2008)
5. Grady, C.L., McIntosh, A.R., Beig, S., Keightley, M.L., Burian, H., Black, S.E.: Evidence from functional neuroimaging of a compensatory prefrontal network in Alzheimer's disease. J. Neurosci 23(3), 986–993 (2003)
6. Haller, S., Nguyen, D., Rodriguez, C., Emch, J., Gold, G., Bartsch, A., Lovblad, K., Giannakopoulos, P.: Individual prediction of cognitive decline in mild cognitive impairment using support vector machine-based analysis of diffusion tensor imaging data. J. Alzheimers Dis. 22(1), 315–527 (2010)
7. Leemans, A., Jeurissen, B., Sijbers, J., Jones, D.K.: ExploreDTI: A graphical toolbox for processing, analyzing, and visualizing diffusion MR data. In: 17th Annual Meeting of Intl. Soc. Mag. Reson. Med., Hawaii, USA, p. 3537 (2009)
8. Machulda, M.M., Senjem, M.L., Weigand, S.D., Smith, G.E., Ivnik, R.J., Boeve, B.F., Knopman, D.S., Petersen, R.C., Jack Jr., C.R.: Functional MRI changes in amnestic and non-amnestic MCI during encoding and recognition tasks. J. Int. Neuropsych. Soc. 15(3), 372–382 (2009)
9. McEvoy, L.K., Fennema-Notestine, C., Roddey, J.C., Hagler Jr., D.J., Holland, D., Karow, D.S., Pung, C.J., Brewer, J.B., Dale, A.M.: Alzheimer disease: Quantitative structural neuroimaging for detection and prediction of clinical and structural changes in mild cognitive impairment. Radiology 251, 195–205 (2009)
10. Sokolova, M.V., Japkowicz, N., Szpakowicz, S.: Beyond accuracy, F-score and ROC: A family of discriminant measures for performance evaluation. In: Sattar, A., Kang, B.-h. (eds.) AI 2006. LNCS (LNAI), vol. 4304, pp. 1015–1021. Springer, Heidelberg (2006)
11. Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., Joliot, M.: Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. Neuroimage 15, 273–289 (2002)
12. Van Dijk, K.R.A., Hedden, T., Venkataraman, A., Evans, K.C., Lazar, S.W., Buckner, R.L.: Intrinsic functional connectivity as a tool for human connectomics: Theory, properties and optimization. J. Neurophysiol 103, 297–321 (2010)
13. Wang, Z., Jia, X., Liang, P., Qi, Z., Yang, Y., Zhou, W., Li, K.: Changes in thalamus connectivity in mild cognitive impairment: Evidence from resting state fMRI. Eur. J Radiol. (2011)

Connectivity-Informed fMRI Activation Detection

Bernard Ng[1], Rafeef Abugharbieh[1], Gael Varoquaux[2],
Jean Baptiste Poline[2], and Bertrand Thirion[2]

[1] Biomedical Signal and Image Computing Lab, UBC, Canada
[2] Parietal Team, Neurospin, INRIA Saclay-Ile-de-France, France
bernardyng@gmail.com

Abstract. A growing interest has emerged in studying the correlation structure of spontaneous and task-induced brain activity to elucidate the functional architecture of the brain. In particular, functional networks estimated from resting state (RS) data were shown to exhibit high resemblance to those evoked by stimuli. Motivated by these findings, we propose a novel generative model that integrates RS-connectivity and stimulus-evoked responses under a unified analytical framework. Our model permits exact closed-form solutions for both the posterior activation effect estimates and the model evidence. To learn RS networks, graphical LASSO and the oracle approximating shrinkage technique are deployed. On a cohort of 65 subjects, we demonstrate increased sensitivity in fMRI activation detection using our connectivity-informed model over the standard univariate approach. Our results thus provide further evidence for the presence of an intrinsic relationship between brain activity during rest and task, the exploitation of which enables higher detection power in task-driven studies.

Keywords: activation detection, connectivity prior, fMRI, resting-state.

1 Introduction

The standard approach for analyzing functional magnetic resonance imaging (fMRI) data involves comparing each brain voxel independently against an expected response to estimate the likelihood of activation [1]. However, accumulating evidence suggests that brain function is also mediated through the interactions between brain regions in what is referred to as functional connectivity [2]. Although incorporation of functional connectivity may provide activation models that better reflect the nature of brain activity, few existing methods have been designed for such purposes [3]. Instead, research efforts have largely focused on modeling the spatial structure of fMRI data through spatial regularization [4-7]. These methods indirectly account for local voxel interactions, but long-range interactions are completely neglected.

The discovery of structure in ongoing brain activity in the absence of external stimulus has ignited enormous research interest [8-10]. Many detected resting-state (RS) networks were found to exhibit high resemblance to those engaged during task performance [9]. These findings suggest potential relationships between brain activity during task and rest, which may serve as an additional source of information for detecting stimulus-evoked activation. In particular, task-based fMRI data typically

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 285–292, 2011.

display rather low signal-to-noise ratio (SNR), especially in patients due to difficulties in performing certain tasks [10]. Since acquiring RS data requires minimal task demands and RS data are less susceptible to behavioural confounds [10], extracting connectivity priors from RS data to inform activation detection may enhance detection sensitivity, which is especially beneficial for studying diseased populations.

In this paper, we propose a novel generative model for integrating connectivity and task-evoked responses under a unified analytical framework. Assuming brain regions displaying functional correlations at rest are more likely to co-activate during task, incorporating RS-connectivity information should improve task activation detection. To learn the connectivity structure from RS data, we employ and compare graphical LASSO (GL) [11] and the oracle approximating shrinkage (OAS) technique [12]. The resulting connectivity information is then used as a prior on the task activation effects. Unlike most existing generative models [4-7] that employ either approximate inference or sampling methods for parameter estimation, our model has the distinct advantage of permitting exact closed-form solutions for both the posterior activation effect estimates and the model evidence. We apply our model on data from a cohort of 65 subjects undergoing a variety of experimental conditions and show increased sensitivity in detecting group activation over the standard univariate method.

2 Connectivity-Informed Activation Model

Motivated by recent studies that showed high resemblance between functional networks detected during rest and task [9], we propose to integrate RS-connectivity and task-evoked responses under a unified generative model. Specifically, let Y be an $d \times n$ matrix containing the intensity time courses of d voxels or d brain regions of a subject. Our proposed model can be summarized as follows:

$$Y \sim N(AX, V_1) = \frac{1}{\left|2\pi V_1\right|^{\frac{n}{2}}} \exp\left(-\frac{1}{2} tr\left((Y - AX)^T V_1^{-1} (Y - AX) \right) \right), \tag{1}$$

$$A \sim MN(M, V_2, K) = \frac{\left|K\right|^{\frac{d}{2}}}{\left|2\pi\alpha^{-1} V_2\right|^{\frac{m}{2}}} \exp\left(-\frac{1}{2} \alpha \cdot tr\left((A - M)^T V_2^{-1} (A - M) K \right) \right), \tag{2}$$

where X is an $m \times n$ matrix with m regressors along the rows for modeling stimulus-invoked response and confounds [1]. A is the $d \times m$ activation effect matrix to be estimated (Section 2.1). V_1 is the $d \times d$ covariance matrix of Y, V_2 is an $d \times d$ covariance matrix modeling the correlations between the activation effects of the d brain regions, and K is an $m \times m$ covariance matrix modeling the correlations between the experimental conditions. $MN(M,V_2,K)$ denotes the matrix normal distribution, which serves as the conjugate prior of (1) [13] with α controlling the degree of influence of this prior on A (Section 2.3). To ensure that our estimate of A is invariant to affine transformations on Y and X, we set M to $0_{d \times m}$ and K to XX^{T} [13]. Setting M to $0_{d \times m}$ also ensures that the activation effect estimates will not be biased towards non-zero values, which could induce false detections. V_1 and V_2 are assumed to be known. Compared to the model in [13] where V_1 and V_2 are assumed to be equal, we show that exact

closed-form solutions for the posterior estimate of A and the model evidence can be derived even for the more general case with V_1 and V_2 being distinct. Permitting distinct V_1 and V_2 accounts for how Y and A might have different correlation structures. Also, we hypothesize that brain regions functionally correlated at rest are more likely to co-activate during task performance. We thus set V_2 to the covariance estimates learned from RS data (Section 2.2). We assume $V_1 = I_{d\times d}$ as conventionally done for analytical simplicity, and defer learning V_1 from data for future work.

2.1 Posterior Activation Effects Estimation

To estimate A, we first derive the joint distribution of Y and A by taking the product of (1) and (2) and isolating the terms involving A from those involving Y:

$$\begin{aligned} P(Y,A \mid X,V_1,V_2) &\propto \exp\left(-\frac{1}{2}tr\left((Y-AX)^T V_1^{-1}(Y-AX)+\alpha X^T A^T V_2^{-1} AX\right)\right) \\ &= \exp\left(-\frac{1}{2}tr\left(\left(V_1^{-1}+\alpha V_2^{-1}\right)\left(AXX^T A^T - 2\left(V_1^{-1}+\alpha V_2^{-1}\right)^{-1} V_1^{-1} YX^T A^T\right)+V_1^{-1}YY^T\right)\right) \end{aligned}. \quad (3)$$

Since the terms involving A take a quadratic form, the posterior distribution of A is again a matrix normal distribution, as expected from the conjugacy of (1) and (2). By completing the square, the maximum a posteriori mean of A can be derived:

$$M_A = \left(V_1^{-1}+\alpha V_2^{-1}\right)^{-1} V_1^{-1} YX^T \left(XX^T\right)^{-1}. \quad (4)$$

Computing M_A requires inverting V_1 and V_2, which is unstable if V_1 and V_2 are set as sample covariance estimated from data with more brain regions than time points. In this work, we assume $V_1 = I_{d\times d}$, and employ and compare two state-of-the-art techniques, namely GL and OAS, for obtaining a well-conditioned V_2.

2.2 Functional Connectivity Estimation

***Graphical LASSO*:** Given a sample covariance matrix, S, computed from data that are assumed to follow a centered multivariate Gaussian distribution, one can estimate a well-conditioned sparse inverse covariance matrix, $\hat{\Lambda}$, by minimizing the negative log-likelihood of the data distribution over the space of positive definite (p.d.) matrices while imposing an l_1 penalty on $\hat{\Lambda}$ [11]:

$$\hat{\Lambda} = \arg\min_{\Lambda>0} tr(\Lambda S) - \log\det(\Lambda) + \lambda\|\Lambda\|_1, \quad (5)$$

where $\|\cdot\|_1$ is the element-wise l_1 norm and λ controls the level of sparsity. Enforcing sparsity simplifies interpretation, since $\hat{\Lambda}_{ij} = 0$ implies brain regions i and j are not connected. To optimize (5), we employ the Two-Metric Projection method [14].

***OAS*:** Assume the data for estimating the ground truth covariance, Σ, is generated from a multivariate Gaussian distribution. The most well-conditioned covariance estimate of Σ is $F = tr(S)/d \cdot I_{d\times d}$ [12]. The idea of OAS is to shrink the ill-conditioned sample covariance, S, towards F so that a well-conditioned covariance estimate, $\hat{\Sigma}$, can be obtained. Specifically, OAS optimizes the following cost function [12]:

$$\hat{\rho} = \arg\min_{\rho} E\left[\left\|\hat{\Sigma} - \Sigma\right\|_F^2\right], \ s.t. \ \hat{\Sigma} = (1-\rho)S + \rho F \ , \tag{6}$$

where ρ controls the amount of shrinkage with the optimal value given by [12]:

$$\hat{\rho} = \frac{\left(1-\frac{2}{d}\right)tr\left(S^2\right) + tr^2(S)}{\left(n+1-\frac{2}{d}\right)\left(tr\left(S^2\right) - \frac{tr^2(S)}{d}\right)} . \tag{7}$$

Thus, no parameter selection is required and the inverse covariance matrix can be obtained by inverting $\hat{\Sigma}$ with stable inversion guaranteed.

2.3 Hyper-Parameters Estimation

A critical hyper-parameter in our model is α, which controls the degree of influence of the connectivity prior on the activation effect estimates. To set α, we first derive the model evidence by integrating $P(Y,A|\alpha,V_1,V_2)$ over A:

$$P(Y \mid \alpha, V_1, V_2) = \frac{\left|\left(V_1^{-1} + \alpha V_2^{-1}\right)^{-1}\right|^{\frac{m}{2}}}{\left|2\pi V_1\right|^{\frac{n}{2}}\left|\alpha^{-1}V_2\right|^{\frac{m}{2}}} \exp\left(-\frac{1}{2}tr\left(V_1^{-1}\left(YY^T \ldots - YX^T\left(XX^T\right)^{-1}XY^T V_1^{-1}\left(V_1^{-1} + \alpha V_2^{-1}\right)^{-1}\right)\right)\right) . \tag{8}$$

Since V_2^{-1} estimated from GL or OAS is p.d., V_2^{-1} can be decomposed into $Q\Gamma Q^{\mathrm{T}}$ where Q are the eigenvectors of V_2^{-1} and Γ contains the eigenvalues of V_2^{-1} along the diagonal. By exploiting this property of p.d. matrix and that V_1 is assumed to be $I_{d\times d}$ hence $V_1^{-1} = I_{d\times d} = QI_{d\times d}Q^{\mathrm{T}}$, the log of (8) can be simplified into:

$$\ln P(Y \mid \alpha, V_1, V_2) = -\frac{m}{2}\sum_{i=1}^{d}\left(\ln(1+\alpha\gamma_i) - \ln(\alpha\gamma_i) - \frac{1}{m}\frac{B_{ii}}{1+\alpha\gamma_i}\right) + C \ , \tag{9}$$

$$B = Q^T YX^T\left(XX^T\right)^{-1}XYQ \ , \tag{10}$$

where γ_i is the i^{th} eigenvalue of V_2^{-1}, and terms that do not depend on α, V_1, or V_2 are grouped into C. Since (9) is a single variable function of α, the optimal α at which (9) is maximized can be efficiently determined using any generic optimization routines.

To optimize the choice of λ for the case where V_2^{-1} is estimated using GL, we define a grid of λ values at which the percentage of non-zero elements in V_2^{-1} roughly ranges from 10% to 90%. For each V_2^{-1} associated with a given λ, we determine the optimal α at which (9) is maximized and compute the log model evidence. The λ associated with the largest log model evidence is then taken as the optimal λ.

3 Materials

***Synthetic Data*:** We generated 300 synthetic datasets based on our proposed model. Each dataset consisted of 10 subjects. Each subject's data comprised 100 regions with the first 20 regions set to be mutually correlated and activated across all subjects. The next 20 regions were set to be mutually correlated but not activated to test if our model will falsely declare correlated regions as activated. The remaining regions were not correlated nor activated. To simulate this scenario, we generated 100 signal time courses, where the first 20 time courses were sine functions with Gaussian noise added. The next 20 time courses were cosine functions with Gaussian noise added. The remaining 60 time courses were simply Gaussian noise. The empirical covariance of these signal time courses was taken as the group covariance of the activation effects and the RS networks, Ω^g. A random p.d. matrix was added to Ω^g to introduce inter-subject variability into each subject's covariance matrix, Ω^i. The degree of variability was controlled by restricting the Kullback-Leibler divergence of $N(0,\Omega^g)$ and $N(0,\Omega^i)$ to be ~1.5. For each subject i, 100 RS time courses of length N_t were drawn from $N(0,\Omega^i)$. We set N_t to 25 to emulate the typical situation where the number of regions exceeds N_t. To generate task data, samples of A were drawn from (2) with V_2 set to Ω^i. To simulate activation, the means of the first 20 regions in (2) were set to a small positive number, δ, that depended on the SNR, δ^2/σ^2. Gaussian noise was added to δ to further introduce inter-subject variability. The resulting A were then used to draw samples of Y from (1) with V_1 set to $\sigma^2 I$ and X being boxcar functions convolved with the hemodynamic response function [1]. Three SNRs were tested: 0.25, 0.5, and 0.75, with 100 synthetic datasets generated at each SNR level.

***Real Data*:** fMRI data were collected from 65 healthy subjects at multiple imaging centers. Each subject performed 10 language, computation, and sensorimotor tasks over a period of ~5 min (140 brain volumes) similar to those in [15]. RS data of ~7 min duration (187 brain volumes) were also collected. 3T scanners from multiple manufacturers were used for acquiring the data with TR = 2200 ms, TE = 30 ms, and flip angle = 75°. Standard preprocessing, including slice timing correction, motion correction, temporal detrending, and spatial normalization, were performed on the task-based data using the SPM8 software. Similar preprocessing was performed on the RS data except a band-pass filter with cutoff frequencies at 0.01 to 0.1 Hz was applied to isolate the signal of interest [8]. Signals from cerebrospinal fluid and white-matter voxels were regressed out from the gray-matter voxels.

To ensure stable sparse inverse covariance estimation using GL [11], we reduced the dimension of the data by grouping the voxels into 1000 parcels. Specifically, we concatenated the RS voxel time courses across subjects and applied the parcellation technique of [16] to generate a group parcellation map. Each subject's brain images (in normalized space) were then parcellated using the group parcel labels. The mean voxel time courses of each parcel from rest and task were taken as the input to our model. To account for scanner variability across imaging centers, we normalized the parcel time courses by subtracting the mean and dividing by the standard deviation.

4 Results and Discussion

***Validation*:** For validation, we compare the sensitivity of our model with connectivity prior estimated from OAS and GL against the ridge regression model (i.e. V_2 set to I) and the standard univariate model [1] in detecting group activation. We denote these models as OAS-CM, GL-CM, R-UM, and S-UM. Since the ground truth activated brain regions are unknown for real data, we employ the max-t permutation test [17] to enforce strict control on false positive rate (FPR) so that we can safely base our validation on the number of parcels detected. We note that for each permutation, the *entire* activation effect map of each subject is multiplied by 1 or -1 chosen at random. Hence, the spatial covariance structure of the activation effect maps is preserved. Also, we emphasize that our validation criterion is independent of the criterion used for optimizing the model parameters, which mitigates bias from being introduced.

***Synthetic Data*:** The mean receiver operating characteristics curves averaged over the synthetic datasets are shown in Fig. 1. At all FPR and SNR levels, OAS-CM and GL-CM achieved higher true positive rates (TPR) than R-UM and S-UM, thus confirming that given the activation effects are inherently correlated, our model can exploit this information to improve group activation detection. We note that there is less need for imposing a prior (which introduces a bias) at higher SNR since more signals are available to estimate A, hence the decrease in sensitivity for OAS-CM and GL-CM. Also, the higher sensitivity achieved by OAS-CM compared to GL-CM might be due to how GL tends to produce unstable covariance estimates for large-scale problems.

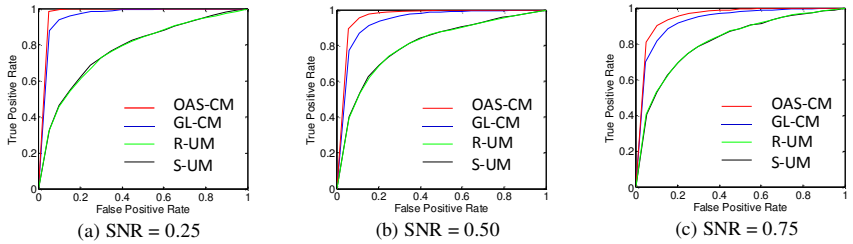

(a) SNR = 0.25 (b) SNR = 0.50 (c) SNR = 0.75

Fig. 1. Synthetic data results. OAS-CM (red) and GL-CM (blue) achieved higher TPR for all FPR and SNR levels than R-UM (green) and S-UM (black).

***Real Data*:** To test the generality of our model, we examined 21 contrasts between the 10 experimental conditions. Contrasts included computation vs. sentence processing task, auditory vs. visual task among others. Shown in Fig. 2(a) is the percentage of parcels detected with significant activation differences averaged over contrasts for p-value thresholds ranging from 0 to 0.5. Incorporating a connectivity prior using OAS-CM and GL-CM outperformed R-UM and S-UM even under the simplifying yet common assumption that $V_1 = I_{d\times d}$. We note that adding a shrinkage prior to control overfitting, as in the case of R-UM, only improved detection mildly. Thus, our results suggest an intrinsic relationship between the correlation structure of activation effects and RS-connectivity, and this relationship is consistent across subjects, hence the

improved group activation detection. Qualitatively, incorporating a RS-connectivity prior resulted in more detections of bilateral activation in brain regions implicated for the contrasts examined, examples of which are shown in Fig. 2(b) and (c).

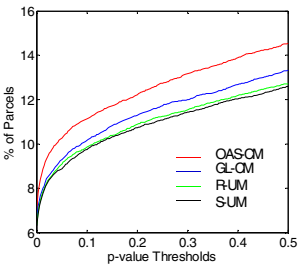

(a)

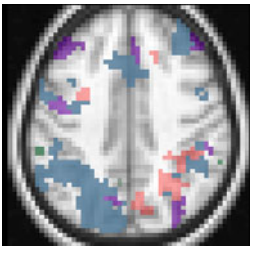

(b) p-value threshold = 0.01, corrected

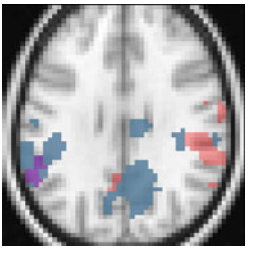

(c) p-value threshold = 0.01, corrected

Fig. 2. Real data results. (a) % of parcels with significant activation differences averaged across contrasts vs. *p*-value thresholds. (b) Parcels detected by contrasting computation against sentence processing task, and (c) auditory against visual task. Red = detected by only OAS-CM. Purple = detected by both OAS-CM and GL-CM. Blue = detected by all methods.

5 Conclusions

We proposed a novel generative model for integrating connectivity and stimulus-induced response under a single analytical framework. Our model permits exact closed-form solutions for the posterior activation effect estimates and the model evidence without resorting to approximate inference or computationally-expensive sampling methods. On real data, we demonstrated that integrating a RS-connectivity prior improves sensitivity in detecting group activation. Our results thus support that the correlation structure of task activation effects is pertinent to RS connectivity, and this relationship is common across subjects. The flexibility of our model permits other V_1 and V_2 to be easily examined. In particular, integrating an anatomical connectivity prior estimated from diffusion MRI data would be a promising direction to explore. We expect that our model can similarly improve intra-subject activation detection, which enables more refined subject-specific analysis for studying patient populations.

Acknowledgements. Jean Baptiste Poline was partly funded by the IMAGEN project, which receives research funding from the E. U. Community's FP6, LSHM-CT-2007-037286. This manuscript reflects only the authors' views and the Community is not liable for any use that may be made of the information contained therein.

References

1. Friston, K.J., Holmes, A.P., Worsley, K.J., Poline, J.B., Frith, C.D., Frackowiak, R.S.J.: Statistical Parametric Maps in Functional Imaging: A General Linear Approach. Hum. Brain Mapp. 2, 189–210 (1995)

2. Rogers, B.P., Morgan, V.L., Newton, A.T., Gore, J.C.: Assessing Functional Connectivity in the Human Brain by fMRI. Magn. Reson. Imaging 25, 1347–1357 (2007)
3. Ng, B., Hamarneh, G., Abugharbieh, R.: Detecting Brain Activation in fMRI Using Group Random Walker. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. 6362, pp. 331–338. Springer, Heidelberg (2010)
4. Descombes, X., Kruggel, F., von Cramon, D.Y.: Spatio-Temporal fMRI Analysis Using Markov Random Fields. IEEE Trans. Med. Imaging 17, 1028–1039 (1998)
5. Penny, W.D., Trujillo-Barreto, N.J., Friston, K.J.: Bayesian fMRI Time Series Analysis with Spatial Priors. NeuroImage 24, 350–362 (2005)
6. Harrison, L.M., Penny, W.D., Asburner, J., Trujillo-Barreto, N.J., Friston, K.J.: Diffusion-based Spatial Priors for Imaging. NeuroImage 38, 677–695 (2007)
7. Woolrich, M.W., Jenkinson, M., Brady, J.M., Smith, S.M.: Fully Bayesian Spatio-temporal Modeling of fMRI Data. IEEE Trans. Med. Imaging 23, 213–231 (2004)
8. Fox, M.D., Raichle, M.E.: Spontaneous Fluctuations in Brain Activity Observed with Functional Magnetic Resonance Imaging. Nat. Rev. Neurosci. 8, 700–711 (2007)
9. Smith, S.M., Fox, P.T., Miller, K.L., Glahn, D.C., Fox, P.M., Mackay, C.E., Filippini, N., Watkins, K.E., Toro, R., Laird, A.R., Beckmann, C.F.: Correspondence of the Brain's Functional Architecture During Activation and Rest. Proc. Natl. Acad. Sci. 106, 13040–13045 (2009)
10. Fox, M.D., Greicius, M.: Clinical Applications of Resting State Functional Connectivity. Front. Syst. Neurosci. 4, 19 (2010)
11. Friedman, J., Hastie, T., Tibshirani, R.: Sparse Inverse Covariance Estimation with the Graphical LASSO. Biostats. 9, 432–441 (2008)
12. Chen, Y., Wiesel, A., Eldar, Y.C., Hero, A.O.: Shrinkage Algorithms for MMSE Covariance Estimation. IEEE Trans. Sig. Proc. 58, 5016–5029 (2010)
13. Minka, T.P.: Bayesian Linear Regression. Technical report, MIT Media Lab (2001)
14. Schmidt, M., Fung, G., Rosales, R.: Fast Optimization Methods for L1 Regularization: A Comparative Study and Two New Approaches. In: Kok, J.N., Koronacki, J., Lopez de Mantaras, R., Matwin, S., Mladenič, D., Skowron, A. (eds.) ECML 2007. LNCS (LNAI), vol. 4701, pp. 286–297. Springer, Heidelberg (2007)
15. Pinel, P., Thirion, B., Meriaux, S., Jober, A., Serres, J., Le Bihan, D., Poline, J.B., Dehaene, S.: Fast Reproducible Identification and Large-scale Databasing of Individual Functional Cognitive Networks. BioMed. Central Neurosci. 8, 91 (2007)
16. Thirion, B., Flandin, G., Pinel, P., Roche, A., Ciuciu, P., Poline, J.B.: Dealing with the Shortcomings of Spatial Normalization: Multi-subject Parcellation of fMRI Datasets. Hum. Brain Mapp. 27, 678–693 (2006)
17. Nichols, T., Hayasaka, S.: Controlling the Familywise Error Rate in Functional Neuroimaging: a Comparative Review. Stat. Methods Med. Research 12, 419–446 (2003)

A Stochastic Linear Model for fMRI Activation Analyses

Leigh A. Johnston[1,2], Maria Gavrilescu[3], and Gary F. Egan[4,5]

[1] Electrical & Electronic Engineering, University of Melbourne, & NICTA Victorian Research Laboratory, Australia
[2] Howard Florey Institute, Florey Neuroscience Institutes, Australia
[3] Defence Science and Technology Orginisation, Australia
[4] Centre for Neuroscience, University of Melbourne, Australia
[5] Monash University, Australia

Abstract. Purpose: The debate regarding how best to model variability of the hemodynamic response function in fMRI data has focussed on the linear vs. nonlinear nature of the optimal signal model, with few studies exploring the deterministic vs. stochastic nature of the dynamics. We propose a stochastic linear model (SLM) of the hemodynamic signal and noise dynamics to more robustly infer fMRI activation estimates. **Methods:** The SLM models the hemodynamic signal by an exogenous input autoregressive model driven by Gaussian state noise. Activation weights are inferred by a joint state-parameter iterative coordinate descent algorithm based on the Kalman smoother. **Results:** The SLM produced more accurate parameter estimates than the GLM for event-design simulated data. In application to block-design experimental visuo-motor task fMRI data, the SLM resulted in more punctate and well-defined motor cortex activation maps than the GLM, and was able to track variations in the hemodynamics, as expected from a stochastic model. **Conclusions:** We demonstrate in application to both simulated and experimental fMRI data that in comparison to the GLM, the SLM produces more flexible, consistent and enhanced fMRI activation estimates.

Keywords: fMRI, hemodynamic signal, Kalman filter, state space model.

1 Introduction

The mapping from neuronal activation to measured blood oxygenation level dependent (BOLD) signal in fMRI involves a complex interplay between physiological and physical processes, most of which are yet to be fully understood [7]. The hemodynamic response function (HRF) is known to vary across subjects, sessions, scans, and brain regions [6]. Approaches to dealing with this variability in HRF have, with few exceptions, focussed on parameterised, deterministic models of the HRF, eg. [14]. Parameterised models may appear at first glance to be stochastic, given that estimated parameters are assumed to be drawn from prior distributions. However, given knowledge of the parameters, the resultant HRFs are deterministic. A notable exception to the deterministic component

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 293–301, 2011.

models are the techniques originated by Ciuciu et al. [3], in which smooth HRFs are ensured by placing a Gaussian prior on the norm of the second derivative of the HRF, with estimates enforced to start and end at zero.

In contrast to deterministic BOLD signal models, we propose a stochastic linear model (SLM) of the BOLD signal, in which noise is used to drive the state dynamics directly. Unlike the methods of Ciuciu et al., our model is stochastic and parameterised, with the BOLD signal estimated voxelwise along with the activation weight. The SLM is a state space model, similar in form to the bilinear dynamical systems (BDS) proposed by [12] and recently extended by [10]. The key difference between the SLM and BDS is that BDS models the supposed neuronal activity via stochastic state dynamics, which are then convolved with HRF kernels to form the BOLD signal. Given the vast difference in temporal scale between fMRI observations and neuronal activity, the SLM does not attempt to separate neuronal activity from the hemodynamic response. Rather, the SLM is a fully stochastic model of the hemodynamic BOLD signal.

The primary focus of the current paper is to present the SLM as a stochastic variant of the general linear model (GLM), and to compare the robustness, statistical significance and consistency of activation estimates derived from the SLM and GLM approaches, rather than to explicitly estimate a correlate of neuronal activity.

2 Methods

The Stochastic Linear Model: Let $k = 1, \ldots, T$ be the discrete time index of repetition time (TR) intervals. The system input is a binary stimulus sequence, u_k, representing task-on/task-off, and system output is the observed voxelwise BOLD signal, y_k. The SLM is a state space model consisting of an ARX state equation modelling the hemodynamic signal, x_k, and an observation equation:

$$\begin{aligned} x_k &= ax_{k-1} + bu_k + w_k, \quad w_k \sim N(0, \sigma_w^2), \\ y_k &= x_k + v_k + r_k + e_k, \quad v_k \sim AR(c, \sigma_n^2),\ r_k \sim MA(d, \sigma_\xi^2),\ e_k \sim N(0, \sigma_e^2). \end{aligned} \tag{1}$$

For full generality, we consider additive autoregressive (AR), moving average (MA) and Gaussian white noise processes in the observation equation.

The SLM can be represented in canonical state space form by embedding the ARMA noise processes in the state vector. For the MA noise process, state embedding occurs via defining a dummy variable, $\tilde{r}_k = d\xi_k$. Let $\mathbf{x}_k = [x_k, v_k, r_k, \tilde{r}_k]'$ be the state vector and $\mathbf{w}_k = [w_k, n_k, \xi_k, d\xi_k]'$ be the noise vector. Then the SLM canonical state space form is,

$$\mathbf{x}_k = A\mathbf{x}_{k-1} + \mathbf{b}u_k + \mathbf{w}_k, \qquad \mathbf{w}_k \sim N(0, Q), \tag{2}$$

$$y_k = \mathbf{g}'\mathbf{x}_k + e_k, \qquad e_k \sim N(0, \sigma_e^2), \tag{3}$$

$$A = \begin{bmatrix} a & 0 & 0 & 0 \\ 0 & c & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} b \\ 0 \\ 0 \\ 0 \end{bmatrix}, \quad Q = \begin{bmatrix} \sigma_w^2 & 0 & 0 & 0 \\ 0 & \sigma_n^2 & 0 & 0 \\ 0 & 0 & \sigma_\xi^2 & d\,\sigma_\xi^2 \\ 0 & 0 & d\,\sigma_\xi^2 & d^2\,\sigma_\xi^2 \end{bmatrix}, \quad \mathbf{g} = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 0 \end{bmatrix}.$$

The full state vector, $\mathbf{x}_k$, and the ARMA noise processes, v_k and r_k, are estimated from the noisy BOLD signal, y_k.

There are three key differences between the SLM and the GLM. Firstly, the SLM dynamics of the neurovascular signal, x_k, are driven by a state noise process, w_k, whereas the GLM neurovascular signal dynamics are deterministic. Secondly, unlike the GLM, the SLM does not contain an explicit activation parameter, β, or hemodynamic response function, h_k. This permits a more flexible modelling of the neurovascular signal where the HRF may be unknown. In order to map the strength of activation, we define an SLM activation weight following in Eq. (4). The third key difference is that in the GLM, attempts are made to mitigate the effect of observation noise through pre-whitening or pre-colouring of the measured BOLD signals [4]. In contrast to these noise suppression procedures, the SLM noise processes are explicitly estimated, through embedding them in the state dynamics.

SLM Activation Estimation: While the activation weight, β^{GLM}, is explicit in a GLM formulation, there is no analogous parameter in the SLM state or observation equation, Eq. (1). Rather we define the activation weight to be

$$\beta^{SLM} \triangleq \frac{b}{1-a}, \tag{4}$$

where a and b are SLM parameters in Eq. (1). This definition is derived directly from the SLM state equation, Eq. (1), by noting that x_k contains stimulus terms in an arithmetic sequence due to the autoregression. We therefore define the activation weight to be the limiting sum of the arithmetic sequence, Eq. (4). Activation weights for higher order AR models can be defined similarly using the characteristic polynomial in the denominator of (4).

State and parameter estimation: The proposed SLM algorithm estimates the noise processes explicitly, rather than suppressing their effect through pre-whitening. The SLM parameters $\Theta = \{a, b, c, d, \sigma_e^2, \sigma_w^2, \sigma_n^2, \sigma_\xi^2\}$ are unknown, as are the state vectors, $\mathbf{x}_k$. Given that both states and parameters are unknown, we seek to estimate the MAP estimates of the states and parameters,

$$\{X, \Theta\}^{MAP} = \arg\max_{X,\Theta} p(X, \Theta | y_1, \ldots, y_T) \tag{5}$$

$$\text{where} \quad p(X, \Theta | y_1, \ldots, y_T) = \prod_{k=1}^{T} \frac{1}{\sqrt{2\pi\sigma_e^2}} \exp\left(-\frac{(y_k - \mathbf{g}'\mathbf{x}_k)^2}{2\sigma_e^2}\right) \times \tag{6}$$

$$\frac{1}{(2\pi)^2 |Q|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}(\mathbf{x}_k - A\mathbf{x}_{k-1} - \mathbf{b}u_k)' Q^{-1} (\mathbf{x}_k - A\mathbf{x}_{k-1} - \mathbf{b}u_k)\right)$$

This MAP optimisation is intractable, and therefore a suboptimal estimation algorithm must be applied. We have chosen iterative coordinate descent algorithm [9], known to converge to a local minimum of the posterior, in which a Kalman smoother and parameter updating are consecutively applied to the system, starting from initial parameter estimates. The computational time of

the iterative estimation is less than 1 second per voxel time-series (for a scan length of 240 volumes), implemented in Matlab and run on a standard 2GHz PC. The algorithm is entirely parallelisable as each voxel analysis is independent of others.

Simulated fMRI Data: Event design data was generated by the balloon model, through random sampling over a realistic range of parameters [5]. SNR was varied by change in neuronal efficacy, ε. Parameters were generated in each of 50 runs according to

$$T = 150s, TR = 2s, a_1 = 1, a_2 = 3.37, \alpha \sim N(0.33, 0.1), E_0\ 0.3 + 0.3U(0,1),$$
$$\tau_s \sim N(1.54, 0.2), \tau_f \sim N(2.46, 0.2), \tau_0 \sim N(1, 0.2), V_0/\sigma_e^2 = 5$$

An estimated activation weight was computed for each run based on the peak of the noiseless BOLD signal, for comparison with $\beta_G LM$ and $\beta_S LM$. Mean square error between the noiseless BOLD signal and the estimated signal from the GLM and SLM methods was similarly computed.

Experimental fMRI Visuomotor Task Data: Three healthy controls were scanned (3T Siemens TRIO, MCRI Melbourne) while performing a visuomotor task, 222 EPI images, in-plane: 3.125x3.125mm^2, TR=1.6s, FA=90°, TE=20ms, 24 axial slices of 5mm with 0.5mm gap. The visuomotor task was a simple block design of alternating periods (30s each). Images were motion corrected and spatially smoothed with a 6mm isotropic Gaussian kernel. The supplementary motor area (SMA) was expertly delineated using landmarks in each subject's native space. The SLM's iterative coordinate descent algorithm was applied for 20 iterations, determined to be sufficient to allow the parameter estimates to converge. For comparison, a GLM using the SPM canonical HRF and temporal derivatives was implemented.

3 Results

Simulations: Simulated data that has been generated using the Balloon model, rather than the actual models underlying SLM or GLM, avoids the problem of bias toward a particular algorithm in the results. The Balloon model, being based on physiology, is also a more realistic model from which to generate data. Varying the 'efficacy' parameter, ε, in the Balloon model [5], varies the cerebral blood flow and hence BOLD signal. Fig. 1 demonstrates that for the simulated event design data, the SLM achieves more accurate estimates of activation weights than the GLM, which consistently underestimates the signal level. The SLM returned a lower MSE than the GLM for almost all efficacies (SNRs), excepting the extremely low SNR regime; as expected the SLM will tend to overfit the signal and match to noise in very low SNR BOLD signals that are dominated by noise.

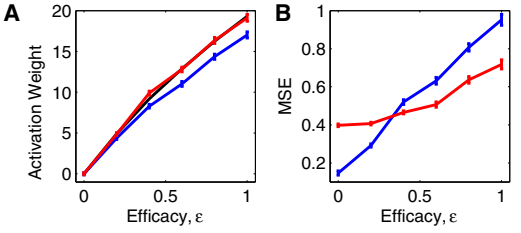

Fig. 1. Performance comparison between SLM and GLM on simulated event design data, simulated for six signal efficacies, $\varepsilon \in [0, 1]$. **A.** Ground-truth activation weights (black) vs SLM (red) and GLM (blue). Where not visible, black line underlies SLM red line. **B.** MSE between ground-truth time-series and each of SLM (red) and GLM (blue) estimates. All estimates display mean ± standard error bars over 50 runs.

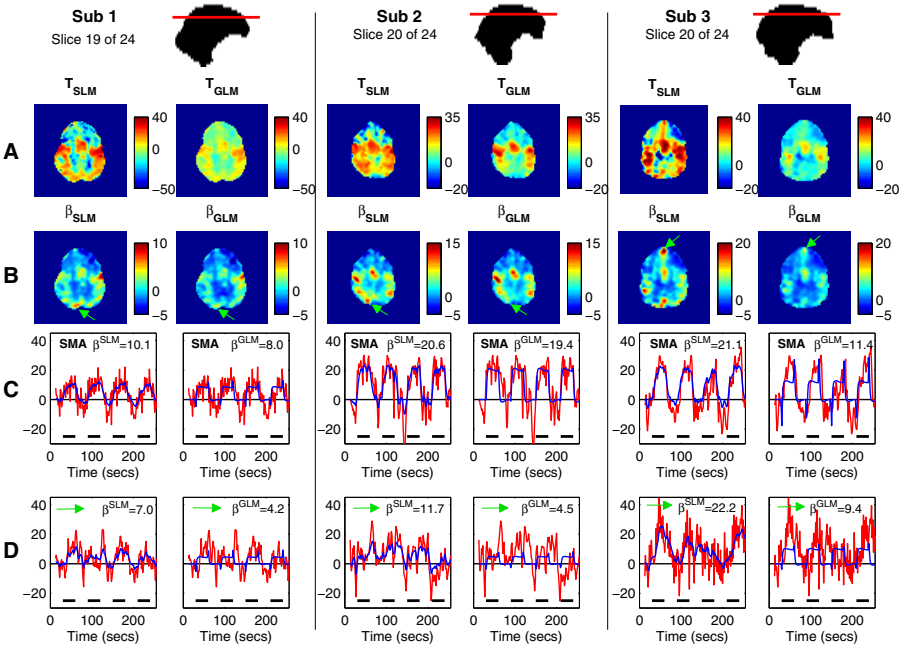

Fig. 2. Estimation of BOLD signal dynamics by SLM and GLM across 3 subjects for a visuomotor block design task. **A.** T statistics. **B.** Activation estimates. **C.** SMA time-series (red) and estimates (blue). **D.** Time-series (red and estimates (blue) for voxels labelled by green arrows in **B**.

Experimental Results: Fig. 2A and B depict the T statistic and estimated activation maps for a representative slice in each subject containing the motor regions as identified manually by an expert. T statistics for the SLM were computed as detailed in App. A. The T statistic maps display a greater range of values with particularly stronger activation in expected motor areas. The SLM activation maps tend to be more punctate than those of the GLM. Fig. 2C displays the observed BOLD signal in the SMA and the corresponding SLM and GLM estimates. For comparison purposes, the same display is presented in Fig. 2D for a second voxel chosen in each subject at a region determined to be activated more strongly by the SLM than the GLM method. These voxels are indicated by green arrows in Fig. 2B. Comparison of the time-series estimates demonstrate that the SLM is more consistent in assignment of activation weights than the GLM, in the following ways: In Subject 1 SMA, the SLM and GLM activation weights are similar (10.1 and 8.0 respectively as indicated on the figures). The GLM fails to capture the shape of the comparison voxel's time-series in Fig. 2D however, resulting in a halving of β^{GLM}, while the SLM determines a reasonable activated time-series shape. The same comment applies to Subjects 2 and 3; The SLM activation weight between SMA and the comparison voxel are reasonable given the relative time-series shapes, while the GLM fails to be able to account for any activation in the comparison voxels. Indeed, in Subject 3, the GLM estimate in the SMA is a poor fit to the data that results in a halved activation weight ($\beta^{GLM} = 11.4$) compared to the SLM SMA ($\beta^{SLM} = 21.1$).

4 Discussion and Conclusion

The use of a state noise to drive the system dynamics is fundamentally different to the inclusion of AR noise in the observation equation, as in [14,13,11]. The model proposed by [1] is similar to the SLM, however the ARX model is presented as an observation equation and is thus reduces to a deterministic model in high SNR.

The SLM is implicitly Bayesian, in that the state equation is equivalently a prior on the signal dynamics, and the resultant estimation algorithm optimises in a maximum *a posteriori* sense for the best state and parameter estimates given the observed data. Our approach differs to Bayesian models in the literature, for example [14,10], in that we have not placed prior distributions governed by hyperparameters on the signal and noise parameters. Such prior distributions can be added at the cost of computational expense. The purpose of the current model is to provide an alternative to the GLM without the use of prior distributions on the parameters.

The stochastic nature of our hemodynamic signal model results in the ability to reshape the input sequence, u_k, to better fit the observed BOLD signal in each voxel. Overfitting of data by models with many parameters is always a

concern. Our simulations and experimental results have demonstrated that the SLM refines the GLM estimates, suggesting that the SLM parameterisation well represents the state and observation processes rather than making overfit inferences. What has previously been described as the nonlinearity of the HRF [2] is now placed in the domain of deterministic vs stochastic. It is of interest to compare in future the SLM with methods in which parameterised, deterministic HRFs are estimated [14], or those in which the HRF is constrained to be a smoothly varying function [3]. Similarly, in future work we will compare the stochastic linear model with nonlinear extended balloon models of the BOLD signal [8] and their estimation strategies, and alternative stochastic modelling approaches such as BDS [10,12].

We have established the stochastic linear model as an alternative to the general linear model for mapping activation, hemodynamic dynamics and noise processes in BOLD fMRI experiments. The SLM provides an inherently flexible model that makes use of uncertainty to drive the estimated hemodynamic signal dynamics, thus accounting for non-canonical responses. Unlike other stochastic models of the BOLD signal, the SLM does not attempt to separate out neuronal activity from the hemodynamics, a decision reflecting the lack of available ground truth at the level of neuronal activity for fMRI experiments.

A SLM Significance Testing

The SLM activation weights are calculated according to Eq. (4), from parameter estimates of a and b at the M^{th} iteration, $a^{(M)}$ and $b^{(M)}$, which are themselves the result of a least squares optimisation of the objective function. Therefore the variances of the parameter estimates are known.

T-statistics for $a^{(M)}$ and $b^{(M)}$ are formed by dividing the estimates by their respective standard deviations. The activation estimate, β^{SLM}, is a ratio between estimates $b^{(M)}$ and $(1 - a^{(M)})$. While it is known that the ratio of two Gaussian random variables produces a Cauchy-distributed random variable, the distribution of the ratio of non-zero-mean T-distributed estimates is considerably more involved. Through simulations, we have empirically derived the following expression for the standard deviation of β^{SLM},

$$\mathrm{std}(\beta^{SLM}) = \mathrm{std}(b^{(M)})\ (1 - a^{(M)}), \tag{7}$$

that is, interestingly, independent of standard deviation of $a^{(M)}$. This expression enables transformation of β^{SLM} to a t-statistic as per $a^{(M)}$ and $b^{(M)}$ above.

To validate the empirical relationship in Eq. (7), we present the results of simulations of the SLM state equation, Eq. (1), for a range of a and σ_w^2. Changes in b modulate the mean without affecting variances, therefore this is chosen to be $b = 1$ without loss of generality. The stimulus u_k is from the experimental visuomotor data, thus $T = 215$. 5000 runs were simulated per parameter pair.

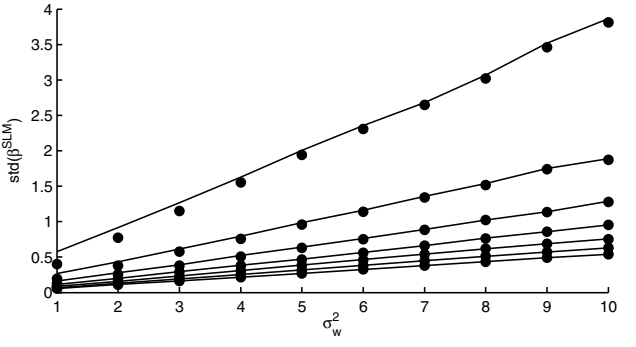

Fig. 3. Simulations validating the empirical relationship in Eq. (7). Standard deviation of β^{SLM} calculated at each parameter $a \in \{-0.75, -0.5, -0.25, \ldots, 0.75\}$ (data sets bottom to top of plot) from 5000 simulation runs (dots) and via Eq. (7) (solid lines).

Fig. 3 depicts the standard deviation of the β^{SLM} estimates along with the empirical relationship in Eq. (7) for each parameter pair, in close agreement. Skewness and kurtosis of the β^{SLM} estimates were concordant with a Gaussian.

References

1. Baraldi, P., Manginelli, A.A., Maieron, M., Liberati, D., Porro, C.A.: An ARX model-based approach to trial by trial identification of fMRI-BOLD responses. NeuroImage 37, 189–201 (2007)
2. Birn, R.M., Saad, Z.S., Bandettini, P.A.: Spatial heterogeneity of the nonlinear dynamics in the FMRI BOLD response. NeuroImage 14, 817–826 (2001)
3. Ciuciu, P., Poline, J.B., Marrelec, G., Idier, J., Pallier, C., Benali, H.: Unsupervised robust nonparametric estimation of the hemodynamic response function for any fMRI experiment. IEEE TMI 22(10), 1235–1251 (2003)
4. Friston, K.J., Josephs, O., Zahahn, E., Holmes, A.P., Rouquette, S., Poline, J.B.: To smooth or not to smooth? Bias and efficiency in fMRI time-series analysis. NeuroImage 12, 196–208 (2000)
5. Friston, K.J., Mechelli, A., Turner, R., Price, C.J.: Nonlinear responses in fMRI: The Balloon model, Volterra kernels, and other hemodynamics. NeuroImage 12, 466–477 (2000)
6. Handwerker, D.A., Ollinger, J.M., D'Esposito, M.: Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. NeuroImage 21, 1639–1651 (2004)
7. Heeger, D.J., Ress, D.: What does fMRI tell us about neuronal activity? Nature Reviews Neuroscience 3, 142–151 (2002)
8. Johnston, L.A., Duff, E.P., Mareels, I., Egan, G.F.: Nonlinear estimation of the BOLD signal. NeuroImage 40, 504–514 (2008)
9. Luenberger, D.G.: Linear and Nonlinear Programming, 2nd edn. Addison Wesley, Reading (1984)

10. Makni, S., Beckmann, C., Smith, S., Woolrich, M.: Bayesian deconvolution fMRI data using bilinear dynamical systems. NeuroImage 42, 1381–1396 (2008)
11. Makni, S., Ciuciu, P., Idier, J., Poline, J.: Joint detection-estimation of brain activity in fMRI using an autoregressive noise model. In: IEEE ISBI, pp. 1048–1051 (2006)
12. Penny, W., Ghahramani, Z., Friston, K.J.: Bilinear dynamical systems. Phil. Trans. R. Soc. B 360(1457), 983–993 (2005)
13. Purdon, P.L., Solo, V., Weisskoff, R.M., Brown, E.N.: Locally regularized spatiotemporal modeling and model comparison for functional MRI. NeuroImage 14, 912–923 (2001)
14. Woolrich, M.W., Jenkinson, M., Brady, J.M., Smith, S.M.: Fully Bayesian spatiotemporal modeling of FMRI data. IEEE TMI 23(2), 213–231 (2004)

Computing the Shape of Brain Networks Using Graph Filtration and Gromov-Hausdorff Metric

Hyekyoung Lee[1,2,3], Moo K. Chung[2,6,7], Hyejin Kang[1,3], Boong-Nyun Kim[5], and Dong Soo Lee[1,3,4]

[1] Department of Nuclear Medicine,
[2] Department of Brain and Cognitive Sciences,
[3] Institute of Radiation Medicine, Medical Research Center,
[4] WCU Department of Molecular Medicine and Biopharmaceutical Sciences,
[5] Department of Neuropsychiatry, Seoul National University, College of Medicine, Seoul, Korea
[6] Department of Biostatistics and Medical Informatics,
[7] Waisman Laboratory for Brain Imaging and Behavior, University of Wisconsin, Madison, WI 53706, USA
mkchung@wisc.edu

Abstract. The difference between networks has been often assessed by the difference of global topological measures such as the clustering coefficient, degree distribution and modularity. In this paper, we introduce a new framework for measuring the network difference using the Gromov-Hausdorff (GH) distance, which is often used in shape analysis. In order to apply the GH distance, we define the *shape of the brain network* by piecing together the patches of locally connected nearest neighbors using the *graph filtration*. The shape of the network is then transformed to an algebraic form called the single linkage matrix. The single linkage matrix is subsequently used in measuring network differences using the GH distance. As an illustration, we apply the proposed framework to compare the FDG-PET based functional brain networks out of 24 attention deficit hyperactivity disorder (ADHD) children, 26 autism spectrum disorder (ASD) children and 11 pediatric control subjects.

1 Introduction

The functional and anatomical connectivity studies based on graph theory have provided new understanding of human brain [1,2]. The characteristic of the brain network is quantified by the global topological measures such as clustering coefficient, characteristic path length and modularity [1,3]. The network comparison is then performed by determining the difference between these topological measures. Each measure reflects different topological characteristic of the brain network. For example, the clustering coefficient and characteristic path length are related with the small-worldness, the assortativity and betweenness are related with the scale-freeness and the modularity is related with the community

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 302–309, 2011.

structure [4,1]. These measures give us a clue for whether the networks have similar topological properties. However, it is unclear which measure is appropriate for network comparison. Instead of trying to find one particular characteristic of network at a given scale, one can also look at the overall change of topological features through persistent homology [5,6,7]. In the persistent homology, the topological features such as the connected components and circles of the network are tabulated in terms of the algebraic form known as Betti numbers. The network difference is then often measured using the bottleneck distance which basically ignores the geometric information of network nodes.

In this paper, we propose a radically different computational framework for determining network difference. Instead of trying to model the topological features of networks, we first define the *shape of network* using the topological concept called the *graph filtration*. The graph filtration is a new graph simplification technique that iteratively build a nested subgraphs of the original graph. The algorithm simplifies a complex graph by piecing together the patches of locally connected nearest nodes. The process of graph filtration can be shown to be mathematically equivalent to the single linkage hierarchical clustering and dendrogram construction. Once the shape of network is defined, we transform the shape into an algebraic form called the *single linkage matrix*. The single linkage matrix is subsequently used in computing the network difference using the Gromov-Hausdorff (GH) metric. The GH metric is a deformation-invariant dissimilarity measure often used in matching deformable shapes [8,9]. The GH metric was never used in measuring the distance between brain networks before.

The proposed method is applied in differentiating functional brain networks with 96 regions of interest (ROIs) extracted from FDG-PET data for 24 attention-deficit hyperactivity disorder (ADHD), 26 autism spectrum disorder (ASD) and 11 pediatric controls (PedCon). Numerical experiments show that the graph filtration framework can differentiate the populations better than most known graph theoretic approaches and the recently popular persistent homology framework. The methodological contributions of this paper are:

(1) We propose a new geometric framework for defining the shape of networks using graph filtration. We introduce the concept of graph filtration and show that it is equivalent to the single linkage hierarchical clustering and dendrogram construction. This implies that there is a mapping from any complex networks to dendrograms.
(2) We determine the distance between networks using the Gromov-Hausdorff metric for the first time. The framework is then used in determining the brain network difference.
(3) We demonstrate that our framework outperforms most of graph theoretic measures and the recently popular persistent homology framework.

2 Main Ideas

The main problem we are trying to solve is to compare and quantify the brain network differences in ADHD, ASD and PedCon populations. We start with briefly introducing the correlation-based brain network construction.

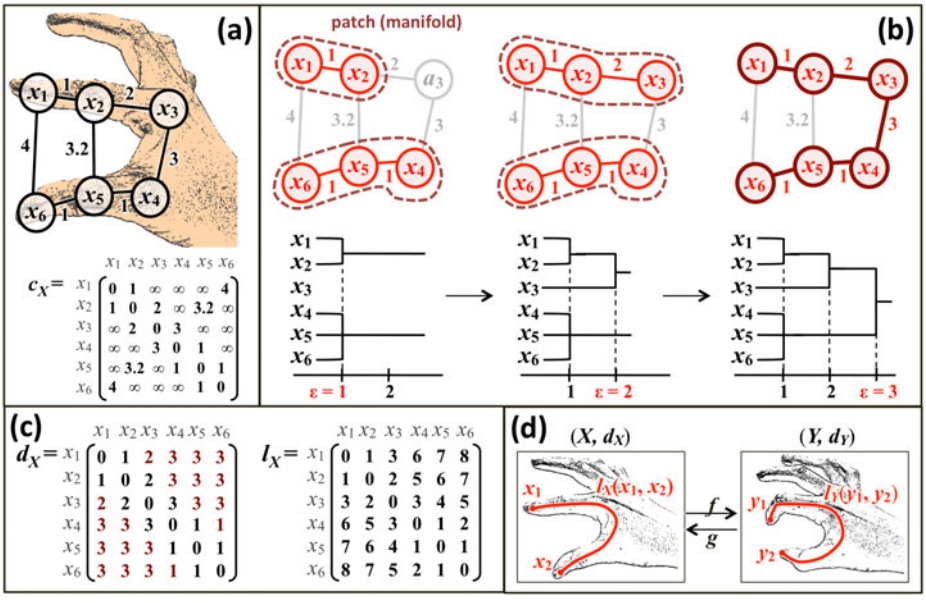

Fig. 1. (a) An example of shape representation using a network of nodes $X = \{x_1, \ldots, x_6\}$ and the distance c_X. The pair (X, c_X) defines the hand. (b) Graph filtration algorithm for representing the graph (X, c_X). (c) The resulting shape can be equivalently represented as the single linkage matrix d_X and the geodesic distance matrix l_X. (d) A deformable hand where d_X and l_X are invariant.

Brain Network Construction. Suppose FDG-PET measurements are obtained in p selected ROIs in n subjects. Each ROI serves as a node in the brain network. Let $X = \{x_1, \cdots, x_p\}$ be the collection of such nodes. Let f_i be the FDG-PET measurement at the node x_i modeled as a random variable. The measurement f_i are assumed to be distributed with mean zero and the covariance $\boldsymbol{\Sigma} = [\sigma_{ij}] \in \mathbb{R}^{p \times p}$. The correlation between f_i and f_j is given by

$$\text{corr}(f_i, f_j) = \frac{\sigma_{ij}}{\sqrt{\sigma_{ii}\sigma_{jj}}}.$$

We can define the metric between the nodes x_i and x_j through the correlation:

$$c_X(x_i, x_j) = 1 - \text{corr}(f_i, f_j).$$

Then the brain network can be represented as the metric space (X, c_X).

Shape of Brain Network. One can characterize the deformable shapes in images using a collection of nodes and the mapping between the deforming nodes. In deformation-invariant shape matching frameworks [8,9], we can identify an open and bended hands as equivalent by establishing the correspondence

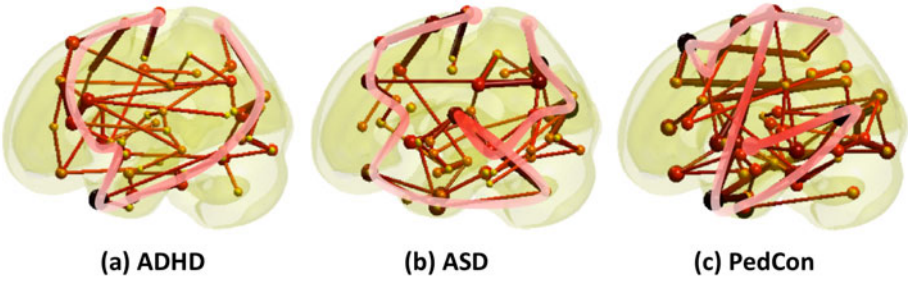

Fig. 2. The shapes of brain networks at the end of the graph filtration (a) ADHD, (b) ASD and (c) PedCon

between nodes (Fig. 1 (d), see below for details). Unlike shapes in images, the shape of brain network is difficult to define and visualize since it is not determined by the Euclidean distance between the nodes, but the correlation between measurements on the nodes. In this paper, we define the *shape of the network* by piecing together patches of locally connected nearest neighbor nodes in an iterative fashion as illustrated in Fig. 1 (b).

The brain network can be viewed as the weighted graph (X, c_X) consisting of the collection of nodes X and the distance c_X. We start with $\epsilon = 0$ and increase the ϵ value at each iteration. The value of ϵ is taken discretely from the smallest $c_X(x_i, x_j)$ to the largest $c_X(x_i, x_j)$. We connect two nodes x_i and x_j if $c_X(x_i, x_j) < \epsilon$. By increasing ϵ, more connected edges are allowed and larger patches are generated. If two nodes are already connected directly or indirectly *via* other intermediate nodes in smaller ϵ values, we do not connect them. For example, in Fig. 1 (b), we do not connect x_2 and x_5 at $\epsilon = 3.2$ since they were already connected through other nodes at $\epsilon = 3$. When ϵ is larger than any distance $c_X(x_i, x_j)$, the iteration terminates since the graph does not change anymore. Suppose G_j corresponds to the graph obtained at the j-th iteration with $\epsilon = \epsilon_j$. Then for $\epsilon_1 < \epsilon_2 < \epsilon_3 < \cdots$, the algorithm generates the sequence of graphs, $G_1 \subset G_2 \subset G_3 \subset \cdots$. Such a sequence of nested graphs is called a *graph filtration* in algebraic topology [5,6]. In this fashion, we define the shapes of the brain network as a sequence of nested subgraphs (Fig. 2).

Single Linkage Matrix. The graph filtration exactly corresponds to the single linkage hierarchical clustering as demonstrated in Fig. 1 (b). The equivalence to the graph filtration and the dendrogram is self-evident. The linking of two nodes corresponds to the linking of leaves in the dendrogram. Increasing the ϵ value in the graph filtration corresponds to increasing the height of the dendrogram.

In the hierarchical clustering, the distance between patches of nodes $\mathcal{C}_1$ and $\mathcal{C}_2$ is given by the distance between the closest members in $\mathcal{C}_1$ and $\mathcal{C}_2$:

$$d_X(\mathcal{C}_1, \mathcal{C}_2) = \min_{x_1 \in \mathcal{C}_1} \min_{x_2 \in \mathcal{C}_2} c_X(x_1, x_2).$$

For example, when $\epsilon = 3$, the distance between the two patches $\{x_1, x_2, x_3\}$ and $\{x_4, x_5, x_6\}$ is given by the distance between x_3 and x_4. Thus, we can represent the shape of brain network as the *single linkage matrix*, where the elements are the single linkage distances between nodes.

Gromov-Hausdorff Distance. After representing the shapes of brain networks, we need to compute the distance between the networks for quantification. Given two metric spaces (X, d_X) and (Y, d_Y), the Gromov-Hausdorff Distance (GH) distance between X and Y is defined as [10,8]:

$$d_{GH}(X,Y) = \inf_{\substack{f:X\to Y \\ g:Y\to X}} \frac{1}{2}\max\left(\mathcal{F}(f), \mathcal{G}(g), \mathcal{H}(f,g)\right), \tag{1}$$

$$\text{where } \mathcal{F}(f) = \sup_{x_1,x_2\in X} |d_X(x_1,x_2) - d_Y(f(x_1), f(x_2))|,$$

$$\mathcal{G}(g) = \sup_{y_1,y_2\in Y} |d_X(g(y_1), g(y_2)) - d_Y(y_1,y_2)|,$$

$$\mathcal{H}(f,g) = \sup_{x\in X, y\in Y} |d_X(x, g(y)) - d_Y(f(x), y)|.$$

We used the single linkage distance for d_X and d_Y. Note that the single linkage distance does not satisfy the triangle inequality but satisfies [11]

$$\max(d_X(x_1,x_2), d_X(x_2,x_3)) \geq d_X(x_1,x_3).$$

In our problem, all the nodes in X and Y are in the fixed locations, thus, the mapping functions f and g are simply given as $f(x_i) = y_i$ and $g(y_i) = x_i$ and Eq. (1) is discretized as [9,12]

$$d_{GH}(X,Y) = \frac{1}{2}\max_{\forall i,j} |d_X(x_i,x_j) - d_Y(y_i,y_j)|.$$

3 Experimental Results

Data Description. The data consists of 24 ADHD (19 boys, mean age: 8.2 $\pm$ 1.6 years), 26 ASD (24 boys, mean age: 6.0 $\pm$ 1.8 years) and 11 PedCon (7 boys, mean age: 9.7 $\pm$ 2.5 years). PET images were preprocessed using Statistical Parametric Mapping (SPM) package. After spatial normalization to the standard template space, mean FDG uptake within 96 ROIs were extracted. The values of FDG uptake were globally normalized to the individuals total gray matter mean count.

Comparison of the Connectivity Matrix. The distance matrices obtained from correlation c_X and single linkage matrices d_X are shown in Fig. 3. The group difference is more evident in the single linkage matrices. The maximum

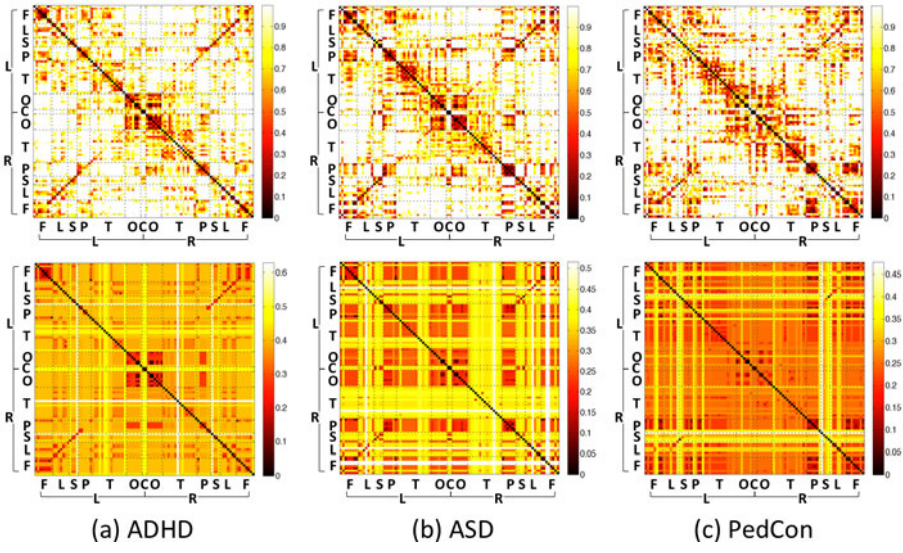

Fig. 3. The correlation-based distance c_X (top) and single linkage matrix d_X (bottom) for (a) ADHD, (b) ASD and (c) PedCon. In each connectivity matrix, the upper-left and the lower-right 48 ROIs are from left and right hemispheres, respectively. The order of ROIs of the left and the right hemispheres are horizontally and vertically symmetric, thus, the diagonal terms from the top-right to the bottom-left represents bilateral symmetry of brain.

single linkage distances of ADHD, ASD and PedCon are 0.62, 0.51, 0.48. The most regions in ADHD are weakly connected except a few strongly connected regions within the occipital (O) and left frontal (F) regions and between the right and the left frontal regions [13,14]. On the other hand, PedCon network is well-connected in the whole brain regions. In ASD, the connection is segmented according to lobes and temporal (T) asymmetry is obviously visible [15,14].

Performance against other Network Measures. We estimated single linkage matrices of 24 ADHD, 26 ASD and 11 PedCon jackknifed resampled data sets and estimated the network differences using 8 different measures including the GH distance, bottleneck distance, assortativity, centrality, clustering coefficient, characteristic path length, small-worldness and modularity (Fig. 4) [6,1,3].

After constructing the distance matrices, we divided the networks into 3 clusters using the hierarchical clustering and evaluated the clustering accuracy by comparing the assigned labels with the true labels. The clustering accuracies of GH distance, characteristic path length and small-worldness are all 100 %. However, the distance between the groups is much larger than the distance within the groups in the GH metric, i.e. $|w - b| = 0.49$ in (Fig. 4(a)), indicating the superior performance of the GH-metric.

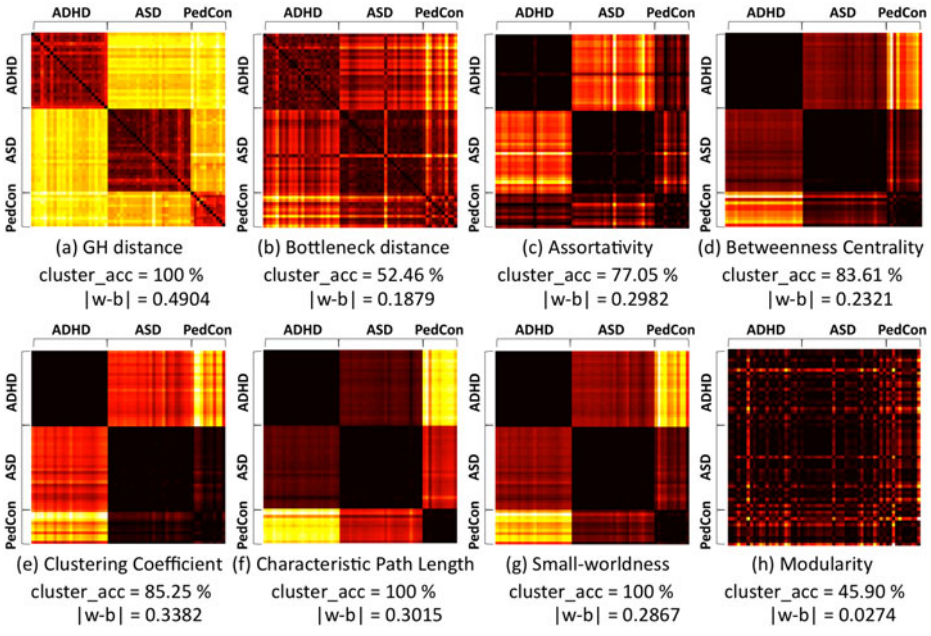

Fig. 4. Network comparison using various network measures: (a) GH distance, (b) bottleneck distance, (c) assortativity, (d) centrality, (e) clustering coefficient, (f) characteristic path length, (g) small-worldness and (h) modularity

4 Conclusions

We presented a novel framework for computing the distance on networks. Using the graph filtration, we defined the shape of the network as a sequence of nested subgraphs. The graph filtration is then transformed into an algebraic form called the single linkage matrix. The single linkage matrices were demonstrated to differentiate the group differences in the ADHD, ASD and PedCon populations. The distance between different single linkage matrices is quantified using the Gromov-Hausdorff metric. The Gromov-Hausdorff metric was validated against other global network measures from graph theory and persistent homology: bottleneck distance, assortativity, centrality, clustering coefficient, characteristic path length, small-worldness and modularity. The GH metric outperforms all of them in terms of the clustering accuracy and the difference between the within- and the between-group distance.

Acknowledgments. This work was supported by the grant R31-2008-000-10103-0 from the WCU project of MEST and NRF, by the grant M103KV010016-08K2201-01610 from Brain Research Center of the 21st Century Frontier Research Program funded by MST, by the grant NRF-2009-351- D00026 from NRF and by NAP of KRCF. The work of M. K. Chung was supported by the WCU Grant from the Government of Korea.

References

1. Rubinov, M., Sporns, O.: Complex network measures of brain connectivity: Uses and interpretations. NeuroImage 52, 1059–1069 (2010)
2. Sporns, O., Tononi, G., Kotter, R.: The human connectome: a structural description of the human brain. PLoS Computational Biology 1, e42 (2005)
3. Bullmore, E., Sporns, O.: Complex brain networks: graph theoretical analysis of structural and functional systems. Nature Reviews Neuroscience 10, 186–198 (2009)
4. Li, L., Alderson, D., Doyle, J.C., Willinger, W.: Towards a theory of scale-free graphs: Definition, properties, and implications. Internet Mathematics 2, 431–523 (2005)
5. Adler, R., Bobrowski, O., Borman, M., Subag, E., Weinberger, S.: Persistent homology for random fields and complexes. ArXiv e-prints (2010)
6. Edelsbrunner, H., Harer, J.: Persistent homology - a survey. Contemporary Mathematics 453, 257–282 (2008)
7. Lee, H., Chung, M.K., Kang, H., Kim, B.N., Lee, D.: Discriminative persistent homology of brain networks. In: IEEE International Symposium on Biomedical Imaging (ISBI), Chicago, IL (2011)
8. Mémoli, F.: Gromov-hausdorff distances in euclidean spaces. In: Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment (CVPR Workshop, NORDIA 2008) (2008)
9. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Efficient computation of isometry-invariant distances between surfaces. SIAM Journal on Scientific Computing 28, 1812–1836 (2006)
10. Kalton, N.J., Ostrovskii, M.: Distances between banach spaces. Forum Mathematicum 11, 17–48 (1999)
11. Carlsson, G., Mémoli, F.: Characterization, stability and convergence of hierarchical clustering methods. Journal of Machine Learning Research 11, 1425–1470 (2010)
12. Mémoli, F., Sapiro, G.: A theoretical and computational framework for isometry invariant recognition of point cloud data. Foundations of Computational Mathematics 5, 313–347 (2005)
13. Kraina, A., Castellanos, F.: Brain development and adhd. Clinical Psychology Review 26, 433–444 (2006)
14. Uddin, L.Q., Menon, V.: Introduction to special topic resting-state brain activity: Implications for systems neuroscience. Frontiers in Systems Neuroscience 4, 37 (2010)
15. Minshew, N., Williams, D.: The new neurobiology of autism: Cortex, connectivity, and neuronal organization. Arch. Neurol., 945–950 (2007)

Model-Driven Harmonic Parameterization of the Cortical Surface

Guillaume Auzias[1], Julien Lefèvre[1,2], Arnaud Le Troter[1], Clara Fischer[3], Matthieu Perrot[3], Jean Régis[4], and Olivier Coulon[1]

[1] LSIS Lab, UMR CNRS 6168, Marseille, France
[2] Université de la Méditerranée, Marseille, France
[3] Neurospin center, CEA, Gif-sur-Yvette, France
[4] INSERM U751, Marseille, France

Abstract. In the context of inter-subject brain surface matching, we present a parameterization of the cortical surface constrained by a model of cortical organization. The parameterization is defined via an harmonic mapping of each hemisphere surface to a rectangular planar domain that integrates a representation of the model. As opposed to previous conformal mapping methods we do not match folds between individuals but instead optimize the fit between cortical sulci and specific iso-coordinate axis in the model. Experiments on both hemispheres of 34 subjects are presented and results are very promising.

Keywords: Surface matching, harmonic mapping, cortical organization, parameterization.

1 Introduction

The rise of surface-based methods for neuroimaging data analysis in the past 10 years has brought a new range of methods in the area of multi-subject comparisons: inter-subject surface matching. Such methods allow to define a common referential in which to perform group studies and inter-subject comparisons. The difficulty comes from the fact that, contrary to volume-based methods, the domain on which the data is defined is different from one subject to another. If one wants to register geometrical information or anatomical landmarks on the cortical surface between different subjects it is necessary to define a common domain on which to perform the registration. This is why existing methods map cortical surfaces to a standard geometry such as a sphere before registration. This mapping is either isometric or conformal [5]. Registration after or during the mapping generally optimize the matching of geometrical information (e.g. convexity [14]) or anatomical landmarks [10,12,13,8] but rarely embeds explicit information about the nature of the variability or stability of the cortical anatomy across subjects.

Nevertheless, the idea of an intrinsic organization of the cortical surface around which variability occurs has arisen [1,3,4], and it has been shown that a model of such organization can be used to perform cortical localization and implicitly

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 310–317, 2011.

solve the problem of inter-subject cortical surface matching [3,2]. In particular in [2] a model is presented that embeds the concept of stable sulcal landmarks conjointly with the notion of organization of these landmarks through specific hypothesis of orientation and alignment of sulci. In this method, the authors propose a scheme of cortical organization, stating that a set of folds that correspond to the first folding location during antenatal life are stable across individuals, and that these folds are organized according to two orthogonal directions and two poles, the insular pole and the cingular pole (Fig.1). This model defines a natural spherical parameterization of the cortical surface, and in the same paper the authors propose an implementation of this model that computes a parameterization constrained by the model for any given cortical mesh and therefore provide a cortical localization and an implicit inter-subject matching. Essentially, this leads to a coordinate system in which the folds that are part of the model always have the same coordinate (longitude or latitude, specified by the model, Fig.1), with a smooth interpolation of the coordinate fields between the constraints. The drawbacks of this approach are that the two coordinate fields are computed independently, the method is therefore unable to theoretically guarantee the integrity of the coordinate system, and that there is no control over the isometric or conformal properties of the resulting mapping.

In this paper we integrate the assumptions given by model presented in [2] and propose a new method for mapping this model onto an individual cortical surface. The relevance of the anatomical hypothesis underlying the model will be investigated in future work. Indeed, we define a mapping that explicitly minimizes angular distortions while matching cortical folds with the model, and introduces a coupling of the two coordinate fields. The originality compared to previous similar harmonic mapping methods (e.g. [10,12]) lies in the fact that folds are not matched to a specific target (e.g. the same fold across subjects) but instead are matched to an iso-coordinate axis of the 2D coordinate system.

In the next section we present our method, and in section 3 the results of our algorithm applied to a set of subjects are presented and discussed.

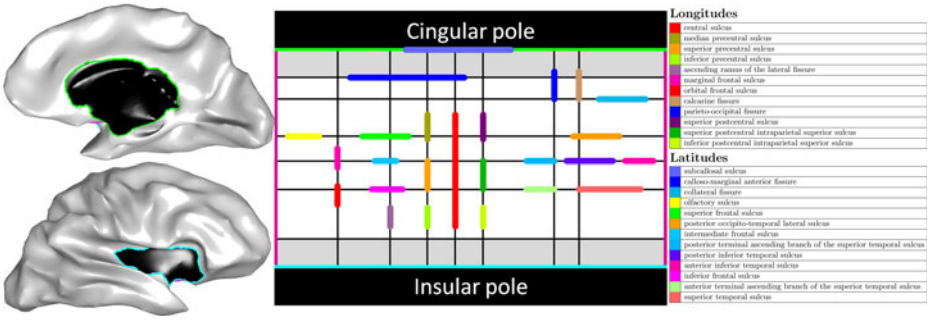

Fig. 1. Left: insular and cingular poles shown on an inflated hemisphere. Right: flat representation of the sulcus-based model of cortical organization.

2 Method

We detail here our method to define a parameterization of the cortical mesh under the sulcal constraints defined on Fig.1. We formalize the parameterization of the cortex as follows. Let us denote $S_T \in \mathbf{R}^3$ the cortical mesh. The parameterization of S_T consists in finding a suitable domain $\Omega \subset \mathbf{R}^2$ and a piecewise linear mapping $h : S_T \rightarrow \Omega$ that is linear on each triangle in S_T and continuous. Such a mapping is uniquely determined by the images $u = (u_1, u_2) = h(v) = (h_1(v_1, v_2, v_3), h_2(v_1, v_2, v_3)) \in \mathbf{R}^2$.

The first step is the definition of the parameterization domain, presented in the next section.

2.1 Parameterization Domain

In contrast with previous approaches, the geometrical assumptions given by our anatomical model impose some constraints on sulcal landmarks but also on the insular and cingular poles, the later being defined around the corpus callosum. While the cortex can be represented as a surface which topology is spherical, i.e. closed without self intersection, the tessellation corresponding to an hemisphere remains open as a hole is formed by the cut around the corpus callosum, i.e. the cingular pole. When a parameterization is achieved on a spherical domain, this hole can be artificially closed but the closing influences the resulting coordinate system. We alternatively suggest to consider the open hemi-cortical representation for fitting our model.

The hemi-spherical cortical surface can then be subdivided into 3 anatomical patches : the insular pole, the cingular pole and the rest of the neocortex. The cingular pole, i.e. the corpus callosum, does not need to be parameterized, as it is not cortex. The parameterization of the insular pole is quite straightforward as its geometry is very close to a plane and is not detailed in the present contribution. Both poles can be segmented automatically (e.g. as presented in [2]), so we focus on the parameterization of the rest of the cortex which is the most challenging.

At this point, the neocortex is represented as a surface with two holes that has the same topology as a cylinder (Fig.2). A cut in the surface linking the two poles is mandatory to obtain the desired topology. We suggest to cut the neocortex following the shortest geodesic path between the two poles. This artificial cut is robustly located as illustrated on Fig.2.

We now suggest to decompose the parameterization of the neocortex into two steps:

- We first define a mapping $f : S_T \rightarrow \Omega$ from the surface of the neocortex onto a 2D rectangle, which corresponds to the integration of the constraints associated to the poles;
- and then include the sulcal constraints, which is thus reduced to a purely 2D deformation problem resulting in a second mapping $g : \Omega \rightarrow \Omega$.

The parameterization mapping h is then defined as the composition of f and g: $h = g \circ f$.

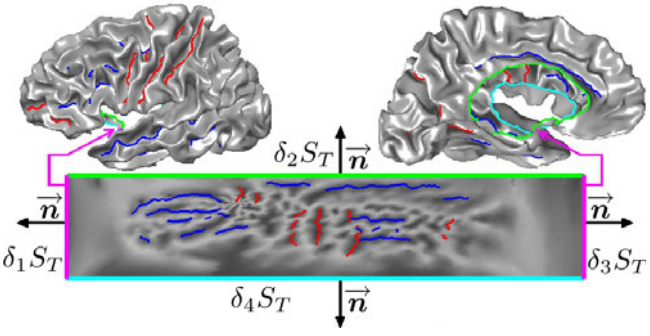

Fig. 2. Each cortical surface is mapped onto a rectangle. Boundaries and sulcal constraints are colored: latitude constraints in blue, longitude constraints in red, insular pole in cyan, cingular pole in green and cut between poles in pink. Mean curvature of the original surface is shown in gray and $\vec{n}$ is the normal of the boundary.

2.2 Unconstrained Harmonic Mapping of the Neocortex onto a Rectangle

Given the intrinsic orthogonal organization of the cortical surface [1,3] the mapping between the neocortex and the rectangular domain should minimize angular distortions. Such mapping has been intensively studied and is denoted as conformal or harmonic mapping. Several research groups have reported work on conformal mapping from the cortical surface to a sphere [11] possibly with additional sulcal landmarks [12]. Hurdal et al. [9] reported a discrete mapping approach that uses circle packing to produce flattened images of cortical surfaces on the sphere, the Euclidean plane, or the hyperbolic plane.

Here, we extend the method introduced to the computer graphics community by Eck et al.[6] which consists in approximating a harmonic map using a finite element method based on P1 basis functions. This technique consists in two steps.

1. First fix the boundary mapping, i.e. fix $f_{\delta S_T} = f^0$, by mapping the boundary δS_T homeomorphically to some convex polygon in the plane. Here this parameterization domain is defined as a planar rectangle.
2. Find the piecewise linear mapping $f : S_T \rightarrow \Omega$ which minimizes the Dirichlet energy subject to the Dirichlet boundary condition $f_{\delta S_T} = f^0$:

$$E_D(f) = 1/2 \int_{S_T} ||\nabla_{S_T} f||^2 \tag{1}$$

Fixing the boundary remains however a major drawback in the context of our model: the conformality of the mapping is not guaranteed and the orthogonality of coordinate axis is lost near the boundary. We thus adapted this procedure by relaxing the boundary constraints such that the points on the boundary can move along the coordinate parallel to the boundary and only the four points in the corners remain fixed (see Fig.2 and Eq.2).

Derivating the Dirichlet energy $E_D(f)$ gives $\nabla E_D(f) = -\Delta_{S_T} f$ and its minimization reduces to solving independently the two linear system of equations which are Poisson equations with mixed boundary conditions:

$$\begin{cases} \Delta_{S_T} f_1 = 0, \in S_T, \\ f_1 = f_1^0 \ on \ \delta_1 S_T \bigcup \delta_3 S_T \\ \nabla f_1 . \overrightarrow{\boldsymbol{n}} = 0 \ on \ \delta_2 S_T \bigcup \delta_4 S_T \end{cases} \quad \begin{cases} \Delta_{S_T} f_2 = 0, \in S_T, \\ \nabla f_2 . \overrightarrow{\boldsymbol{n}} = 0 \ on \ \delta_1 S_T \bigcup \delta_3 S_T \\ f_2 = f_2^0 \ on \ \delta_2 S_T \bigcup \delta_4 S_T \end{cases} \tag{2}$$

where $\delta S_T = \delta_1 S_T \bigcup \delta_2 S_T \bigcup \delta_3 S_T \bigcup \delta_4 S_T$ is the domain boundaries as shown on Fig.2.The existence and uniqueness of the solution has been proved (e.g. in [7]) and Δ_{S_T} can be discretized with the harmonic weights defined in [6]. Inverting the two matrices involved in those two linear equations can be achieved in few seconds with efficient linear algebra libraries.

2.3 Harmonic Mapping under Orthogonal Constraints

We now define two sets of constraints K_{lat} and K_{lon} corresponding to the sulcal fundi specified in our anatomical model (as latitudes and longitudes respectively). Contrary to other methods, we want to align those sulcal fundi onto specific coordinates (longitude or latitude) and not to other landmarks.

The constrained mapping $g : \Omega \rightarrow \Omega$ that minimizes the angle distortion is defined as the optimum of the following equation under the same boundary conditions as in Eq.2:

$$E(g) = E_D(g) + \lambda E_P(g) = \frac{1}{2} \int_\Omega ||\nabla_{S_T} g||^2 + \frac{\lambda}{2} \int_\Omega ||g(u) - p(g(u))||^2 du \tag{3}$$

where p is defined as: $p : \begin{pmatrix} u_1 \\ u_2 \end{pmatrix} \rightarrow \begin{pmatrix} u_1^{lon} \\ u_2 \end{pmatrix}$ if $u \in K_{lon}$, p: $\begin{pmatrix} u_1 \\ u_2 \end{pmatrix} \rightarrow \begin{pmatrix} u_1 \\ u_2^{lat} \end{pmatrix}$ if $u \in K_{lat}$, and p =Id otherwise. $E_P(g)$ measures the distance between each sulcal fundus and the corresponding target longitude (resp. latitude) coordinate value u_1^{lon} (resp. u_2^{lat}) via the projection p. The derivative of the constraint energy $E_P(g)$ can be computed as:

$$\nabla E_P(g) = 2(I - D_g^* p)(g - p(g)) \tag{4}$$

where $D_g^* p$ is the adjoint operator applied to the derivative matrix of p, i.e.

$$D_g^* p(u) = \begin{pmatrix} \frac{\partial p_1}{\partial u_1} & \frac{\partial p_1}{\partial u_2} \\ \frac{\partial p_2}{\partial u_1} & \frac{\partial p_2}{\partial u_2} \end{pmatrix}^t = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \text{ if } u \in K_{lon}, D_g^* p(u) = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \text{ if } u \in K_{lat} \tag{5}$$

and $D_g^* p = Id$ otherwise. We thus obtain

$$\nabla E(g) = -\Delta_{S_T} g + \lambda (I - D_g^* p)(g - p(g)) \tag{6}$$

and the energy is then minimized using a classical iterative gradient descend scheme. Note that the projection of a point onto the corresponding axis $p(g(u))$ depends on its current location $g(u)$ which is updated between two iterations. This original data-driven term $E_P(g)$ introduces a coupling between the two coordinate fields, latitudes and longitudes.

3 Results

We applied this parameterization process with $\lambda = 1$ to the left and right cortical surfaces of 34 subjects, for which sulci were identified manually by an expert after an automatic extraction of cortical surfaces using the BrainVisa[1] T1 processing pipeline on T1-weighted MR images of the subjects. Fig.3 presents the mapping on the rectangular domain of all sulcal fundi that are part of the model for all 34 subjects. Top row shows the result of the unconstrained conformal mapping presented in section 2.2 and bottom row shows the results of the constrained mapping presented in section 2.3. It is visible that the mapping without sulcal

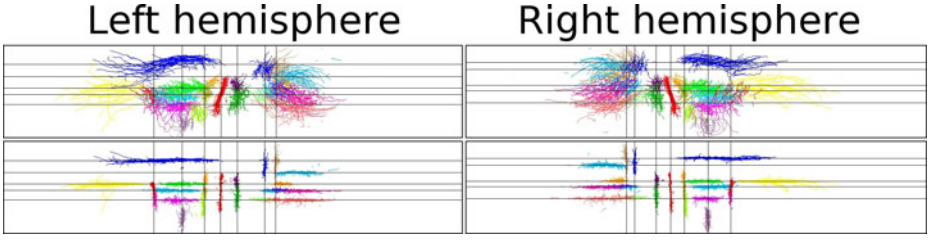

Fig. 3. Sulcal constraints of the 34 subjects mapped onto the rectangular domain with unconstrained mapping (top row) and with the model-driven mapping (bottom row). The sulci are colored according to the model shown on Fig.1

constraints shows a reasonable orthogonality of "longitude" and "latitude" sulci, which tends to advocate for the notions of orthogonal principal direction fields introduced in [1,3], and illustrates the pertinence of the model presented in [2]. On the bottom row of Fig.3, the effects of the model-driven harmonic mapping are very clear. Sulci have been well aligned on the axis of the model. The implicit inter-subject matching performed by our method can also be observed since sulci show a good alignment across subjects. Residual variability is also visible, often at the extremity of sulcal fundi for which the assumption of alignment is sometime corrupted. It is also understood that this assumption is subject to variability and cannot be systematically observed due to the complexity of the folding process during growth, which is subject to many factors. The inter-subject matching can also be checked on Fig.4, in which the mean curvature has been averaged across all subjects and is represented on a single cortical surface. The major sulcal and gyral structures are well preserved: deep blue areas indicate that the dispersion of sulci has been reduced, and red parts show that the crown of gyri are matched across subjects although they are not part of the model. The improved alignment of gyri shows that the coordinates are well interpolated between sulcal constraints and that the harmonic properties of the mapping help to align a number of gyri that are also subject to the orthogonal organization. The density of the model constraints also shows its influence for instance in

[1] http://brainvisa.info

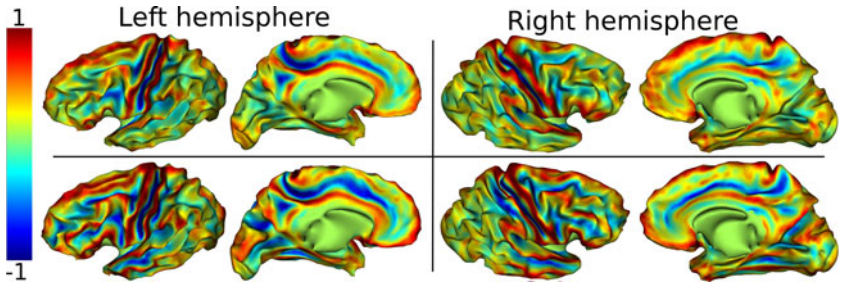

Fig. 4. Mean curvature averaged across 34 subjects. Top row: after unconstrained planar mapping; bottom row: after mapping constrained by the model.

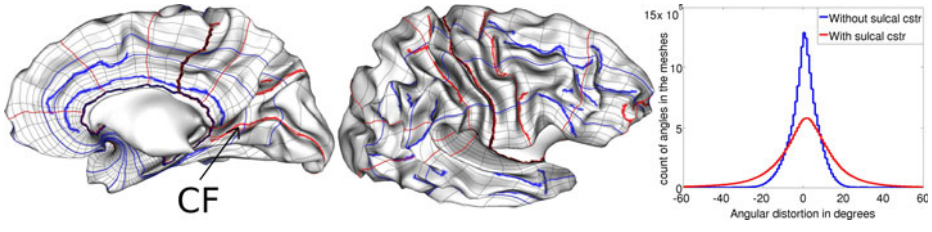

Fig. 5. The resulting parameterization, shown as a coordinate grid on the slightly smoothed cortical surface of one hemisphere. In blue (resp. red), the sulci and axes of the rectangular domain corresponding to latitude (resp. latitudes) constraints of the model. The histogram of the angular distortion through all 34 left and right cortical surfaces before (blue) and after (red) sulcal alignment is shown on the right.

the parietal lobe where the model lacks information [2] and the alignment of curvatures is not as good.

The adaptation of the resulting coordinate grid to local anatomy can be seen for one subject on Fig.5. Iso-coordinate axes follow the local geometry and comply with the sulcal constraints. See e.g. the Calcarine Fissure (CF) which is a good predictor of the position of primary visual functional areas. Finally, Fig.5 also shows the control over angular distortions, as indicated by their distribution through the 34 left and right cortical surfaces before and after the constrained mapping.

4 Conclusion

We have proposed here an implementation of the cortical model presented in [2], performed via a mapping to a planar domain. The harmonic aspect of our mapping allows to explicitly take into account the directional nature of the model, based on the notion of orthogonality of sulcal axis. As opposed to previous harmonic approaches, we do not match folds across subjects or between subjects and a template, but instead we optimize the fit between cortical sulci and specific

iso-coordinate axis in the model. Results show a good alignment of structures, sulci and gyri, across subjects. The process should prove useful for cortical localization and surface matching in the context of surface-based group analysis of functional and anatomical studies, although more work remains to be done, in particular to compare our method to reference methods such as [14].

References

1. Régis, J., Mangin, J.-F., Ochiai, T., Frouin, V., Rivière, D., Cachia, A., Tamura, M., Samson, Y.: Sulcal roots generic model: a hypothesis to overcome the variability of the human cortex folding patterns. Neurol. Med. Chir. 45, 1–17 (2005)
2. Clouchoux, C., Rivière, D., Mangin, J.-F., Operto, G., Régis, J., Coulon, O.: Model-driven parameterization of the cortical surface for localization and inter-subject matching. NeuroImage 50, 552–566 (2010)
3. Toro, R., Burnod, Y.: Geometric atlas: modeling the cortex as an organized surface. NeuroImage 20, 1468–1484 (2003)
4. Lohmann, G., Von Cramon, Y., Colchester, A.: Deep sulcal landmarks provide an organizing framework for human cortical folding. Cereb. Cortex 18, 1415–1420 (2008)
5. Desbrun, M., Meyer, M., Alliez, P.: Intrinsic parameterizations of surface meshes. Computer Graphics Forum 21, 209–218 (2002)
6. Eck, M., DeRose, T., Duchamp, T., Hoppe, H., Lounsbery, M., Stuetzle, W.: Multiresolution analysis of arbitrary meshes. In: 22nd Conference on Computer Graphics and Interactive Techniques SIGGRAPH 1995, pp. 173–182. ACM, New York (1995)
7. Allaire, G.: Analyse numérique et optimisation. Éditions de l'École Polytechnique (2005)
8. Lui, L., Thiruvenkadam, S., Wang, Y., Thompson, P., Chan, T.: Optimized Conformal Surface Registration with Shape-based Landmark Matching. SIAM Journal on Imaging Sciences 3, 52–78 (2010)
9. Hurdal, M.K., Stephenson, K.: Discrete conformal methods for cortical brain flattening. Neuroimage 45, 86–98 (2009)
10. Joshi, A., Shattuck, D., Thompson, P., Leahy, R.: Surface-constrained volumetric brain registration using harmonic mappings. IEEE Trans. Med. Imag. 26, 1657–1669 (2007)
11. Haker, S., Angenent, S., Tannenbaum, A., Kikinis, R., Sapiro, G., Halle, M.: Conformal surface parameterization for texture mapping. IEEE Trans. Vis. Comput. Graphics 6, 181–189 (2000)
12. Wang, Y., Lui, L.M., Chan, T.F., Thompson, P.M.: Optimization of brain conformal mapping with landmarks. In: Duncan, J., Gerig, G. (eds.) MICCAI 2005. LNCS, vol. 3750, pp. 675–683. Springer, Heidelberg (2005)
13. Tosun, D., Rettmann, M.E., Prince, J.L.: Mapping techniques for aligning sulci across multiple brains. Med. Image Anal. 8, 295–309 (2004)
14. Fischl, B., Sereno, M.I., Tootell, R.B., Dale, A.M.: High-resolution intersubject averaging and a coordinate system for the cortical surface. Hum. Brain Mapp. 8, 272–284 (1999)

Assessing Regularity and Variability of Cortical Folding Patterns of Working Memory ROIs

Hanbo Chen[1], Tuo Zhang[2], Kaiming Li[2], Xintao Hu[2], Lei Guo[2], and Tianming Liu[1]

[1] Department of Computer Science and Bioimaging Research Center,
The University of Georgia, Athens, GA, USA
[2] School of Automation, Northwestern Polytechnical University, Xi'an, China

Abstract. Cortical folding patterns are believed to be good predictors of brain cytoarchitecture and function. For instance, neuroscientists frequently apply their domain knowledge to identify brain Regions of Interests (ROIs) based on cortical folding patterns. However, quantitative mapping of cortical folding pattern and brain function has not been established yet in the literature. This paper presents our initial effort in quantification of the regularity and variability of cortical folding pattern features for working memory ROIs identified by task-based fMRI, which is widely accepted as a standard approach to localize functionally-specialized brain regions. Specifically, we used a set of shape attributes for each ROI base on multiple resolution decomposition of cortical surfaces, and described the meso-scale folding pattern via a polynomial-based approach. We also applied brain atlas label distribution as a global-scale description of ROI folding pattern. Our studies suggest that there is deep-rooted regularity of cortical folding patterns for certain working memory ROIs across subjects, and folding pattern attributes could be useful for the characterization, recognition and prediction of ROIs, if extracted and applied in a proper way.

Keywords: ROI, folding pattern, prediction.

1 Introduction

The relationship between brain anatomy and its function has been of keen interest for years. Extensive neuroscience research suggests that the cortical folding process and axongenesis process are closely coupled [1]. Since the structural connectivity pattern has been shown to be good predictor of function [2], it is reasonable to believe that cortical folding is a good predictor of brain function as well. For instance, the cortical folding patterns are shown to be good predictors of the brain cytoarchitecture [3]. However, quantitative mapping of cortical folding pattern and brain function has not been adequately studied yet in the literature. The accomplishment of such a goal entails the availability of both quantitative cortical folding pattern descriptors and accurate localizations of functionally-specialized brain regions.

This paper presents our initial effort in quantitative mapping of the relationship between cortical folding patterns and brain function, particularly in quantification of the regularity and variability of folding pattern features for working memory ROIs. First, we identify brain regions involved in the working memory system by using

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 318–326, 2011.

task-based fMRI. Then, we introduce two groups of descriptive cortical folding pattern features for these ROIs in multiple scales. One is to present the global-scale folding pattern by using the widely used MNI (Montreal Neurological Institute) atlas labels. Another is to present the meso-scale folding pattern by using the recently published polynomial models [4]. Compared with other folding pattern descriptors [5, 6, 7], the advantages of the polynomial model not only lie in its meso-scale view which is more suitable for ROIs, but also its compactness for describing the symmetric shape patterns and its soundness for classifying the cortical shapes into eight primitive folding patterns [8]. To effectively extract folding pattern features for the ROIs, the cortical surface is decomposed into multiple resolutions [7] and attributes are generated on the multi-resolution surfaces. In order to assess the regularity and variability of the shape attributes for each ROI in the working memory network, we performed a set of statistical and computational analysis for ROIs within a population. The experimental results show that there exists deep-rooted regularity of cortical folding patterns for certain ROIs. Finally, we performed ROI prediction based on the extracted multi-scale folding attributes, and our results show that only certain ROIs with consistent folding patterns across individuals can be predicted accurately.

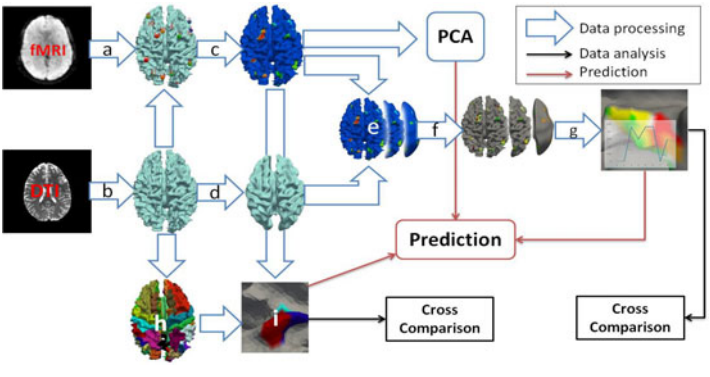

Fig. 1. The flowchart of our approaches. (a) Obtain working memory network ROIs from task-based fMRI. (b) Reconstruct cortical surface from DTI data. (c) Map ROIs onto cortical surface. (d) Decompose cortical surface into multi-resolution surfaces. (e) Map ROIs on multi-resolution surfaces. (f) Compute folding patterns. (g) Generate meso-scale attributes for each ROI. (h) Atlas labeling. (i) Generate global-scale attribute for each ROI.

2 Methods

2.1 Overview

The flowchart of the proposed algorithmic pipeline is outlined in Fig. 1. The folding pattern of human cerebral cortex is a multi-scale concept [9]. Hence, we extract shape features for each ROI base on both folding patterns of multi-resolution surfaces and MNI atlas labels, and then assess the regularity and variability of these features. The cortical surface reconstructed from DTI data [10] is decomposed into multi-resolutions by the spherical wavelet algorithm approach [7]. The volumetric working

memory ROIs were obtained from task-based fMRI, and then are mapped onto the reconstructed surface, as well as other decomposed multi-resolution surfaces. Then, the folding patterns of the ROIs are computed by the polynomial models [4], and the feature vectors of each ROI is then obtained and analyzed via the methods illustrated in Fig. 1 separately. Finally, ROI prediction based on these folding pattern features is performed and a PCA (principal component analysis) model is introduced to constrain the spatial relationship among ROIs.

2.2 Data Acquisition and Pre-processing

Seventeen healthy students were recruited to participate in this study. Each participant performed a modified version of the OSPAN task (3 block types: OSPAN, Arithmetic, and Baseline) while fMRI data was acquired. DTI scans were also acquired for each participant. FMRI and DTI scans were acquired on a 3T GE Signa scanner. Acquisition parameters are as follows: fMRI: 64x64 matrix, 4mm slice thickness, 220mm FOV, 30 slices, TR=1.5s, TE=25ms, ASSET=2; DTI: 128x128 matrix, 2mm slice thickness, 256mm FOV, 60 slices, TR=15100ms, ASSET=2, 3 B0 images, 30 optimized gradient directions, b-value=1000. Each participant's fMRI data was analyzed using FSL FEAT. Individual activation map reflecting the OSPAN (complex span) contrast was identified. Totally, 15 activated ROIs were recognized and applied for this study (Table 1). These ROIs were mapped on the cortical surface and on the corresponding decomposed surfaces, as shown in Fig. 2. The size of ROI is defined by the number of neighborhood vertex rings, which is set as 5 in this paper.

Table 1. The names and coresponding IDs of 15 recognized working memory ROIs

Region name	Region ID		Region name	Region ID	
	Left	Right		Left	Right
Insula	1	10	Precuneus	6	13
Medial Frontal Cortex	2		Superior Frontal Gyrus	7	14
Occipital Pole	3		Inferior Parietal Lobule	8	15
Paracingulate Gyrus	4	12	Dorsolateral Prefrontal Cortex		9
Precentral Gyrus	5		Lateral Occipital Gyrus		11

2.3 Multi-scale Surface Decomposition

In this paper, the over-complete spherical wavelets algorithm [7] is applied to decompose the original cortical surface reconstructed from DTI data into a cascade of 7 lower resolutions (Fig. 2). The lower the resolution is, the smoother and less convoluted the surface will be. In this spherical wavelet transform process, a set of wavelet coefficient samples on the subdivided icosahedron grid are obtained from the input shape described as a set of mesh vertices, which corresponds directly to the resolution level of the spherical wavelet transform. After the transformation, each vertex has a set of wavelet coefficients representing its position in lower resolutions. The total number of vertices and the indices of vertices do not change during this decomposition process [7]. Thus, the correspondence of a vertex remains in all resolutions, which is illustrated by the same ROI colors in Fig. 2. More details of the over-complete spherical wavelets algorithm are referred to [7].

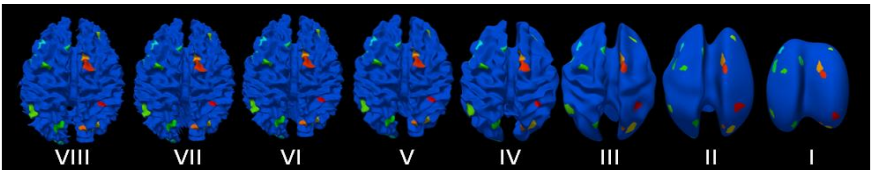

Fig. 2. An example of multi-resolution decomposition for a cortical surface. The ROIs are mapped onto all resolutions respectively encoded by certain colors.

2.4 Folding Pattern Descriptors

2.4.1 Meso-Scale Attributes Based on Polynomial Folding Descriptors

Based on the above surface decomposition, the folding pattern attributes of each ROI are extracted on multi-resolutions, respectively. To compute the folding pattern of a surface patch, a parametric folding descriptor using polynomials [4] is applied:

$$Z = aX^2 + bY^2 + cX^3 + dY^3 \tag{1}$$

where, a and b describe the mirror symmetric components of the patch along x axis and y axis respectively, while c and d represent the rotational symmetric components [4]. Then, a model-driven method similar to that in [4] is applied to classify the combination of the symmetric components into eight primitive patterns of peak, pit, ridge, valley, saddle ridge, saddle valley, flat, and inflection [8]. The attribute vector of ROI on each surface resolution is defined as the normalized histogram of eight patterns within ROI. Thus, there is an 8-dimensional attribute vector for each ROI on each resolution of surface representation. Hence, we have a 64-dimensional meso-scale folding pattern attribute vector for each ROI.

2.4.2 Global-scale Folding Attribute Based on MNI Labels

The HAMMER tool [12] was used to label the cortical surface into MNI labels. Then, the global-scale attribute vector of ROI based on MNI atlas labels is defined as the normalized histogram of MNI labels within each ROI as follows.

$$V_i = \frac{\left(v_i^{h_1}, v_i^{h_2}, \cdots, v_i^{h_n}\right)}{\left|v_i^{h_1}, v_i^{h_2}, \cdots, v_i^{h_n}\right|} \tag{2}$$

where $v_i^{h_k}$ is the number of vertices with MNI label h_k within ROI i. The size of the vector is dependent on the combination of MNI label within an ROI.

2.5 Folding Pattern Feature Analysis

The cosine similarity has been widely applied as a measurement of similarity between feature vectors, and is adopted here. Denote R_m^p as ROI #m of subject p and A_m^p as the attribute vector of R_m^p, the average similarity S_{mn} between 2 ROIs m and n is defined to measure the feature similarity for each ROI pair within a group as follows:

$$S_{mn} = \frac{\sum_{p=1}^{N-1}\sum_{q=p+1}^{N} F(R_m^p, R_n^q)}{C_N^2} \tag{3}$$

$$F(R_m^p, R_n^q) = \frac{A_m^p \cdot A_n^q}{\left\|A_m^p\right\|\left\|A_n^q\right\|} \tag{4}$$

where N is the number of subjects.

Considering the length of global-scale attribute can be different, when comparing the similarity between them, the attribute vectors are normalized to the same length. Elements that represent certain MNI labels are viewed as the weights of attribute vector in certain direction. For those elements not shared between attributes, they are set to zero in order to unify the length of attribute vectors.

2.6 ROI Prediction

An ROI prediction framework is developed based on the cortical folding features described above. Before performing the ROI prediction, all of the subjects are aligned to the same template space (a randomly selected subject in the dataset) by using the linear image registration algorithm FSL FLIRT [13]. Then, the ROI prediction is formulated and solved by minimizing the energy function defined below:

$$E = \lambda E_{int} + (1-\lambda)E_{ext} \tag{5}$$

where internal term E_{int} is the spatial constraint of ROIs' locations which regularizes the search of ROI in a certain space, and the external term E_{ext} describes the distance between folding pattern feature vectors, and λ trades off these two terms.

As there exists certain consistence in the spatial distribution pattern of ROIs, an ROI coordinate PCA model is introduced to constrain the spatial relationship among ROIs. E_{int} is defined as the reconstruction error when projecting the candidate ROIs into the sub-space represented by the PCA model. The external energy for one candidate group is formulated as:

$$E_{ext} = \frac{1}{n}\sum_{i=1}^{i=n}\left(\beta E_{global}^i + (1-\beta)E_{meso}^i\right) \tag{6}$$

where E_{global}^i, E_{meso}^i are the distances between global and meso scale folding pattern feature vectors of candidate ROI #i and those in the training set, and β is the trade-off. The optimization is solved by search over the whole space.

3 Results

3.1 Meso-scale Folding Pattern Analysis

Quantitative assessments of the regularity and variability of the folding pattern attribute vectors for each ROI are summarized in Fig. 3. It is evident in Fig. 3(a) that the meso-scale attributes of certain ROIs including ROI #1, #2, and #7, are less

variable in that the first principal components can account over 85% of the variances. However, the attributes of other ROIs, including ROI #6, have much more variability and the first principal components can only account for 68% of the variance. These results demonstrate that there is deep-rooted regularity of cortical folding patterns for specific working memory ROIs such as ROI #1 and #7 across subjects, and the folding pattern attributes could be potentially used as predictors of functional ROIs.

The similarities between attribute vectors of each pair of ROIs across 17 subjects were calculated in different resolution using the methods in section 2.5 and shown in Fig. 3(b). It can be seen from Fig. 3(b) that, in the first and second resolutions of surface, the 'flat' folding pattern dominates over others and the attribute vectors are almost the same in these two resolutions. However, starting with the third resolution, the attribute vectors are much more distinctive among ROIs which indicates that the distinctiveness of folding pattern attributes of ROIs is dependent on the resolutions of surface representations. If we look at the diagonals in the matrices in Fig. 3(b), some ROIs such as ROI #2, #3, #7, and #13, have higher similarity, while others such as ROI #5 and #11 have lower similarity. This result partly reflects the similar conclusion made by PCA analysis. An interesting finding in Fig. 3(b) is that ever since the fourth resolution, 4 dark lines can be obviously observed in the matrices, which are highlighted by blue arrow in Fig. 3(b.VIII). As dark represents low similarity, it means that the attribute vectors of ROI #1 and ROI #10 are very different from those of other ROIs and share unique patterns. An intriguing fact is that these two ROIs are left and right insular respectively. We used a classifier based on the support vector machine (SVM) [11] to differentiate these two ROIs from the rest by the folding attributes. The classification test was performed in a leave-one-out fashion. We obtained a relatively high true positive rate of 85.29% in test and a low false negative rate of 2.26%. This promising result suggests that folding patterns can characterize certain brain ROIs (e.g., insular) and can provide distinctive morphological signatures for ROI recognition.

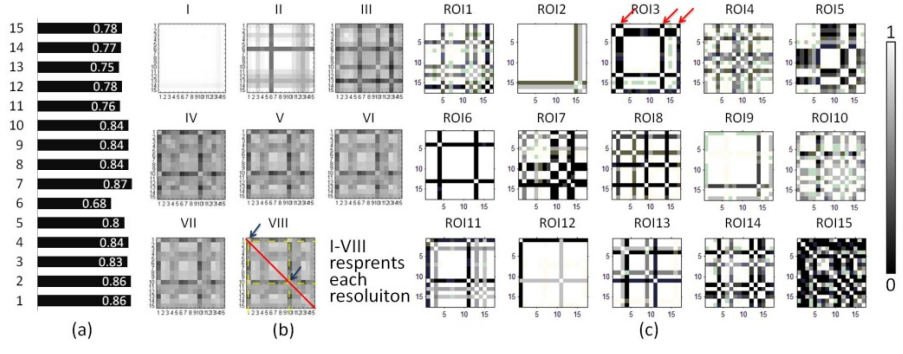

Fig. 3. (a): Histograms of percentages of variance that the first principal components can account for after applying PCA on meso-scale feature vectors of each ROI. (b): Average similarity S_{mn} between meso-scale feature vectors of 15 ROIs on each resolution. Each line represents an ROI. Color bar is on the right. (c): Similarity of global-scale ROI attributes between 17 subjects. Each line represents a single subject.

3.2 Global-Scale Folding Pattern Analysis

The similarities of global-scale attribute vectors defined in Eq. (2) of each ROI across 17 subjects were calculated using the methods in section 2.5. The similarity of attribute of each ROI is shown in Fig. 3(c). It can be seen from Fig. 3(c) that, some ROIs such as ROI #2, #9, and #12, have higher similarities across subject while others such as ROI #15, are more variable across subjects. It is interesting that the ROI #2 and #9 are again more consistent across subjects in terms of meso-scale folding pattern features, as shown in Fig. 3(a-b). This result further indicates that the level of regularity and variability of functional ROI folding patterns is very ROI-specific.

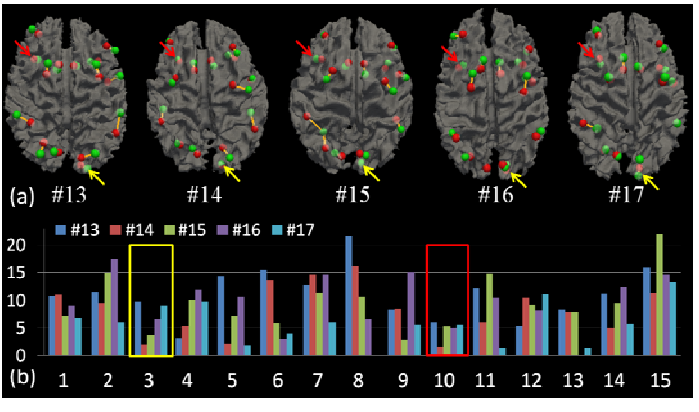

Fig. 4. (a) Prediction results. The center of benchmark ROIs and predicted ROIs are presented by green and red bubbles respectively and connected by yellow line for each pair. ROI #3 and #10 are highlighted by yellow and red arrows accordingly. (b) Euclidean distances between 15 predicted ROIs and benchmark ROIs for 5 subjects. The horizontal axis indexes ROI. ROI #10 is highlighted with red for its good prediction result, in which the average error is 4.7 mm.

3.3 ROI Prediction

We used 12 subjects from #1-#12 as training dataset to obtain the prior model of ROIs. Then, the prediction framework described in Section 2.6 was performed on the rest of 5 subjects (Fig. 4(a)). The prediction result was measured by Euclidean distance between predicted ROI and the benchmark ROI provided by task-based fMRI, as shown in Fig. 4(b). Among all of these 15 ROIs, the prediction accuracy is the best for ROI #10 (average prediction error: 4.7 mm), which significantly improved the initialized ROI positions by linear registration (average initialization error: 5.4 mm). This can be explained by the relatively consistent meso-scale and global-scale folding patterns for ROI #10 across subjects. Though the prediction results of ROI #3 are better for subjects from #14 - #16, but worse for subjects #13 and #17. This is partially caused by the uniqueness of global scale features of these subjects, as highlighted by red arrows in Fig. 3(b). This result suggests that stratification of ROIs and subjects into sub-groups based on their folding patterns

could facilitate ROI prediction within homogeneous subgroups and potentially improve ROI prediction accuracy. For other ROIs, our methods can not significantly improve the initialized ROI positions in that the folding pattern features of these ROIs are not consistent in either meso or global scale across individuals, as shown in Fig. 3.

It should be noted that although our ROI prediction results in Fig. 4 are preliminary, it suggests the possibility of using folding patterns to predict certain functional ROIs without the availability of task-based fMRI data. In many clinical applications in which there is no task-based fMRI data available, it would be very helpful to accurately predict certain functional ROIs based only on the folding pattern features that can be extracted from widely available structural MRI data.

4 Conclusion

We presented a novel framework to assess folding patterns of ROIs that are identified by working memory task-based fMRI. The regularity and variability of these attributes have been assessed within and across ROIs, as well as within and across subjects. The major conclusions of our work are as follows. (1) There are deep-rooted regularities of cortical folding patterns for specific working memory ROIs, e.g., insular, across subjects. However, the intrinsic relationship between the regularity of cortical folding patterns and brain connectivity and function remains to be elucidated in the future. (2) Folding pattern attributes, if extracted and applied in a proper way, could be very useful for the characterization, recognition and prediction of brain ROIs, e.g., ROI #10 of insular in the working memory network in Fig. 4. In the absence of fMRI data, this ROI prediction capability based on widely available MRI data could be helpful in localizing functional brain regions in clinical applications.

References

1. Van Essen, D.C.: A tension-based theory of morphogenesis and compact wiring in the central nervous system. Nature 385(6614), 313–318 (1997)
2. Sporns, O., Chialvo, D.R., Kaiser, M., Hilgetag, C.C.: Organization, development and function of complex brain networks. Trends Cogn. Sci. 8(9), 418–425 (2004)
3. Passingham, R.E., Stephan, K.E., Kötter, R.: The anatomical basis of functional localization in the cortex. Nat. Rev. Neurosci. 3(8), 606–616 (2002)
4. Zhang, T., Guo, L., Li, G., Nie, J., Liu, T.: Parametric representation of cortical surface folding based on polynomials. In: Yang, G.-Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (eds.) MICCAI 2009. LNCS, vol. 5762, pp. 184–191. Springer, Heidelberg (2009)
5. Zilles, K., Armstrong, E., Schleicher, A., Kretschmann, H.J.: The human pattern of gyrification in the cerebral cortex. Anat. Embryol. (Berl.) 179, 173–179 (1988)
6. Yu, P., Yeo, B.T., Grant, P.E., Fischl, B., Golland, P.: Cortical Folding Development Study based on Over-Complete Spherical Wavelets. In: IEEE Workshop on MMBIA, pp. 1–8 (2007)
7. Yeo, B.T., Yu, P., Grant, P.E., Fischl, B., Golland, P.: Shape analysis with overcomplete spherical wavelets. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part I. LNCS, vol. 5241, pp. 468–476. Springer, Heidelberg (2008)

8. Besl, P.J., Jain, R.C.: Segmentation Through Variable-Order Surface Fitting. IEEE Transactions on Pattern Analysis and Machine Intelligence 10, 167–192 (1988)
9. Li, K., Guo, L., Li, G., Nie, J., Faraco, C., Cui, G., Zhao, Q., Miller, L.S., Liu, T.: Gyral folding pattern analysis via surface profiling. NeuroImage 52(4), 1202–1214 (2010)
10. Liu, T., Li, H., Wong, K., Tarokh, A., Guo, L., Wong, S.T.: Brain tissue segmentation based on DTI data. NeuroImage 38, 114–123 (2007)
11. Chang, C.-C., Lin, C.-J.: LIBSVM: a library for support vector machines. Software (2001), http://www.csie.ntu.edu.tw/cjlin/libsvm
12. Shen, D., Davatzikos, C.: HAMMER: hierarchical attribute matching mechanism for elastic registration. IEEE Trans. Med. Imaging 21(11), 1421–1439 (2002)
13. FMRIB Software Library, http://www.fmrib.ox.ac.uk/fsl/index.html

Conformal Metric Optimization on Surface (CMOS) for Deformation and Mapping in Laplace-Beltrami Embedding Space

Yonggang Shi[1], Rongjie Lai[2], Raja Gill[3], Daniel Pelletier[4], David Mohr[5], Nancy Sicotte[6], and Arthur W. Toga[1,⋆]

[1] Lab of Neuro Imaging, UCLA School of Medicine, Los Angeles, CA, USA
[2] Dept. of Mathematics, University of Southern California, Los Angeles, CA, USA
[3] Dept. of Neurology, UCLA School of Medicine, Los Angeles, CA, USA
[4] Department of Neurology, Yale School of Medicine, New Haven, CT, USA
[5] Department of Preventive Medicine, Northwestern University, Feinberg School of Medicine, Chicago, IL, USA
[6] Cedar Sinai Medical Center, Los Angeles, CA, USA

Abstract. In this paper we develop a novel technique for surface deformation and mapping in the high-dimensional Laplace-Beltrami embedding space. The key idea of our work is to realize surface deformation in the embedding space via optimization of a conformal metric on the surface. Numerical techniques are developed for computing derivatives of the eigenvalues and eigenfunctions with respect to the conformal metric, which is then applied to compute surface maps in the embedding space by minimizing an energy function. In our experiments, we demonstrate the robustness of our method by applying it to map hippocampal atrophy of multiple sclerosis patients with depression on a data set of 109 subjects. Statistically significant results have been obtained that show excellent correlation with clinical variables. A comparison with the popular SPHARM tool has also been performed to demonstrate that our method achieves more significant results.

1 Introduction

Surface mapping is an important technique in studying brain morphometry and has the potential of pinpointing atrophy in various pathologies [1]. While many methods were proposed for the modeling and mapping of anatomical surfaces [2–6], there is still a lack of general, yet feature sensitive, methods for the characterization and mapping of 3D anatomical surfaces. To overcome this challenge, one promising technique that is receiving increased interests is to use Laplace-Beltrami (LB) eigenfunctions as modeling tools of anatomical structures [7–10]. In this work, we propose a novel, and general approach for surface analysis with LB eigenfunctions by optimizing conformal metrics on surfaces.

⋆ This work was supported by NIH grants 5P41RR013642, R01-MH059708, and 5R01MH080892-03, and a DoD grant W81XWH-10-1-0882.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 327–334, 2011.

The eigenfunctions of the LB operator can be considered as the extension of the Fourier basis onto 3D surfaces. The critical difference is that the LB eigenfunctions depend on surface geometry, and are invariant up to isometry. This robustness makes them ideal for intrinsic modeling of anatomical surfaces across population. For general shape classification, the LB eigenvalues were proposed as a DNA-like signature [7]. For detailed analysis of surface geometry, the eigenfunctions provide more information and have been applied to smoothing [8, 11], feature extraction [9, 12], and surface mapping [10].

The metric optimization method proposed in this work is based the embedding of surfaces into the Hilbert space l^2 with their LB eigen-systems [13]. This embedding is scale and pose invariant, and provides a general framework for intrinsic surface analysis. A histogram feature was developed for shape classification with the LB embedding [13]. One important result is that the surface is still a manifold in the embedding space and a rigorous distance measure between embedded manifolds was proposed [14]. The main contribution of this work is that we develop a general approach for surface deformation in the high-dimensional embedding space, which can be applied to various shape analysis tasks such as surface mapping. By iteratively optimizing conformal metrics on a surface, we can evolve its LB eigenvalues and eigenfunctions, and realize its deformation in the embedding space. With this novel technique, we develop an intrinsic approach for surface mapping and demonstrate its application in mapping hippocampal atrophy of multiple sclerosis (MS) patients with depression. Statistically significant results are obtained that show excellent correlation with clinical measures of depression. We also compare our method with the popular SPHARM tool [5] and demonstrate our method is able to achieve more significant mapping results.

2 Conformal Metric Optimization and LB Embedding

In this section, we introduce LB embedding with conformal metrics and develop numerical schemes with finite element methods on triangular meshes. For metric optimization, we derive the derivatives of eigenvalues and eigenfunctions with respect to the weight function in the conformal metric.

2.1 LB Embedding with Conformal Metrics

Let $(\mathcal{M}, g)$ be a genus-zero Riemannian surface where the metric g is the standard metric induced from R^3. For a function $f : \mathcal{M} \to \mathbb{R}$, the LB operator on $\mathcal{M}$ with the metric g is defined as:

$$\Delta^g_{\mathcal{M}} f = \frac{1}{\sqrt{G}} \sum_{i=1}^{2} \frac{\partial}{\partial x_i} (\sqrt{G} \sum_{j=1}^{2} g^{ij} \frac{\partial f}{\partial x_j}) \tag{1}$$

where (g^{ij}) is the inverse matrix of $g = (g_{ij})$ and $G = \det(g_{ij})$. Because the spectrum of $\Delta^g_{\mathcal{M}}$ is discrete, its eigen-system is defined as

$$\Delta^g_{\mathcal{M}} f_n = -\lambda_n f_n \quad (n = 0, 1, 2, \cdots) \tag{2}$$

where λ_n and f_n are the n-th eigenvalue and eigenfunction, respectively. Using the LB eigen-system, an embedding $I^g_{\mathcal{M}} : \mathcal{M} \to l^2$ was proposed in [13]:

$$I^g_{\mathcal{M}}(x) = (\frac{f_1(x)}{\sqrt{\lambda_1}}, \frac{f_2(x)}{\sqrt{\lambda_2}} \cdots, \frac{f_n(x)}{\sqrt{\lambda_n}}, \cdots) \quad \forall x \in \mathcal{M}. \tag{3}$$

Note that the embedding is not unique because of the sign ambiguities in the eigenfunction, i.e., both f_n and $-f_n$ are the n-th eigenfunction. For practical applications, the sign ambiguities can be resolved by anatomical priors or simply searching through all 2^N possible sign combinations when up to N eigenfunctions are used for numerical implementation.

The embedding $I^g_{\mathcal{M}}(\mathcal{M})$ is scale and pose invariant, and automatically aligns surfaces in the l^2 space for the intrinsic analysis of anatomical surfaces [14]. On the other hand, non-isometric shape differences remain intact in the embedding space and affect further analysis. As an example, we show in Fig. 1(a) and (d) two hippocampal surfaces $\mathcal{M}_1$ and $\mathcal{M}_2$ with different degree of bending. Such non-isometric shape differences lead to different eigenfunctions as shown in Fig. 1(b) and (d). Using closest point matching in the embedding space, which is approximated with the first 10 eigenfunctions in this example, we can project the mesh structure of $\mathcal{M}_1$ onto $\mathcal{M}_2$, and vice versa. The impact of non-isometric shape differences can be clearly seen in the large distortions in the projected mesh structures that are plotted in Fig. 1(c) and (f).

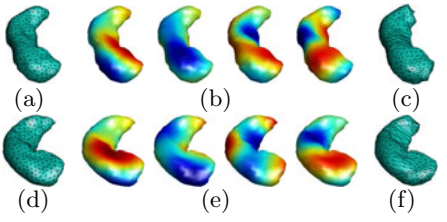

Fig. 1. Impact of non-isometric shape differences. (a)(d) Two surfaces $\mathcal{M}_1$ and $\mathcal{M}_2$. (b)(e) The 4,5,6,7-th eigenfunctions on the two surfaces. (c) (f) Projection of $\mathcal{M}_2$ onto $\mathcal{M}_1$, and $\mathcal{M}_1$ onto $\mathcal{M}_2$ with closest point matching in the embedding space.

A class of conformal metrics on $\mathcal{M}$ are denoted as $\hat{g} = \omega g$, where $\omega : \mathcal{M} \to \mathbb{R}^+$. To achieve better surface matching in the embedding space, we propose in this work to optimize the Riemannian metric in the class of metrics conformally equivalent to g. By iteratively perturbing the weight function, we can realize surface deformation in the embedding space, and minimize non-isometric shape differences. The existence of such weight functions are theoretically guaranteed since all genus-zero surfaces are conformally equivalent. Following (1), the LB operator with the conformal metric is $\Delta^{\hat{g}}_{\mathcal{M}} = \frac{1}{\omega}\Delta^g_{\mathcal{M}}$, and the eigen-system of the weighted operator is

$$\Delta^{\hat{g}}_{\mathcal{M}} f = -\lambda f. \tag{4}$$

Using the relation between $\hat{g}$ and g, we have the weak form of (4):

$$\int_{\mathcal{M}} \nabla^g_{\mathcal{M}} f \nabla^g_{\mathcal{M}} \eta d\mathcal{M} = \lambda \int_{\mathcal{M}} \omega f \eta d\mathcal{M} \quad \forall \eta : \mathcal{M} \to \mathbb{R} \tag{5}$$

where $\nabla^g_{\mathcal{M}}$ is the gradient operator on $\mathcal{M}$ with the standard metric g, and η is a test function.

For numerical computation, we represent $\mathcal{M}$ as a triangular mesh with L vertices $\mathcal{V} = \{v_i | 1 \leq i \leq L\}$. At each vertex v_i, we denote its barycentric coordinate function as ϕ_i and represent the weight function as $\omega = \sum_{j=1}^{L} \omega_j \phi_j$, and the eigenfunction as $f = \sum_{k=1}^{L} \beta_k \phi_k$. By choosing the test function $\eta = \phi_i (1 \leq i \leq L)$, we convert the weak form into its matrix form:

$$Q\beta = \lambda \overline{U}(\omega)\beta. \tag{6}$$

The elements of the matrix Q are defined as $Q_{ik} = \int_{\mathcal{M}} < \nabla\phi_i, \nabla\phi_k > d\mathcal{M}$. The matrix $\overline{U}(\omega)$ is a function of ω with its elements defined as $\overline{U}_{ik}(\omega) = \sum_{j=1}^{L} \omega_j U_{ijk}$, where $U_{ijk} = \int_{\mathcal{M}} \phi_i \phi_j \phi_k d\mathcal{M}$. By solving (6), we can compute the LB embedding $I_{\mathcal{M}}^{wg}$ under the new metric $\hat{g} = wg$.

2.2 Eigen-Derivatives

To realize surface deformation in the embedding space, we derive the derivatives of the eigenvalues and eigenfunctions with respect to the weight function ω.

Let λ_n and f_n denote the n-th eigenvalue and eigenfunction of the LB operator under the conformal metric ωg. We compute the derivative with respect to ω_j, the j-th component of ω, on both sides of (6) and have:

$$Q\frac{\partial f_n}{\partial \omega_j} = \frac{\partial \lambda_n}{\partial \omega_j}\overline{U} f_n + \lambda_n \frac{\partial \overline{U}}{\partial \omega_j} f_n + \lambda_n \overline{U}\frac{\partial f_n}{\partial \omega_j} \tag{7}$$

where the elements of $\frac{\partial \overline{U}}{\partial \omega_j}$ are defined as $[\frac{\partial \overline{U}}{\partial \omega_j}]_{ik} = U_{ijk}$. Pre-multiplying both sides with f_n^T, we obtain:

$$\frac{\partial \lambda_n}{\partial \omega_j} = -\lambda_n f_n^T \frac{\partial \overline{U}}{\partial \omega_j} f_n \tag{8}$$

because $f_n^T \overline{U} f_n = 1$ and $f_n^T (Q - \lambda_n \overline{U}) = 0$.

To compute the derivative of the eigenfunction, we need to solve

$$(Q - \lambda_n \overline{U})\frac{\partial f_n}{\partial \omega_j} = F_n \tag{9}$$

where $F_n = -\lambda_n f_n^T \frac{\partial \overline{U}}{\partial \omega_j} f_n \overline{U} f_n + \lambda_n \frac{\partial \overline{U}}{\partial \omega_j} f_n$. Because $Q - \lambda_n \overline{U}$ is singular, we follow [15] and write $\frac{\partial f_n}{\partial \omega_j} = \mu_{nj} + c_{nj} f_n$ with the constraint that the p-th component of u_{ij} is zero, where p is the index of the component that has the largest magnitude in f_n. This is realized by setting the p-th component of F_n as zero and the p-th row and column of $(Q - \lambda_n \overline{U})$ as zero except the diagonal term, which is set to one. Equation (9) then becomes

$$\begin{bmatrix} [Q - \lambda_n \overline{U}]_{11} & 0 & [Q - \lambda_n \overline{U}]_{12} \\ 0 & 1 & 0 \\ [Q - \lambda_n \overline{U}]_{21} & 0 & [Q - \lambda_n \overline{U}]_{22} \end{bmatrix} \mu_{nj} = \begin{bmatrix} [F_n]_1 \\ 0 \\ [F_n]_2 \end{bmatrix} \tag{10}$$

where $[F_n]_1$ is the 1 to $(p-1)$-th components of F_n, and $[F_n]_2$ is the $p+1$ to the end of the vector F_n. Assuming there is no multiplicity at λ_n [15], this problem is non-singular, and we can solve it to obtain μ_{nj}. To compute c_{nj}, we use the condition that $f_n^T \overline{U} f_n = 1$. By taking derivatives on both sides, we have $\frac{\partial f_n}{\partial \omega_j} \overline{U} f_n = 0$ and get $c_{nj} = -\mu_{nj}^T \overline{U} f_n$. This completes the solution for $\frac{\partial f_n}{\partial \omega_j}$.

3 Surface Mapping via Deformation in Embedding Space

In this section, we demonstrate the application of metric optimization by applying it to compute surface maps in the embedding space. Let $(\mathcal{M}_1, g_1)$ and $(\mathcal{M}_2, g_2)$ denote two surfaces, and ω_1 and ω_2 the weight functions on them, respectively. The eigenvalues and eigen-functions of $(\mathcal{M}_m, \omega_m g_m)(m = 1, 2)$ are denoted as $\lambda_{m,n}$ and $f_{m,n}$. To match these two surfaces in the embedding space, we minimize the following energy function with respect to the conformal metrics:

$$E(\omega_1, \omega_2) = \frac{1}{S_1}\int_{\mathcal{M}_1} d_1^2(\mathbf{x}) d\mathcal{M}_1 + \frac{1}{S_2}\int_{\mathcal{M}_2} d_2^2(\mathbf{x}) d\mathcal{M}_2 + \xi \sum_{i=1}^{2}\int_{\mathcal{M}_i} \|\nabla\omega_i\|^2 d\mathcal{M}_i. \quad (11)$$

The distances are defined as $d_1(x) = \min_{y\in\mathcal{M}_2} \|I_{\mathcal{M}_1}^{\omega_1 g_1}(x) - I_{\mathcal{M}_2}^{\omega_2 g_2}(y)\|_2$ and $d_2(x) = \min_{y\in\mathcal{M}_1} \|I_{\mathcal{M}_2}^{\omega_2 g_2}(x) - I_{\mathcal{M}_1}^{\omega_1 g_1}(y)\|_2$, where $I_{\mathcal{M}_1}^{\omega_1 g_1}$ and $I_{\mathcal{M}_2}^{\omega_2 g_2}$ are LB embeddings chosen to minimize the distances among all possible sign combinations [14]. The third term in the energy encourages smoothness in the weight functions and ξ is a regularization parameter.

To find the optimal metrics that minimize the energy, we iteratively update the weight functions in the gradient descent direction to deform the surfaces in the embedding space. We represent each surface as a triangular mesh $\mathcal{M}_m = (\mathcal{T}_m, \mathcal{V}_m)$ for $m = 1, 2$. For both surfaces, we use the first N eigenfunctions to approximate the embedding. At each iteration, we denote $\mathbf{u}_1(\mathcal{V}_1) = A\mathcal{V}_2$, and $\mathbf{u}_2(\mathcal{V}_2) = B\mathcal{V}_1$ as the closest point maps that minimizes d_1 and d_2 in the embedding space, where A and B are interpolation matrices. Using these two maps, we can write the energy at the current iteration in discrete form:

$$\begin{aligned} E(\omega_1, \omega_2) = \sum_{n=1}^{N} \Bigg(& \frac{1}{S_1}\left(\frac{f_{1,n}}{\sqrt{\lambda_{1,n}}} - \frac{f_{2,n}(\mathbf{u}_1)}{\sqrt{\lambda_{2,n}}}\right)^T U_1 \left(\frac{f_{1,n}}{\sqrt{\lambda_{1,n}}} - \frac{f_{2,n}(\mathbf{u}_1)}{\sqrt{\lambda_{2,n}}}\right) \\ & + \frac{1}{S_2}\left(\frac{f_{2,n}}{\sqrt{\lambda_{2,n}}} - \frac{f_{1,n}(\mathbf{u}_2)}{\sqrt{\lambda_{1,n}}}\right)^T U_2 \left(\frac{f_{2,n}}{\sqrt{\lambda_{2,n}}} - \frac{f_{1,n}(\mathbf{u}_2)}{\sqrt{\lambda_{1,n}}}\right)\Bigg) + \xi(\omega_1^T Q_1 \omega_1 + \omega_2^T Q_2 \omega_2) \end{aligned} \quad (12)$$

where U_m and Q_m are matrices defined in (6) with uniform weight, i.e., the standard metric. Using the eigen-derivatives with respect to the weight functions, we can derive the gradient flows for the weight functions as follows:

$$\begin{aligned} \frac{\partial E}{\partial \omega_1} = 2\sum_{n=1}^{N}\Bigg[& \frac{1}{S_1}\left(\frac{1}{\sqrt{\lambda_{1,n}}}\frac{\partial f_{1,n}}{\partial \omega_1} - \frac{\partial \lambda_{1,n}}{\partial \omega_1}\frac{(f_{1,n})^T}{2\sqrt[3/2]{\lambda_{1,n}}}\right) U_1 \left(\frac{f_{1,n}}{\sqrt{\lambda_{1,n}}} - \frac{Af_{2,n}}{\sqrt{\lambda_{2,n}}}\right) \\ & - \frac{1}{S_2}\left(\frac{\partial f_{1,n}}{\partial \omega_1}\frac{B^T}{\sqrt{\lambda_{1,n}}} - \frac{\partial \lambda_{1,n}}{\partial \omega_1}\frac{(Bf_{1,n})^T}{2\sqrt[3/2]{\lambda_{1,n}}}\right) U_2 \left(\frac{f_{2,n}}{\sqrt{\lambda_{2,n}}} - \frac{Bf_{1,n}}{\sqrt{\lambda_{1,n}}}\right)\Bigg] + 2\xi Q_1 \omega_1 \end{aligned} \quad (13)$$

$$\begin{aligned} \frac{\partial E}{\partial \omega_2} = 2\sum_{n=1}^{N}\Bigg[& \frac{1}{S_2}\left(\frac{1}{\sqrt{\lambda_{2,n}}}\frac{\partial f_{2,n}}{\partial \omega_2} - \frac{\partial \lambda_{2,n}}{\partial \omega_2}\frac{(f_{2,n})^T}{2\sqrt[3/2]{\lambda_{2,n}}}\right) U_2 \left(\frac{f_{2,n}}{\sqrt{\lambda_{2,n}}} - \frac{Bf_{1,n}}{\sqrt{\lambda_{1,n}}}\right) \\ & - \frac{1}{S_1}\left(\frac{\partial f_{2,n}}{\partial \omega_2}\frac{A^T}{\sqrt{\lambda_{2,n}}} - \frac{\partial \lambda_{2,n}}{\partial \omega_2}\frac{(Af_{2,n})^T}{2\sqrt[3/2]{\lambda_{2,n}}}\right) U_1 \left(\frac{f_{1,n}}{\sqrt{\lambda_{1,n}}} - \frac{Af_{2,n}}{\sqrt{\lambda_{2,n}}}\right)\Bigg] + 2\xi Q_2 \omega_2 \end{aligned} \quad (14)$$

Starting from a pair of embeddings $I_{\mathcal{M}_1}^{\omega_1 g_1}$ and $I_{\mathcal{M}_2}^{\omega_2 g_2}$ that achieve the minimum energy among 2^N possible sign combinations, we iteratively deform the embeddings by optimizing the metrics in the gradient descent direction following above

equations. Note that the search through 2^N combinations only needs be done once in the first iteration to resolve sign ambiguities in eigenfunctions. After that, we can resolve the sign ambiguities efficiently by comparing correlations between corresponding eigenfunctions in consecutive iterations because only small perturbations are introduced in one iteration. Once the iterative process converges, we obtain $\mathbf{u}_1$ and $\mathbf{u}_2$ as the maps between these two surfaces.

4 Experimental Results

In this section, we present experimental results on hippocampal surface mapping to demonstrate the application of our method in brain imaging research. In the first experiment, we present detailed results on the mapping of two surfaces. The robustness and clinical relevance of our method are demonstrated in the second experiment on a clinical dataset of 109 subjects. A comparison with the popular SPHARM tool is presented in the third experiment.

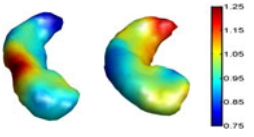

Fig. 2. The optimized weight function in the conformal metric of the two surfaces

4.1 Mapping Results of Two Surfaces

In this experiment, we apply our metric optimization method to the two surfaces in Fig. 1(a) and (d). The parameters are $N = 20$, $\xi = 0.1$ and 150 iterations of metric optimization according to equations (13) and (14) are used to obtain the final maps $\mathbf{u}_1$ and $\mathbf{u}_2$. The computational time is around 15 minutes on a PC.

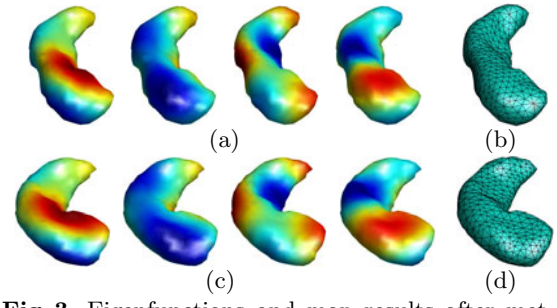

Fig. 3. Eigenfunctions and map results after metric optimization. (a) Eigenfunctions of $\mathcal{M}_1$. (b) $\mathbf{u}_2(\mathcal{M}_2)$. (c) Eigenfunctions of $\mathcal{M}_2$. (d) $\mathbf{u}_1(\mathcal{M}_1)$.

The optimized weight functions of these two surfaces are shown in Fig. 2, where corresponding regions exhibit complimentary metric deformations to account for non-isometric differences of these parts. Under the conformal metrics, the corresponding eigenfunctions of the two surfaces are shown in Fig. 3(a) and (c). We can see they are almost identical and agree much better than the ones in Fig. 1. The quality of the maps can be visualized by the mesh quality of $\mathbf{u}_1(\mathcal{M}_1)$ and $\mathbf{u}_2(\mathcal{M}_2)$ in Fig. 3(b) and (d). Compared with the meshes in Fig. 1(c) and (f), we can see the mesh structures are much more uniform.

4.2 Hippocampal Mapping in MS Patients with Depression

In this experiment, we demonstrate the robustness of our method by applying it to a clinical study of hippocampal atrophy in MS patients with depression. Using the Center for Epidemiologic Studies-Depression (CES-D) scale as

the measure for depression, the 109 female subjects in this study are split into two groups: low depression (CES-D$\leq$ 20) and high depression (CES-D$>$ 20). To study group differences, the right hippocampi are mapped with our metric optimization method to an atlas surface, which is the right hippocampus of one randomly selected subject. Using the computed maps, we project the mesh structure of the atlas onto all hippocampal surfaces to establish one-to-one correspondences across subjects. At each corresponding triangle of the 109 surfaces, a one-sided t-test is applied using a thickness measure [10] to test the hypothesis that MS patients with high depression have more severe hippocampal atrophy.

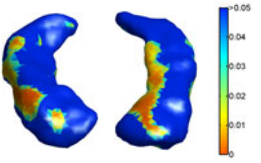

Fig. 4. Top (left) and bottom (right) views of the thickness p-value map from our method

As shown in the significance map of p-values in Fig. 4, the highlighted regions indicate larger atrophy occurs in the right hippocampus of MS patients with high depression. To correct for multiple comparisons, we applied 1000 permutation tests and an overall p-value of 0.017 is obtained, which means the overall significance of the thickness map. To further validate the clinical relevance of the thickness map, we test the correlation of the thickness measure and CES-D scores at each triangle. The correlation coefficients are plotted in Fig. 5(a) with the significance map of the correlation in Fig. 5(b), which indicates highlighted regions in Fig. 4 match excellently to regions with significant negative correlations between thickness and CES-D scores. This shows that patients with more severe depressive symptoms in deed have more hippocampal atrophy in the highlighted regions detected by our method.

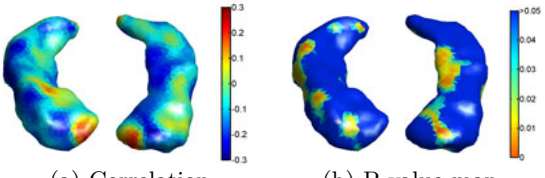

(a) Correlation. (b) P-value map.

Fig. 5. The correlation between CESD and thickness, and its p-value map

4.3 Comparison with SPHARM

In this experiment, we apply the popular SPHARM tool to map the same group of hippocampal surfaces in the second experiment. We adopt the suggested parameters for hippocampus in the manual of SPHARM [5]. Using the correspondences established by the SPHARM maps, the same one-sided t-tests are applied at each triangle to test for group differences among the 109 surfaces. The resulting significance map of p-values is plotted in Fig. 6. To correct for multiple comparisons, we also apply 1000 permutation tests and the overall p-value is 0.07. By comparing the results in Fig. 4, 5 and 6, we can see that our method achieves better performance by detecting more group differences that correlate well with clinical variables.

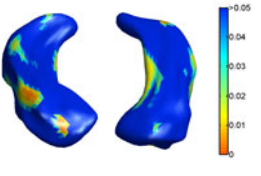

Fig. 6. Top (left) and bottom (right) views of the thickness p-value map from the SPHARM tool

5 Conclusions

In this work we proposed a general approach for surface deformation in LB embedding space by optimizing conformal metrics on surfaces. In future work, we will minimize bias in statistical analysis by developing group-based techniques. We will also conduct more extensive comparisons with existing methods.

References

1. Thompson, P., Hayashi, K., de Zubicaray, G., Janke, A., Rose, S., Semple, J., Hong, M., Herman, D., Gravano, D., Doddrell, D., Toga, A.: Mapping hippocampal and ventricular change in Alzheimer disease. NeuroImage 22(4), 1754–1766 (2004)
2. Joshi, S., Miller, M.: Landmark matching via large deformation diffeomorphisms. IEEE Trans. Imag. Process. 9(8), 1357–1370 (2000)
3. Davies, R., Twining, C., Allen, P., Cootes, T., Taylor, C.: Shape discrimination in the hippocampus using an MDL model. In: Taylor, C.J., Noble, J.A. (eds.) IPMI 2003. LNCS, vol. 2732, pp. 38–50. Springer, Heidelberg (2003)
4. Yeo, B., Sabuncu, M., Vercauteren, T., Ayache, N., Fischl, B., Golland, P.: Spherical demons: Fast surface registration. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part I. LNCS, vol. 5241, pp. 745–753. Springer, Heidelberg (2008)
5. Styner, M., Oguz, I., Xu, S., Brechbühler, C., Pantazis, D., Levitt, J., Shenton, M., Gerig, G.: Framework for the statistical shape analysis of brain structures using SPHARM-PDM. The Insight Journal (2006)
6. Yushkevich, P., Zhang, H., Gee, J.: Continuous medial representation for anatomical structures. IEEE Trans. Med. Imag. 25(12), 1547–1564 (2006)
7. Reuter, M., Wolter, F., Peinecke, N.: Laplace-Beltrami spectra as Shape-DNA of surfaces and solids. In: Computer-Aided Design, vol. 38, pp. 342–366 (2006)
8. Qiu, A., Bitouk, D., Miller, M.: Smooth functional and structural maps on the neocortex via orthonormal bases of the Laplace-Beltrami operator. IEEE Trans. Med. Imag. 25(10), 1296–1306 (2006)
9. Shi, Y., Lai, R., Krishna, S., Sicotte, N., Dinov, I., Toga, A.: Anisotropic Laplace-Beltrami eigenmaps: Bridging Reeb graphs and skeletons. In: Proc. MMBIA, pp. 1–7 (2008)
10. Shi, Y., Morra, J., Thompson, P., Toga, A.: Inverse-consistent surface mapping with Laplace-Beltrami eigen-features. In: Prince, J.L., Pham, D.L., Myers, K.J. (eds.) IPMI 2009. LNCS, vol. 5636, pp. 467–478. Springer, Heidelberg (2009)
11. Seo, S., Chung, M., Vorperian, H.: Heat kernel smoothing using Laplace-Beltrami eigenfunctions. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. 6363, pp. 505–512. Springer, Heidelberg (2010)
12. Reuter, M.: Hierarchical shape segmentation and registration via topological features of Laplace-Beltrami eigenfunctions. Int'l Journal of Computer Vision 89(2-3), 287–308 (2010)
13. Rustamov, R.M.: Laplace-beltrami eigenfunctions for deformation invariant shape representation. In: Proc. Eurograph. Symp. on Geo. Process. pp. 225–233 (2007)
14. Lai, R., Shi, Y., Scheibel, K., Fears, S., Woods, R., Toga, A., Chan, T.: Metric-induced optimal embedding for intrinsic 3D shape analysis. In: Proc. CVPR, pp. 2871–2878 (2010)
15. Nelson, R.: Simplified calculation of eigenvector derivatives. AIAA Journal 14(9), 1201–1205 (1976)

Area-Preserving Surface Flattening Using Lie Advection

Guangyu Zou[1,⋆], Jiaxi Hu[2], Xianfeng Gu[3], and Jing Hua[2]

[1] Innovisgroup, Inc., China
[2] Wayne State University, USA
[3] State University of New York at Stony Brook, USA
gyzou@innovisgroup.com

Abstract. In this paper, we propose a novel area-preserving surface flattening method, which is rigorous in theory, efficient in computation, yet general in application domains. Leveraged on the state-of-the-art flattening techniques, an infinitesimal area restoring diffeomorphic flow is constructed as a Lie advection of differential 2-forms on the manifold, which yields strict equality of area elements between the flattened and the original surfaces at its final state. With a surface represented by a triangular mesh, we present how an deterministic algorithm can be faithfully implemented to its continuous counterpart. To demonstrate the utility of this method, we have applied our method to both the cortical hemisphere and the entire cortex. Highly complied results are obtained in a matter of seconds.

Keywords: Brain mapping, area-preserving flattening, Lie advection.

1 Introduction

Given the fact that many anatomical surfaces are intrinsically 2D and highly undulated, flattening techniques constitute a major means of visualizing pathologies that are deeply buried within the folds [4]. Flattening (parametrization) also enables many procedures on a regular parameter domain, therefore resulting in more efficient and stable computation. However, existing flattening methods, including conformal mapping [2,6,5], typically suffer from severe, unpredictable area distortion when dealing with extruding shapes that contains rich, complex features, which largely impedes the capture of anatomical characteristics as well as other associated imaging modalities on a planar domain. Although it is known that surface cuts can effectively reduce the distortion of final flattening, such practice is typically not preferred, since neighborhoods on opposite sides of a cut will become far apart in the final flattened representation.

The technique presented in this paper theoretically ensures a uniform sampling of the surface on the parameter domain. Each patch retains exactly the same area when flattened to 2D. Based on the mathematical advance, a series of analysis tasks regarding neuronal density, activation extent, cortical thickness,

⋆ Corresponding Author.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 335–342, 2011.

among others, can be instead invoked in 2D where more compact data structures, more efficient discretization schemes and faster data access are available. Towards the same goal, a handful of respectable efforts have been made [3,4,8,9]. However, the flattening results typically correspond only to local minima with respect to certain objective functionals [3,4], and lack provable guarantee of area preservation. Pons et al. [8] designed a tangential motion according to a given normal motion in a cortical inflating procedure, which preserves area exactly, but is not capable of flattening a 3D surface into 2D. A mass (e.g., area/volume) preserving mapping was explicitly sought in $\mathbf{R}^n$ by a gradient decent method to the Monge-Kantorovich functional in [9], implemented on a regular Cartesian grid. It is not clear how this method can be extended to a general manifold.

In contrast, our method computes global strict area-preserving (A.P.) flattening of arbitrary 2-manifolds using *Lie advection*, a concept from classical mechanics. To our best knowledge, this is the first work that employs Lie advection as a tool to manipulate area changes in the context of surface flattening. Besides a general framework, our method also allows an efficient, yet accurate discretization scheme that is motivated by preserving the original geometric and algebraic structures of the continuous model in the limit, therefore rendering better numerical fidelity.

A similar idea was mathematically sketched in [1]. The discussion was restricted to a spherical domain. Despite similar concepts, our method is deterministic and derived for arbitrary surface manifolds. In the remainder of this paper, we take brain surface as an acting example, while the presented method is principally applicable to general surfaces.

2 Basic Idea

Starting with an arbitrary initial diffeomorphism from an given surface to the desired domain, e.g., a conformal parametrization [7,5], we can subsequently evolve it to an A.P. alternative as follows. Suppose M and N are two differentiable 2-manifolds, associated by a diffeomorphism $f : M \to N$. The area element of a surface is a differential 2-form. Let ω_i, $i = M, N$ be the area form of M and N, respectively. The pullback of ω_N under f is a differential 2-form on M, denoted as $f^*(\omega_N)$. Suppose M and N have the same area integral after an appropriate scaling, that is, $\int_M \omega_M = \int_N \omega_N$. Computing an A.P. map $\mu : M \to N$ now is equivalent to finding a diffeomorphism $\varphi : M \to M$, such that $\varphi^*(\omega_M) = f^*(\omega_N)$. Therefore, μ is given by $f \circ \varphi^{-1}$. To accomplish this, we first linearly interpolate a 2-form over time:

$$\omega_t = (1-t)\omega_M + tf^*(\omega_N), \quad t \in [0,1]. \tag{1}$$

Note that $\omega_0 = \omega_M$ and $\omega_1 = f^*(\omega_N)$. In the following, we will design a one parameter family of diffeomorphisms, such that the corresponding flow deforms the area element in the same fashion as ω_t.

More specifically, consider a smooth surface M with a smooth vector field V on it. Given any point $p \in M$, there exist a unique integral curve $\gamma(t)$ of V passing through it, such that

$$\begin{cases} \frac{d\gamma_p(t)}{dt} = V(\gamma_p(t)), \\ \gamma_p(0) = p. \end{cases} \tag{2}$$

A one parameter family of diffeomorphisms (which are also automorphisms) ϕ_t, parameterized by $t \in [0, 1]$, can be defined on M as

$$\phi_t(p) = \gamma_p(t). \tag{3}$$

We want $\phi_t^*(\omega_0) = \omega_t$. Substituting p with ω_t in Eq. (3) and computing time derivative at $t = 0$, we get

$$\frac{d\phi_t^*\omega}{dt}\Big|_{t=0} = \omega_M - f^*(\omega_N), \tag{4}$$

which, by definition, is the *Lie derivative* of ω_t with respect to V. Hence, the central equation to solve is

$$\mathcal{L}_{V(t)}\omega_t = \omega_M - f^*(\omega_N), \tag{5}$$

where $\mathcal{L}_{V(t)}$ denotes the Lie derivative with respect to $V(t)$. Intuitively, $\mathcal{L}_{V(t)}$ estimates the change of ω_t along the flow of $V(t)$. By using Cartan's formula $\mathcal{L}_V = di_V + i_V d$, where i_V denotes the interior product with respect to $V(t)$, we have

$$d(i_V\omega_t) + i_V d\omega_t = \omega_M - f^*(\omega_N). \tag{6}$$

Since ω_t is a 2-form, $d\omega_t = 0$ on M. Hence, we have

$$d(i_V\omega_t) = \omega_M - f^*(\omega_N). \tag{7}$$

By definition, we can write $\omega_M = h_1 du \wedge dv$ and $f^*(\omega_N) = h_2 du \wedge dv$. h_1 and h_2 are the scaling factors from a mutual parameter domain to M and N, respectively. Now let

$$\Delta g = h_2 - h_1, \tag{8}$$

where Δ denotes the Laplacian-Beltrami differential operator. We substitute Eq. 8 into Eq. 7. Using the fact that M and N are 2-manifolds, $V(t)$ can be solved as

$$V(t) = \frac{1}{(1-t)h_1 + th_2}\nabla g. \tag{9}$$

Note that, as implied by Eq. (9), $V(t)$ varies both in magnitude and direction over time. we need to solve $V(t)$ at each time. g is essentially a harmonic scalar field on M. As a result, the corresponding gradient vector field ∇g is guaranteed by construction to be highly smooth and free of extraneous critical points. Time integration of $V(t)$ yields a diffeomorphism. Finally, the A.P. mapping is given by $f \circ \phi_{t=1}^{-1} : M \to N$.

3 Algorithm

In practice, the surface is represented by a triangular mesh. The position of vertex v_i is denoted by $\mathbf{v}_i \in \mathbf{R}^3$. The 1-ring neighbors of v_i is denoted as $N_1(i)$. For a vertex v_i, the associated surface patch is chosen to be the barycentric finite volume, denoted as A_i.

We first assume the target domain is a unit square, which is commonly used to parameterize topological disks. Later, we will show how this framework can be extended to the unit sphere to flatten closed genus zero surfaces in spherical geometry. The square boundary condition is set up as follows: first, the boundary is isometrically mapped to that of the unit square D; next, a discrete conformal map of the interior is computed in the least squares sense [7]. Fig. 1(a) and 1(b) show the lateral and mesial views of a cortical hemisphere, respectively. The initial (conformal) flattening is illustrated in Fig. 1(c), the area distortion of which is color encoded in Fig. 1(d). Notice how some lateral cortical patterns suffer from the intense geometric stretch, which greatly impairs inspection of these regions. Note that, since our A.P. surface flattening method is independent of the initial mapping, other surface parametrization methods can also be equally employed in order to achieve specific functionality relevant in applications. In

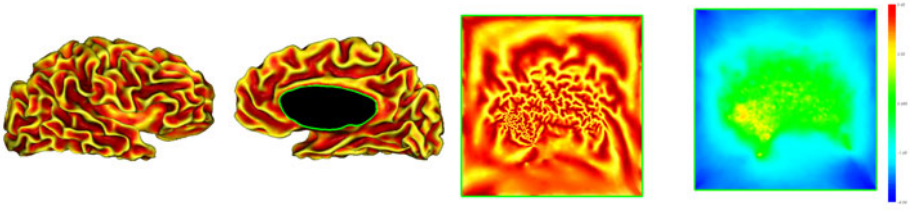

(a) Lateral view (b) Mesial view (c) Conformal map (d) Area distortion

Fig. 1. Initial conformal flattening. (a) and (b) show the lateral and mesial views of a cortical hemisphere. (c) is the initial (conformal) flattening on a unit square. The area distortion is color-coded in (d).

the following, we describe the essential steps of the algorithm in the order they occur in the procedure.

3.1 Solving $\Delta g = h_2 - h_1$

For an area-preserving flattening, h_1 is always 1, whereas h_2 is the model area/parameter area ratio at play. When the total areas of the model and the domain are not equal, h_2 is subject to a normalization to make them equal. The result will therefore be of relative area preservation.

Given a continuous function on the surface, its discrete version is represented as a vector, defined on the vertex set V. To solve Eq. (8) on a triangular mesh, Δg is estimated at v_i in the same manner as [3]. Note that this discrete approximation requires acute angles. For obtuse ones, a proper remeshing procedure should be employed. As such, Eq. (8) can be written as a linear system: $Lx = b$,

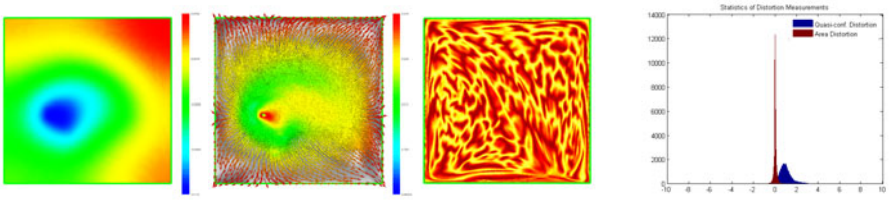

(a) $g(t=0)$ (b) $\nabla g(t=0)$ (c) A.P. map$(t=1)$ (d) Area/angle distortion

Fig. 2. Area-preserving flattening on a unit square. The solved function g is shown in (a) and its gradient vector field in (b). After time integration of a dynamic diffeomorphic flow following the vector field, the proposed A.P. flattening is obtained in (c). (d) shows the histograms of area and angle distortion metrics.

where $x = \boldsymbol{g}$, $b = \boldsymbol{h}_2 - \boldsymbol{h}_1$, L represents the coefficient matrix of the discrete Laplace-Beltrami operator. As L is sparse, Eq. (8) can be solved efficiently in linear time. Fig. 2(a) shows the solved function g.

3.2 Computing ∇g

Now that the g is obtained, we can proceed to compute the corresponding gradient vector field on the triangulated domain. We consider a face f_{ijk} with its three corners lying at $\mathbf{v}_i$, $\mathbf{v}_j$, $\mathbf{v}_k$ in $\mathbf{R}^3$. Also, let $\mathbf{n}$ be a unit normal vector perpendicular to the plane spanned by f_{ijk}. By assuming linear interpolation within each triangle, the gradient vector can be easily computed by solving the 3×3 linear system:

$$\begin{bmatrix} \mathbf{v}_j - \mathbf{v}_i \\ \mathbf{v}_k - \mathbf{v}_j \\ \mathbf{n} \end{bmatrix} \nabla g = \begin{bmatrix} g_j - g_i \\ g_k - g_j \\ 0 \end{bmatrix}, \tag{10}$$

for which a closed-form solution exists. To obtain a unique vector at each vertex, ∇g at vertex v_i is defined as an average of the gradients of the adjacent faces, weighted by the incident angle α_{jk}^i of each face f_{ijk} at v_i. The resulting vector field is shown in Fig. 2(b).

3.3 Time Integration of $V(t)$

Recall that the Lie derivative is defined as the instantaneous change of forms evaluated at $\phi_t(x)$, which is a dynamic definition (See Eq. (4)). In fact, $V(t)$ only coincides with the flow vector field at $t = 0$. In other cases, $V(t)$ needs be transported accordingly. By an simple analytical integration, we have

$$\int_0^1 V(t)dt = \frac{\ln h_2 - \ln h_1}{h_2 - h_1} \nabla g. \tag{11}$$

When $h_1 = 1$, we have $\lim_{h_2 \to 1} \frac{\ln h_2 - \ln h_1}{h_2 - h_1} = 1$, which means that, when h_2 is sufficiently close to 1, the displacement vector field can be properly approximated by ∇g. Thus, the area-correcting process is divided into K sequential

steps. With each step, the area element is only modified by an small amount δh towards the target setting such that the overall area adjustment is equal to $K\delta h$. More specifically, we let $h_2 = 1 + \delta h$ and $h_1 = 1$ as the input of the analytical integration (Eq. (11)), and the result gives the corresponding displacement for the area change of δh. This procedure is discretized into 50 steps for the results shown. By choosing step length carefully, we can guarantee that the time integration converges. The final A.P. flattening of the cortical hemisphere is shown in Fig. 2(c). Fig. 3(d) presents the statistics of the area (darkred) and angle (blue) distortions; the employed metrics will be explained in Section 4. All boundary vertices can either be fixed, or evolved by the covariant derivative of $V(t)$ along ∂D. Empirically, the latter one gives better approximation to the continuous case, but less robust to degenerate mesh triangles.

3.4 Genus Zero Surface Flattening

The entire brain surface is often modeled as a topological sphere, i.e., a closed surface of genus zero, thus is preferred to be flattened on a unit sphere ($\mathbf{S}^2$) without any topological changes. Our method can easily adapt to this case as well. We first compute its initial mapping using the method described in [5]. The Lie advection flowing to the area preservation is then performed in the tangent spaces of the spherical domain. In practice, whenever a vertex moves out of the unit sphere, it is pulled back by $\tilde{v}_i = v_i/|v_i|$. Fig. 3(a) shows the entire cortex, while its spherical conformal mapping is shown in 3(b). Its A.P. flattening is shown in Fig. 3(c) and the corresponding histograms about the area (darkred) and angle (blue) distortions is presented in 3(d).

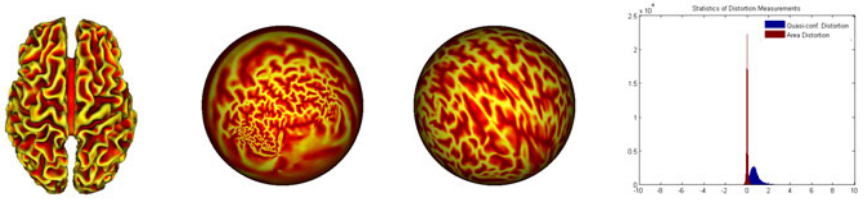

(a) $g(t=0)$ (b) $\nabla g(t=0)$ (c) A.P. map$(t=1)$ (d) Area/angle distortion

Fig. 3. Area preserving mapping of closed genus zero surface. (a) is the original brain surface model. Its spherical conformal mapping and subsequent A.P. flattening are shown in (b) and (c), respectively. Similarly, the statistics of the result is shown in (d).

4 Results

All brain surface models are reconstructed on the gray matter/white matter interface of 3D MRI brain volumes. To verify that the area elements (associated with the vertice) are preserved globally independent of the triangulation, we build up our distortion measures on their dual cells–triangular faces. To be specific, we examine both the area distortion and the quasi-conformal distortion

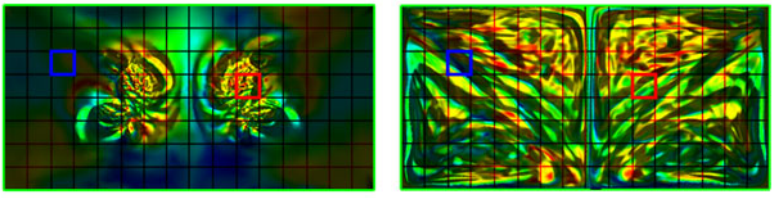

(a) Conformal flattening result (b) A.P. flattening result

Fig. 4. Application of quantitative surface-based analytics. The color visualizes the integrated PET data. (a) shows the flattening result via conformal mapping, where most anatomical and functional meaningful areas are squeezed in a few square elements. In contrast, the cortical surface is more evenly sampled under A.P. flattening, as shown in (b).

per face over the mesh. The area distortion metric Υ and the quasi-conformal distortion metric Λ are computed respectively as follows:

$$\Upsilon = \ln(\gamma_{max} \cdot \gamma_{min}), \quad \Lambda = \ln \frac{\gamma_{max}}{\gamma_{min}}, \tag{12}$$

where $(\gamma_{max}, \gamma_{min})$ are the larger and smaller eigenvalues of the Jacobian of the affine transformation that maps the domain triangle to the surface, normalized in such a way that the total area of the surface equals that of the domain. In both cases, a value of 0 indicates no distortion at all, while distortions can deviate on both sides.

In Fig. 2(d) and Fig. 3(d), the statistics of the area distortion (darkred) and the quasi-conformal ones (blue) are illustrated for both examples, respectively. From the distribution of metric Υ, we can see that the areas of triangles are clearly well preserved. Because for general manifolds, no isometric mapping exists except for a few special cases, quasi-conformal distortion is expected for A.P. maps, as indicated by the distribution of Λ.

In terms of runtimes, the underlying mesh of the cortical hemisphere model is composed of 48,287 faces and the A.P. flattening was obtained in 57 sec; the model of the entire cortex is with 99,736 faces and it took 122 sec. All experiments were conducted on an Intel T6600 2.20GHz laptop with 3GB RAM. Generally, the cost linearly depends on the size of the mesh and the number of discretized steps needed for the desired accuracy.

For a visual analytical framework that is designed to integrate various complementing neuroimaging modalities from multiple sources and for quantitative analyses, A.P. flattening exhibits unique advantages. In Fig. 4, the brain surface is sliced open along the medial plane at the bottom, without passing any significant anatomical features. The entire brain surface is then mapped to a rectangle with a height/width ratio of 1:2. To capture the statistics of abnormal brain activity, a sufficient number of sampling points should be grouped at certain resolution. Each unit forms a cortical element. Since the surface of the brain is uneven and varies across human subjects, subdividing it into a set of equal

geometric elements is nontrivial. With A.P. flattening, we can easily define finite homotopic cortical elements on the parametric domain with simple isotropic grid. Comparing to other "quasi" area-preserving/conformal mapping, each element based on an A.P. flattening accounts for an identical amount of portion in the original brain surface. With these well-defined equiareal cortical elements, a variety of functional patterns can be readily quantified on a per element basis.

5 Conclusion

In this paper, we have presented a special surface flattening methodology that is strictly area-preserving. Given an arbitrary initial parametrization, an A.P. result can be efficiently and uniquely obtained via the Lie advection of area forms along the domain. As our method only depends on the intrinsic surface geometry, it is insensitive to different discretized representations, including variations in triangulation and/or resolution. Our implementation strives to preserve the analytical ingredients of the computation as far as possible. As a result, our method is highly efficient and stable in getting consistent results, as opposed to an optimization approach. While conformal maps are known to be powerful for shape analysis, in a number of other scenarios, such as studies of neuronal density, cortical functional activation extent, and cortical thickness variations, where accurate modality sampling and statistical sensitivity are critical, areal preservation is highly preferable. Extensive comparison between our method and existing approaches in clinical practice will be conducted in the near future.

References

1. Angenent, S., Haker, S., Tannenbaum, A., Kikinis, R.: On area preserving mappings of minimal distortion. In: Djaferis, T., Schick, I. (eds.) System Theory: Modeling, Analysis, and Control. Kluwer, Holland (1999)
2. Angenent, S., Haker, S., Tannenbaum, A., Kikinis, R.: On the laplace-beltrami operator and brain surface flattening. IEEE Trans. Med. Imag. 8, 700–711 (1999)
3. Desbrun, M., Meyer, M., Alliez, P.: Intrinsic parameterizations of surface meshes. Comput. Graph. Forum 21, 209–218 (2002)
4. Fischl, B., Sereno, M., Dale, A.: Cortical surface-based analysis II: Inflation, flattening, and a surface-based coordinate system. NeuroImage 9(2), 195–207 (1999)
5. Gu, X., Wang, Y., Chan, T., Thompson, P., Yau, S.: Genus zero surface conformal mapping and its application to brain surface mapping. IEEE Trans. Med. Imag. 23, 949–958 (2003)
6. Hurdal, M., Stephenson, K.: Discrete conformal methods for cortical brain flattening. NeuroImage 45, S86–S98 (2009)
7. Lévy, B., Petitjean, S., Ray, N., Maillot, J.: Least squares conformal maps for automatic texture atlas generation. ACM Trans. Graph. 21(3), 362–371 (2002)
8. Pons, J., Keriven, R., Faugeras, O.: Area preserving cortex unfolding. In: Barillot, C., Haynor, D.R., Hellier, P. (eds.) MICCAI 2004. LNCS, vol. 3216, pp. 376–383. Springer, Heidelberg (2004)
9. Zhu, L., Haker, S., Tannenbaum, A.: Flattening maps for the visualization of multibranched vessels. IEEE Trans. Med. Imaging 24(2), 191–198 (2005)

Non-parametric Population Analysis of Cellular Phenotypes

Shantanu Singh[1], Firdaus Janoos[1], Thierry Pécot[1], Enrico Caserta[3], Kun Huang[2], Jens Rittscher[4], Gustavo Leone[3], and Raghu Machiraju[1]

[1] Dept. of Computer Science and Engg., The Ohio State University, U.S.A.
[2] Dept. of Biomedical Informatics, The Ohio State University, U.S.A.
[3] Dept. of Molecular Genetics, The Ohio State University, U.S.A.
[4] General Electric Global Research Center, U.S.A.

Abstract. Methods to quantify cellular–level phenotypic differences between genetic groups are a key tool in genomics research. In disease processes such as cancer, phenotypic changes at the cellular level frequently manifest in the modification of cell population *profiles*. These changes are hard to detect due the ambiguity in identifying distinct cell phenotypes within a population. We present a methodology which enables the detection of such changes by generating a phenotypic signature of cell populations in a data–derived feature–space. Further, this signature is used to estimate a model for the *redistribution* of phenotypes that was induced by the genetic change. Results are presented on an experiment involving deletion of a tumor–suppressor gene dominant in breast cancer, where the methodology is used to detect changes in nuclear morphology between control and knockout groups.

Keywords: Microscopy Image Analysis, Cell Nucleus, Shape Analysis.

1 Introduction

Gene targeting methods have elucidated the roles of several genes in disease processes such as cancer. In these experiments, a genetic perturbation such as a gene knockout is introduced in a model organism and the effects on the genetic as well as phenotypic makeup of the organism are examined. Recent advances in microscopy coupled with automated image analysis tools have enabled researchers to *quantify* a broad range of such phenotypic alterations at the cellular level [8]. In these studies, the phenotype of interest, such as cell morphology, is quantified using image features across a population of cells, and differences between normal and genetically perturbed populations are examined. In the event that a clear hypothesis can be posed based on a priori knowledge, it is possible to extract specific cellular features of interest [2] and investigate this hypothesis using standard statistical tests.

In the absence of such prior knowledge, investigations rely on exploratory data analysis techniques, making the detection of phenotypical changes more challenging. In addition, some of the changes affect only a certain subpopulation

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 343–351, 2011.

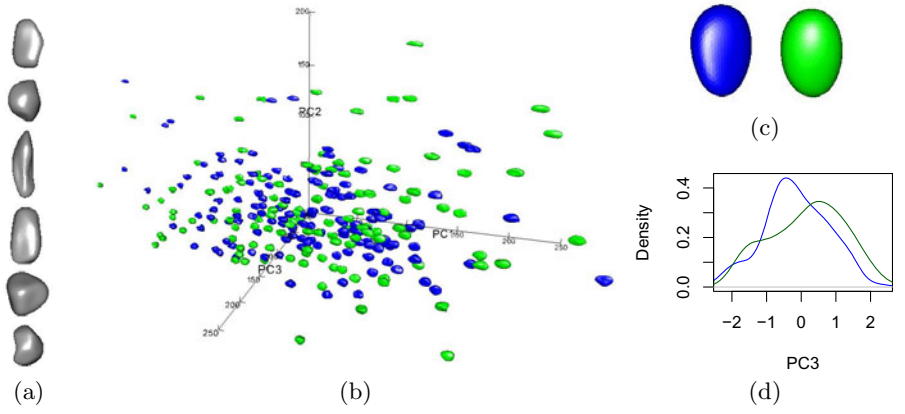

Fig. 1. Detecting changes in cell populations. (a) Examples of fibroblast nuclei from the study (b) Scatter plot of fibroblast nuclei obtained from a control group (Wild-type) and genetically modified group (*PTEN*-deleted) plotted in top three principal components (PC) of the shape space; no evidence of separation between the groups (c) Mean nuclear shapes of the two groups; difference between means is not statistically significant (Sec. 4.1) (d) Estimated (marginal) densities of the two groups in the third PC; *distributional* differences between the groups are statistically significant (Sec. 4.1).

of the cells, resulting in very subtle differences between groups. In this case it is no longer sufficient to analyze the global statistic of a specific cellular feature to detect the presence of an effect. In particular, during the onset of a disease such as the early stage cancer, small perturbations in the system are not easily detectable using existing methods. The quantification of such changes through the analysis of alterations in the population as a whole is the focus of this paper.

In this paper a methodology is presented that is applied to detect the phenotypic effect of deletion of a tumor–suppressor gene (*PTEN*) in the context of breast cancer. Deletion of *PTEN* in fibroblasts has been shown to increase the predisposition to cancer in certain mouse models [11]. Of specific interest are changes at the cellular level in the early stages of cancer development which provide an important indication of what precursor activities occur prior to tumor metastasis. Motivated by the seminal role played by the cell nucleus in cancer [13], nuclear morphology has been used in this study as a proxy for the cell phenotype. The study focuses on detecting the effect of *PTEN* deletion on the nuclear shape of fibroblasts in an exploratory manner. The following data analysis approach, illustrated in Fig. 1, explicates the guiding principles of the proposed method.

A spherical harmonics–based shape parameterization was used to represent nuclear morphology (Sec. 3.1) and PCA was used to reduce dimensionality. A plot of the data in the top three principal components is shown in Fig. 1b. No evidence of a discernible separation between the genotype groups was observed. The differences in the mean shapes (Fig. 1c) were not observed to be statistically

significant based on a non–parametric test for differences between means (Sec. 4.1). However, the presence of an effect was observed when testing for differences in the *shape* of the distributions of the two populations (Fig. 1d). Note that the shape of the distribution essentially represents the phenotypic *population profile*, a characterization of the mix of the nuclear phenotypes present in a population. By using a non–parametric test for difference in distributions between the two groups, a difference between the marginal distributions in the third PC was observed at a significance of $p = 0.002$ (See Sec. 4.1 for further details).

The principles behind this analysis approach are summarized as follows. *First*, the feature set used is *agnostic* in that we did not seek specific structures (such as in [2]) to measure the cell phenotype. Rather, the biologically relevant modes of variation were learned from data using dimensionality reduction, which were used as the feature set. *Second*, the cell population exhibits a significant degree of phenotypic heterogeneity as seen in Fig. 1a. The change in the population profile corresponds to a change in the *composition* of this heterogeneity, which is detected through the analysis. In summary, the above approach enables the detection of phenotypic changes in *population profiles* in a *data–driven* manner.

In this paper, the proposed methodology builds on these principles to further estimate an *interpretable model* for the changes that have occurred. This is achieved by modeling the *redistribution* of cellular phenotypes that putatively occurred due to the induced perturbation. The formulation of the problem is equivalent to finding the Earth Mover's Distance (EMD) [10] in a data–derived feature space. The rest of the paper is organized as follows. A review of related work is presented in Section 2. The proposed methodology is discussed in Section 3 and results on breast cancer datasets are presented in Section 4. We conclude with a discussion in Section 5.

2 Related Work

Several methods for modeling cellular morphology have been proposed in image cytometry literature. A physically–based deformation model of a cell was proposed in [4] which served as an atlas to compare sub–cellular features. A multi–modal approach incorporating diverse cellular features was proposed in [6] that enabled simultaneous segmentation and classification of cells. A statistical methodology for the analysis of sub-cellular features was proposed in [1] that used descriptors from spatial statistics to characterize the internal structure of cells. The above approaches focused on the characterization of individual cells. A computational anatomy framework was used in [9] to quantify shape differences in cells by analyzing the deformation fields in mapping between the samples. This approach enables group analysis on cell populations in an agnostic framework, similar to the approach proposed in this paper. In contrast to the above approaches, this paper presents a methodology that enables the detection of subtle variations between population profiles as a whole, and additionally provides an interpretable model of differences between them.

3 Mathematical Approach

A non–parametric shape representation is adopted in order to support an unbiased exploratory study of nuclear phenotypes. Let $\psi : \mathcal{Z} \rightarrow \mathcal{X}$ denote the shape representation, where $\mathcal{Z}$ is the domain of nuclear shapes and $\mathcal{X}$ is the domain of shape features. A low–dimensional feature–space $\tilde{\mathcal{X}}$ corresponding to biologically plausible nuclear shapes is estimated by learning a transformation $\phi_D : \mathcal{X} \rightarrow \tilde{\mathcal{X}}$ where $D \subset \mathcal{X}$ is a large set of nuclear shapes. Given a normal cell population $P \subset \mathcal{Z}$ and a perturbed one $Q \subset \mathcal{Z}$, the set of most likely transitions that explain the apparent redistribution of phenotypes is estimated. A process similar to computing the Earth Mover's Distance metric [10] is used to measure the total cost of the transformation (Sec. 3.2).

3.1 Nuclear Morphology Representation

Nuclear morphology is modeled using a spherical harmonic (s.h.) representation of its surface [3]. The surface is mapped to a sphere by a distortion–minimizing transformation resulting in the representation of the original surface through the set of coordinate functions $[x(\phi,\theta), y(\phi,\theta), z(\phi,\theta)]$ on the unit sphere. These functions are represented in the s.h. basis as $x(\phi,\theta) \triangleq \sum_{l=0}^{\infty} \sum_{m=-l}^{l} c_{l,m}^{x} Y_{l,m}(\phi,\theta)$ where $Y_{l,m}(\phi,\theta)$ are the s.h. basis functions of degree l and order m and $c_{l,m}^{x}$ is the corresponding coefficient [3]. The functions $y(\phi,\theta)$ and $z(\phi,\theta)$ are represented similarly. The resulting set of s.h. coefficients form the shape vector, with dimensionality $3(L+1)^2$, where L is the maximum degree of the harmonics considered.

Previous *in vivo* studies indicate that changes in morphology have low–dimensional modes of variation [5]. Linear dimensionality reduction methods are well suited to estimate these modes of variation [7]. Motivated by these studies, PCA is subsequently used to estimate a low–dimensional subspace of biologically relevant nuclear shapes. The resulting linear transform is represented as ϕ_D^{ϵ}, where ϵ is the fraction of variance that is discarded, resulting in a k–dimensional embedding.

3.2 Estimating Redistribution of Phenotypes

The differences between the distributions of P and Q represent the differences in the phenotypic profiles of the two populations of cells. This difference is modeled in terms of the putative *redistribution* of phenotypes that transformed P to Q. To estimate this, the optimal "moves" required to transform one population to another is computed under a cost model $d : \mathcal{Z} \times \mathcal{Z} \rightarrow \mathbb{R}$. The choice of d is taken to be the Euclidean distance in ϕ_D^{ϵ} since this subspace captures the principal modes in which cell shapes vary [7]. It is noted here that the above is a formulation of the earth mover's distance [10] and can be solved using linear programming as follows.

Denote $P_{\phi} = \phi_D^{\epsilon} \circ \psi(P)$ where $\psi(P)$ corresponds to the shape representation and ϕ_D^{ϵ} to the low–dimensional embedding. P_{ϕ} and Q_{ϕ} are converted to

empirical normalized histograms in the form $\hat{P}_\phi \triangleq \{(s_1, p_1), (s_1, p_1), \ldots, (s_1, p_n)\}$ and $\hat{Q}_\phi \triangleq \{(t_1, q_1), (t_1, q_1), \ldots, (t_1, q_m)\}$ where the pairs (x, c) indicate that $x \in \tilde{\mathcal{X}}$ is a bin center and c is the normalized frequency of points in the bin. The bin centers are obtained by vector quantization. The optimal solution is found in terms of the "flow" $R \triangleq (r_{ij})$ that minimizes the total cost $\sum_{i=1}^{m}\sum_{j=1}^{n} r_{ij}d_{ij}$ where $d_{ij} \triangleq d(s_i, t_j)$ and R obeys the constraints (i) $r_{ij} \geq 0$, $1 \leq i \leq m$, $1 \leq j \leq n$ (ii) $\sum_{i=1}^{m} r_{ij} \leq p_j$, $1 \leq j \leq n$ (iii) $\sum_{j=1}^{n} r_{ij} \leq q_i$, $1 \leq i \leq m$ (iii) $\sum_{i=1}^{m}\sum_{j=1}^{n} r_{ij} = min(\sum_{i=1}^{m} p_i, \sum_{j=1}^{n} q_j)$.

Given the optimal solution $R^* \triangleq (r^*_{ij})$, the putative redistribution is obtained by *removing* r_{ij} instances of phenotype s_i from $\hat{P}_\phi$ and *adding* an equal number of instances of phenotype t_j to $\hat{Q}_\phi$. Table 2 provides an illustrative example for this interpretation. Thus by the above formulation, we are able to *interpret* the difference between the two populations in terms of the most likely phenotypic transitions that occurred. Further, the difference between the two distributions is measured by the total cost of the transformation given by $D(P_\phi, Q_\phi) \triangleq \sum_{i=1}^{m}\sum_{j=1}^{n} r^*_{ij}d_{ij} / \sum_{i=1}^{m}\sum_{j=1}^{n} r^*_{ij}$.

In experiments with small sample sizes, histogram estimation in the dimensionality of the embedded space may result in unstable estimates. In this event, the following strategy is used to retain the dimensions in which the marginal distributional differences are (statistically) most significant. Denote $\phi_i(\cdot) \triangleq [\phi(\cdot)]_i$. The statistic $D(P_{\phi_i}, Q_{\phi_i})$ is computed for $i \in \{1 \ldots k\}$ and its significance is obtained through permutation testing. The bases are ranked by increasing order of p-value, and a cutoff p-value is used to select the top l features.

4 Results

The method was applied to the knockout study discussed in Section 1 to determine the phenotypic changes resulting from the selective deletion of the *PTEN* gene. The study focused on the morphological changes in fibroblasts, the cell type in which the gene was selectively deleted.

Data Collection and Feature Extraction. Tissue sections were collected from two–month old female mice belonging to control (P) and *PTEN*-deleted (Q) groups, with three mice per genotype group. The sections were stained with DAPI, a fluorescent marker for DNA, to identify cell nuclei. Endogenously expressed cell-specific fluorescence was used to detect fibroblasts. 3D images of the tissue were acquired with a confocal microscope at an in-plane resolution $0.31\mu m$ and axial resolution of $0.5\mu m$. Nuclei were segmented using Otsu thresholding followed by morphological closing to fill holes in the volume. Segmentation errors were manually corrected using a semi-automatic segmentation tool [12]. For each nucleus, s.h. coefficients were computed to the fifth degree resulting in a 108-dimensional feature–vector. In all, a random subset of 125 fibroblast nuclei from each group were selected for the study.

4.1 Statistical Tests

A set of statistical tests were performed on the cell populations to test for differences between the *PTEN*-deleted and control groups as follows. Tests were performed in *each dimension* of (a) the s.h. feature–space $\mathcal{S}_{sh}$ and (b) the principal components feature–space $\mathcal{S}_{pc}$, obtained by reducing dimensionality of $\mathcal{S}_{sh}$ keeping 95% of the total variation. Test for difference between the group means was performed using the Student's t test (ST). The Kolmogorov-Smirnov test (KS) was used to test for difference between the distributions of the groups. Permutation testing was performed over 20,000 iterations. In each iteration, the labels of the data were randomly permuted and the relevant test statistic was computed to generate the null distribution. Results from these tests are reported in Table 1. The columns $p^{(1)}$ through $p^{(5)}$ list the five most significant *uncorrected* (w.r.t. multiple comparisons) p-values in ascending order. The index of the component is shown in parentheses. The p-values that are significant *after* multiple test correction using false discovery rate method are underlined. For both, $\mathcal{S}_{sh}$ and $\mathcal{S}_{pc}$, ST did not report an effect in any of the features, an indicator that the heterogeneity of the cellular population precludes the use of a simple test of difference between means to identify differences. The KS test on the other hand, reported a statistically significant difference in the third PC. None of the components in $\mathcal{S}_{sh}$ reported significance in the KS test due to the effect of multiple comparison correction over 108 comparisons.

Table 1. Permutation testing results

Feat	Test	$p^{(1)}$	$p^{(2)}$	$p^{(3)}$	$p^{(4)}$	$p^{(5)}$
$\mathcal{S}_{sh}$	ST	0.177 (17)	0.209 (27)	0.387 (12)	0.513 (35)	0.551 (31)
$\mathcal{S}_{sh}$	KS	0.011 (12)	0.021 (23)	0.210 (17)	0.319 (31)	0.332 (27)
$\mathcal{S}_{pc}$	ST	0.191 (2)	0.262 (1)	0.291 (3)	0.655 (7)	0.706 (10)
$\mathcal{S}_{pc}$	KS	0.002 (3)	0.050 (4)	0.181 (10)	0.237 (5)	0.286 (7)

4.2 Estimating Redistribution of Phenotypes

While the KS test established the presence of a difference between the distributions, there was no interpretation provided by the test about *how* the phenotype transformation occurred. Such an interpretation is provided by using the method described in Section 4, where the most likely redistribution that explains the differences is modeled as follows.

Estimating Embedding from Data. A collection of 1600 fibroblast nuclei collected across several images in the same animal model were used to obtain the biologically relevant modes of variation using PCA. By using a larger data set, a larger gamut of the biological variability is captured and results in a better characterization of the natural modes of variation. Keeping 95% of the total variation resulted in an 11-dimensional feature-space. The representation of the two populations in this feature space were computed.

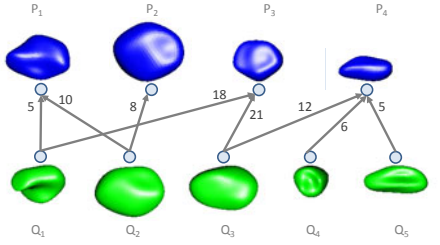

	P_1	P_2	P_3	P_4
Q_1	**4.8(1.2)**	3.1(1.2)	**18.2(1.8)**	0.3(0.9)
Q_2	**10.1(2.1)**	**7.8(1.9)**	0.4(0.1)	0.3(0.1)
Q_3	1.1(0.6)	1.2(0.3)	**21.2(2.4)**	**12.1(1.8)**
Q_4	0.8(0.3)	2.1(0.8)	1.3(0.4)	**5.8(2.9)**
Q_5	0.4(0.1)	1.8(0.4)	0.3(0.1)	**6.9(1.2)**

Table 2. Estimating the redistribution of phenotypes. The nuclei $P_1 ... P_4$ and $Q_1 ... Q_5$ represent the cluster centers within the groups Wild-type and *PTEN*-deleted respectively. Edge weights correspond to total percentage of the optimal flow from one phenotype subgroup to another. Edges with weight less than five are not shown for clarity. The table shows the bootstrap estimates of edges weights demonstrating the stability of the solution (1 s.d. shown in brackets). Dominant flows are highlighted.

Estimating Redistribution. Since the sizes of the datasets in the experiment were small ($n = 125$/group), the dimensionality was further reduced as described in Sec. 3.2 in order to get stable estimates of histograms. The sorted plot of the negative log p-values are shown in the inset. A sharp knee was observed at PC4 and was used as the criterion for selecting PC3 and PC4 for the rest of the analysis. Next, the empirical histograms $\hat{P}_\phi$ and $\hat{Q}_\phi$ in this space — representing the population profiles — were computed by adaptive binning, for which centroids were obtained using using k-means clustering. A total of $n_p = 4$ and $n_q = 5$ centroids were selected for the two groups based on the AIC criterion. The centroids for P and Q are visualized in Table 2 (left). This solution is visualized through the directed graph shown in the figure. In order to establish stability of the result, bootstrap estimates of the edges weights were estimated over 1000 iterations. It is observed that phenotypes P_4, small elongated nuclei were transformed to Q_3, significantly larger nuclei and to Q_4, smaller and rounded ones, while a certain amount of them remained relatively unaltered Q_5. Further, a significant fraction of rounded nuclei P_4 acquired a curved morphology Q_1. This model of redistribution thus provides an interpretation of the phenotypic shifts that occurred, giving additional insights into the nature of the genetic perturbation.

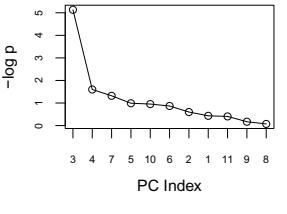

5 Discussion

In this paper, a methodology was presented to quantify differences between two cell populations in terms of their phenotypic profiles and an explanatory model by which the differences can be interpreted was provided. Using an example of perturbations induced by the deletion of a tumor–suppressor gene in the

mouse model, it was first shown that the phenotypic differences between control and perturbed fibroblast cell populations can be characterized in terms of their differences in distribution in a data–derived phenotype space. The proposed methodology was further used to estimate the putative *redistribution* of nuclear phenotypes. This model of redistribution provides a hypothesis about the nature of changes induced by genetic perturbations in the animal system.

The experiment described was conducted in the context of understanding the role of the tumor microenvironment in breast cancer [11]. A characterization of the changes that take place in fibroblasts — cells that play a major role in this microenvironment — can lead to the generation of new hypothesis about cancer progression. While nuclear morphology was used as a proxy for cell phenotype in this study, the proposed methodology can be applied to any phenotypic characterization of a cell, so long as an appropriate cost model d can be estimated for the same. In future studies, we plan to integrate more morphological features including cell image texture and cellular context for further characterization of the effects of genetic perturbation in the tumor microenvironment.

References

1. Andrey, P., et al.: Statistical Analysis of 3D Images Detects Regular Spatial Distributions of Centromeres and Chromocenters in Animal and Plant Nuclei. PLoS Computational Biology 6(7) (July 2010)
2. Boland, M.V., Murphy, R.F.: A neural network classifier capable of recognizing the patterns of all major subcellular structures in fluorescence microscope images of HeLa cells. Bioinformatics 17(12), 1213–1223 (2001)
3. Brechbühler, C.: Parametrization of Closed Surfaces for 3-D Shape Description. Computer Vision and Image Understanding 61(2), 154–170 (1995)
4. Gladilin, E., Goetze, S., Mateos-Langerak, J., Van Driel, R., Eils, R., Rohr, K.: Shape normalization of 3D cell nuclei using elastic spherical mapping. Journal of Microscopy 231(Pt. 1), 105–114 (2008)
5. Keren, K., Pincus, Z., Allen, G.M., Barnhart, E.L., Marriott, G., Mogilner, A., Theriot, J.A.: Mechanism of shape determination in motile cells. Nature 453(7194), 475–480 (2008)
6. Lin, G., Chawla, M.K., Olson, K., Barnes, C.A., Guzowski, J.F., Bjornsson, C., Shain, W., Roysam, B.: A multi-model approach to simultaneous segmentation and classification of heterogeneous populations of cell nuclei in 3D confocal microscope images. Cytometry Part A 71(9), 724–736 (2007)
7. Pincus, Z., Theriot, J.A.: Comparison of quantitative methods for cell-shape analysis. Journal of Microscopy 227(Pt. 2), 140–156 (2007)
8. Rittscher, J.: Characterization of Biological Processes through Automated Image Analysis. Annual Review of Biomedical Engineering 12, 315–344 (2010)
9. Rohde, G.K., Ribeiro, A.J.S., Dahl, K.N., Murphy, R.F.: Deformation-based nuclear morphometry: capturing nuclear shape variation in HeLa cells. Cytometry Part A 73(4), 341–350 (2008)

10. Rubner, Y., Tomasi, C., Guibas, L.J.: The Earth Movers Distance as a Metric for Image Retrieval. International Journal of Computer Vision 40(2), 99–121 (2000)
11. Trimboli, A.J., Fukino, K., de Bruin, A., Wei, G., Shen, L., Tanner, S.M., Creasap, N., Rosol, T.J., Robinson, M.L., Eng, C., Ostrowski, M.C., Leone, G.: Direct evidence for epithelial-mesenchymal transitions in breast cancer. Cancer Research 68(3), 937–945 (2008)
12. Yushkevich, P.A., Piven, J., Hazlett, H.C., Smith, R.G., Ho, S., Gee, J.C., Gerig, G.: User-guided 3D active contour segmentation of anatomical structures: significantly improved efficiency and reliability. NeuroImage 31(3), 1116–1128 (2006)
13. Zink, D., Fischer, A.H., Nickerson, J.A.: Nuclear structure in cancer cells. Nature reviews. Cancer 4(9), 677–687 (2004)

Vertex-Wise Shape Analysis of the Hippocampus: Disentangling Positional Differences from Volume Changes

Hosung Kim[1], Tommaso Mansi[2], Andrea Bernasconi[1], and Neda Bernasconi[1]

[1] Department of Neurology and McConnell Brain Imaging Center, Montreal Neurological Institute and Hospital, McGill University, Montreal, Quebec, Canada
[2] Siemens Corporate Research, Image Analytics and Informatics, Princeton, NJ, USA
{khs001,neda,andrea}@bic.mni.mcgill.ca,
tommaso.mansi@siemens.com

Abstract. Hippocampal atrophy and developmental positional variants may co-occur in various neurological disorders. We propose a surface-based framework to analyze independently volume and positioning. After extracting the spherical harmonics combined with point distribution models (SPHARM-PDM) from manual labels, we computed displacement vectors between individual surfaces and the template. Then, we computed surface-based Jacobian determinants (SJD) from these vectors to localize volume changes. To analyze positional variants, we constructed a mean meridian axis (MEMAX), inheriting the shape-constrained point correspondences of SPHARM, on which we compute local curvatures and position vectors. We validated our method on synthetic shapes, and a large database of healthy subjects and patients with temporal lobe epilepsy. Our comprehensive analysis showed that in patients atrophy and positional changes co-occurred at the level of the posterior hippocampus. Indeed, in this region, while SPHARM-PDM showed mirrored deformations, SJD detected atrophy, and shape analysis of MEMAX unveiled medial positioning due to bending.

Keywords: surface, shape analysis, hippocampus, medial axis, Jacobian determinant, epilepsy.

1 Introduction

The most frequent drug-resistant epilepsy is temporal lobe epilepsy (TLE) related to hippocampal sclerosis, which generally appears as atrophy on MRI [1]. Detecting hippocampal atrophy is clinically relevant as patients have about 70% chance of becoming seizure free after surgery [2].

As atrophy reflects a single morphological characteristic, it may be insufficient to fully describe pathology. Indeed, in addition to atrophy, developmental anomalies, secondary to incomplete folding of the fetal hippocampus, may increase the susceptibility to TLE [3]. Such anomalies have been reported in other brain disorders in which the pathogenic mechanisms involve the temporal lobe, such as schizophrenia [4] and autism [5]. On MRI, hippocampal developmental anomalies appear as atypical

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 352–359, 2011.

shape and positioning of the hippocampus [3]. The complexity of the brain anatomy may hamper their visual identification, particularly when they co-occur with atrophy.

Advanced surface-based shape models, including deformation-based mapping and spherical harmonics description combined with point distribution models (SPHARM-PDM), have allowed localization of morphological changes that may not be readily visible [6], and successfully identified hippocampal pathology in TLE [7] and schizophrenia [8]. The displacement metric provided by these approaches, however, does not allow differentiating volume changes from positional differences. Thus, a biologically meaningful interpretation of findings may be difficult, particularly when both morphological characteristics co-exist. Moreover, the concomitant presence of atrophy and positional changes, by increasing inter-subject shape variability, may reduce statistical sensitivity to detect subtle pathology.

A surface-based Jacobian metric has been proposed as a robust technique to assess volume changes [9]. Shape models based on the medial axis (*i.e.*, the central path of a 3D object), on the other hand, allow quantifying local positional variations [8, 10]. Nevertheless, as these methods rely on coarse-scale sampling, they may lack sensitivity to assess high-dimensional features.

We propose a surface-based framework to analyze independently volume and positioning. We measured surface-based Jacobian determinants as in [9]. We construct a medial axis model that inherits the shape-constrained point correspondence of SPHARM-PDM, and quantifies fine-scale local positional vectors and curvature. We validated sensitivity and specificity using synthetic shapes, and assessed the pathological variations of the hippocampus by applying our method to a large database of healthy subjects (n=46) and TLE patients (n=78).

2 Methods

2.1 Spherical Harmonic Description and Point Distribution Model (SPHARM- PDM)

The hippocampal meshes are first extracted from manual labels and are parameterized using SPHARM-PDM, an area-preserving and distortion-minimizing spherical mapping that ensures point-wise correspondence through an icosahedron subdivision [6]. A template is then constructed by averaging all the individual SPHARM-PDM surfaces after rigid-body alignment. Individual surfaces are again rigidly aligned to this template to avoid any bias. Finally, vertex-wise displacement vectors are calculated between each individual and the template. The surface normal components of these vectors are used to quantify inward and outward deformations.

2.2 Surface-Based Analysis of Volume Changes

Volume changes are quantified on each vertex of the SPHARM-PDM surface as described in [9]. Briefly, we apply the heat equation to interpolate the vertex-wise displacement vectors within the volume domain enclosed by the surface boundary. The Jacobian determinant is then calculated from this interpolated vector field and projected back onto the SPHARM-PDM surface. Finally, after substracting 1 from the Jacobian determinant, growth ($J>0$) or shrinkage ($J<0$) of a unit-size cube is quantified along the surface-normal direction.

2.3 Construction of the Medial Axis

Conventional approaches create the central path of a 3D object by joining points that are equidistant from the two closest facing points on the object's boundary [8, 10-11]. To provide inter-subject point correspondence, all individual skeletons are deformed into a common spatial frame by a thin plate registration [11]. Registration errors and smoothing in the resampling step may introduce biases in the subsequent analysis steps. In our approach, using the latitudes and longitudes parameters of SPHARM-PDM, we parameterize the prime meridians of the object's two hemispheres (**Fig 1A, B**). Then we create a central route by performing pair-wise averaging of the prime meridians. This allows the route to inherit the shape-constrained correspondence of the SPHARM, thus minimizing biases, mainly those related to registration.

2.3.1 Geometric Parameters of SPHARM

The spherical harmonic basis functions Y_l^m of degree l and order $-l \le m \le l$ are defined by the spherical coordinates θ, ϕ as:

$$Y_l^m(\theta,\phi) = \sqrt{\frac{2l+1(l-m)!}{4\pi(l+m)!}} P_m^l(\cos\theta) e^{im\phi} \qquad Y_l^{-m}(\theta,\phi) = (-1)^m Y_l^{m*}(\theta,\phi) \tag{1}$$

In **(1)**, Y_l^{m*} denotes the complex conjugate of Y_l^m and P_l^m the associated Legendre polynomials [6]. Summation of the spherical harmonics from 1 to L degree approximates a given structural boundary to a surface. A vertex point $\vec{s}$ on **S** is then uniquely determined by:

$$\vec{s}(x,y,z) = \vec{s}(x(\theta,\phi), y(\theta,\phi), z(\theta,\phi)) = \sum_{l=1}^{L} \sum_{m=-l}^{l} \vec{c}_l^m Y_l^m(\theta,\phi) \tag{2}$$

The 3D weight vectors $\vec{c}_l^m$ are computed through least-square minimization. θ, ϕ define latitudes and longitudes, respectively (yellow and green in **Fig. 1A, B**). Accordingly, $\vec{s}(0, \pi/2)$ and $\vec{s}(\pi, \pi/2)$ indicate the poles, and $\vec{s}(\pi/2, \phi)$ the equator [6].

2.3.2 Mean Meridian Axis (MEMAX)

The functions $\vec{s}(\theta,0)$ and $\vec{s}(\theta,\pi)$ in **Eq. (2)**, describe a set of points that constitute the prime meridian (blue in **Fig. 1A, B**). The prime meridian points on a given hemisphere **s**= $\vec{s}(0 \le a \le \pi/2, 0 \; or \; \pi)$ and their counterparts on its opposite side $\hat{\mathbf{s}}$=$\vec{s}(\pi/2 \le \pi - a \le \pi, 0 \; or \; \pi)$ are averaged in a pair-wise fashion, thus yielding a mean meridian axis **M**, henceforth called MEMAX (red in **Fig. 1C**) as the skeleton of a shape.

For shape analysis, we sample K points on **M**. We define **M**(k), where k={k | 1, 2, .., K}, using equidistant subdivisions of θ, that is:

$$\begin{aligned} \mathbf{M}(\mathrm{k}) &= \left\{ \vec{s}\left(\frac{\pi}{2} - \frac{\pi}{2} \times \frac{2k}{K}, 0\right) + \vec{s}\left(\frac{\pi}{2} + \frac{\pi}{2} \times \frac{2k}{K}, 0\right) \right\} / 2, \quad \text{if } 1 \le k \le K/2 \\ &= \left\{ \vec{s}\left(\frac{\pi}{2} \times \frac{2(k-K/2)}{K}, \pi\right) + \vec{s}\left(\frac{\pi}{2} + \frac{\pi}{2} \times \frac{2(k-K/2)}{K}, \pi\right) \right\} / 2, \quad \text{if } K/2 \le n \le K \end{aligned} \tag{3}$$

2.4 Assessment of Local Positioning

To assess shifting, group analysis is performed by comparing the location of each sample point on **M** using Hotelling T^2 metric [12] as:

$$T^2 = (\mu_1 - \mu_2)'(\sum{}_1 \frac{1}{n_1} + \sum{}_2 \frac{1}{n_2})^{-1}(\mu_1 - \mu_2) \quad (4)$$

In this equation, μ_1, μ_2 are the group mean vectors, $\sum_1$, $\sum_2$ the group-separated covariance matrices and n_1, n_2 the number of samples in groups 1 and 2.

The curvature at a given point (x,y,z) of a 3D curve is generally computed using the Frenet-Serret formula. This measurement, however, provides only the degree of curving without distinguishing convexity from concavity, thus requiring additional visual inspection to fully understand its nature. Instead, we propose to measure the signed curvature varying along the normal directions of a surface defined by the prime meridian or the equator. We create an open surface **C** where its boundaries are defined by the prime meridian **s**(k) and $\hat{\mathbf{s}}$(k). Since the MEMAX **M**(k) is a geometric mean of the two prime meridians, **M**(k) is placed in the middle **s**(k) and $\hat{\mathbf{s}}$(k) on the surface **C** (**Fig. 1D** left). We then calculate the Gaussian curvature κ_{Gauss} at each sample point **M**(k) on **C** (**Fig. 1D** right). To differentiate the convexity (+) from the concavity (-), signs are computed using the normal vector of the k_{th} facet on **C** $\mathbf{n}(\mathbf{C}_k)=(\mathbf{M}(k+1)-\mathbf{M}(k)) \times (\mathbf{s}(k) - \hat{\mathbf{s}}(k))$, as:

$$Sign(\mathbf{n}(\mathbf{C}_k)) = Sign\left(\mathbf{n}(\mathbf{C}_k) \bullet \mathbf{n}(\mathbf{C}_{k+1}) / \left\|\mathbf{n}(\mathbf{C}_k)\right\| \cdot \left\|\mathbf{n}(\mathbf{C}_{k+1})\right\|\right) \quad (5)$$

The signed curvature is finally determined as:

$$K = Sign(\mathbf{n}(\mathbf{C}_k)) \, |\kappa_{Gauss}| \quad (6)$$

In the same fashion, we create the surface **C**' on which boundaries are defined by the equator and place **M**(k) on **C**'. Then, we compute curvature K' to measure changes along the normal to **C**'.

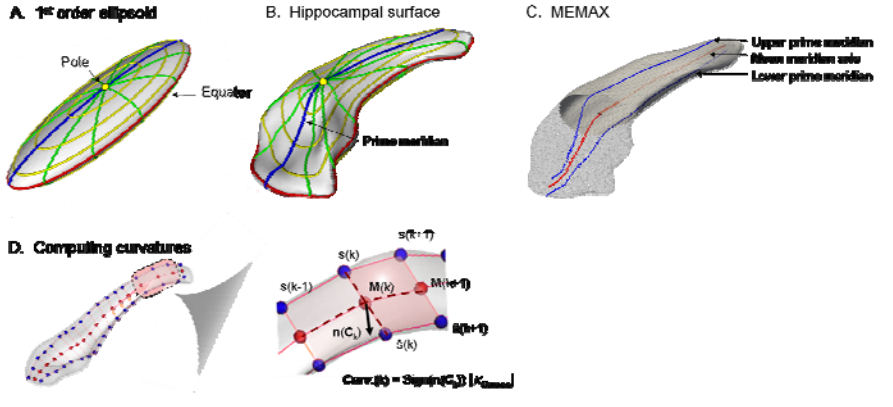

Fig. 1. Modeling of the mean meridian axis (MEMAX; A-C) and computing point-wise curvature (D). See text for details.

3 Experiments and Results

3.1 Synthetic Data

We created a rounded cylindrical volume with a resolution of 0.25 x 0.25 x 0.25 mm^3 and applied (**Fig. 2A-C**): *i)* bending, *ii)* local shifting (4.0 mm), and *iii)* simultaneous local shrinking (0.5 mm) and shifting (4.0 mm). We extracted the SPHARM-PDM surfaces and calculated the displacement vectors. Then, at each vertex, we computed the surface-normal component of the displacement vector (SNV) and the surface-based Jacobian determinant (SJD) as a metric of local volume change. The MEMAX were extracted from all shapes. Finally, we quantified point-wise positioning of the variants compared to the reference.

The SNV showed mirrored inward/outward deformations on facing surfaces reflecting bending/shifting (**Fig. 2D**, top and middle). In case of simultaneous local shrinking and shifting, SNV did not enable differentiating shrinkage from shifting (**Fig. 2D**, bottom). Contrary to SNV, SJD was not influenced by bending/shifting (**Fig. 2E**, top and middle) and accurately localized the 0.5 mm local shrinkage (**Fig. 2E**, bottom). Analysis of MEMAX detected bending and local shifting only (**Fig. 2F**), and was robust to shrinking (**Fig. 2F**, middle and bottom)**.**

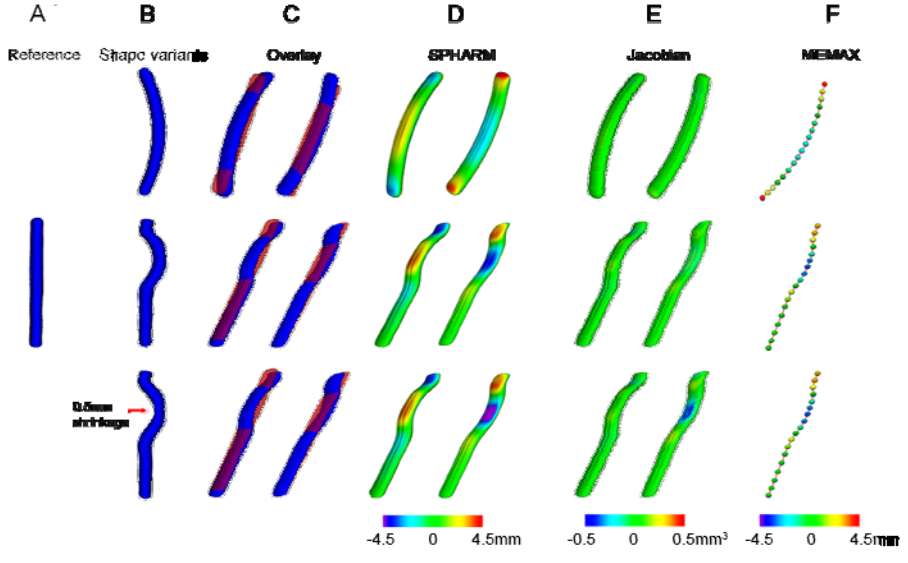

Fig. 2. Analysis of synthetic data. The -/+ signs indicate the direction of displacement (in **D**), volume (in **E**), and shifting (in **F**) with respect to the reference.

3.2 Shape Analysis of the Hippocampus in Temporal Lobe Epilepsy

3.2.1 Subjects and Image Pre-processing

We studied 78 drug-resistant TLE patients (age: 36±10). Diagnosis and lateralization of the focus into left (LTLE; n=34) and right (RTLE; n=44) were determined by

video-EEG and MRI evaluation. The control group consisted of 46 age- and sex-matched (age: 32±12) healthy subjects. Our Ethics Committee approved this study.

High-resolution 1mm-isotropic 1.5T MRIs were linearly registered to a stereotaxic space. An expert manually segmented the hippocampi in all subjects. Using z-score normalization based on healthy controls' distribution, all patients showed hippocampal atrophy (i.e. z<-2) ipsilateral to the seizure focus.

3.2.2 Optimization of Sampling Scale Along MEMAX

We first extracted the MEMAX of the right hippocampi in controls and patients with RTLE. We then evaluated the effect of varying the total number of sample points (Ns = 2-36) on the sensitivity to detect positional changes (*i.e.*, position vectors and Gaussian curvatures). For each individual and each metric, we averaged all sample points. We then performed group-wise Hotelling T^2 test on the mean position vectors and Student's t-test on the mean curvatures. Results are shown in **Fig. 3**.

Comparing groups, we detected shifting in TLE with Ns ranging from 14 to 34 (T^2>3.6) and bending with Ns ranging from 14 to 22 (t>2.82). The mean inter-point distance spanned from 2.1 to 3.3 mm. The sampling along MEMAX that provided the most significant group differences was Ns = 18 (T^2=3.71; t=2.85). Thus, for the subsequent analyses, we used 18 points, with a mean inter-point distance of 2.7 mm.

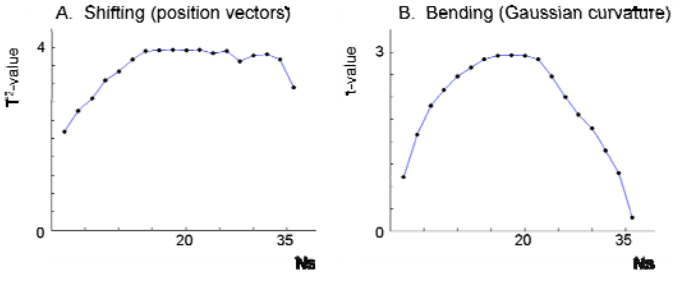

Fig. 3. Relationship between sampling points (Ns) along MEMAX and sensitivity in detecting shifting (**A**) and bending (**B**) in TLE patients

3.2.3 Clinical Analysis of Local Hippocampal Shape Changes in TLE

We performed vertex-wises group comparisons of surface-normal component of the displacement vector (SNV) and the surface-based Jacobian determinant (SJD), and point-wise comparisons of the Gaussian curvatures and positional vectors. We used the false discovery rate (FDR) to correct for multiple comparisons [13]. To localize changes, we schematically outlined hippocampal subfields on the surface template according to histological parcellations [7].

Results are presented in **Fig. 4 A** and **B**. Ipsilateral to the seizure focus, both SNV and SJD revealed areas of inward deformation and atrophy, respectively, in the hippocampal head and body (FDR<0.001). In the tail, where SNV detected *mirrored* inward/outward deformations, SJD on the other hand unmasked atrophy. In addition, SJD detected small areas of CA1 inward deformations (FDR<0.05) in the contralateral hippocampus.

Curvature analysis of the mean meridian axis (MEMAX) detected ipsilateral bending towards the mid-sagittal plane (LTLE: 1 sample point, FDR=0.04; RTLE: 2

sample points, FDR<0.05; **Fig. 4D**), at the transition between hippocampal body and tail. We also found ipsilateral medial shifting at the level of the tail (LTLE: 3 points, FDR<0.01; RTLE: 6 points, FDR<0.05; **Fig. 4E**), which overlapped with the region of *mirrored* inward/outward deformations detected by SNV (**Fig. 4A**).

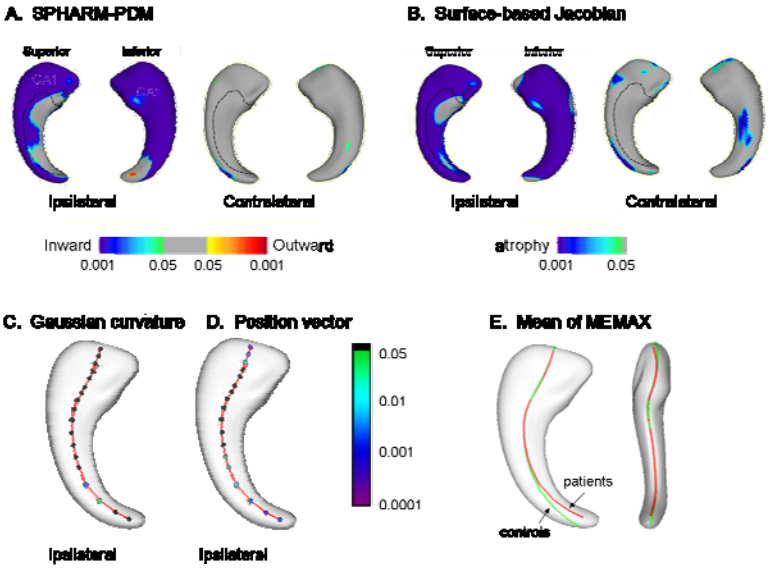

Fig. 4. Comparisons between TLE and controls for SPHARM-PDM (**A**), Surface-based Jacobian (**B**), Gaussian curvature (**C**), position vectors (**D**). The geometric means of MEMAX (**E**) is shown to ease the interpretation of C/ D. Color-scales show FDR-corrected p-values.

4 Discussion and Conclusion

We previously proposed [9] to quantify local volume changes by computing SJD. In case of co-occurrence of volume and positional anomalies this method failed to quantify the latter. Here we unambiguously disentangled volume from positional changes by quantifying independently both morphological characteristics, *i.e.* by calculating SJD and Gaussian curvatures on the mean meridian axis (MEMAX).

In agreement with histology [14], SJD detected atrophy in TLE encompassing the CA1 subfield. Moreover, our analysis showed that atrophy and positional changes co-occurred in the hippocampal tail. In this region, while SPHARM-PDM showed mirrored deformations, SJD detected atrophy, and shape analysis of MEMAX unveiled medial positioning due to bending.

In conventional medial models, a small number of sample points (*e.g.*, 8) are used mainly to reduce registration errors [8, 10-11]. This relatively coarse sampling may reduce sensitivity to subtle changes. By using geometric parameters of SPHARM [6], MEMAX allowed shape-constrained correspondence without a registration step. We found that modulating the number of sample points influenced the sensitivity to detect positional changes, with the largest changes at a finer scale of 18 points.

A single longitudinal central path would not allow measuring thickness nor retrieve the original shapes of an object. Our model was specifically designed to quantify positional changes. To achieve this, conventional medial representations would require additional processing steps, such as pruning [10-11], to eliminate undesirable secondary branches producing irregular skeleton topologies.

Medial positioning is a feature of hippocampal developmental abnormalities [3] that may predispose to TLE. Disentangling atrophy from positioning anomalies may provide new insights in the pathogenesis of a variety of other brain disorders, such as schizophrenia and autism, in which these morphometric characteristics coexist.

References

1. Berkovic, S.F., Andermann, F., Olivier, A., Ethier, R., Melanson, D., Robitaille, Y., Kuzniecky, R., Peters, T., Feindel, W.: Hippocampal sclerosis in temporal lobe epilepsy demonstrated by magnetic resonance imaging. Ann. Neurol. 29(2), 175–182 (1991)
2. Schramm, J., Clusmann, H.: The surgery of epilepsy. Neurosurgery 62(suppl. 2), 463–481 (2008); discussion 481
3. Bernasconi, N., Kinay, D., Andermann, F., Antel, S., Bernasconi, A.: Analysis of shape and positioning of the hippocampal formation: An mri study in patients with partial epilepsy and healthy controls. Brain 128(Pt. 10), 2442–2452 (2005)
4. Connor, S.J., Ng, V., McDonald, C., Schulze, K., Morgan, K., Dazzan, P., Murray, R.M.: A study of hippocampal shape anomaly in schizophrenia and in families multiply affected by schizophrenia or bipolar disorder. Neuroradiology 46(7), 523–534 (2004)
5. Salmond, C.H., Ashburner, J., Connelly, A., Friston, K.J., Gadian, D.G., Vargha-Khadem, F.: The role of the medial temporal lobe in autistic spectrum disorders. European Journal of Neuroscience 22(3), 764–772 (2005)
6. Styner, M., Oguz, I., Xu, S., Brechbuhler, C., Pantazis, D., Gerig, G.: Statistical shape analysis of brain structures using spharm-pdm. In: MICCAI 2006 Opensource Workshop (2006)
7. Hogan, R.E., Wang, L., Bertrand, M.E., Willmore, L.J., Bucholz, R.D., Nassif, A.S., Csernansky, J.G.: MRI-based high-dimensional hippocampal mapping in mesial temporal lobe epilepsy. Brain 127(8), 1731–1740 (2004)
8. Styner, M., Pantazis, D., Gerig, G.: Boundary and medial shape analysis of the hippocampus in schizophrenia. Med. Image Anal. 8(3), 197–203 (2004)
9. Kim, H., Besson, P., Colliot, O., Bernasconi, A., Bernasconi, N.: Surface-based vector analysis using heat equation interpolation: A new approach to quantify local hippocampal volume changes. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part I. LNCS, vol. 5241, pp. 1008–1015. Springer, Heidelberg (2008)
10. Joshi, S., Pizer, S., Fletcher, P.T., Yushkevich, P., Thall, A., Marron, J.S.: Multiscale deformable model segmentation and statistical shape analysis using medial descriptions. IEEE Trans. Med. Imaging 21(5), 538–550 (2002)
11. Styner, M., Gerig, G., Lieberman, J., Jones, D., Weinberger, D.: Statistical shape analysis of neuroanatomical structures based on medial models. Medical Image Analysis 7(3), 207–220 (2003)
12. Seber, G.F.: Multivariate observations, vol. 686. John Wiley & Sons Inc., Chichester (1984)
13. Benjamini, Y., Hochberg, Y.: Controlling the false discovery rate. J. Royal Stat. Soc. 57(1), 289–300 (1995)
14. Blumcke, I., Thom, M., Wiestler, O.D.: Ammon's horn sclerosis: A maldevelopmental disorder associated with temporal lobe epilepsy. Brain Pathol. 12(2), 199–211 (2002)

Localized Component Analysis for Arthritis Detection in the Trapeziometacarpal Joint

Martijn van de Giessen[1,2,5], Sepp de Raedt[1,⋆], Maiken Stilling[3], Torben B. Hansen[3], Mario Maas[4], Geert J. Streekstra[5], Lucas J. van Vliet[1], and Frans M. Vos[1,4]

[1] Quantititative Imaging Group, Delft University of Technology, The Netherlands
[2] Division of Image Processing, Leiden University Medical Center, The Netherlands
[3] Dept. of Orthopaedics, Holstebro Regional Hospital, Denmark
[4] Dept. of Radiology, AMC Amsterdam, The Netherlands
[5] Dept. of Biomed. Engineering and Physics, AMC Amsterdam, The Netherlands

Abstract. The trapeziometacarpal joint enables the prehensile function of the thumb. Unfortunately, this joint is vulnerable to osteoarthritis (OA) that typically affects the local shape of the trapezium. A novel, local statistical shape model is defined that employs a differentiable locality measure based on the weighted variance of point coordinates per mode. The simplicity of the function and the smooth derivative enable to quickly determine localized components for densely sampled surfaces.

The method is employed to assess a set of 60 trapezia (38 healthy, 22 with OA). The localized components predominantly model regions affected by OA, contrary to shape variations found with PCA. Furthermore, identification of pathological trapezia based on the localized modes of variation is improved compared to PCA.

1 Introduction

The prehensile biomechanical function of the thumb is determined for a large part by the complex saddle-shaped trapeziometacarpal (TMC) joint (Fig. 1a). Unfortunately, the unique motion of the TMC joint renders trapezium and metacarpal 1 vulnerable to impairment by osteoarthritis (OA), e.g. [7]. Depending on the progression of the disease, the symptoms range from discomfort to the complete inability to manipulate objects. Initially, OA manifests as local, destructive shape changes of the adjacent joint surfaces [4] that result in joint space narrowings and affect motion patterns. Such shape changes are particularly difficult to diagnose using conventional X-ray images. We hypothesize that a model of known shape variations can aid in detecting pathological changes.

A common approach to describe shape variations is to use active shape models (ASM) [3]. An ASM models coordinate distributions of surface points on example shapes by a mean shape and orthogonal modes of variation using principal component analysis (PCA). However, PCA typically describes global shape

⋆ M. van de Giessen and S. de Raedt contributed equally to this work.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 360–367, 2011.

variations. It has been observed, mainly in segmentation applications, that these global descriptions are not always appropriate to model local variations in shape, e.g. [8]. A well-known alternative to PCA is independent component analyis (ICA). Although ICA tends to describe shape variations in a more localized way [10], the ICA components are non-orthogonal. Since a shape representation using non-orthogonal components is not unique, this complicates a statistical comparison between healthy and pathological shapes.

Two methods that do preserve the orthogonality between components are VARIMAX [6] and sparse component analysis (SPCA) as proposed in [2]. Both techniques aim to minimize the number of point coordinates that are affected by each mode. However, there is not an explicit optimization for spatially coherent variations. To improve on this aspect, Alcantara et al. [1] proposed a method for localized component analysis (LoCA) in hippocampi. This method adapts the cost function from [2] with a term that penalizes concurrent movement of distant surface points. A limitation, though, is that the proposed localization term in [1] does not have a continuous derivative. This makes the method difficult to optimize using gradient descent methods and thereby sensitive to local minima.

In this work we propose a new method for localized component analysis with a penalty function for non-locality that has continuous 1st and higher order derivatives. Using this function a localized statistical model is defined for a set of 38 healthy and 22 pathological trapezia. The local shape variations are used in the design of a classifier for identification of pathological trapezia.

2 Materials and Methods

The method proposed in this work uses a point distribution model (PDM) as used in ASMs [3]. However, instead of using the principal components of variation, determined by PCA, the components are optimized to explicitly model local variations. Therefore a novel, differentiable locality measure is derived.

2.1 Active Shape Models

A PDM assumes a set of example shapes, represented by points for which the correspondence relations to a point on any of the other shapes have been established. It comprises a linear model that describes how the surface points $\mathbf{x}_i$ of shape S_i deviate from the mean $\bar{\mathbf{x}}$ of a set of example shapes:

$$\mathbf{x}_i = \bar{\mathbf{x}} + P\mathbf{b}_i \tag{1}$$

where $\mathbf{x}_i = \begin{bmatrix} x_1\ y_1\ z_1 \cdots x_N\ y_N\ z_N \end{bmatrix}^T$ contains the concatenated coordinates of N points on surface S_i, $\bar{\mathbf{x}}$ describes the mean shape, $P = \begin{bmatrix} \mathbf{p}_1\ \mathbf{p}_2 \cdots \mathbf{p}_M \end{bmatrix}$ contains the M modes of variation as columns $\mathbf{p}_j, j = 1, \ldots, M$ and $\mathbf{b}_i$ is a vector with weightings for each mode. When the columns of P are orthogonal, $\mathbf{b}_i$ uniquely describes surface S_i.

Thus, a set of surfaces is described with the mean coordinate vector $\bar{\mathbf{x}}$ and a matrix B where each column $\mathbf{b}_i$ describes a different example surface. The

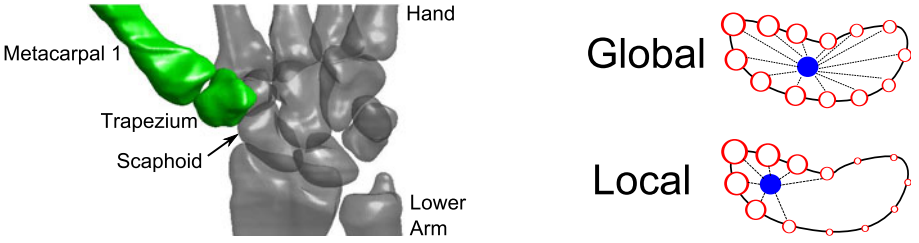

Fig. 1. (a) The TMC joint in a left-handed wrist. (b) Schematic overview of the weighted variance as a measure for locality. The radius of the red circles indicates the magnitude of the variation of the surfaces points in a mode. The blue dots are the weighted means $\mathbf{m}_j$. The sum of the lengths of the dashed lines squared, weighted with the squared magnitudes of the variations is the localization measure in (6).

variance per mode σ_j^2 (for mode j) is a measure for the 'importance' of a mode and is estimated as the variance of each row of B. It is customary to compute the relative variance per mode as $v_j = \sigma_j^2 / \sum_{k=1}^{M} \sigma_k^2$. PCA is often implemented using an eigenvalue or singular value decomposition, but the same P can also be obtained by rotating the components and thereby minimizing [2]:

$$C_{\text{pca}}(P) = \sum_{j=1}^{M} -v_j \log(v_j) \tag{2}$$

A common strategy for such a minimization is to pairwise rotate components in the plane spanned by them until C_{pca} is minimal [2,1]. This procedure ensures that the rotating components remain orthogonal to all other components. The formulation (2) has the advantage that it can be extended to sparse or local component analysis, where a trade-off between compactness of the model (PCA) and other criteria is needed such as sparsity or locality.

2.2 Localized Component Analysis

For a localized description of shape variation within a set the following cost function is minimized (similar to [2,1]):

$$C(P) = (1-\lambda) C_{\text{pca}}(P) + \lambda C_{\text{loca}}(P) \tag{3}$$

where λ is a parameter to balance between compactness (C_{pca}) and locality (C_{loca}). The latter will be derived in this section.

Each column $\mathbf{p}_j$ of P is normalized ($\|\mathbf{p}_j\| = 1$) and contains the relative displacements of all surface points compared to $\bar{\mathbf{x}}$. The relative displacement per surface point (in 3D) is thus given by each triplet of elements of vector $\mathbf{p}_j$:

$$\mathbf{p}_j = [\mathbf{d}_{j,1}; \cdots; \mathbf{d}_{j,N}] \tag{4}$$

Here $\mathbf{d}_{j,k}$ is the 3-element vector with the displacement of surface point k. The relative displacement amplitudes of all points are thus given by $d_{j,k} = \|\mathbf{d}_{j,k}\|$. The newly derived locality measure in this work uses these displacement measures as weights in a variance analysis of the coordinates of the affected points in the following manner.

Intuitively, one might observe that points with large relative displacements should be close together in a description of local variations (Fig. (1b)). In other words, the distribution of points that have a high value of $d_{j,k}$ should be compact. Therefore, a measure for the compactness of large variations is the variance of the coordinates weighted by $d_{j,k}$. The most likely estimate of the set of varying shapes is given by the mean shape and therefore the weighted variance is computed on the mean cloud $\bar{\mathbf{x}}$. By denoting the coordinate vector for each point in $\bar{\mathbf{x}}$ as $\mathbf{z}_k = \begin{bmatrix} x_k & y_k & z_k \end{bmatrix}^T$, the weighted variance vector for the j-th mode is given by

$$\mathbf{s}_j = \sum_{k=1}^{n} d_{j,k}^2 \left(\mathbf{z}_k - \mathbf{m}_j\right)^2 \text{ where } \mathbf{m}_j = \sum_{k=1}^{N} d_{j,k}^2 \mathbf{z}_k \tag{5}$$

is the weighted mean point. The weighted mean $\mathbf{m}_i$ can be interpreted as the point around which the displacement concentrates. Note that this point does not have to be positioned on the shape surface. The cost function C_{loca} is then composed from the weighted variance vectors as

$$C_{\text{loca}}(P) = \sum_{j=1}^{M} \sum_{l=1}^{d} s_{j,l} \tag{6}$$

where d is the number of dimensions and $s_{j,l}$ is the l-th element of $\mathbf{s}_j$. We dub this method Localized Variance Component Analysis (LVCA).

To obtain localized components, $C(P)$ is minimized using sequential pairwise rotations of components as in [2]. This split in pairwise rotations is allowed because both $C_{\text{pca}}(P)$ and $C_{\text{loca}}(P)$ are summations in which the terms that depend on non-rotated components remain constant. A crucial aspect of our method is that (5) is by definition a smooth function that has a continuous derivative. Therefore (3) can be minimized efficiently with a gradient-descent minimization. We employ the Boyden-Fletcher-Goldfarb-Shanno algorithm.

3 Experiments

The advantages of a localized statistical shape model over a PCA model are demonstrated in two datasets. Initially, the effect of the parameter λ is studied in an artificial set of facial contours. Subsequently, the shape changes between healthy trapezia and trapezia with OA are assessed and used for classification.

3.1 Data

Faces. The 'faces dataset' was obtained from the IMM Face Database [9] that consists of 240 images of faces annotated with 58 points each. The focus was on the 120 full frontal faces (the others presenting lateral views). Translations, rotations and scale differences were removed with a Procrustes analysis.

Trapezia. Two datasets of CT scans were separately acquired: a set of healthy trapezia and a set with trapezia affected by osteoarthritis.

The normal dataset contains 38 CT scans with a mean age of 51 (21-70) years. All CT scans were obtained as a result of bilateral scanning over a period of two years at the Academic Medical Center, The Netherlands. Individuals in this set had no complaints (pain, dysfunction) regarding included wrists and all wrist bones were fully developed (epiphyseal lines closed). An experienced musculoskeletal radiologist (MM) deemed all included wrists as healthy and excluded any wrists showing signs of: fractures, arthritis, fragments and fusions.

The OA group contained 22 CT scans from patients with a mean age of 58 (40-72) years. The CT scans were consecutively acquired over a period of one year from Holstebro Regional Hospital and Herning Regional Hospital, Denmark. The dataset contained preoperative scans from patients who received a TMC joint prosthesis. The patients had all been referred to the department of hand surgery due to thumb pain and dysfunction. They were all graded for TMC OA using the Eaton Glickel Classication system by an experienced hand surgeon (TBH) using the anteroposterior (AP) radiographs [4]. The group was comprised of 6 stage II and 18 stage III patients with TMC OA in addition to a single patient with stage I and a single patient with stage IV TMC OA.

All scans were resampled to have an isotropic voxel size of $0.3 \times 0.3 \times 0.3$ mm^3. The voxel sizes in the original scans slightly varied, but were almost isotropic with voxel volume differences in the order of 10%. The trapezia in both sets were segmented using a level set based method, with the same parameter settings as in [5]. 10^4 points per bone surface, corresponding between the bones, were established using non-rigid registration similar to [5].

3.2 Faces

The modes of variations determined by the proposed localized component analysis method (LVCA) were compared to those found by PCA and the method proposed by Alcantara [1] (LoCA). Alcantara's method needs a user-defined function that describes the locality. Therefore, we chose a non-normalized Gaussian: $\exp\left(-q^2/2\sigma^2\right)$, with $\sigma = 0.02$, where q is the distance between points and the face coordinates are between 0 and 1 in horizontal and vertical direction. For both LoCA and LVCA, the influence of the parameter λ was studied.

The 10 modes of variation with the largest variance per mode are shown in Fig. 2. The PCA model (effectively $\lambda = 0$ in (3)) clearly distinguishes itself from the other models by the global variations. For the other extreme case ($\lambda = 1$), both LoCA and LVCA render variations that are too sparse, which is observable as several modes only perturb a single point. Clearly, such extremely localized components are not very informative. This behavior arises because the dataset contains more shapes (120) than point coordinates ($2 \times 58 = 116$). Reversely, if there are less example shapes than point coordinates, then by definition any mode of variation must affect more than one point. Setting $\lambda < 1$ allows for a trade-off between local and principal components. This can be seen in the rows

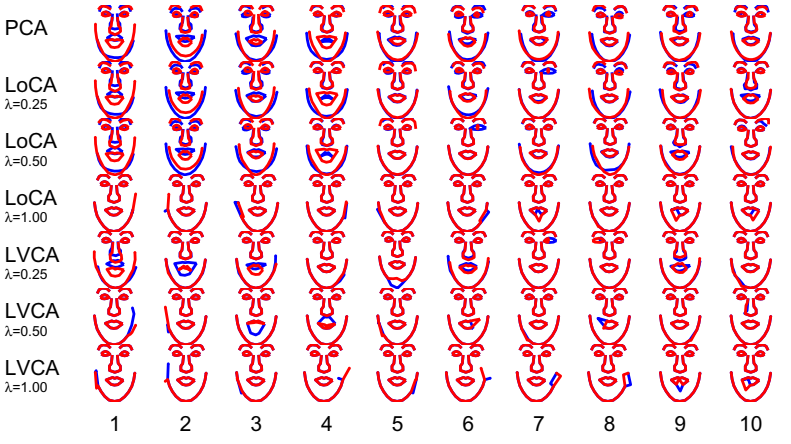

Fig. 2. The 10 modes of variation with the highest variance in the IMM faces data for PCA, localized components as in [1] (LoCA) and as proposed (LVCA), both for $\lambda = 1, 0.5, 0.25$. For each mode the deviations from the mean shape at +/- 3 standard deviations are shown.

with $\lambda = 0.25, 0.5$ for both LoCA and LVCA of Fig. 2, which show variations that are locally concentrated (e.g. around the mouth, nose or eyes), but leave structures further away unaffected. However, LoCA easily gets stuck in local minima, which is demonstrated in modes 1 to 4 for $\lambda = 0.25, 0.5$, which are (almost) the same as the corresponding PCA modes. Computing components was approximately 130 times faster for LVCA than for LoCA.

3.3 Trapezia

The trapezia experiment compared PCA and the proposed LVCA to identify shape changes between healthy and pathological shapes. Clearly, the number of point coordinates (3×10^4) is vastly larger than the number of example trapezia (60). To obtain the most localized components possible in this set, λ in (3) was set to 1. The coefficients of the modes of variations for the PCA and LVCA model are denoted as $\mathbf{b}_i^{\text{PCA}}$ and $\mathbf{b}_i^{\text{LVCA}}$, respectively. Using the coefficients with the largest discrepancy between healthy and pathological bones, a linear logistic classifier is trained. The number of features used is optimized using forward feature selection. The performance of feature selection and classifier training were tested by a 50 times repeated cross validation.

Fig. 3a demonstrates that for PCA the coefficients of the healthy and pathological trapezia lie (in general) closer together than for LVCA. This is to be expected, as the changes of the bone surface due to osteoarthritis are mainly local. In Fig. 4 the 10 modes with the largest differences between healthy and pathological coefficients are shown, in which the color reflects the values $d_{j,k} \times \sigma_j$ (See (4)), i.e. the average deviation (in mm). Clearly, the modes of the PCA

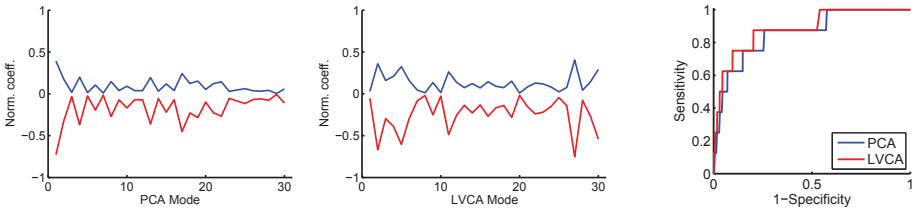

Fig. 3. (a) Normalized means of $\mathbf{b}_i^{\mathrm{PCA}}$ and $\mathbf{b}_i^{\mathrm{LVCA}}$ (i.e. expressed in terms of the corresponding standard deviations (SD)) of the first 30 modes for PCA (left) and LVCA (right), for healthy (blue,up) and pathological shapes (red,down). Split healthy and pathological SDs are just below 1. (b) Average ROC curves (over 50 experiments) for a logistic linear classifier using $\mathbf{b}_i^{\mathrm{PCA}}$ and $\mathbf{b}_i^{\mathrm{LVCA}}$ as features.

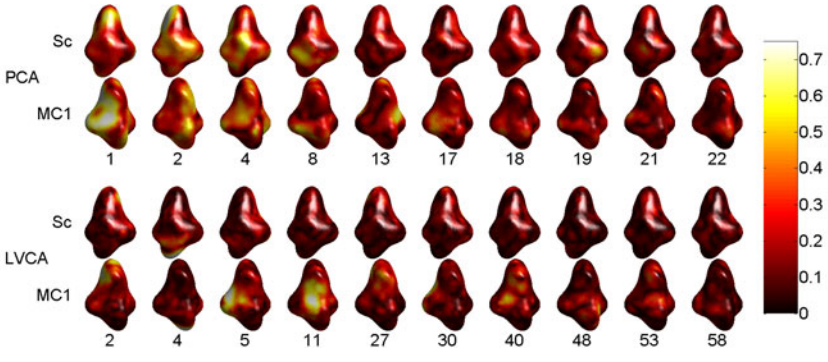

Fig. 4. The 10 modes of variation, sorted by decreasing variance, for which OA trapezia differ the most from healthy trapezia, for PCA and LVCA. Both sides of the trapezium adjacent to the scaphoid (Sc) and metacarpal 1 (MC1) are shown. Intensity increases with variation amplitude (in mm).

model affect the shape in a more global way than the local modes. Particularly the local modes 2, 4 (at the bottom edge), 27 and 30 show the locations of osteophytes, whereas modes 5 and 11 show the flattening of the surface of the TMC joint. The four last local modes in Fig. 4 (40, 48, 53, 58) contain very small variations, which are mainly noise. Observe that the global characteristics of the PCA modes make that they cannot be related to specific bone surface locations that are affected by osteoarthritis. The average receiver operating curves (ROC) for the classifiers are shown in Fig. 3b. These curves show that the LVCA components give a slightly better trade-off between sensitivity and specificity than the PCA components. We hypothesize that this must be attributed to the higher sensitivity of LVCA to local deviations as the classification improvement mainly concerns trapezia with small osteophytes and that deviate as in LVCA mode 11.

4 Discussion

We proposed a new method for efficient modeling of local shape variations. A novel, differentiable locality measure was defined, which effectively equals a 3D weighted variance of the point coordinates per mode. The simplicity of the function and the smooth derivative enable to quickly determine localized components for densely sampled surfaces. The LVCA models based on the IMM faces dataset clearly show that the locality of the modes depends on the trade-off parameter λ. For (relatively) small sets of densely sampled shapes, such as the trapezia, the number of components in the trainingset is sufficiently limited to obtain clinically meaningful local shape variations for $\lambda = 1$. Larger, sparser training sets require a $\lambda < 1$. Estimating LVCA components was shown to be less sensitive to local minima and 2 orders of magnitude faster than LoCA components in the IMM faces set. The LVCA modes corresponded to trapezoid surfaces variations where changes are expected from clinical practice. Reversely, for PCA components the variations were global. Particularly, the area around the facet adjacent to the scaphoid bone was (incorrectly) detected to contain important variations. Classification of pathological trapezia showed a consistent improvement with LVCA coefficient features, compared to PCA.

References

1. Alcantara, D.A., Carmichael, O., Harcourt-Smith, W., Sterner, K., Frost, S.R., Dutton, R., Thompson, P., Delson, E., Amenta, N.: Exploration of shape variation using localized components analysis. IEEE T. Pattern Anal. 31(8), 1510–1516 (2009)
2. Chennubhotla, C., Jepson, A.: Sparse PCA - Extracting multi-scale structure from data. In: ICCV 2001, vol. 1, pp. 641–647 (2001)
3. Cootes, T.F., Taylor, C.J., Cooper, D.J., Graham, J.: Active shape models - their training and application. Comput. Vis. Image Und. 61(1), 38–59 (1995)
4. Eaton, R.G., Glickel, S.Z.: Trapeziometacarpal osteoarthritis. staging as a rationale for treatment. Hand Clin. 3(4), 455–471 (1987)
5. Van de Giessen, M., Foumani, M., Streekstra, G.J., Strackee, S.D., Maas, M., van Vliet, L.J., Grimbergen, C.A., Vos, F.M.: Statistical descriptions of scaphoid and lunate bone shapes. J. Biomech. 43(8), 1463–1469 (2010)
6. Kaiser, H.: The VARIMAX criterion for analytic rotation in factor analysis. Psychometrika 23(3), 187–200 (1958)
7. North, E.R., Rutledge, W.M.: The trapezium-thumb metacarpal joint: the relationship of joint shape and degenerative joint disease. Hand 15(2), 201–206 (1983)
8. Shen, D., Herskovits, E.H., Davatzikos, C.: An adaptive-focus statistical shape model for segmentation and shape modeling of 3-D brain structures. IEEE T. Med. Imaging 20(4), 257–270 (2001)
9. Stegmann, M.B., Ersbøll, B.K., Larsen, R.: FAME – a flexible appearance modelling environment. IEEE Trans. Med. Imaging 22(10), 1319–1331 (2003)
10. Suinesiaputra, A., Frangi, A.F., Kaandorp, T.A., Lamb, H.J., Bax, J.J., Reiber, J.H.C., Lelieveldt, B.P.F.: Automated detection of regional wall motion abnormalities based on a statistical model applied to multislice short-axis cardiac MR images. IEEE T. Med. Imaging 28(4), 595–607 (2009)

Geometric Correspondence for Ensembles of Nonregular Shapes

Manasi Datar[1], Yaniv Gur[1], Beatriz Paniagua[2], Martin Styner[2], and Ross Whitaker[1]

[1] Scientific Computing and Imaging Institute, University of Utah, USA
[2] University of North Carolina at Chapel Hill, USA

Abstract. An ensemble of biological shapes can be represented and analyzed with a dense set of point correspondences. In previous work, optimal point placement was determined by optimizing an information theoretic criterion that depends on relative spatial locations on different shapes combined with pairwise Euclidean distances between nearby points on the same shape. These choices have prevented such methods from effectively characterizing shapes with complex geometry such as thin or highly curved features. This paper extends previous methods for automatic shape correspondence by taking into account the underlying geometry of individual shapes. This is done by replacing the Euclidean distance for intrashape pairwise particle interactions by the geodesic distance. A novel set of numerical techniques for fast distance computations on curved surfaces is used to extract these distances. In addition, we introduce an intershape penalty term that incorporates surface normal information to achieve better particle correspondences near sharp features. Finally, we demonstrate this new method on synthetic and biological datasets.

1 Introduction

A well established method for performing statistics on an ensemble of shapes is to compare configurations of corresponding landmarks placed on the individual shapes. In recent years, several methods have proposed an automatic placement of landmarks in a way that captures statistical properties of an ensemble [1,2]. The method of Cates et al [1] uses a formulation of ensemble entropy to deploy a dense set of landmarks, or *particles*, which assign correspondences between shapes within a population. The positions of the particles on the shape surfaces are optimized using a variational framework that tries to find a balance between model simplicity via minimum entropy, and geometric accuracy of the surface representation. However, medical or biological objects shapes are often composed of sharp features and regions of high curvature. In such cases, nearby particles in the ambient space may be separated by a large distance on the object's surface (see Fig. 1). Thus, the Euclidean distance measure cannot capture correctly the underlying local geometry and prevents the method from producing a faithful shape representation. This limitation reflects a failure of Euclidean distance

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 368–375, 2011.

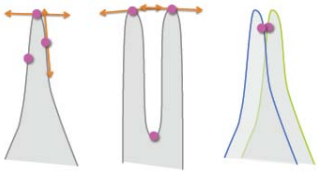

Fig. 1. Points near sharp features (left) are not able to achieve good distributions with Euclidean distance, because they do not lie in the same tangent space, which is necessary for movement that is constrained to the surfaces. Points may be nearby and interact (center) even though they sample very different parts of the surface. Points on nearby features (right) on different shapes (blue and green) can come into incorrect correspondence if the system considers only distance.

to account for the intrinsic distances between points on the surface, suggesting geodesic distance as a better choice. However, geodesic distances are generally not computable in closed form, and interparticle interactions are part of the inner loop of an interactive optimization process. Thus, the computational burden of geodesics are prohibitive. This paper makes several contributions that enable better modeling of ensembles composed of shapes with a complex geometric structure. First, we incorporate geodesic distance measures into the framework proposed by Cates, et al [1]. While accurate geodesic distance computation is unwieldy for implicit surfaces, very fast methods exist to compute geodesic distances between vertices of 3D mesh representations of shapes. Thus, we propose precomputing all pairwise distances on a somewhat fine 3D mesh of an input surface and interpolate, as required, in the process of optimizing intrashape particle interactions. To address mismatches of correspondences on highly curved features across different shapes, we introduce an intershape penalty that accounts for the behavior of normals on highly curved geometry. Hence, the second contribution of our paper is to integrate this intershape penalty term into the variational framework for model optimization given in [1] to improve particle correspondences near sharp features. As a final contribution we demonstrate the use of a correspondence-based method for the analysis of highly curved (or nonregular) shapes—the left ventricle myocardium of the human heart—which has, so far not feasible with point correspondences.

2 Background

In the following section we provide a brief overview the particle-system correspondence optimization method as proposed in [1]. The general strategy of this method is to represent correspondences as point sets that are distributed across an ensemble of similar shapes by minimizing an objective function that quantifies the entropy of the system. We also describe an efficient, fine-grained algorithm for solving the eikonal equation on triangular meshes, as proposed by Fu et al [3]

Correspondence Optimization. Let us define a surface as a smooth, closed manifold of codimension one, which is a subset of $\mathbb{R}^d$ (e.g., $d = 3$ for volumes). We sample the surface $\mathcal{S} \subset \mathbb{R}^d$ using a discrete set of N points that are considered as random variables $\mathbf{Z} = (\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_N)^T, \mathbf{X} \in \mathbb{R}^d$ drawn from a probability density function (PDF), $p(\mathbf{X})$. We denote a realization of this PDF with lower case, and thus we have $\mathbf{z} = (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N)^T$, where $\mathbf{z} \in \mathcal{S}^N$. We refer to the positions $\mathbf{x}$ as *particles*, and to a set of particles as a *particle system*. The amount of information encoded in this random sampling is, in the limit, the differential entropy of the PDF, given by $H[\mathbf{X}] = -E\{\log p(\mathbf{X})\}$, where $E\{\cdot\}$ is the expectation. Approximating the expectation by the sample mean, we have $H[\mathbf{X}] \approx -\frac{1}{Nd}\sum_i \log p(\mathbf{x}_i)$. To determine the probability of a particle's position, $p(\mathbf{x}_i)$, [1] uses a nonparametric Parzen-window density estimation given by a mixture of multivariate, isotropic Gaussian kernels with standard deviation σ that determines the strength of particles interaction with N neighbouring particles within the defined window. An ensemble comprised of M surfaces, $\mathcal{E} = \mathbf{z}^1, \ldots, \mathbf{z}^M$ can be described by a $Nd \times M$ matrix of particle positions $P = (\mathbf{x}_j^k)$, where $k = 1, \ldots, M$ and $j = 1, \ldots, N$. Let $\mathbf{z}^k \in \mathbb{R}^{Nd}$ be an instance of a random variable $\mathbf{Z}$, then, the combined ensemble and shape cost function is defined by

$$Q = H(\mathbf{Z}) - \sum_k H(P^k) \tag{1}$$

This cost function is composed of two interacting terms. The first term produces a compact distribution of samples in shape space, while the second term provides uniformly-distributed correspondence positions on the shape surfaces, to achieve a faithful shape representation. The optimization process of this cost function is defined via gradient descent as described in [1].

Fast Geodesic Distance Computation. The use of Euclidean distance between particles in the Parzen-window density estimation in [1] requires that nearby particles interact in the local tangent plane of the surface. However, it is not the case for thin structures with high curvature, such as the one illustrated in Fig. 1 (left). To address this, we replace the Euclidean distance in the kernel by the *geodesic* interparticle distance. However, this modification demands a large number of pairwise geodesic distance computations. Such computations are not feasible without the recent developments in fast, parallel algorithms for solving hyperbolic partial differential equations (PDEs) as well as extremely fast SIMD hardware in the form of graphics processors (GPUs). The distance between each point a on the surface and every other point, is given by the solution to the eikonal equation $|\nabla u| = 1$, as discussed in [4], using the boundary condition $u(a) = 0$. The computation of distances to many thousands of points on large ensembles of shapes is feasible only if the eikonal equation can be solved in a small fraction of a second. The fast iterative method (FIM) [4] for regular grids is not worst-case optimal, but is extremely efficient on parallel, SIMD architectures, such as GPUs. Here, we use an extension of the FIM for triangular meshes [3]. This algorithm computes, for instance, distances between nodes on a mesh with thousands of vertices in less than 30 seconds on a GPU.

3 Methodology

The input to the shape correspondence system is a collection of implicitly defined surfaces. For this paper, the input surfaces are binary segmentations, and we use the preprocessing, initialization, and particle optimization pipeline described in [1]. Here we describe the integration of the geodesic distance for interparticle interactions and the surface normal based penalty term for intershape correspondence into the framework described in Sec. 2.

Particle Position Optimization using Geodesic Distances. A triangulation of each input surface is generated using the algorithm described in [5]. An example triangulation is shown in Fig. 2(a) along with the corresponding synthetic shape. The numerical technique for fast distance computation on 3D triangulated surfaces described at the end of Sec. 2 is then used to precompute geodesic distances between each vertex and all other vertices within a prescribed distance, $d_{\max}$. The parameter $d_{\max}$ is chosen to coincide with the limited range of influence of the Gaussian kernels that control the range of influence of each particle. This truncation results in a sparse, symmetric matrix of geodesic distances. The entries in this matrix are then converted into a fixed point format and stored using a List of Lists (LIL) representation for efficient memory usage and fast access. We call this matrix M_G such that $D_G(v_1, v_2) = M_G[v_1, v_2]$, where $D_G(v_1, v_2)$ is the geodesic distance between vertices v_1, v_2. Geodesic distances between particle positions on the implicit surface can now be computed via a barycentric interpolation scheme described below.

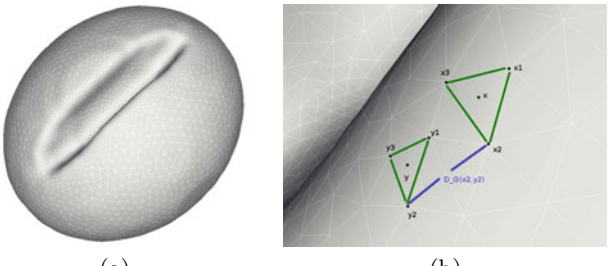

(a) (b)

Fig. 2. (a) An example of a triangle mesh used for geodesic distance computations. (b) Configuration for two-layered interpolation of geodesic distance between arbitrary points: x and y are contained in triangles defined by vertices (x_1, x_2, x_3) and (y_1, y_2, y_3) respectively. The geodesic distances between vertices for all shapes are precomputed on a GPU.

To use this discrete set of distances between particles, which lie in the volume and are constrained to lie on the implicit surface, we interpolate the mesh-vertex distances to the faces of the triangles. This requires two layers of linear interpolation on the faces of the mesh. Let the barycentric coordinates of a point x in a triangle T_x defined by vertices (v_1, v_2, v_3) be given by (α, β, γ) such that

the location of x can be given as $x = \alpha v_1 + \beta v_2 + \gamma v_3$ where $\alpha + \beta + \gamma = 1$. Consequently, any function of x can be interpolated as $f(x) = \alpha f(v_1) + \beta f(v_2) + \gamma f(v_3)$ provided its value is known at all vertices in the mesh. For the case of geodesic distances, the function f is the distance to another arbitrary point y, which can be evaluated on each vertex using this same interpolation scheme for the triangle T_x that contains y. To compute $D_G(x, y)$ in a fast and efficient manner, we first determine the triangle faces on the mesh that contain points x and y, by projecting them onto the nearest face in the mesh. Let these triangles defined by vertices (x_1, x_2, x_3) and (y_1, y_2, y_3), as shown in Fig. 2. Since the geodesic distance is a function defined between every pair of vertices in the mesh, we can approximate the geodesic distance between points x and y as

$$\begin{aligned} D_G(x, y) &\approx \alpha D_G(x_1, y) + \beta D_G(x_2, y) + \gamma D_G(x_3, y), \\ D_G(x_i, y) &\approx \alpha D_G(x_i, y_1) + \beta D_G(x_i, y_2) + \gamma D_G(x_i, y_3). \end{aligned} \tag{2}$$

Each $D_G(x_i, y_i)$ is simply an entry in the matrix M_G as described above. Thus, using this two-layered interpolation scheme, we can approximate geodesic distances between particle positions on the implicit surface. The Gaussian forces of repulsion governing the motion of particles can then be computed as a function of these geodesic distances to improve sensitivity to the underlying geometry.

Correspondence Optimization with Surface Normals. The cost function described in Eq. 1 relies on particle positions to find a balance between a compact ensemble representation and a good distribution of particles on each surface. However, with an ensemble containing highly curved or convoluted surfaces, like those shown in Fig. 1, a reliance on only positional information may lead to incorrect correspondences. To address this shortcoming, we propose the addition of an intershape penalty term based on surface normals to disambiguate correspondences near highly curved features. Thus, we associate with each particle on each surface a pair of d-tuples $(x_i, n_i) \in \mathbb{R}^d \times S^2$, where S^2 is the unit sphere.

We denote the total collection of N normals across M shapes as V. With the assumption that $N > M$. Assuming a Gaussian model with a covariance Σ, we can compute the entropy

$$H(V) \approx \frac{1}{2} \log |\Sigma| \approx \frac{1}{2} \log |\sum_i \sum_k \hat{n}_i^k \cdot (\hat{n}_i^k)^T| \tag{3}$$

For the ith particle on the kth shape, $\hat{n}_i^k = d(n_i^k, \bar{n}_i)$, where $\bar{n}_i$ is the Fréchet mean defined in [6]. Since the normals are points on the Riemannian manifold $\mathcal{M} \epsilon S^2$, $\hat{n}_i^k = Log_{\bar{n}_i}(n_i^k)$ [6]. In the tangent plane $\mathcal{T}_{\bar{n}_i}\mathcal{M}$, we have

$$\hat{n}_i^k = Log_{\bar{n}_i}(n_i^k) = \frac{P_t(n_i^k - \bar{n}_i) \arccos(n_i^k \cdot \bar{n}_i)}{1 - (n_i^k \cdot \bar{n_i})^2} \tag{4}$$

where P_t is the idempotent projection matrix given by $(I - \bar{n}_i \cdot (n_i^k)^T)$. Since Σ will not have a full rank in practice, we implement a regularization similar to that

described in [1] to introduce a lower bound on the eigenvalues. The optimization problem in Eq. 1 can now be reformulated as

$$Q = H(\mathbf{Z}) - \sum_k H(P^k) + H(V) \tag{5}$$

The Riemannian distances are functions of normals; $\hat{n}_i^k = f(n_i^k)$, which in turn are a function of position; $n_i^k = n(x_i^k)$, the gradient descent on $H(V)$ with respect to particle position x_i is given by the chain rule:

$$\frac{\partial H(V)}{\partial x_i^k} = \frac{\partial H(V)}{\partial \hat{n}_i^k} \cdot \frac{\partial f(n_i^k)}{\partial n_i^k} \cdot \frac{\partial n(x_i^k)}{\partial x_i^k} \tag{6}$$

This incremental update gets projected onto the tangent plane of the surface, as part of the algorithm described in Sec. 2, in order to maintain the constraint that particles remain on the surface. As with the geodesic distances, the curvature, $\partial n_i^k / \partial x_i^k$, is precomputed. Here we use the formulation of curvature for the level sets of the volume using finite differences (combined with a Gaussian kernel of standard deviation 1.0). The means of the normals are updated after each full iteration (one update for every particle on every shape).

4 Results and Discussion

This section details experiments designed to illustrate and validate the proposed method. First, we present an experiment with synthetically generated *coffee bean* shapes, that consist of an ellipsoid with a *slot* or indentation, creating a high-curvature feature that would confound the previous approaches. We also present an application to a study of group differences in the left ventricular myocardium.

Synthetic Data
Computational solid geometry methods were used to compute the intersection of a small ellipsoid with axes a, b and c, and a larger ellipsoid with axes A, B and C, to create a *coffee bean* shape. The *slot* was then moved and scaled stochastically, to create a population of 10 *coffee bean* shapes. The position of the *slot* was chosen from a uniform distribution in the range $[-B/3, B/3]$, and its width was sampled from a Gaussian distribution of $\mu = 8$ and $\sigma = 2$. Both, the method in [1] and the proposed method were applied to distribute 1024 correspondences across the ensemble. Both methods identified two dominant modes of variation, with significantly different amount of leakage into smaller modes. These modes are illustrated in Fig. 3 for both the methods, to 2 standard deviations. The proposed method lost 4% of the total variation into smaller modes, compared to 16% lost by the original method. Thus, the proposed method was able to characterize the variation in the population better than the original method, while remaining faithful to the original shape representation (as seen from the reconstructions in Fig. 3).

Application to Group Comparison. We applied the proposed methods to study group differences in the left ventricular myocardium of ischemic patients

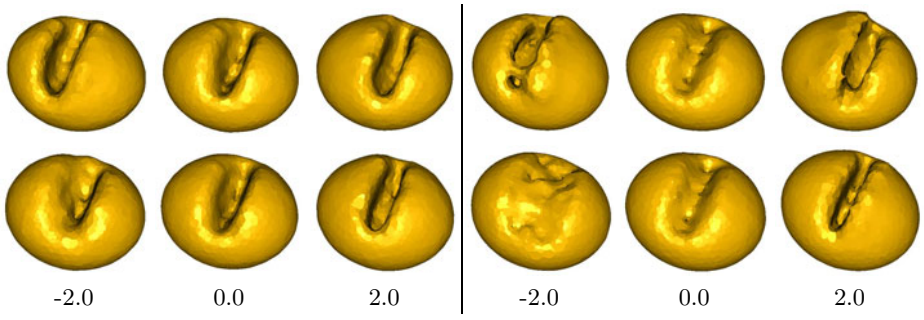

Fig. 3. Mean shape computed from the proposed method (left) and the original method (right), projected onto the first (top) and second (bottom) PCA modes, and ±2 standard deviations

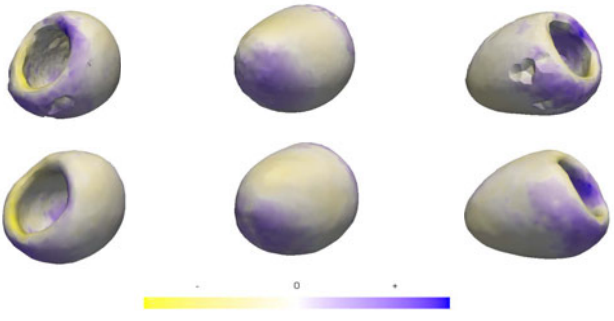

Fig. 4. Visualizing mean differences between normal and ischemic groups (blue denotes expansion and yellow denotes contraction) using [1] (top row) and the proposed method (bottom row)

and non-ischemic controls, using segmented volumes of the left ventricular myocardium at end diastole (ED) as inputs. The proposed method was used to initialize and optimize 1024 correspondences across the ensemble of 21 (12 patients, 9 controls) shapes. We then used *parallel analysis* to project the correspondences into a lower dimensional space determined by choosing an optimal number of basis vectors from principal component analysis (PCA). A standard, parametric Hotelling T^2 test was used to test for group differences, with the null hypothesis that the two groups are drawn from the same distribution. In this case, the hypothesis test results in a highly significant *p-value* of 0.005, with 7 PCA modes chosen by *parallel analysis*.

Fig. 4 (bottom row) shows the differences between the mean shape surfaces for the normal and ischemic groups. To visualize the group differences driving statistical results, we use the linear discriminant vector, rotated from PCA space into the full dimensional shape space, and mapped onto the mean group shape surfaces to give an indication of the significant morphological differences between

groups. The above experiment was also conducted using the method described in [1]. The resulting group differences, visualized in Fig. 4 (top row), were also found to be statistically significant with a *p-value* of 0.005 using the Hotelling T^2 test. However, the shape differences obtained using the proposed method are found to be more consistent with previously published results presented in [7], as compared to those obtained using [1].

5 Conclusion and Future Work

This paper extends the method given by [1] to improve particle distribution and correspondences across an ensemble of highly convoluted surfaces. The first contribution is the inclusion of geodesic distance to compute the intrashape particle interactions, which results in improved sensitivity of the particle distribution to the underlying surface geometry. The second contribution is the introduction of an intershape penalty term based on surface normals, to improve correspondence near sharp features. Results on synthetic and real data indicate that the proposed method provides a practical solution to compute correspondence models of ensembles of highly convoluted surfaces in an efficient and robust manner.

Acknowledgements. This work is supported by the NIH/NCRR Center for Integrative Biomedical Computing - 2P41 RR0112553-12, the NIH/NCBC National Alliance for Medical Image Computing - U54-EB005149 and NSF grant CCF-073222. We are also grateful to Dr. Raimond Winslow at The Center for Cardiovascular Bioinformatics and Modeling, John Hopkins University, for providing left ventricle data for the group comparison experiment.

References

1. Cates, J., Fletcher, P.T., Styner, M., Shenton, M., Whitaker, R.: Shape modeling and analysis with entropy-based particle systems. In: Karssemeijer, N., Lelieveldt, B. (eds.) IPMI 2007. LNCS, vol. 4584, pp. 333–345. Springer, Heidelberg (2007)
2. Davies, R., Twining, C., Allen, P., Cootes, T., Taylor, C.: Shape discrimination in the hippocampus using an MDL model. In: Taylor, C.J., Noble, J.A. (eds.) IPMI 2003. LNCS, vol. 2732, pp. 38–50. Springer, Heidelberg (2003)
3. Fu, Z., Kirby, M., Whitaker, R.: A fast iterative method for solving the eikonal equation on triangulated meshes. SIAM Journal on Scientific Computing (2011) (to appear)
4. Jeong, W., Whitaker, R.: A fast iterative method for eikonal equations. SIAM Journal on Scientific Computing 30(5), 2512–2534 (2008)
5. Meyer, M.D., Georgel, P., Whitaker, R.T.: Robust particle systems for curvature dependent sampling of implicit surfaces. In: Proceedings of the International Conference on Shape Modeling and Applications, pp. 124–133 (June 2005)
6. Fletcher, P., Lu, C., Pizer, S., Joshi, S.: Principal geodesic analysis for the study of nonlinear statistics of shape. IEEE Trans. Med. Imaging 23(8), 995–1005 (2004)
7. Ardekani, S., Weiss, R., Lardo, A., George, R., Lima, J., Wu, K., Miller, M., Winslow, R., Younes, L.: Computational method for identifying and quantifying shape features of human left ventricular remodeling. Ann. Biomed. Eng. 37(6), 1043–1054 (2009)

Hippocampal Surface Mapping of Genetic Risk Factors in AD via Sparse Learning Models

Jing Wan[1,2,⋆], Sungeun Kim[1,⋆], Mark Inlow[1,3], Kwangsik Nho[1], Shanker Swaminathan[1], Shannon L. Risacher[1], Shiaofen Fang[2], Michael W. Weiner[4], M. Faisal Beg[5], Lei Wang[6], Andrew J. Saykin[1,⋆⋆], Li Shen[1,2,⋆⋆], and ADNI

[1] Radiology and Imaging Sciences, Indiana University School of Medicine, IN, USA
[2] Computer and Information Science, Purdue University Indianapolis, IN, USA
[3] Mathematics, Rose-Hulman Institute of Technology, IN, USA
[4] Radiology, Medicine and Psychiatry, UC San Francisco, CA, USA
[5] School of Engineering Science, Simon Fraser University, BC, Canada
[6] Psychiatry and Behavioral Sciences, Northwestern University, IL, USA

Abstract. Genetic mapping of hippocampal shape, an under-explored area, has strong potential as a neurodegeneration biomarker for AD and MCI. This study investigates the genetic effects of top candidate single nucleotide polymorphisms (SNPs) on hippocampal shape features as quantitative traits (QTs) in a large cohort. FS+LDDMM was used to segment hippocampal surfaces from MRI scans and shape features were extracted after surface registration. Elastic net (EN) and sparse canonical correlation analysis (SCCA) were proposed to examine SNP-QT associations, and compared with multiple regression (MR). Although similar in power, EN yielded substantially fewer predictors than MR. Detailed surface mapping of global and localized genetic effects were identified by MR and EN to reveal multi-SNP-single-QT relationships, and by SCCA to discover multi-SNP-multi-QT associations. Shape analysis identified stronger SNP-QT correlations than volume analysis. Sparse multivariate models have greater power to reveal complex SNP-QT relationships. Genetic analysis of quantitative shape features has considerable potential for enhancing mechanistic understanding of complex disorders like AD.

1 Introduction

Recent advances in brain imaging and high throughput genotyping techniques enable new approaches to study the influence of genetic variation on brain structure and function. Existing imaging genetics studies employ summary statistics (e.g., volume, thickness) [7] and detailed voxel-wise measures [8] as phenotypes to discover genetic risk factors. Genetic mapping of hippocampal shape, an

⋆ Equal contribution by J Wan (wanjing@iupui.edu) and S Kim (sk31@iupui.edu).
⋆⋆ Correspondence to L Shen (shenli@iupui.edu) or AJ Saykin (asaykin@iupui.edu).
Data collection and sharing for this project was funded by the Alzheimer's Disease Neuroimaging Initiative (ADNI) (U01 AG024904, http://adni.loni.ucla.edu). This project was also supported in part by NIBIB R03 EB008674, NIA 1RC 2AG036535, CTSI-IUSM/CTR(RR025761), NIA P30 AG10133, and NIA R01 AG19771.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 376–383, 2011.

Table 1. Participant characteristics

Category	HC	MCI	AD	p-value
Gender (M/F)	91/75	184/103	68/61	0.041
Baseline Age (years; Mean±STD)	76.18±4.91	74.99±7.21	75.36±7.78	0.198
Education (years; Mean±STD)	16.20±2.63	15.71±2.98	15.07±3.04	$<$ **0.005**
Handedness (R/L)	155/11	260/27	121/8	0.411

under-explored area, has strong potential as a neurodegeneration biomarker for Alzheimer's disease (AD) and mild cognitive impairment (MCI). The present study investigates genetic effects of top candidate single nucleotide polymorphisms (SNPs) on hippocampal shape features in a large cohort.

Massive univariate analyses are often used in imaging genetics [7,8], and can quickly identify important associations between individual SNPs and imaging quantitative traits (QTs). However, it treats SNPs and QTs as independent units, and overlooks relationships in which multiple SNPs jointly effect multiple QTs. In this work, two multivariate sparse models, the elastic net and sparse canonical correlation analysis, are used to study genetic effects on hippocampal shape and are expected to have greater power to reveal complex SNP-QT relationships. These models could enable discovery of a small set of relevant features which may provide potential surrogate biomarkers for therapeutic trials.

2 Materials and Methods

Magnetic resonance imaging (MRI) and genotype data were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database [7]. ADNI is a landmark investigation sponsored by the NIH and industrial partners designed to collect longitudinal neuroimaging, biological and clinical information from over 800 participants that will track the neural correlates of memory loss from an early stage. Further information can be found at www.adni-info.org. 582 non-Hispanic Caucasian participants (166 Healthy Control (HC), 287 MCI, 129 AD participants) with segmented hippocampal data and quality controlled (QC) genotype data were included in this study (Table 1).

Hippocampal Shape: Hippocampi were segmented from the baseline MRI scans by applying probabilistic-based FreeSurfer and Large Deformation Diffeomorphic Metric Mapping (FS+LDDMM) [3]. This fully-automated segmentation pipeline first uses FreeSurfer subcortical labeling to provide information for initialization, and then employs LDDMM to generate a diffeomorphic transformation so that anatomical structures can be mapped consistently and smoothly. To remove size effect, total intracranial volume (ICV) was adjusted to a constant (i.e., mean ICV of all HCs) and each hippocampus was scaled accordingly. Rigid body transformation was then applied to register each hippocampus to a template (defined as the mean of all HCs) in a least square fashion. Surface signals were extracted as the deformation along the surface normal direction of the

template, and were adjusted for baseline age, gender, education, and handedness using the regression weights derived from the HC participants (Table 1).

Candidate SNPs: The SNP data were genotyped using the Human 610-Quad BeadChip (Illumina, Inc., San Diego, CA). We focused on top AD genetic risk factors, including top 23 SNPs from the AlzGene database [1] as of 09/01/2010, and a SNP from the TOMM40 gene adjacent to the APOE gene. The TOMM40 SNP was included because it was unclear whether the SNP played a unique role in AD or served solely as an APOE marker. Four SNPs were excluded due to failed imputation or quality check. Among the remaining 20 SNPs (Fig. 1(a)), 10 SNPs were available from the ADNI data and 10 SNPs were successfully imputed using MACH v1 [4] and IMPUTE v2 [6] software packages. The QC criteria for the SNP data include (1) call rate check per subject and per SNP marker, (2) gender check, (3) sibling pair identification, (4) the Hardy-Weinberg equilibrium test, (5) marker removal by the minor allele frequency and (6) population stratification. The selected 20 SNPs were numerically coded to test additive genetic effect, i.e., dose dependent effect of the minor allele.

Overall Strategy: For comparative analysis, *multiple regression models* were fit using all 20 SNPs to predict the hippocampal volume (mean of left and right, covaried for age, gender, education, handedness and ICV) and, in addition, the surface signal at each location or vertex on the hippocampal surface. The *elastic net regression* was then applied to identify a small set of relevant SNPs for each surface location. Finally, *sparse canonical correlation analysis* was used to examine more complex relationships between SNP sets and surface regions.

Multiple Regression: Under the additive model, the surface signals are linearly related to the number of minor alleles. This implies, assuming no interactions between SNPs, the multiple regression model $S_{i,j} = \beta_{0,j} + \beta_{1,j}\text{SNP}_{i,1} + \cdots + \beta_{20}\text{SNP}_{i,20} + \epsilon_{i,j}$, where $S_{i,j}$ is the surface signal at vertex j for subject i. The model utility F test was used to test the null hypothesis of no relationship between S_j and the 20 SNPs for the $j = 1, \ldots, 13222$ vertices. Gaussian random field theory (RFT) methods [13], implemented in SurfStat [12], were used to ensure the family-wise error rate did not exceed 0.05. While this procedure can detect any linear relationship between S_j and the SNPs this flexibility comes at the cost of reduced power to detect a relationship between a specific SNP and S_j. Sparse regression methods, which seek to accurately predict the response variable using a minimal number of predictors, address this and other regression shortcomings by integrating variable selection and model estimation.

Elastic Net Regression: The ability of sparse regression methods to detect and model genetic relationships was investigated by estimating the above model at each hippocampal location using elastic net (EN). EN produces sparse solutions by adding a coefficient magnitude penalty to the least squares objective function [14]. More specifically, the EN coefficient estimates minimize the penalized least squares objective function $\text{ElNet}_j(\beta_0, \beta_1, ..., \beta_{20}) = \sum_{i=1}^{n}(S_{i,j} - \hat{S}_{i,j})^2 + \lambda P_\alpha(\beta_1, \ldots, \beta_{20})$, in which $\hat{S}_{i,j} = \beta_{0,j} + \beta_{1,j}\text{SNP}_{i,1} + \cdots + \beta_{20,j}\text{SNP}_{i,20}$

and the penalty $\hat{P}_\alpha(\beta_1, \ldots, \beta_{20}) = \alpha \sum_{k=1}^{20} |\beta_k| + (1-\alpha) \sum_{k=1}^{20} \beta_k^2$ is a convex combination of the L_1 lasso and L_2 ridge penalties. This objective function has two parameters: λ controls the amount of shrinkage; and α adjusts the trade-off between lasso and ridge to capitalize on their strengths and minimize their weaknesses. The preceding regression analysis was duplicated using the Glmnet [2,9] implementation of EN with $\alpha = 0.5$ and λ chosen using 10-fold cross-validation.

Sparse Canonical Correlation Analysis: The surface signals represent samples of a smooth function defined on the hippocampus. Methods which capitalize on the resulting correlation between surface signals at neighboring vertices by modeling the joint relationship between multiple surface signals and SNPs should provide increased power to detect any relationships present [10]. To investigate this possibility for linear relationships, sparse canonical correlation analysis (SCCA) was used. Let $X_i = (\mathrm{SNP}_{i,1}, \mathrm{SNP}_{i,2}, \ldots, \mathrm{SNP}_{i,20})'$ be the vector of the 20 SNPs for subject i and $Y_i = (S_{i,1}, S_{i,2}, \ldots, S_{i,m})'$ be the vector consisting of the surface signals at the $m = 13,222$ vertices. Canonical correlation analysis (CCA) produces linear combinations (canonical variates) $U_j = A_j' Y$ and $V_j = B_j' X$, $j = 1, \ldots, 20$, such that the correlation between U_j and V_j is maximized subject to orthogonality constraints. Two major weaknesses of CCA are that it requires the number of observations n to exceed the combined dimension of Y and X (here 13,242) and that it produces nonsparse A_j and B_j which are difficult to interpret. The SCCA method employed here ameliorates these weaknesses using the penalized matrix decomposition approach [11]. This method maximizes the correlation between U and V subject to the coefficient vector constraints $P_1(A) \leq c_1$ and $P_2(B) \leq c_2$. Here the L_1 penalty $P(A) = \sum_{k=1}^{p} |A(k)|$ was used for both P_1 and P_2. Values for c_1 and c_2 were chosen using Witten and Tibshirani's permutation tuning procedure. The SCCA analyses were computed using the R package PMA (Penalized Multivariate Analysis v.1.0.7.1).

3 Results

In the volumetric analysis of 20 SNPs, only APOE SNP (rs429358) has a significant ($p \leq 0.0004$) effect on the hippocampal volume. The Pearson correlation coefficient between the APOE SNP and hippocampal volume was -0.159.

Fig. 2(a) shows the map of F-statistics of multiple regression (MR). Regions with $F \geq 3.0$ and spatial extent ≥ 2.4 resels have a random field theory adjusted p-value ≤ 0.05. Fig. 2(b) shows the mean of the absolute residuals (fitted errors) over all subjects. The residual map of elastic net (EN) is almost identical to Fig. 2(b), showing similar predictive power between EN and MR.

However, the predictors selected by EN are much more sparse than those of MR (see Fig. 1(a-c)). Combining Fig. 1(c) with (a) and (b), we can extract the coefficient map for a specific SNP and examine localized genetic effects on the surface. Shown in Fig. 1(d-g) are examples of the APOE and TOMM40 SNPs, which elucidate the benefit of sparsity achieved in EN compared to MR. While MR indicates a global effect on the surface (f,g), EN identifies localized regional effects (d,e) and yields useful information for biomarker discovery.

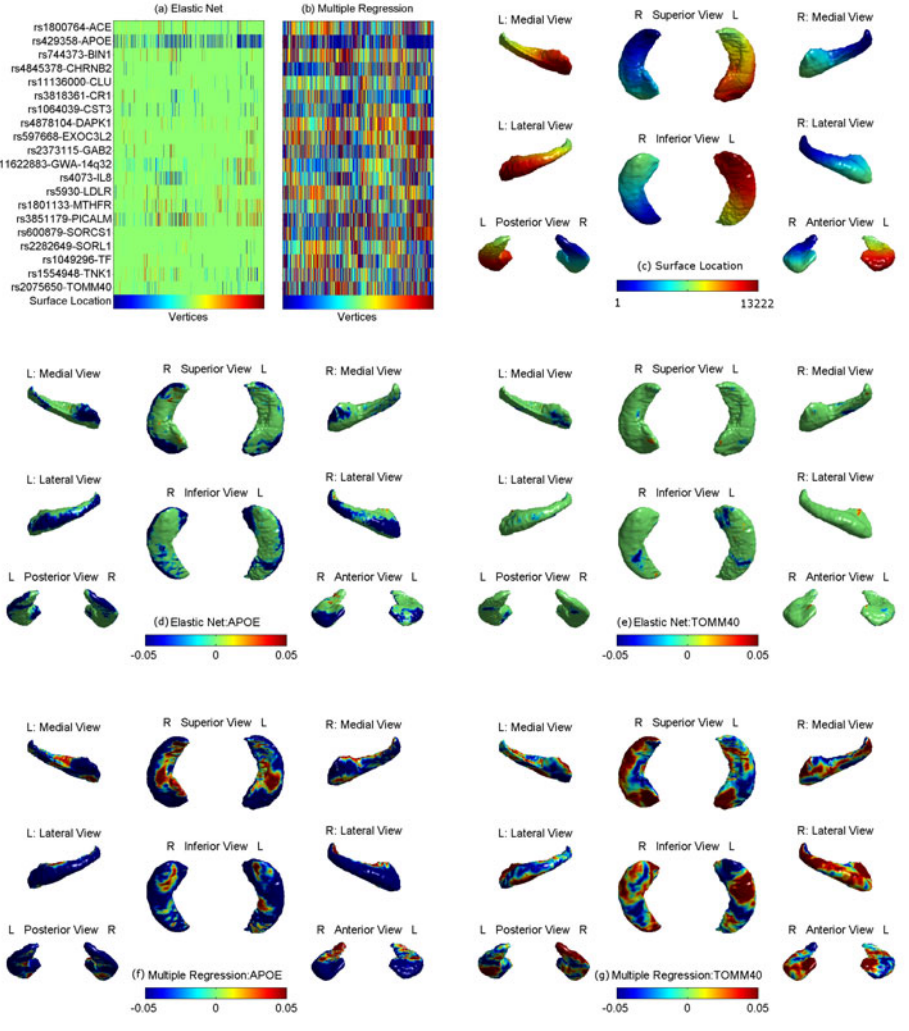

Fig. 1. (a-c) Heat maps of regression coefficients for elastic net (a) and multiple regression (b), where the hippocampal surface location (bottom row in (a,b)) is color-coded and mapped in (c). (d-e) Surface map of genetic effects of the APOE and TOMM40 SNPs estimated by elastic net (d,e) and multiple regression (f,g).

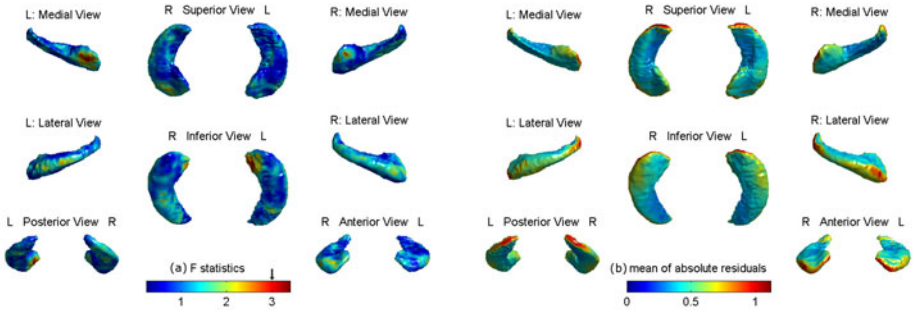

Fig. 2. (a) F-map of multiple regression. (b) Mean of absolution residuals.

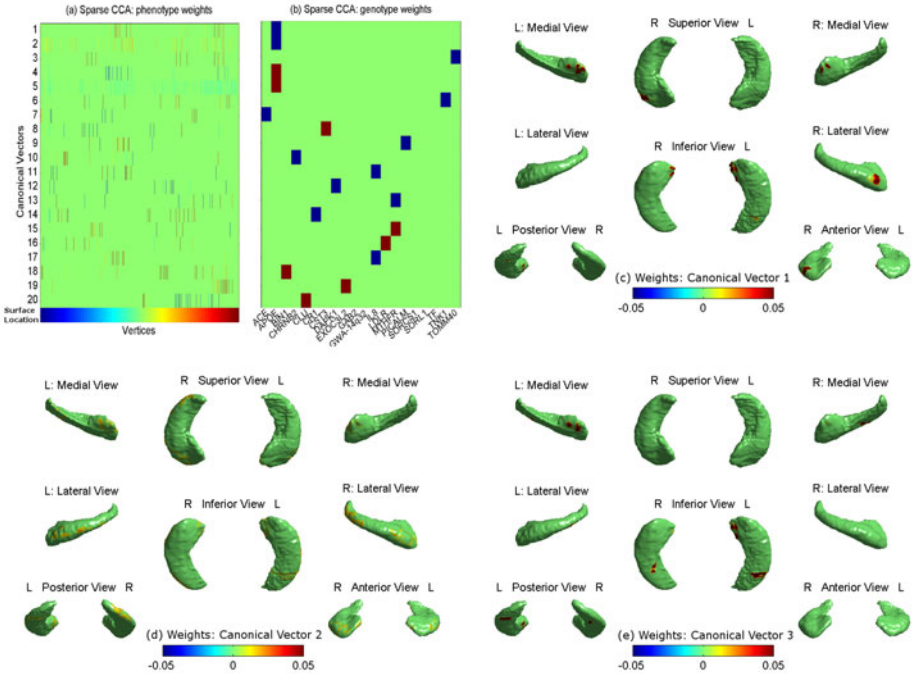

Fig. 3. (a-b) Weights of canonical vectors ordered by descending correlations between surface signals (a) and SNPs (b). (c-e) Surface maps of the top three canonical vectors: the first three rows in (a) are mapped onto the surface.

Fig. 3 shows the results of SCCA. Weights of 20 canonical vectors for vertex-based surface signals (a) and SNPs (b) were color-coded as heatmaps. The top three rows in (a) were mapped onto the hippocampal surface and shown in (c-e), respectively. In (a-b), canonical vector pairs (i.e., corresponding rows in (a-b)) were ordered by descending correlation between surface signals and SNPs; and

the correlation coefficients of all 20 pairs ranged from 0.26 to 0.17 in descending order. This clearly demonstrated the increased power of shape analysis, since the strongest correlation between each of 20 SNPs and hippocampal volume in our volumetric analysis was between the APOE SNP and hippocampal volume with a magnitude of 0.159. This was corroborated by the fact that the maximum absolute correlation between the surface signal and APOE SNP was 0.20 among all vertices and was 0.19 among the vertices with $F \geq 3.0$.

In addition, the parameters for SCCA were automatically tuned by 100 permutations to increase the sparsity and smoothness. As a result, the identified surface locations, correlated with each SNP were more sparse than those for the same SNP from EN (see Fig. 3(a-b) vs Fig. 1(a)). Interestingly, the sparsity was maximized for SNPs, since each canonical SNP vector selected exactly one SNP (Fig. 3(b)), yielding a simple model easy to interpret (i.e., multi-SNP-multi-location associations became single-SNP-multi-location ones).

Fig. 3(c-d) show surface regions related with the APOE SNP (rs429358) at different correlation levels. The correlated vertices in Fig. 3(c-d) have non-zero weights as in Fig. 1(d,f), but they are localized to smaller regions in Fig. 3(c-d). Fig. 3(e) shows surface regions related with the TOMM40 SNP (rs2075650). All vertices with non-zero weights in Fig. 3(e) also have non-zero weights in Fig. 1(e,g). However, compared to Fig. 1(e,g), vertices with non-zero weights in Fig. 3(e) are highly sparse and spatially localized to smaller areas. These two types of patterns are complimentary: the associations derived from EN are multi-SNP-single-location, while those found in SCCA are single-SNP-multi-location.

Five-fold cross-validation of SCCA yielded equally sparse SNP-QT patterns. The most consistent canonical component identified in all five trials is similar to the top finding using the entire data: the genetic vector contains only APOE, and the phenotype vector shows a pattern like Fig. 3(c). Training and testing correlation coefficients are 0.279 ± 0.017 (mean $\pm$ SD) and 0.175 ± 0.068, respectively, while the magnitudes of correlation coefficients between APOE and hippocampal volume in the same data are 0.159 ± 0.012 and 0.163 ± 0.056, respectively.

4 Discussion

Detailed surface mappings of localized genetic effects were identified from our hippocampal shape analysis. Different from existing massive univariate analyses [7,8], this study is among the first to simultaneously use multiple response variables with multiple predictors for analyzing real neurogenomic data [5,10] and may be the first for studying genetic influences on hippocampal morphometry using this paradigm. In our analyses, we combined two promising sparse multivariate models with a typical morphometric method. Investigation of other statistical models (e.g., [10]) and surface metrics, coupled with pathway analyses, will be important future topics to potentially yield new discoveries. As the best known AD genetic risk factor, APOE was the most prominent signal in all of our analyses, which to some extent validated the efficacy of our methods. Replication in independent large samples will be important to confirm the imaging

genetic findings. Genetic analysis of quantitative shape features has considerable potential for examining disease mechanisms from a novel perspective that can inform selection of imaging biomarkers for early detection and therapeutic trials.

References

1. Bertram, L., McQueen, M.B., Mullin, K., Blacker, D., Tanzi, R.E.: Systematic meta-analyses of alzheimer disease genetic association studies: the alzgene database. Nat. Genet. 39(1), 17–23 (2007)
2. Friedman, J., Hastie, T., Tibshirani, R.: Regularization paths for generalized linear models via coordinate descent. J. Stat. Softw. 33(1), 1–22 (2010)
3. Khan, A.R., Wang, L., Beg, M.F.: Freesurfer-initiated fully-automated subcortical brain segmentation in mri using large deformation diffeomorphic metric mapping. Neuroimage 41(3), 735–746 (2008)
4. Li, Y., Willer, C.J., Ding, J., Scheet, P., Abecasis, G.R.: MaCH: using sequence and genotype data to estimate haplotypes and unobserved genotypes. Genet. Epidemiol. 34(8), 816–834 (2010)
5. Liu, J., Pearlson, G., Windemuth, A., Ruano, G., Perrone-Bizzozero, N.I., Calhoun, V.: Combining fmri and snp data to investigate connections between brain function and genetics using parallel ica. Hum. Brain Mapp. 30(1), 241–255 (2009)
6. Marchini, J., Howie, B., Myers, S., McVean, G., Donnelly, P.: A new multipoint method for genome-wide association studies via imputation of genotypes. Nature Genetics 39, 906–913 (2007)
7. Shen, L., Kim, S., Risacher, S.L., Nho, K., Swaminathan, S., West, J.D., Foroud, T., Pankratz, N., Moore, J.H., Sloan, C.D., Huentelman, M.J., Craig, D.W., Dechairo, B.M., Potkin, S.G., Jack Jr., C.R., Weiner, M.W., Saykin, A.J., ADNI: Whole genome association study of brain-wide imaging phenotypes for identifying quantitative trait loci in MCI and AD: A study of the ADNI cohort. Neuroimage 53(3), 1051–1063 (2010)
8. Stein, J.L., Hua, X., Lee, S., Ho, A.J., Leow, A.D., Toga, A.W., Saykin, A.J., Shen, L., Foroud, T., Pankratz, N., Huentelman, M.J., Craig, D.W., Gerber, J.D., Allen, A.N., Corneveaux, J.J., Dechairo, B.M., Potkin, S.G., Weiner, M.W., Thompson, P.: Voxelwise genome-wide association study (vgwas). Neuroimage 53(3), 1160–1174 (2010)
9. Tibshirani, R.: Glmnet, http://www-stat.stanford.edu/~tibs/glmnet-matlab/
10. Vounou, M., Nichols, T.E., Montana, G.: Discovering genetic associations with high-dimensional neuroimaging phenotypes: A sparse reduced-rank regression approach. Neuroimage 53(3), 1147–1159 (2010)
11. Witten, D.M., Tibshirani, R., Hastie, T.: A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. Biostatistics 10(3), 515–534 (2009)
12. Worsley, K.J.: Surfst, http://www.math.mcgill.ca/keith/surfstat
13. Worsley, K.J., Andermann, M., Koulis, T., MacDonald, D., Evans, A.C.: Detecting changes in nonisotropic images. Hum. Brain Mapp. 8(2-3), 98–101 (1999)
14. Zou, H., Hastie, T.: Regularization and variable selection via the elastic net. J. R. Statist. Soc. 67(2), 301–320 (2005)

Euclidean Geodesic Loops on High-Genus Surfaces Applied to the Morphometry of Vestibular Systems

Shi-Qing Xin[1], Ying He[1], Chi-Wing Fu[1], Defeng Wang[2], Shi Lin[2], Winnie C.W. Chu[2], Jack C.Y. Cheng[3], Xianfeng Gu[4], and Lok Ming Lui[5]

[1] School of Computer Engineering, Nanyang Technological University
[2] Department of Imaging and Interventional Radiology, CUHK
[3] Department of Orthopaedics and Traumatology, CUHK
[4] Department of Computer Science, Stony Brook University
[5] Department of Mathematics, The Chinese University of Hong Kong

Abstract. This paper proposes a novel algorithm to extract feature landmarks on the vestibular system (VS), for the analysis of Adolescent Idiopathic Scoliosis (AIS) disease. AIS is a 3-D spinal deformity commonly occurred in adolescent girls with unclear etiology. One popular hypothesis was suggested to be the structural changes in the VS that induce the disturbed balance perception, and further cause the spinal deformity. The morphometry of VS to study the geometric differences between the healthy and AIS groups is of utmost importance. However, the VS is a genus-3 structure situated in the inner ear. The high-genus topology of the surface poses great challenge for shape analysis. In this work, we present a new method to compute exact geodesic loops on the VS. The resultant geodesic loops are in Euclidean metric, thus characterizing the intrinsic geometric properties of the VS based on the real background geometry. This leads to more accurate results than existing methods, such as the hyperbolic Ricci flow method [13]. Furthermore, our method is fully automatic and highly efficient, e.g., one order of magnitude faster than [13]. We applied our algorithm to the VS of normal and AIS subjects. The promising experimental results demonstrate the efficacy of our method and reveal more statistically significant shape difference in the VS between right-thoracic AIS and normal subjects.

1 Introduction

In medical image analysis, surface-based morphometry has been commonly applied for disease analysis. For example, in human brain mapping, neuroscientists are interested in detecting shape changes or abnormalities on brain cortical surfaces for analyzing various brain diseases. Hippocampal surface morphometry has also been an active research area for the analysis of Alzheimer's disease.

Although surface-based shape analysis has been extensively studied, most existing methods can only deal with surfaces with simple topology. For example, both the brain cortical surface and the hippocampal surface are of genus 0. Dealing with high-genus surfaces are apparently much more difficult, due

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 384–392, 2011.

to their complicated topology. One typical example of high-genus anatomical structures is called the vestibular system (VS). The VS is a genus-3 structure situated in the inner ear, which is responsible for detecting head movements and sending postural signals to the brain. The morphometry of VS plays an important role in the analysis of various diseases such as the Adolescent Idiopathic Scoliosis (AIS) disease. AIS is a $3D$ spinal deformity which affects about 4% schoolchildren worldwide. The etiology of AIS is still unclear but believed to be a multi-factorial disease. One popular hypothesis was suggested to be the structural changes in the VS that induce the disturbed balance perception, and further cause the spinal deformity[1,4,9]. Some recent works have also revealed the statistical difference in global morphology of the VS between right-thoracic AIS and normal controls [8,13]. Hence, analyzing the shape of VS is crucial for understanding AIS. However, since the VS has high-genus topology, it poses great challenge for shape analysis, and an effective algorithm is thus needed.

In this work, we proposed a novel algorithm to effectively extract feature landmarks on the VS, for the disease analysis of AIS. The basic idea is to compute exact geodesic loops on the VS. The resultant geodesic loops are in Euclidean metric, thus characterizing the intrinsic geometric properties of the VS based on the real background geometry. This leads to more accurate results than the existing methods, such as the hyperbolic Ricci flow method [13]. Furthermore, our method is fully automatic and highly efficient, e.g., one order of magnitude faster than [13]. To test the effectiveness of our proposed method, we applied the algorithm to the VS of normal and AIS subjects. Experimental results demonstrate the efficacy of our method and reveal more statistically significant shape differences in the VS between right-thoracic AIS and normal subjects.

2 Related Work

Shi et al. [8] proposed to consider the radii of the canals for the global morphology of the VS. The drawback was that the complete geometry of the surface has not been fully analyzed. In [13], Zeng et al. proposed to extract the geodesic spectra of the VS to compare geometric difference between the normal and unhealthy groups. This algorithm requires the computation of uniformization metric, which is computed by discrete Ricci flow or Yamabe flow [5]. Since both curvature flows are highly non-linear PDEs, the computational costs are very high. It is known that embedding surfaces of genus $g \geq 2$ onto the hyperbolic space $\mathbb{H}^2$ could be error prone due to the numerical truncation, the input meshes must come with good quality (i.e., the vertices are distributed uniformly over the surface) and contain small number of vertices (usually a few thousands). Compared to this approach, our method is computationally efficient and numerically stable. Furthermore, the computed geodesic of [13] was based on the hyperbolic metric and hence it was not reflecting the real background geometry of the VS.

In this paper, we suggest computing $2g$ characteristic geodesic loops and use their lengths for the purpose of AIS disease analysis. A discrete geodesic on a polyhedral surface in $\mathbb{R}^3$ is defined to be a polyline that is locally shortest

everywhere [7]. In detail, Mitchell et al. [6] pointed out the requirements of a discrete geodesic l: 1) For each face f, the intersection between f and l must be empty or a line segment; 2) For each edge that l intersects, the entry angle must equal to the departure angle; and 3) For each vertex that l passes through, the two angles passed through by l must be no less than π.

Closed geodesics are also called *geodesic loops*. Though defined on polygonal meshes, discrete geodesics may pass through the face interior, rather than restricted on mesh edges. They are intrinsic to geometry and insensitive to the mesh tessellation and resolution. Therefore, the resultant spectra depend on the actual geometry of the input VS geometry, independent of the triangulation.

To our knowledge, very few algorithms are known for the geodesic loop problem. Wu and Tai [11] proposed the discretized geodesic curvature flow (dGCF) to compute geodesic loops on triangle meshes using a level set formulation. Later, Zhang *et al.* [14] improved dGCF to FGCF (fast geodesic curvature flow) by reducing the problem dimension. Both dGCF and FGCF require numerical solvers and the computation result is sensitive to the user specified parameters (e.g., convergence tolerance). In sharp contrast, our geodesic loop algorithm is parameter-free and does not require any numerical solver. Furthermore, it can deform an arbitrary curve into exact geodesic loop within finite steps. The experimental results show that our method is much faster than dGCF [11] and FGCF [14].

3 Our Algorithm

In this section, we present the algorithm for computing 3 tunnel loops and 3 handle loops for the VS models, which can be paired into $(a_i, b_i), i = 1, 2, 3$. The spectrum determined by their lengths can be used for AIS detection purpose. The exact description of tunnel loops and handle loops is available in [2].

Overview. Given a vestibular system modeled by a triangular mesh of genus $g(= 3)$, we compute the geodesic tunnel and handle loops, which form the homology basis, and then the geodesic spectra with the following steps (see Figure 1):
Step 1: For every mesh vertex v_i, we compute a system of loops: $\mathcal{L}_i = \{l_i^1, l_i^2, ..., l_i^{2g}\}$; and then we obtain $2gn$ loops.
Step 2: Sort the loops of $\mathcal{L}$ by length in the ascending order;
Step 3: Find $2g$ representative loops $l_1, l_2, \cdots, l_{2g}$ from $\mathcal{L}$ such that they form a homology basis;
Step 4: Evolve $l_1, l_2, \cdots, l_{2g}$ into exact closed geodesic loops $\widehat{l_1}, \widehat{l_2}, \cdots, \widehat{l_{2g}}$;
Step 5: Match $\widehat{l_1}, \widehat{l_2}, \cdots, \widehat{l_{2g}}$ into g pairs and compute the geodesic spectra.

Computing a system of loops restricted on mesh edges. Given a closed, oriented surface M of genus g, the system of loops [10] is a set of $2g$ simple loops sharing a common base point. By cutting along a system of loops, M is topologically equivalent to a disk. Erickson and Whittlesey [3] devised an efficient $O(n \log n)$-time algorithm for an 1-skeleton of a surface, i.e., the graph of edges embedded in the polyhedral surface (n is the number of mesh vertices). Given a source vertex s, we compute a system of loops rooted at s by (see Figure 2):

1. Using Dijkstra's algorithm on the embedded graph G, we obtain a shortest path tree T that encodes parent-child relationships rooted at s. Then we set a flag "STOP" for edges on T and a flag "PASS" for other edges.
2. We take each face as a node and define two faces f_1, f_2 to be neighbors if they share a "PASS" edge $e = v_1v_2$. At the same time, we set the weight between f_1 and f_2 to be $\|e\| + d(v_1) + d(v_2)$, where $d(v_i)$ is the shortest distance from s to $v_i, i = 1, 2$. This induces a weighted connected graph G^*.
3. We compute the maximum spanning tree of G^*, say T^*, in $O(n \log n)$ time.
4. There are $2g$ mesh edges that are in neither T nor T^*. For each edge $e = \overline{v_1v_2}$ that belongs to neither T nor T^*, we compute a loop:

$$\Pi_{\overline{v_1v_2}} = \Pi(s, v_1) \bigcup \overline{v_1v_2} \bigcup \Pi(v_2, s),$$

where $\Pi(s, v_1)$ and $\Pi(v_2, s)$ are the shortest paths, from s to v_1 and v_2, respectively, computed by Dijkstra's algorithm.

Identifying handle and tunnel loops. After finding a group of closed curves $\overline{\mathcal{L}}$, serving as generators, from the loop set $\mathcal{L}$, we extract $2g$ curves that can be matched into g pairs. Then, we need to differentiate handle and tunnel loops.

Checking homotopic curves. Among the $2gn$ loops, many are not homotopic to each other. To efficiently check if two curves are not homotopic, we have the following observation: Let $\mathcal{C}_1$ and $\mathcal{C}_2$ be two non-self-intersecting closed curves restricted on mesh edges and they do not intersect. Let $f : V \to \mathbb{R}$ be a smooth function defined on the VS mesh vertices. Then $df : E \to \mathbb{R}$ is a one-form defined on each oriented edge. If the given two curves $\mathcal{C}_1$ and $\mathcal{C}_2$ are homologous, then $\mathcal{C}_1 - \mathcal{C}_2$ bounds a surface patch, say $\partial D = \mathcal{C}_1 - \mathcal{C}_2$. As a result, $\int_{\mathcal{C}_1} df - \int_{\mathcal{C}_2} df = \int_{\mathcal{C}_1 - \mathcal{C}_2} df = \int_{\partial D} df = \int_D ddf = 0$. The last two equations come from the Stoke's theorem and the fact that $d \circ d = 0$. Thus, if $\int_{\mathcal{C}_1} df \neq \int_{\mathcal{C}_2} df$, the two curves are not homologous, which implies that they are not homotopic either.

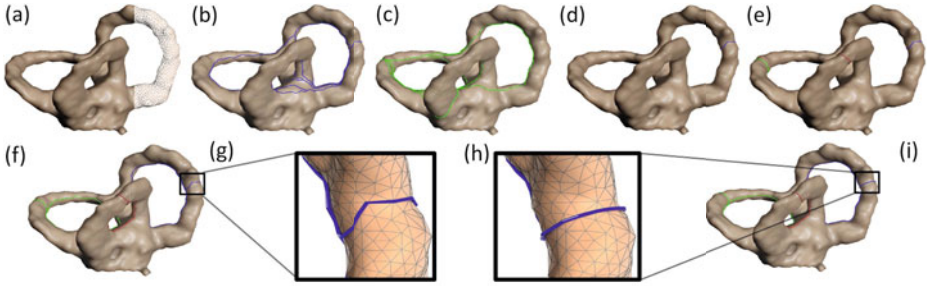

Fig. 1. Algorithm pipeline. Taking a mesh model of a vestibular system as the input (see (a)), we compute a system of loops for each mesh vertex (see (b) and (c) for two examples). As a result, we obtain $2gn$ loops $\mathcal{L}$, where n is the number of mesh vertices and $g = 3$ the genus of VS. In an incremental manner, we compute the optimal homology basis $\mathcal{L}^* \subset \mathcal{L}$ with $2g$ loops by gradually adding the loops from short to long (d-f). Finally, we evolve each curve in $\mathcal{L}^*$ into an exact geodesic loop (i). The close-up views in (g) and (h) reveal the difference between a curved loop and an exact geodesic loop, where the latter has minimal length.

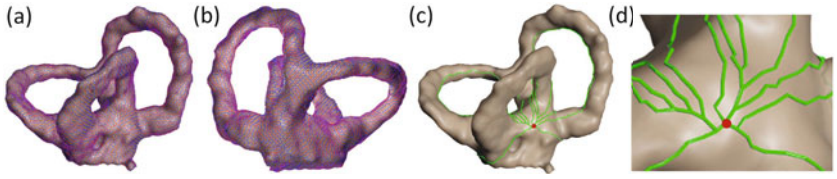

Fig. 2. A system of loop rooted at a given vertex: (a) and (b) show the front and back views of the shortest path tree and the dual graph. (c) shows an example system of loops with $2g$ curves sharing a common base point. (d) shows the close-up view.

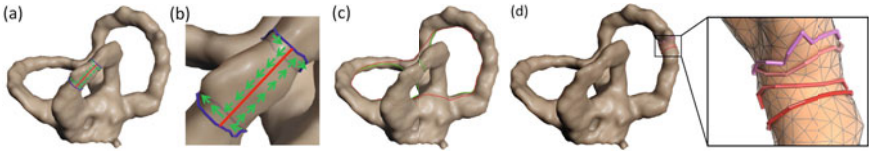

Fig. 3. The two blue loops in (a) are homotopic since one curve can be continuously deformed to the other along the red path (see (b) for the close-up view). The red loop in (c) can be reduced by the three green loops. (d) Our algorithm can deform an arbitrary closed curve (in pink) into an exact geodesic loop (in red) within finite steps. The intermediate results are shown in the close-up view.

The above test can efficiently eliminate a large number of cases where two curves are not homotopic. For curves that pass this test, we use the following approach to test whether they are homotopic. Let $\mathcal{P}(\mathcal{C}_1, \mathcal{C}_2)$ be the shortest path (restricted on mesh edges) between $\mathcal{C}_1$ and $\mathcal{C}_2$ [1] Clearly, if the surface patch generated by $\mathcal{C}_1 \bigcup \mathcal{C}_2 \bigcup \mathcal{P}(\mathcal{C}_1, \mathcal{C}_2)$ is a topological disk, then $\mathcal{C}_1$ and $\mathcal{C}_2$ are homotopic, and thus, we can continuously transform one curve into the other.

Curve reduction. For two homotopic closed curves, we say the longer loop can be reduced by the shorter one; see Figure 3 (a). If the composite loop $l_a \bigcup l_b \bigcup \mathcal{P}(l_a, l_b)$ is homotopic to a loop l_c and $\|l_c\| < \|l_a\|$, we say that l_a can be reduced into l_c by l_b, where $\mathcal{P}(l_a, l_b)$ is the shortest path connecting l_a and l_b. Generally, we can try to reduce a new loop l_i by existing loops $l_j (1 \leq j < i)$ in order. The goal of curve reduction is to transform the new loop such that it is not homotopic to any one of the existing loops and as short as possible. Figure 3 (c) shows an example where the red loop can be completely reduced by the three green loops (in any reduction order).

Generators and homology basis. We first sort the loops in $\mathcal{L}$ by length, extract the shortest one, say l_1, and append it into an empty set $\overline{\mathcal{L}}$. For the second shortest loop l_2, we try to reduce it by the existing loops in $\overline{\mathcal{L}}$ (l_1 in this case). If l_2 cannot be reduced by any loop in $\overline{\mathcal{L}}$, we append it into $\overline{\mathcal{L}}$. In an incremental style, we can extract a group of closed curves from $\mathcal{L}$ that serve as generators that has $2g$ closed curves. We observe that for geometrically simple models, the resulting

[1] The "multi-source-all-destination" Dijkstra's algorithm is applied to find the shortest path between $\mathcal{C}_1$ and $\mathcal{C}_2$ (see Figure 3 (a)).

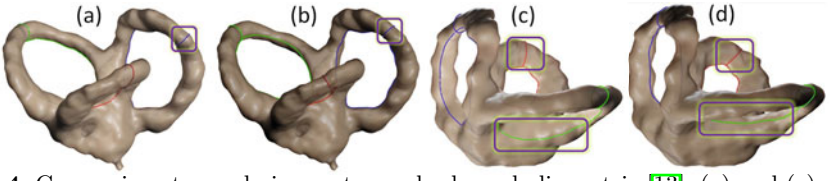

Fig. 4. Comparison to geodesic spectra under hyperbolic metric [13]. (a) and (c) show result of [13]. (b) and (d) show our result. The boxes highlight the difference.

generators computed in such a greedy scheme are exactly an optimal homology basis with g pairs that satisfy (1) each pair of loops has a common point; (2) loops in different pairs have no common points; and (3) the total length is minimum in the 1-skeleton of the input mesh. For complicated geometric models, universal covering space [13] should be considered.

Handle and tunnel loops. A closed oriented surface M embedded in $\mathbb{R}^3$ partitions $\mathbb{R}^3$ into interior sub-space I and exterior sub-space O. Tunnel and handle loops are special homology basis that span a surface in the complement space of the input surface. Specifically, a handle loop (tunnel loop) is null-homologous in I (O) but not in O (I). So we can differentiate handles and tunnels by shrinking the loops a little bit and check if they are inside the mesh.

Deforming loops by shortening the length. Given a closed curve $\mathcal{C}$ restricted on mesh edges, we will deform it into an exact geodesic loop (generally not restricted on mesh edges) that is locally shortest on the polyhedral surface. The two key operations are 1) computing a shortest loop restricted on a closed face sequence and 2) updating the edge sequence. Note that the deformation proceeds on the surface (including the face interior), rather than only on mesh edges. The detailed algorithm is as follows:

1. Find the closed edge sequence Γ that contains $\mathcal{C}$ and build a loop set $\bar{\mathcal{L}} = \emptyset$;
2. Cut Γ along one edge $e = v_1v_2 \in \Gamma$ into an open face sequence Γ' and unfold it onto a 2D plane;
3. Since Γ' encloses a simple polygon $\mathcal{G}$, we can find the shortest path, inside $\mathcal{G}$, between v_1 and its image point v_1', and then add the path into $\bar{\mathcal{L}}$. Similarly, we add the shortest from v_2 into $\bar{\mathcal{L}}$, as well as the image point of v_2 .
4. By traversing all edges in Γ, we can add more loops into $\bar{\mathcal{L}}$;
5. Replace the loop $\mathcal{C}$ with the shortest loop $\mathcal{C}'$ in $\bar{\mathcal{L}}$. (Note: $\mathcal{C}'$ is shorter than $\mathcal{C}$.)
6. If $\mathcal{C}'$ passes through some vertex v in Γ, then update Γ from one side of v to the other side to test whether the new edge sequence can give a shorter loop, and go to Step 2; otherwise, output the resultant loop.

It is clear that the resultant loop is a closed polyline through an alternative sequence of edges and vertices, and it can be proved that the resultant loop is locally shortest everywhere and shorter than the initial loop. The algorithm is highly efficient and outperforms previous algorithms; see [12] for detail.

4 Experimental Results

Image acquisition and Segmentation of VS: Experiments were done on 11 girls with right-thoracic AIS (mean age 15 years old with variance 1.7 years old); (mean Cobb's angle 27.27 degrees with variance 15.62 degrees) and 11 age-matched healthy girls. VS surface meshes were extracted from the T2-weighted MRI scanning of the inner ears using the 1.5T MR Scanner (Sonata, Siemens, Erlangen, Germany) with a quadrature head coil.

Computation of Exact Geodesic Loops: Our proposed algorithm can deform any closed loop to an exact geodesic loop within finite number of steps; we experimented it with a large number of 3D models and the statistical results showed that the complexity is roughly linear to the number of vertices in-between the initial and resultant loops. For the given VS meshes with roughly 13K triangles, our entire algorithm (including computation of the systems of loops, identification of $2g$ homology basis, and deforming them into geodesic loops) takes less than 1 second on a workstation with an Xeon 2.66GHz CPU and 4GB RAM. The hyperbolic Ricci flow approach [13], in sharp contrast, is a highly non-linear PDE with high computational cost. Our results show that our method is at least one magnitude order of faster than [13] (see also Figure 4).

Exact Geodesic Loops for the Morphometry of VS: We have computed the exact geodesic loops on the VS of the healthy and AIS groups. Figure 5 shows some examples, where the upper row are some normal VS models while the lower row are some AIS VS models. The statistics of the geodesic loops (a_i, b_i)'s on the VS from each groups are also as shown in Table 1. The statistical difference in the geodesic loops between groups are also evaluated using t tests. Basically, a larger a_i means the canal is longer, whereas a larger b_i means the canal is thicker. Notice that the exact geodesic loops are computed based on the real background (Euclidean) metric. It is different from [13] in which the lengths of features are computed based on the hyperbolic metric. Hence, our algorithm describes the actual geometry of the VS. From the statistics, AIS tends to have larger a_i's (hence longer canals) and smaller b_i's (hence thinner canals) than the normal groups. In particular, AIS tends to have larger a_1 and a_3 with very high statistical significance (both with P-value < 0.02). This implies the lateral and superior canals are generally longer for AIS subjects. The ratio a_1/b_1 and a_3/b_3 also tend to be larger in the AIS group with high statistical significance ($P = 0.0083$ and $P=0.0321$). It means the conformal modules of the lateral and superior canals are significantly different between groups. This implies once again that the shape in the lateral and superior canals tends to be significantly different between the two groups. Compared to [13], our new algorithm can detect more geometric differences between groups with high statistical significance. Also, the P-value for the detected shape differences is much smaller than that in [13]. It means our extracted feature curves can conclude shape differences with much higher statistical significance. We also ran a multiple-comparison test. The results in the test agree with our findings that the shape in the lateral and superior canals tends to be different more significantly between the AIS and healthy groups.

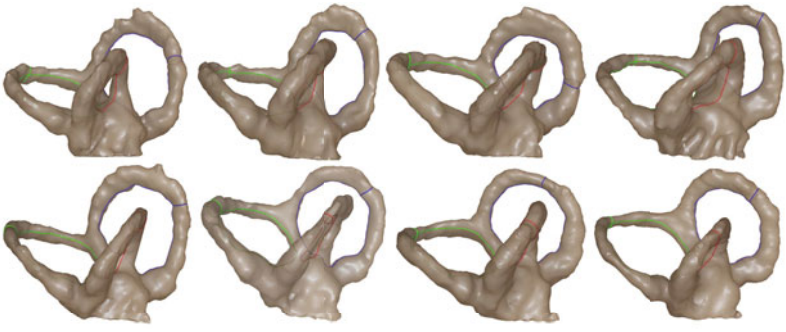

Fig. 5. More results. Row 1: normal subjects; Row 2: AIS subjects

Table 1. Statistics on the exact geodesic loops between the normal and AIS groups

Combination	Mean(Normal)	Mean(AIS)	P-value
a_1	8.7835 ± 1.2119	9.8544 ± 0.6469	**0.0177**
a_2	16.4447 ± 1.3652	17.1932 ± 0.9952	0.1573
a_3	14.8008 ± 1.2991	16.2392 ± 1.1075	**0.0112**
b_1	3.8099 ± 0.3643	3.4912 ± 0.5146	0.1092
b_2	3.0688 ± 0.4108	2.8679 ± 0.3759	0.2455
b_3	3.4521 ± 0.4980	3.2118 ± 0.3699	0.2135
a_1/b_1	2.3185 ± 0.3543	2.8882 ± 0.5390	**0.0083**
a_2/b_2	5.4207 ± 0.6631	6.1018 ± 0.9580	0.0668
a_3/b_3	4.3743 ± 0.7917	5.1234 ± 0.7322	**0.0321**
$a_1 + a_2 + a_3$	40.0290 ± 2.7322	43.2868 ± 2.2105	**0.0060**
$b_1 + b_2 + b_3$	10.3309 ± 0.8356	9.5710 ± 0.8810	0.0510
$a_1 + b_1 + a_2 + b_2 + a_3 + b_3$	50.3599 ± 3.2353	52.8578 ± 2.6241	0.0606

5 Conclusion

This work proposes an effective method that extracts exact geodesic loops on high-genus surfaces to support analysis of feature landmarks on the VS. The extracted loops are computed using the Euclidean metric, and hence can be used as geometric features to detect shape differences between genus-3 VS. Our proposed method is fully automatic and highly efficient. We applied it to the VS of normal and right-thoracic AIS subjects; experimental results reveal geometric differences in the two groups with high statistical significance.

Acknowledgements. This project was partially supported by NRF2008IDM-IDM004-006, AcRF 69/07 and HK RGC grant (ID: 411910).

References

1. Byl, N., Gray, J.: Complex balance reactions in different sensory conditions: adolescents with and without idiopathic scoliosis. Journal of Orthopaedic Research 11(2), 215–227 (1993)

2. Dey, T.K., Li, K., Sun, J.: Computing handle and tunnel loops with knot linking. In: Computer Aided Design, vol. 41, pp. 730–738 (2009)
3. Erickson, J., Whittlesey, K.: Greedy optimal homotopy and homology generators. In: Proc. of ACM-SIAM Symposium on Discrete Algorithms, pp. 1038–1046 (2005)
4. Guo, X., Chau, W., Hui-Chan, C., Cheung, C., Tsang, W., Cheng, J.: Balance control in adolescents with idiopathic scoliosis and disturbed somatosensory function. Spine 31, 437–440 (2006)
5. Jin, M., Luo, F., Gu, X.D.: Computing general geometric structures on surfaces using ricci flow. In: Computer-Aided Design, vol. 39(8), pp. 663–675 (2007)
6. Mitchell, J.S.B., Mount, D.M., Papadimitriou, C.H.: The discrete geodesic problem. SIAM J. Comput. 16, 647–668 (1987)
7. Sharir, A., Baltsan, A.: On shortest paths amidst convex polyhedra. In: Proc. of Symposium on Computational Geometry, SCG 1986, pp. 193–206 (1986)
8. Shi, L., Wang, D., Chu, W., Heng, P., Burwell, R., Cheng, J.: Statistical morphometry of the vestibular system in adolescent idiopathic scoliosis. In: Annual World Congress for Brain Mapping and Image Guided Therapy, Harvard Med. Sch. (2009)
9. Simoneau, M., Lamothe, V., Hutin, E., Mercier, P., Teasdale, N., Blouin, J.: Evidence for cognitive vestibular integration impairment in idiopathic scoliosis patients. BMC Neuroscience 10(102) (2009)
10. de Verdière, E.C., Lazarus, F.: Optimal system of loops on an orientable surface. In: Proc. of the Symposium on Foundations of Computer Science, pp. 627–636 (2002)
11. Wu, C., Tai, X.: A level set formulation of geodesic curvature flow on simplicial surfaces. IEEE Tran. on Vis. and Computer Graphics (TVCG) 16, 647–662 (2010)
12. Xin, S.Q., He, Y., Fu, C.W.: Efficiently computing exact geodesic loops within finite steps. IEEE Tran. on Vis. and Computer Graphics, TVCG (accepted, 2011)
13. Zeng, W., Lui, L., Shi, L., Wang, D., Chu, W., Cheng, J., Hua, J., Yau, S., Gu, X.: Shape analysis of vestibular systems in adolescent idiopathic scoliosis using geodesic spectra. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. 6363, pp. 538–546. Springer, Heidelberg (2010)
14. Zhang, J., Wu, C., Cai, J., Zheng, J., Tai, X.: Mesh snapping: Robust interactive mesh cutting using fast geodesic curvature flow. Computer Graphics Forum (In Proc. of Eurographics 2010) 29(2), 517–526 (2010)

A Statistical Model of Shape and Bone Mineral Density Distribution of the Proximal Femur for Fracture Risk Assessment

Tristan Whitmarsh[1], Karl D. Fritscher[2], Ludovic Humbert[1], Luís Miguel Del Rio Barquero[3], Tobias Roth[4], Christian Kammerlander[4], Michael Blauth[4], Rainer Schubert[2], and Alejandro F. Frangi[1]

[1] Center for Computational Imaging & Simulation Technologies in Biomedicine (CISTIB), Universitat Pompeu Fabra (UPF) and CIBER-BBN, Barcelona, Spain
[2] Institute for Biomedical Image Analysis (IBIA), University for Health Sciences, Medical Informatics and Technology (UMIT), Hall in Tirol, Austria
[3] CETIR Centre Mèdic, Barcelona, Spain
[4] Department for Trauma Surgery and Sports Medicine, Innsbruck Medical University, Innsbruck, Austria

Abstract. This work presents a statistical model of both the shape and Bone Mineral Density (BMD) distribution of the proximal femur for fracture risk assessment. The shape and density model was built from a dataset of Quantitative Computed Tomography scans of fracture patients and a control group. Principal Component Analysis and Horn's parallel analysis were used to reduce the dimensionality of the shape and density model to the main modes of variation. The input data was then used to analyze the model parameters for the optimal separation between the fracture and control group. Feature selection using the Fisher criterion determined the parameters with the best class separation, which were used in Fisher Linear Discriminant Analysis to find the direction in the parameter space that best separates the fracture and control group. This resulted in a Fisher criterion value of 6.70, while analyzing the Dual-energy X-ray Absorptiometry derived femur neck areal BMD of the same subjects resulted in a Fisher criterion value of 0.98. This indicates that a fracture risk estimation approach based on the presented model might improve upon the current standard clinical practice.

1 Introduction

With the rapid advancement of medical imaging technologies, the development of image analysis methods has progressed in accordance. Specifically statistical models have received a great deal of interest and have been used in a wide range of fields for the diagnosis of diseases and detection of symptoms.

Research on statistical models for orthopaedics have mainly focused on modeling the bone shape for reconstructions, which are to be used in therapy planning, and not so much for analysis and diagnosis. In Schuler *et al.* [1], a method is presented to estimate the fracture load using statistical appearance models. This

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 393–400, 2011.

gives an estimation of the femur strength, which could potentially be used to derive a fracture risk. In current clinical practice, Dual-energy X-ray Absorptiometry (DXA) derived femur neck areal Bone Mineral Density (BMD) is used. This measure however is limited by its two-dimensionality while it has been shown that the 3D distribution of the BMD greatly influences the fracture risk [2]. In Li *et al.* [3], a statistical model was constructed of the proximal femur by image registration and the principal components of this model were evaluated for their discriminatory power. This work however only uses statistical information about the BMD distribution, while it has been shown that also the shape has an influence on fracture risk [4].

In this work, we propose a statistical model of the BMD distribution as well as the shape, whereby, by analyzing the input data consisting of Quantitative Computed Tomography (QCT) scans of fracture patients and a control group, the variation that influences the fracture risk most is determined. This in turn can be used to derive a fracture risk for an unseen patient. The underlying shape and density model is constructed using an intensity based registration process on the QCT scans. Principal Component Analysis (PCA) results in a dimensionality reduction whereby Horn's parallel analysis [5] determines the number of modes of variation to be used for further analysis. As a result, each subject in the model can be said to have a specific set of shape and density model parameters. Using the Fisher criterion, the parameters that best separate the fracture and non-fracture group were determined. Over these parameter values, Fisher Linear Discriminant Analysis (FLDA) was applied, which determined the direction in the feature space that best separates the two classes. The resulting variation is analyzed and compared with previous findings on the shape and BMD distribution with respect to the fracture risk. Finally, the class separation resulting from FLDA was compared to the class separation using the femur neck areal BMD of the same subjects, which is the measurement used in clinical routine to assess the fracture risk.

2 Method

2.1 Registrations

The model construction depends on accurate intensity based registrations of the QCT volumes. In order to prevent the pelvis area from resulting in misalignments in the registration process, the pelvis area is semi automatically removed from the volumes. In addition, a thresholding is applied, which removes the soft tissue structures that can negatively affect the registration process.

Figure 1 depicts the model construction pipeline. First, a reference volume is chosen based on its regular shape and BMD quality. In this volume, the bone is manually segmented and a regular mesh is constructed from this segmentation. All volumes are subsequently registered to this reference volume by an intensity based similarity registration. This is followed by a multi-scale intensity based Thin Plate Spline (TPS) registration using the mesh vertices as the control points on the reference image. In order to reduce the computation load of

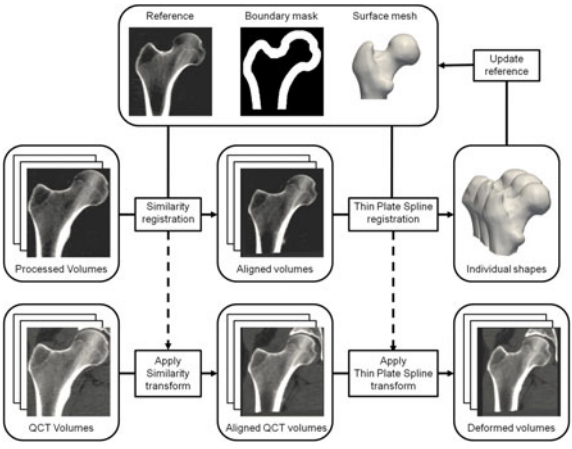

Fig. 1. Model construction pipeline

the TPS registration, the registration is restricted to the region of interest by specifying a mask of the bone boundary in the reference volume (Figure 1). The TPS registrations result in the corresponding landmark locations on the target volumes and thus provides the surface mesh for all similarity aligned volumes. All similarity aligned meshes are subsequently scaled to their original size using the uniform scale value resulting from the similarity transform.

To remove some of the bias from the reference selection, the reference volume (and mask) are deformed to the average of the similarity aligned meshes using a TPS transformation defined by the vertices. In a second iteration, the volumes are then registered onto this updated reference while using the average shape for the TPS registrations.

2.2 Shape and Density Model

The registrations results in the set of aligned patient specific surface meshes with a vertex correspondence between them. PCA is then applied to the vertices and a new shape $\mathbf{s}$ can be expressed as the average shape $\overline{\mathbf{s}}$ and a linear combination of the first m eigenvectors corresponding to the main modes of variation: $\mathbf{s} = \overline{\mathbf{s}} + \sum_{i=1}^{m} \mathbf{p}_i a_i$. Here $\mathbf{p}_i$ is the i-th eigenvector resulting from the singular value decomposition of the covariance matrix and a_i the corresponding scalar coefficient referred to here as the shape model parameter.

To model the BMD distribution, a final iteration of the registration process is performed to deform all volumes to the same average bone shape, whereafter the resulting similarity and TPS transformation is applied to the unprocessed QCT volume for each subject. This results in shape normalized volumes with a voxel correspondence between them. Since the TPS transformation is defined by landmarks on the bone surface only, the TPS interpolation preserves the internal BMD distribution. PCA is then applied to the voxel densities inside the bone so that a new volume $\mathbf{v}$ can be expressed as the average volume $\overline{\mathbf{v}}$ and the

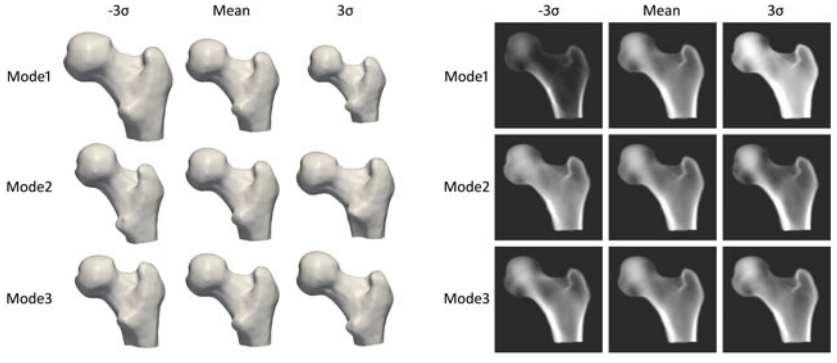

Fig. 2. The mean and the first three modes of variation of the shape model (left) and projections of the density model (right), varying 3 standard deviations (σ)

first n eigenvectors: $\mathbf{v} = \overline{\mathbf{v}} + \sum_{i=1}^{n} \mathbf{q}_i b_i$. Here $\mathbf{q}_i$ is the i-th eigenvector and b_i the corresponding density model parameter. To determine the number of shape model parameters m and density model parameters n to retain, Horn's parallel analysis was used.

2.3 Fisher Linear Discriminant Analysis

After the construction of the shape and density model, the model parameter vectors of the femora in both fracture and control group are used as input for FLDA in order to determine the vector in the feature space that best separates these two groups.

Not all parameters resulting from the dimensionality reduction will contribute to a good class separation. Therefore, first a feature selection based on the Fisher criterion is performed. The Fisher criterion for feature k is defined as: $F(k) = \frac{\mathbf{S}_B^k}{\mathbf{S}_W^k}$, where $\mathbf{S}_B^k$ and $\mathbf{S}_W^k$ are the k-th diagonal element of $\mathbf{S}_B$ and $\mathbf{S}_W$, respectively. Here $\mathbf{S}_B$ and $\mathbf{S}_W$ are the between-class and within-class scatter matrices and are defined as $\mathbf{S}_B = (\mathbf{m}_1 - \overline{\mathbf{m}})(\mathbf{m}_1 - \overline{\mathbf{m}})^T + (\mathbf{m}_2 - \overline{\mathbf{m}})(\mathbf{m}_2 - \overline{\mathbf{m}})^T$ and $\mathbf{S}_W = \sum_{\mathbf{x} \in C_1} (\mathbf{x} - \mathbf{m}_1)(\mathbf{x} - \mathbf{m}_1)^T + \sum_{\mathbf{x} \in C_2} (\mathbf{x} - \mathbf{m}_2)(\mathbf{x} - \mathbf{m}_2)^T$. C_1 and C_2 are the feature sets of the fracture and non-fracture group, $\mathbf{m}_1$ and $\mathbf{m}_2$ are their respective means and $\overline{\mathbf{m}}$ the overall mean. Each feature is analyzed for their class separation by calculating this Fisher criterion. The features are subsequently sorted and the features beyond the converging point (nearing zero) are discarded.

FLDA is applied to the remaining features to determine the vector that best separates the fracture and control group. FLDA finds the vector $\mathbf{w}$ that maximizes the total scatter while minimizing the within scatter of the two classes. This is achieved by maximizing the Rayleigh quotient: $J(\mathbf{w}) = \frac{\mathbf{w}^T S_B \mathbf{w}}{\mathbf{w}^T S_W \mathbf{w}}$.

The resulting direction in the feature space is visualized by generating an instance of the model at the two extremes of this vector corresponding to 3 and -3 standard deviation from mean. Three standard deviations are chosen to

constrain them to statistically valid instances. For both this high and low fracture risk extreme, the density model instance is deformed towards the shape model instance, again using a TPS transformation defined by the shape vertices. These visualizations allow for the assessment of the combined variations of shape and BMD distribution that influences the fracture risk, and these are subsequently compared to current knowledge of hip fracture predictors.

2.4 Data

A dataset of CT scans of the pelvis area of 58 fracture patients was collected at the Universitätsklinik für Radiodiagnostik in Innsbruck using the GE LightSpeed VCT Multi Slice CT device (GE Healthcare, Chalfont St. Giles, UK). The patients were all female with an average age of 79 ± 10 years and all had suffered a proximal femur fracture. From this set the contralateral un-fractured femurs were taken, which can be justified by research on proximal femur symmetry [6], and this set defined the high fracture risk group. In addition, the European forearm phantom [7] was included into every scan acquisition to convert the CT scans to QCT volumes.

Besides the fracture group, a group of 58 (all female) patients was collected at the CETIR Medical Center (Barcelona, Spain) using the Philips Gemini GXL 16 system (Philips Healthcare, Best, The Netherlands). These patients had a lower average age of 55 ± 12 years, to represent a control group with a low fracture risk and a normal BMD distribution. The CT scans were calibrated using the Mindways calibration phantom (Mindways Software Inc., Austin, TX, United States).

Different studies have shown that the scan device has little influence on the density calibration [8]. However, the European forearm phantom calibrates the volumes to a hydroxyapatite (HA) density, whereas the Mindways phantom is constructed of K_2HPO_4. Previous research shows that one calibration material is highly correlated to another and can be converted using a linear transformation [9]. To get the conversion formula for the phantoms in this work, both phantoms were scanned together using the Siemens scanner. This resulted in the following conversion formula: $y = 1.1053x - 17.788$ ($R^2 = 0.9998$). This way, the scans with the European forearm phantom were calibrated to relate to K_2HPO_4 densities.

Femur neck BMD measurements were performed on the DXA scans of all subjects. For the fracture group patients, the Hologic Discovery W bone densitometry system (Hologic Inc., Bedford, MA, United States) was used while for the control group the scans were performed using the GE Healthcare's Lunar iDXA scanner (GE Healthcare, Chalfont St. Giles, UK). The femur neck areal BMD values were subsequently converted to standardized BMD (sBMD) [10].

3 Results

Figure 2 shows the first three modes of variation of the resulting shape and density model. Horn's parallel analysis resulted in 10 shape model parameters

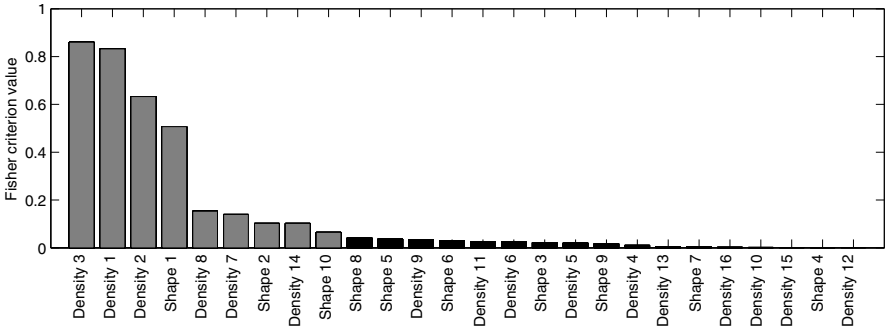

Fig. 3. A bar graph of the features with the corresponding Fisher criterion value, sorted in descending order. The gray bars indicate the features selected for further analysis while the features with the black bars were discarded.

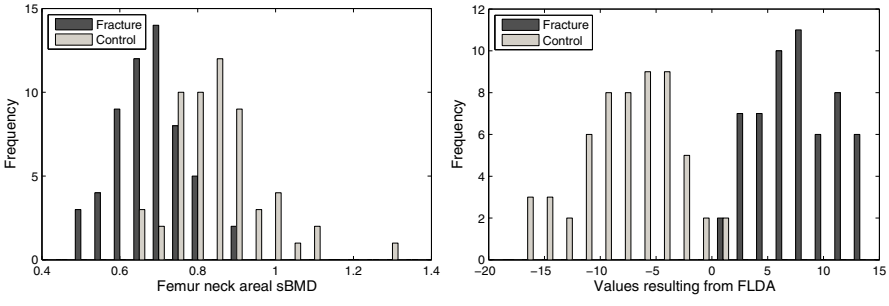

Fig. 4. The histograms corresponding to the femur neck areal sBMD (left) and the values resulting from FLDA (right) for both fracture and control group.

and 16 density model parameters to form the feature space. The feature selection method using the Fisher criterion resulted in 3 shape model parameters and 6 density model parameters to be used, as shown in Figure 3. Applying FLDA to the values of these parameters of the subjects in the model resulted in a class separation with a Fisher criterion value of 6.70. In comparison, the femur neck areal sBMD values separates the fracture and control group with a Fisher criterion value of 0.98.

For each subject in the model, the corresponding shape and density model parameters can be projected onto the vector resulting from FLDA. This represents a single value measure for the fracture risk. In Figure 4, the histogram corresponding to these values are shown for both fracture and control group. The same histogram is given for the femur neck areal sBMD values as a comparison.

In Figure 5, the high and low fracture risk extremes with respect to the mean are visualized, which show the difference in shape and BMD distribution between high and low fracture risk femurs as described by the model.

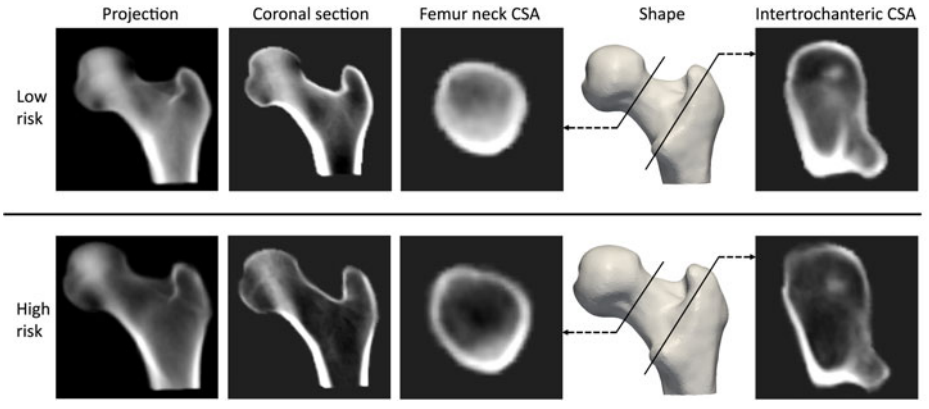

Fig. 5. Model instances at the low fracture risk (top) and high fracture risk (bottom) extreme. From left to right: the projection, the coronal cross section, the femur neck cross-sectional area (CSA), the shape model instance and the intertrochanteric CSA.

4 Discussion and Conclusions

This work presents a shape and density model of the proximal femur for fracture risk estimation. A dimensionality reduction and feature selection method is used to extract the shape and density model parameters that best separate the fracture and control group. FLDA is subsequently used to find the direction in this parameter space of greatest separation.

In Figure 5 it can be seen that the volumetric BMD and the cortical thickness is one of the main discriminators between fracture and non-fracture femora according to the presented model, which is in accordance with recent findings [2]. The visualization of the femur neck and intertrochanteric cross-sectional area (CSA) also reflects previous observations where a significantly smaller cortical CSA, and a larger trabecular CSA is associated with a higher fracture risk [4]. Regarding the shape, a larger hip axis length is associated with a higher fracture incidence [4], which is also reflected in the presented model. The neck-shaft angle has been shown to be significantly larger for fracture patients [4], and this is captured by the model through the second shape model parameter (Figure 2 and 3). These similarities indicate that the presented model captures the differences in shape and BMD distribution between high and low fracture risk subjects and can thus be used in a computer aided diagnosis system.

The Fisher criterion value resulting from FLDA shows that the proposed model better separates the fracture and control group than DXA derived femur neck areal sBMD with a value of 6.70 as opposed to 0.98. The improved class separation can also be seen in the histograms of Figure 4. This indicates that a model-based fracture risk estimation approach might improve upon the current standard clinical practice.

Acknowledgments. The research leading to these results has received funding from the ERDF Operational Programme of Catalonia 2007-2013, through the 3D-FemOs project (exp.VALTEC09-02-0012) and the EU FP7 programme under grant agreement nr 269909 (MySPINE project). A.F. Frangi is partially funded by the ICREA-Academia programme and T. Whitmarsh is supported by the Instituto de Salud Carlos III through PFIS predoctoral grant (FI09/00757).

References

1. Schuler, B., Fritscher, K.D., Kuhn, V., Eckstein, F., Link, T.M., Schubert, R.: Assessment of the individual fracture risk of the proximal femur by using statistical appearance models. Medical Physics 37(6), 2560–2571 (2010)
2. Bousson, V., Adams, J., Engelke, K., Aout, M., Cohen-Solal, M., Bergot, C., Haguenauer, D., Goldberg, D., Champion, K., Aksouh, R., Vicaut, E., Laredo, J.: In vivo discrimination of hip fracture with quantitative computed tomography: Results from the prospective European Femur Fracture Study (EFFECT). Journal of Bone and Mineral Research 26(4), 881–893 (2010)
3. Li, W., Kornak, J., Harris, T., Lu, Y., Cheng, X., Lang, T.: Hip fracture risk estimation based on principal component analysis of QCT atlas: a preliminary study. In: Medical Imaging 2009: Biomedical Applications in Molecular, Structural, and Functional Imaging, p. 72621M. SPIE, San Jose (2009)
4. Ito, M., Wakao, N., Hida, T., Matsui, Y., Abe, Y., Aoyagi, K., Uetani, M., Harada, A.: Analysis of hip geometry by clinical CT for the assessment of hip fracture risk in elderly Japanese women. Bone 46(2), 453–457 (2010)
5. Horn, J.: A rationale and test for the number of factors in factor analysis. Psychometrika 30(2), 179–185 (1965)
6. Faulkner, K., Genant, H., McClung, M.: Bilateral comparison of femoral bone density and hip axis length from single and fan beam DXA scans. Calcified Tissue International 56, 26–31 (1995)
7. Ruegsegger, P., Kalender, W.A.: A phantom for standardization and quality control in peripheral bone measurements by PQCT and DXA. Physics in Medicine and Biology 38(12), 1963–1970 (1993)
8. Fujii, Y., Tsunenari, T., Tsutsumi, M., Miyauchi, A., Yamada, H., Fukase, M., Yoshimoto, Y., Okuno, Y., Kusakabe, H., Miyoshi, K., Fukunaga, M., Morita, R., Fujita, T.: Quantitative computed tomography: Comparison of two calibration phantoms. Journal of Bone and Mineral Metabolism 6, 17–20 (1988)
9. Suzuki, S., Yamamuro, T., Okumura, H., Yamamoto, I.: Quantitative computed tomography: comparative study using different scanners with two calibration phantoms. British Journal of Radiology 64(767), 1001–1006 (1991)
10. Genant, H.K., Grampp, S., Glüer, C.C., Faulkner, K.G., Jergas, M., Engelke, K., Hagiwara, S., van Kuijk, C.: Universal standardization for dual X-ray absorptiometry: Patient and phantom cross-calibration results. Journal of Bone and Mineral Research 9(10), 1503–1514 (1994)

Estimation of Smooth Growth Trajectories with Controlled Acceleration from Time Series Shape Data

James Fishbaugh, Stanley Durrleman, and Guido Gerig

Scientific Computing and Imaging Institute
University of Utah
Salt Lake City, Utah

Abstract. Longitudinal shape analysis often relies on the estimation of a realistic continuous growth scenario from data sparsely distributed in time. In this paper, we propose a new type of growth model parameterized by acceleration, whereas standard methods typically control the velocity. This mimics the behavior of biological tissue as a mechanical system driven by external forces. The growth trajectories are estimated as smooth flows of deformations, which are twice differentiable. This differs from piecewise geodesic regression, for which the velocity may be discontinuous. We evaluate our approach on a set of anatomical structures of the *same* subject, scanned 16 times between 4 and 8 years of age. We show our acceleration based method estimates smooth growth, demonstrating improved regularity compared to piecewise geodesic regression. Leave-several-out experiments show that our method is robust to missing observations, as well as being less sensitive to noise, and is therefore more likely to capture the underlying biological growth.

1 Introduction

The study of time dependent shapes is an emerging field in Computational Anatomy, with potential applications to early brain development, aging studies, or the analysis of evolving pathologic structures. As longitudinal data becomes more widely available, the need for computer models of anatomical evolution becomes increasingly important. Two approaches have been followed so far: the first consists in computing a realistic growth scenario from cross-sectional time-series data, like in [4,10,6,3]. The second approach involves estimating several individual growth trajectories and combining them with a framework for 4D registration between growth trajectories or 4D atlas construction, to statistically analyze the growth variability within a population, like in [13,8,14,7,9].

In any case, the methods rely greatly on the estimation of growth models from time series data, which are sparsely distributed in time. Growth models provide a tool to generate shapes at any instant in time (within the interval defined by the data), offering us the opportunity to continuously measure shape properties. This is in contrast to using sparse measurements such as volume or circumference

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 401–408, 2011.

for 1D regression absent the shape information. The problem can be stated as "temporal shape regression" and can be solved by purely descriptive statistical methods like the extension of kernel regression to Riemannian manifolds [4], or by generative statistical models which define a parameterized family of realistic growth models and the one which best fits the actual data is estimated based on a regularized least-square criterion [8,7,2]. We favor this last approach, since it makes explicit the assumptions which drive the estimation of growth trajectories and therefore enables the inclusion of realistic biological priors to constrain the estimation.

The growth model in [7] is based on a continuous flow of diffeomorphisms, with piecewise geodesics interpolating between shapes. This method estimates continuous non-linear growth between shapes, but does not guarantee differentiable growth as the speed of evolution is discontinuous at observation time-points. Our work is motivated by the assumption that the evolution of biological tissue is inherently *smooth* in time. If we consider the growth of biological tissue as a mechanical system driven by external forces, then the evolution of any particle on an anatomical surface is continuous with continuous derivative and therefore does not change direction *instantaneously*, as observed in the growth model estimated from [7].

Temporal smoothness can be enforced via smooth interpolation between 3D deformations estimated at discrete time-points, using B-splines or polynomial interpolation as in [11,1]. However, these approaches are not based on the inference of a generic growth model, which captures the dynamics of the shape changes over time.

Based on these considerations, we propose a new growth model parameterized by acceleration, rather than velocity as in the large deformation setting of [12]. The estimated acceleration could be considered an indication of the forces which drive the growth of the anatomical structures. From this parameterization, we gain one order of differentiability and guarantee that shape evolution is smooth in both space and time. We further deviate from the large deformations framework by introducing a new regularization term which accounts for the total amount of acceleration. As a consequence, our model does not constrain the flow of deformation between shapes to be geodesic, or close to a geodesic path. By contrast, the approach in [16] estimates twice differentiable trajectories as random perturbations of geodesic paths.

The evaluation of our new methodology on real anatomical surfaces reveals the differences between our approach and piecewise geodesic regression. Our regression yields a twice differentiable evolution with improved regularity, thus discarding more noise from the data to fit a more realistic growth trajectory. Also, we demonstrate that volume measurements taken out of our 3D shape regression are compatible with a 1D regression of these measurements, whereas piecewise geodesic regression appears to overfit. Lastly, we show via leave-several-out experiments that our method better interpolates between data and is therefore more robust to missing observations. This suggests a greater ability to capture the underlying growth of the anatomical structures.

2 Shape Regression Parameterized by Acceleration

The problem of longitudinal shape regression involves inferring a continuous shape evolution from a discrete set of shapes S_{t_i} observed at time t_i. Shape evolution is modeled as the continuous deformation of a baseline shape S_0, formally defined as $R_t = \phi_t(S_0)$ where R_t corresponds to S_0 having undergone the transformation ϕ_t with t varying continuously within the time interval. The time-varying deformation ϕ_t is a general transformation from $\mathbb{R}^N$ to $\mathbb{R}^N$ with $\phi_0(S_0) = S_0$. The baseline shape is deformed over time to closely match the observed shapes ($R_{t_i} \sim S_{t_i}$) while the rigidity of the deformation is controlled via a regularity term. This leads to a variational problem in the form of a trade off between fidelity to data and regularity. For measuring shape similarity, we follow the work of [15], modeling shapes as currents.

We define the acceleration field $a(x,t)$ at point x and time t as a vector field of the form

$$a(x,t) = \sum_{i=1}^{N} K^V(x, x_i(t))\alpha_i(t) \tag{1}$$

where x_i are the shape points carrying a point force vector α_i, and $K^V(x,y) = exp(-\|x-y\|^2/\lambda_V^2)$ is a Gaussian kernel of dimension mass^{-1} with standard deviation λ_V controlling the spatial extent at which the acceleration field varies.

The time-varying point force vectors $\alpha_i(t)$ parameterize a flow of deformation $\phi_t(x_i(t))$ by the integration of the 2nd-order ODE $\ddot{\phi}_t(x_i(t)) = a(x_i(t),t)$ with initial position $x_i(0)$ and initial velocity $\dot{x}_i(0)$. The initial positions of the particle are assumed to be fixed at the vertices of the baseline shape, while the initial velocity of the particles have to be determined by the algorithm.

Let $\mathbf{x}(t)$, $\mathbf{a}(t)$, and $\boldsymbol{\alpha}(t)$ be the concatenation of the $x_i(t)$'s, $a_i(t)$'s, and the $\alpha_i(t)$'s. This parameterization leads to the specific regression criterion

$$E(\dot{\mathbf{x}}(0), \boldsymbol{\alpha}(t)) = \sum_{t_i} \|\phi_{t_i}(\mathbf{x}(0)) - \mathbf{x}(t_i)\|^2_{W^*} + \gamma \int_0^T \|\mathbf{a}(t)\|^2_V \, dt \tag{2}$$

where $\|\cdot\|_{W^*}$ is the norm on currents and regularity is defined as $\|\mathbf{a}(t)\|^2_V = \boldsymbol{\alpha}(t)K^V(\mathbf{x}(t),\mathbf{x}(t))\boldsymbol{\alpha}(t)$, interpreted as the 'total amount of acceleration', measured using the norm in the reproducing kernel Hilbert space defined by the interpolating kernel [5].

3 Description of the Algorithm

We implement an adaptive step size gradient descent algorithm. The gradient of the criterion (2) with respect to force vectors and initial velocity is written as

$$\nabla_{\alpha_{i(t)}} E(t) = 2\gamma\alpha_i(t) + \eta_i^{\dot{x}}(t) \quad \text{and} \quad \nabla_{\dot{x}_i(0)} E = \eta_i^{\dot{x}}(0) \tag{3}$$

where variables $\eta_i^x(t)$ and $\eta_i^{\dot{x}}(t)$ satisfy coupled ODEs shown in Appendix A.

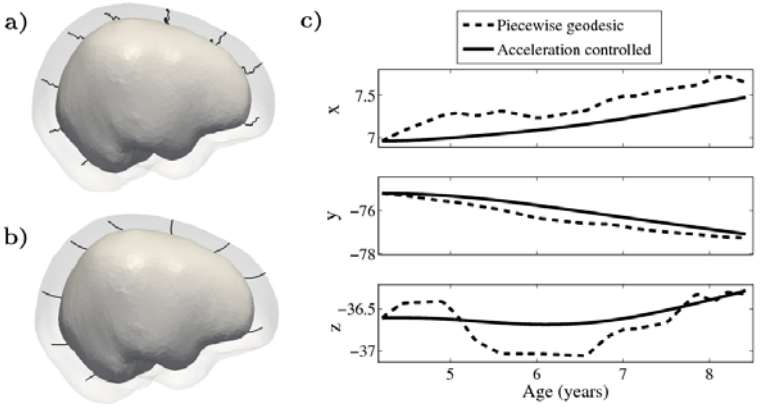

Fig. 1. a) and **b)** Shape evolution from baseline (solid) to final configuration (transparent) using a model based on piecewise geodesics (a) and our method (b) with point trajectories for selected particles displayed as black lines. **c)** The path of a point on the forebrain is decomposed into coordinates. Growth is estimated using 15 target shapes, highlighting the speed discontinuities present in the piecewise geodesic evolution.

During each iteration of gradient descent, the trajectories of shape points are computed by solving the 2nd-order ODE $\ddot{\phi}_t(x_i(t)) = a(x_i(t), t)$ using a Verlet integration scheme. The auxiliary variables $\eta_i^x(t)$ and $\eta_i^{\dot{x}}(t)$ are computed using an Euler method with prediction/correction. Eventually we compute the gradients given in equation (3). The algorithm may be started with zero initial velocity and force, though we notice faster convergence when initial velocity is determined by geodesic diffeomorphic registration between the baseline and first target shape as in [15].

4 Experiments

To evaluate our method, we use longitudinal image data from a child that has been scanned 16 times between four and eight years of age. The MRI data is first rigidly aligned to establish a common reference frame. The intracranial volume and lateral ventricles are segmented from each image using an EM based tissue classification algorithm and a level-set based active contour segmentation tool.

We estimate the evolution of the intracranial surface using a regression model based on the piecewise geodesic flow of diffeomorphisms as in [7]. The standard deviation of the Gaussian kernel controlling deformation is set to 50 mm, roughly 30% of the diameter of the baseline intracranial surface. For the scale of currents we use 20 mm, with a regularity weight of 0.1. Finally, time is discretized in increments of 0.0425 years. We also produce a growth trajectory using our proposed method with the same parameter settings as above except we weight regularity by 0.01 (the two weighted terms cannot be compared since they have different 'physical' dimension). The parameters were tuned empirically to produce regressions of comparable quality with both methods.

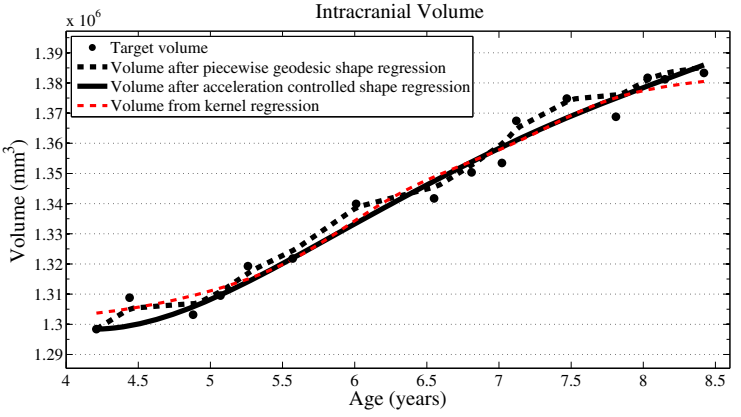

Fig. 2. Volume measurements derived from our growth model are consistent with a kernel regression ($\sigma = 0.5$) performed on the sparse volume measurements. Our model describes the continuous evolution of *shape* and volume is measured after regression.

Shape evolution is considerably smoother using our proposed regression model as compared to the piecewise geodesic model. This is particularly evident in the trajectories of the shape points across time, a subset of which are shown in Fig. 1. It is an important distinction that the trajectories estimated by our method are *not* a smoothing of the piecewise geodesic method. Rather, the trajectories are the result of fundamentally different assumptions on the underlying model which results in a more realistic estimation of growth.

The smoothness constraints imposed by our model limit the shape variation we can capture over short time periods. Consequently, we investigate the accuracy of our model by examining how closely we match the target data: our estimated growth scenario decreases the initial sum of squared residual by 148%, compared to a 153% decrease from the piecewise geodesic method. While our method does not come as close to matching the target data, this suggests that our method is less sensitive to noise and less likely to overfit.

Next, we investigate the application of our model to the study of measurements derived from shape. Here we obtain a continuous non-linear model of volume, shown in Fig. 2. The results are consistent with a $1D$ regression model, such as kernel regression, applied to the sparse volume measurements. However, we have focused our modeling efforts on capturing the evolution of shape, with continuous volume measurements resulting naturally from the estimated growth. In addition, the piecewise geodesic method appears to be overfitting, producing unrealistic volume measurements, further suggesting that our method is more robust in the presence of noisy data.

Finally, we consider the evolution of the lateral ventricles, which exhibit considerably more complexity than the intracranial surface. The horns of the segmented lateral ventricles are as thin as a few millimeters, making regression particularly challenging. As with the intracranial volume, ventricle growth is estimated using a piecewise geodesic model and our acceleration based model. The

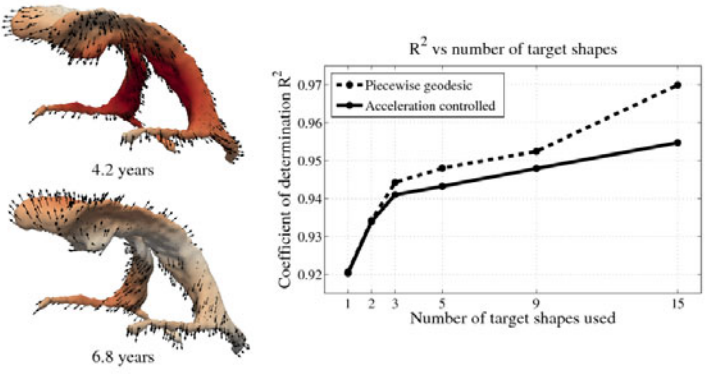

Fig. 3. Left: Snapshots from a continuous shape evolution of lateral ventricles estimated by our regression model. Acceleration vectors are displayed on the surface, with color denoting magnitude. **Right:** The impact of the number of target shapes on R^2.

scale of deformation is set at 6 mm, the scale of currents to 2 mm, and regularity is weighted by 0.1 and 0.01, respectively.

The impact of missing data is examined by performing leave-several-out experiments, the results of which are summarized in Fig. 3. In all experiments, selected target shapes were chosen as uniformly across time as possible. Our method demonstrates robustness with respect to the number of target shapes, with only minimal increase in the coefficient of determination R^2 when using more than 3 targets. This suggests that our method captures the underlying growth with limited data, as additional target data does not greatly alter the estimation. In contrast, piecewise geodesic regression is more influenced by additional target data and is therefore likely to overfit.

5 Conclusion

We have introduced a new 2nd-order regression model for estimating smooth evolution from a collection of time dependent shape data. This is based on a new way of parameterizing growth based on acceleration rather than velocity. We show on real anatomical data that, compared to the standard piecewise geodesic model, our method is less sensitive to noise introduced during segmentation and is robust to missing data, and is therefore more likely to characterize the underlying biological growth. The evolution of volume extracted after shape regression was shown to be compatible with a 1D regression on the observed volume measurements. Our method may be improved by additionally solving for initial positions of the shape points as in [7], to address the apparent underestimation of initial volume in Fig. 2.

Note that the new concept introduced in this paper has been implemented for 3D-surface data modeled as currents but can be easily adapted to a variety of other data and metrics. Future work will focus on the interpretation of the estimated acceleration in terms of external forces exerted on the biological tissue. This will enable the addition of more biological and mechanical priors.

Acknowledgments. This work was supported by NIH grant RO1 HD055741 (ACE, project IBIS) and by NIH grant U54 EB005149 (NA-MIC).

References

1. Castillo, E., Castillo, R., Martinez, J., Shenoy, M., Guerrero, T.: Four-dimensional deformable image registration using trajectory modeling. Physics in Medicine and Biology 55, 305–327 (2010)
2. Craene, M.D., Camara, O., Bijnens, B.H., Frangi, A.F.: Large diffeomorphic FFD registration for motion and strain quantification from 3D-US sequences. In: Ayache, N., Delingette, H., Sermesant, M. (eds.) FIMH 2009. LNCS, vol. 5528, pp. 437–446. Springer, Heidelberg (2009)
3. Datar, M., Cates, J., Fletcher, P., Gouttard, S., Gerig, G., Whitaker, R.: Particle based shape regression of open surfaces with applications to developmental neuroimaging. In: Yang, G.Z., Hawkes, D. J., Rueckert, D., Noble, J.A., Taylor, C. J. (eds.) MICCAI 2009. LNCS, vol. 5762, pp. 167–174. Springer, Heidelberg (2009)
4. Davis, B., Fletcher, P., Bullitt, E., Joshi, S.: Population shape regression from random design data. In: Proc. of ICCV, pp. 1–7 (October 2007)
5. Durrleman, S.: Statistical models of currents for measuring the variability of anatomical curves, surfaces and their evolution. Thèse de sciences (phd thesis). Université de Nice-Sophia Antipolis (March 2010)
6. Durrleman, S., Pennec, X., Trouvé, A., Ayache, N., Braga, J.: Comparison of the endocast growth of chimpanzees and bonobos via temporal regression and spatiotemporal registration. In: MICCAI-STIA Workshop (2010)
7. Durrleman, S., Pennec, X., Trouvé, A., Gerig, G., Ayache, N.: Spatiotemporal atlas estimation for developmental delay detection in longitudinal datasets. In: Yang, G.-Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (eds.) MICCAI 2009. LNCS, vol. 5761, pp. 297–304. Springer, Heidelberg (2009)
8. Ehrhardt, J., Werner, R., Schmidt-Richberg, A., Schulz, B., Handels, H.: Generation of a mean motion model of the lung using 4D-CT image data. In: Proc. of Eurographics Workshop on VCBM, pp. 69–76 (2008)
9. Hart, G., Shi, Y., Zhu, H., Sanchez, M., Styner, M., Niethammer, M.: DTI longitudinal atlas construction as an average of growth models. In: MICCAI-STIA Workshop (2010)
10. Khan, A., Beg, M.: Representation of time-varying shapes in the large deformation diffeomorphic framework. In: Proc. of ISBI, pp. 1521–1524 (2008)
11. Metz, C., Klein, S., Schaap, M., van Walsum, T., Niessen, W.: Nonrigid registration of dynamic medical imaging data using nd+t b-splines and a groupwise optimization approach. Medical Image Analysis 15(2), 238–249 (2011)
12. Miller, M.I, Trouvé, I., Younes, L.: On the metrics and Euler-Lagrange equations of Computational Anatomy. Annual Review of Biomedical Engineering 4, 375–405 (2002)
13. Perperidis, D., Mohiaddin, R.H., Rueckert, D.: Spatio-temporal free-form registration of cardiac MRI sequences. Medical Image Analysis 9(5), 441–456 (2005)
14. Peyrat, J.M., Delingette, H., Sermesant, M., Pennec, X., Xu, C., Ayache, N.: Registration of 4D Time-Series of Cardiac Images with Multichannel Diffeomorphic Demons. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part II. LNCS, vol. 5242, pp. 972–979. Springer, Heidelberg (2008)

15. Vaillant, M., Glaunès, J.: Surface matching via currents. In: Christensen, G.E., Sonka, M. (eds.) IPMI 2005. LNCS, vol. 3565, pp. 381–392. Springer, Heidelberg (2005)
16. Vialard, F., Trouvé, A.: Shape splines and stochastic shape evolutions: A second-order point of view. Quarterly of Applied Mathematics To Appear

A Differentiation of the Regression Criterion

Using matrix notation, we denote the current state of the system of shape points by the vector $\mathbf{X}(t) = (\mathbf{x}(t), \dot{\mathbf{x}}(t))^t$ concatenating position and velocity of every point. The state of the system is evolved by the following differential equation:

$$\dot{\mathbf{X}}(t) = F(\mathbf{X}(t), \boldsymbol{\alpha}(t)) = \begin{pmatrix} \dot{\mathbf{x}}(t) \\ \ddot{\mathbf{x}}(t) = \mathbf{K}(\mathbf{x}(t), \mathbf{x}(t))\boldsymbol{\alpha}(t) \end{pmatrix} \tag{4}$$

with initial condition $\mathbf{X}(0) = \mathbf{X}_0 = (\mathbf{x}_0, \dot{\mathbf{x}}_0)^t$.

We now rewrite (2) as $E(\mathbf{X}(t)) = \sum_{t_i} A(\mathbf{X}(t_i)) + \gamma \int_0^T L(\mathbf{X}(t), \alpha(t))dt$. Let δE be a variation of the criterion E with respect to a variation $\delta\boldsymbol{\alpha}(t)$ of the impulse vectors $\boldsymbol{\alpha}(t)$, which induces a variation of the state variable $\mathbf{X}(t)$:

$$\delta E = \sum_{t_i} \left(d_{\mathbf{X}(t_i)} A_i\right) \delta\mathbf{X}(t_i) + \gamma \int_0^T (\partial_{\mathbf{X}(t)} L(t))\delta\mathbf{X}(t) + (\partial_{\boldsymbol{\alpha}(t)} L(t))\delta\boldsymbol{\alpha}(t)dt \tag{5}$$

The ODE in (4) shows that these variations $\delta\mathbf{X}(t)$ satisfy a linear inhomogeneous ODE. The method of variation of parameters gives the solution

$$\delta\mathbf{X}(t) = R_{0t}\delta\mathbf{X}_0 + \int_0^T R_{ut}\partial_{\boldsymbol{\alpha}(u)} F(u)\delta\boldsymbol{\alpha}(u)\mathbf{1}_{\{u \leq t_i\}}du \tag{6}$$

where $R_{ut} = \exp\left(\int_u^t \partial_{\mathbf{X}(s)} F(s)ds\right)$ and $\mathbf{1}_{\{t \leq t_i\}} = 1$ if $t \leq t_i$ and 0 otherwise.

Plugging this equation into (5) leads to:

$$\nabla_{\boldsymbol{\alpha}} E(t) = \partial_{\boldsymbol{\alpha}(t)} L(t)^t + \partial_{\boldsymbol{\alpha}(t)} F(t)^t \boldsymbol{\eta}(t) \quad \text{and} \quad \nabla_{\mathbf{X}_0} E = \boldsymbol{\eta}(0) \tag{7}$$

where we denote the auxiliary variable $\boldsymbol{\eta}(t)$ as

$$\boldsymbol{\eta}(t) = \sum_i \nabla_{\mathbf{X}_{t_i}} A_i \mathbf{1}_{\{t \leq t_i\}} + \int_t^T \partial_{\mathbf{X}} L(u)^t + \partial_{\mathbf{X}} F(u)^t \boldsymbol{\eta}(u)du \tag{8}$$

From now on, we decompose the vectors into 2 blocks (the x-component and the $\dot{x}$-component). Due to the definition of A, L and F, we have $\nabla_{\mathbf{X}(t_i)} A_i = (\nabla_{\mathbf{x}_i} A_i \;\; \mathbf{0})^t$, $\partial_{\mathbf{X}} L = (\gamma\boldsymbol{\alpha}^t(\partial_1 + \partial_1)(\mathbf{K}(\mathbf{x}, \mathbf{x})\boldsymbol{\alpha}) \;\; \mathbf{0})^t$, $\partial_{\boldsymbol{\alpha}} L = 2\gamma\boldsymbol{\alpha}^t\mathbf{K}(\mathbf{x}, \mathbf{x})$, $\partial_{\mathbf{X}} F = \begin{pmatrix} \mathbf{0} & \mathbf{1} \\ (\partial_1 + \partial_2)\mathbf{K}(\mathbf{x}, \mathbf{x})\boldsymbol{\alpha} & \mathbf{0} \end{pmatrix}$ and $\partial_{\boldsymbol{\alpha}} F^x = (\mathbf{0} \;\; \mathbf{K}(\mathbf{x}, \mathbf{x}))$.

Therefore, the gradient of the regression criterion with respect to the L^2 metric given in (7) is now equal to: $\nabla_{\boldsymbol{\alpha}} E(t) = \mathbf{K}(\mathbf{x}(t), \mathbf{x}(t))\left(2\gamma\boldsymbol{\alpha}(t) + \boldsymbol{\eta}^{\dot{x}}(t)\right)$, where we have decomposed the auxiliary variable $\boldsymbol{\eta}$ into $\boldsymbol{\eta} = (\boldsymbol{\eta}^x, \boldsymbol{\eta}^{\dot{x}})$.

The matrix $\mathbf{K}(\mathbf{x}(t), \mathbf{x}(t))$ is precisely the Sobolev metric induced by the kernel on the set of L^2 functions, so the gradient is given in coordinates as in (3).

Minimization of Intra-Operative Shaping of Orthopaedic Fixation Plates: A Population-Based Design

Habib Bou-Sleiman[1,⋆], Lucas E. Ritacco[2], Lutz-Peter Nolte[1], and Mauricio Reyes[1]

[1] Institute for Surgical Technology and Biomechanics, University of Bern, Stauffacherstrasse 78, 3014 Bern, Switzerland
{habib.bousleiman, mauricio.reyes}@istb.unibe.ch
[2] Hospital Italiano de Buenos Aires, Gascón 450, Buenos Aires, Argentina

Abstract. In this paper we present a new population-based method for the design of bone fixation plates. Standard pre-contoured plates are designed based on the mean shape of a certain population. We propose a computational process to design implants while reducing the amount of required intra-operative shaping, thus reducing the mechanical stresses applied to the plate. A bending and torsion model was used to measure and minimize the necessary intra-operative deformation. The method was applied and validated on a population of 200 femurs that was further augmented with a statistical shape model. The obtained results showed substantial reduction in the bending and torsion needed to shape the new design into any bone in the population when compared to the standard mean-based plates.

Keywords: Orthopaedic implant design, population-based analysis, bone fixation plate.

1 Introduction

Bone fixation plates are commonly used in orthopaedic surgeries to preserve, maintain, and help restore the original anatomy of a diseased or fractured bone. The success of reconstructive and corrective interventions heavily relies on the proper design and application of the implants. Current trends tend to offer fixation plates that are pre-contoured to the specific target location in which they are supposed to function [1,2]. Pre-contouring is commonly based on the average anatomy of the target population, or on a template bone considered as a representative of that population. To date, the available pre-contoured plates are not capable of providing an optimal fit to all operated patients. This is mainly due to the differences in morphology of the human skeleton between and within different populations. Factors such as age, gender, and ethnic origin play an important role in defining the morphology of the bones [3].

⋆ Corresponding author.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 409–416, 2011.

Furthermore, fixation plates are manually shaped during the surgery in order to adapt them to the patient-specific anatomy, a delicate and time-consuming procedure that is prone to high inaccuracies [4]. Such procedures require wide exposure of the bone [2] and longer surgical time, thus increasing the risks and costs of the intervention. The development of locked internal fixators and anatomically pre-shaped plates reduced the need of plate bending [1,2,5,6], however, accurate patient-specific implant pre-contouring is still not accessible.

In [7], the authors propose a method to design plates by searching through the parametric space of a statistical shape model of the bones. They search for the most significant population-specific shape variability patterns that affect the fit of the fixation plate using criteria based on surface-to-surface distances. They finally propose manual changes to the current design and prove that the new plate fits more bones from the population according to the same design criteria.

In this paper, we present a new population-based automatic approach to the design of orthopaedic fixation plates. We propose a design that minimizes the maximum amount of intra-operative manual bending and torsion that are to be applied to the plate. The proposed deformation model and metrics go in agreement with the type of deformation the surgeon applies during the intervention. The design criteria are more clinically significant than those used in [7]. The benefit of such method is two-fold. Explicitly, less mechanical stresses are applied to the plate and therefore a better long-term mechanical stability is expected. Whereas the lesser required shaping implicitly indicates a better pre-operative surface-to-surface fit.

2 Methods

2.1 Experimental Data

A population of 200 segmented computed tomography (CT) datasets with varying image parameters and scanners was used in this work. Point distribution models (PDM) of the left femur were generated and aligned. Dense point-to-point correspondence was computed using an image-based log-domain demons registration with a polyaffine regularization [8].

Two patches of interest have been selected, outlined, and extracted from the initial datasets by an experienced orthopaedic surgeon. The patches are located on the distal medial and the distal lateral sections of the femur. Fig. 1a-b shows these two patches.

Using the pre-established point-to-point correspondence, the vertices of the chosen patches were propagated throughout the PDM in order to obtain 200 samples of each type of patch that are subsequently realigned. In the present case, each distal medial patch is composed of 1,343 surface points whereas its lateral counterpart is composed of 1,625 points. These patches are considered as both bone surfaces and potential new designs of the contact surface of the fixation plate. The two terms *plates* and *patches* will be used interchangeably throughout the rest of this paper.

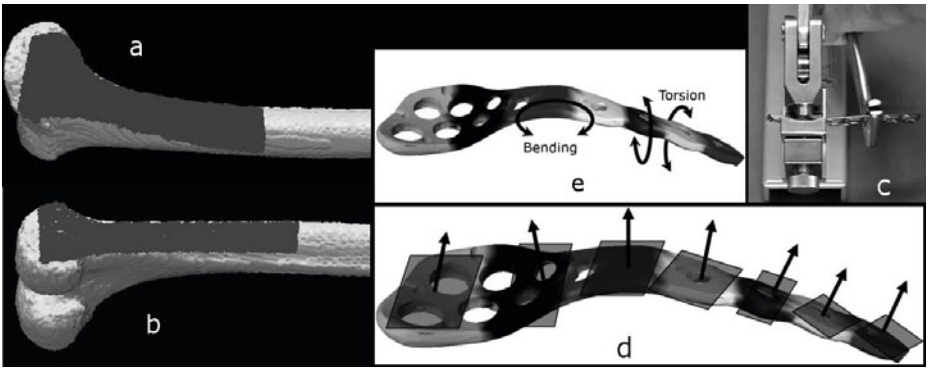

Fig. 1. Selection of the (a) medial and (b) lateral distal patches of the femur. (c) Intra-operative bending of a fixation plate. (d) Example of how the plate can be divided into discrete sections. A plane is fitted to each individual section and represented by its normal vector. (e) Illustration showing the directions of bending and torsion.

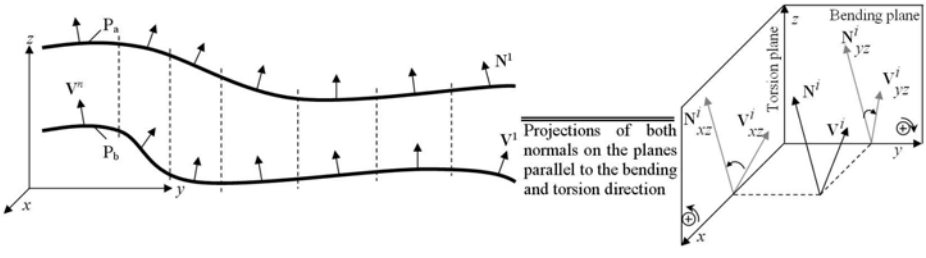

Fig. 2. Illustration of how the method computes the bending and torsion angles required to shape one patch (P_a) into another (P_b). Both components are measured relative to the main axis of the patches, assuming the main axis is aligned with the y-axis of the coordinate system. In this example, the bending component is the angle between the projections of the normals onto the yz-plane, whereas the torsion component is the angle between the projections of the normals onto the xz-plane.

2.2 Plate Discretization and Representation

We propose to first represent the contact surface of a fixation plate (or the bone patch) by dividing it into discrete sections lengthwise and fitting a plane to each section. The number of sections depends on the size of the implant and anatomical location. It should be small enough to capture the anatomical features, but large enough not to be affected by local noise in the surface. The size of the individual sections should also be within the operational limits of the common shaping tools (see Fig. 1c). Each plane is then represented by a vector normal to it. An example of how the plate can be divided is shown in Fig. 1d.

Algorithm 1. Computation of bending and torsion

In: Two sets of normal vectors $\mathbf{N}$ and $\mathbf{V}$ to the planes constituting the bone patches P_a and P_b

Out: Maximum bending and torsion angles, D_Ω and D_Θ, needed to shape P_a into P_b

1: Initialize $\Omega^0_{P_a,P_b} = 0;\ \Theta^0_{P_a,P_b} = 0;\ R^0_\Omega = 0;\ R^0_\Theta = 0$
2: **for** each section $i \in [1 \text{ to } n]$
3: $\quad \Omega^i_{P_a,P_b} \leftarrow \alpha_{\mathbf{N}^i_{yz},\mathbf{V}^i_{yz}} - R^i_\Omega$
4: $\quad \Theta^i_{P_a,P_b} \leftarrow \alpha_{\mathbf{N}^i_{xz},\mathbf{V}^i_{xz}} - R^i_\Theta$
5: $\quad$ **for** each $j \in [(i+1) \text{ to } n]$
6: $\quad\quad R^j_\Omega \leftarrow R^j_\Omega + \Omega^i_{P_a,P_b}$
7: $\quad\quad R^j_\Theta \leftarrow R^j_\Theta + \Theta^i_{P_a,P_b}$
8: $\quad$ **end for**
9: **end for**
10: $D_\Omega\,(P_a, P_b) \leftarrow \max\left\{\Omega^i_{P_a,P_b} \mid i = 1 \ldots n\right\}$
11: $D_\Theta\,(P_a, P_b) \leftarrow \max\left\{\Theta^i_{P_a,P_b} \mid i = 1 \ldots n\right\}$

2.3 Measurement of Bending and Torsion

The key formulation of our solution is to examine the effect of shaping each section of the implant on the rest of the sections and compute the amount of deformation needed to shape one patch to another based on currently available shaping tools. The angles between two corresponding normals to the planes are split into two independent components, namely bending and torsion, both relative to the main axis of the patch. Fig. 2 illustrates how the patches are discretely represented by a set of normal vectors and how the bending and torsion angles are measured.

The method used to measure how much bending (D_Ω) and torsion (D_Θ) are required to shape one patch into another is formulated in Algorithm 1, where n is the number of sections per patch, α is an angle between two vectors, Ω and Θ are the vectors of bending and torsion angles, and R_Ω and R_Θ are the vectors of residual bending and torsion angles, both respectively. Letting P_a be the patch to be contoured to P_b, the algorithm starts from one end of the patches, say the proximal end, and sequentially goes along the main axis until the other end is reached. This process is analogous to drum pressing the plate against a template until it acquires its shape. It also mimics what the surgeon would do during the intervention. It starts by initializing the angles to zero, indicating an alignment of the first normals without any deformation of the patch or plate (line 1). The algorithm then proceeds by iterating over two main steps. The first computes the bending and torsion required to align two corresponding normal vectors while taking into account residual angles from previous steps (lines 3-4). The second is an update step that computes the vector of residual angles and

consecutively assigns its elements to the remaining normal vectors in the chain of sections along the patch (lines 6-7). This update step describes the effect of applying a local deformation on the rest of the plate. Our interest is to compute the maximum needed local deformation. Thus the highest values in the lists of bending and torsion are extracted and stored for later processing (lines 10-11). These values are to be later minimized since they are the maximal forces that a plate design will undergo during the surgery.

2.4 Finding the Optimal Template from a Population

By applying the analysis described above to the whole population, one is able to measure pair-wise maximal bending and torsion angles needed to shape any one patch into another. Our goal is to find the patch that simultaneously minimizes both components among the studied population. Consistent with the notation presented in Algorithm 1, we propose to optimize our metrics using Eq. (1) through Eq. (3). In Eq. (3), *NewDesign* is the index of the patch that requires the least amount of physical shaping, and p the size of the population. The weighting factors ω_Ω and ω_Θ are used to impose optional non-equal significances on both components. In our experiments, we used unity weights indicating equal contribution of bending and torsion to the design criterion.

$$d^a_\Omega = \max\{D_\Omega(P_a, P_b) \mid b = 1 \dots p\} \; ; \tag{1}$$

$$d^a_\Theta = \max\{D_\Theta(P_a, P_b) \mid b = 1 \dots p\} \; ; \tag{2}$$

$$NewDesign = \underset{a}{\operatorname{argmin}} \left\{ \sqrt{\omega_\Omega {d^a_\Omega}^2 + \omega_\Theta {d^a_\Theta}^2} \mid a = 1 \dots p \right\} . \tag{3}$$

2.5 Further Optimization

In order to avoid falling into local minima and to ensure that the method finds the optimal plate design, we further examined the three best ranked patches from the previous steps. We generated a statistical shape model [9] of the initial populations and retained the first 42 principal components capturing 99% of the natural shape variability. The parameters of the three patches within the statistical space were recovered using a least-squares approximation $\mathbf{c} = \mathbf{Q}^+\mathbf{y}$, where $\mathbf{c}$ is the vector of shape parameters, $\mathbf{Q}$ the matrix of eigenvectors scaled by the corresponding eigenvalues, and $\mathbf{y}$ the vector of coordinates of the surface points. Latin hypercube sampling [10] was used to generate 100 instances in the statistical vicinity of each of the best choices. Latin hypercube sampling was preferred over uniform or random dense sampling since it offers a similar space coverage with a substantially lower number of samples. The same search described earlier was applied again to the augmented populations.

2.6 Finding an Adequate Number of Sections

The number of sections per plate or bone patch is a fundamental parameter in the design process. Care must be taken while defining this parameter since it directly

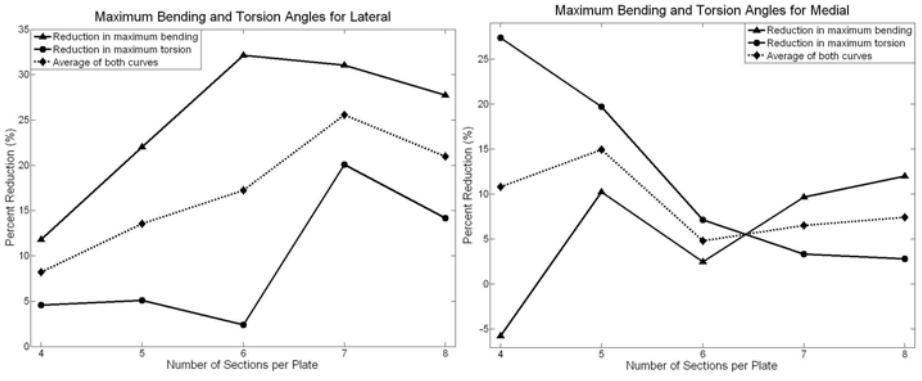

Fig. 3. Reduction in the maximum bending and torsion angles needed to shape the proposed design to any instance in the population, for lateral (*left*) and medial (*right*) plates. The reduction is relative to the results obtained for the standard design and plotted against the number of sections used to divide the bone patches or plates. The average curve has its peak at seven sections for the lateral case, whereas five sections resulted in a higher improvement for the medial dataset.

affects the quality of the final design. For each one of the considered anatomical sites, we tested the algorithm above using a varying number of sections per plate. We recorded the amount of improvement[1] over the standard mean-based designs. Ideally, the goal is to pick the number that yields the highest improvement in both bending and torsion components. We used the average values of both components as a selection criterion.

3 Results

The method described above was separately applied to the sets of bone patches extracted from the medial and lateral distal femur. The first step was to identify the optimal number of sections that divide each instance of the population. Fig. 3 shows the results of the corresponding tests, where seven and five dividing sections yielded the highest improvement for the lateral and medial cases, respectively.

Following the initial step, the configuration of the discretization pattern that yielded the best improvement was retained and applied to the whole population as well as to the standard design. Fig. 4 plots the obtained results and highlights the difference between the maximum required deformation for the standard plates and that for the enhanced designs. Table 1 shows a comparison between the standard plates and the proposed improved design. It also compares the highest pre-operative mean surface distances (MSD) measured for every case. The results indicate that the intra-operative deformation required

[1] Improvement is measured as the percent reduction in maximum bending and torsion from the values recorded for the standard design.

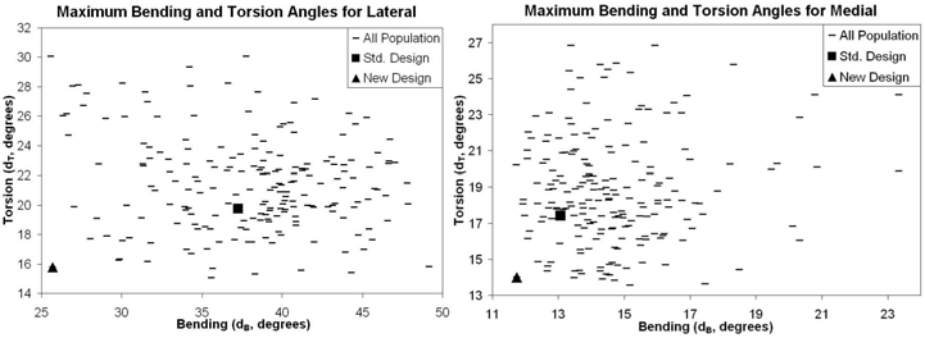

Fig. 4. Maximum bending and torsion required to shape any instance of the population (−) to all other instances, for lateral (*left*) and medial (*right*) plates. The new design (▲) resulted in lower bending and torsion angles than the standard design (■).

to fit the standard plate to any potential patient exceeds that required by the new design. Consequently, this implies that it is easier to contour the new plate design to the patient-specific anatomy.

4 Discussion

In this paper we presented a novel population-based method to design pre-contoured orthopaedic fixation plates. The method minimizes the maximum deformation (bending and torsion) required to fit an implant to the patient-specific anatomy. The presented deformation model is consistent with the actual intra-operative deformation and complies with the current shaping tools. The results indicate that the population mean is not the optimal design, neither the MSD is the optimal design metric. The highest possible pre-operative MSD (preceding intra-operative shaping) between the plate and any bone surface is also reduced (medial) or almost unchanged (lateral) with the new design. The three-dimensional surface representation of the bone patch designated as *NewDesign* can be used as a template to manufacture better pre-contoured implants.

We intend to apply the method presented herein to populations of different anatomical sites. Different population characteristics can also be included in the analysis such as gender, race, age group, etc. This would divide the population

Table 1. Comparison between the standard mean-based design and the new design proposed in this paper. The values refer to the highest measured angles and distances.

	Lateral Plate			**Medial Plate**		
	Bending	**Torsion**	**MSD**	**Bending**	**Torsion**	**MSD**
Standard Design	37.25°	19.74°	23.15mm	13.07°	17.41°	17.03mm
New Design	25.27°	14.00°	23.22mm	11.74°	13.22°	14.36mm
Improvement	**32.16%**	**29.08%**	**-0.30%**	**10.18%**	**24.07%**	**15.68%**

into subgroups that can be used to design population-specific implants. The population could also be divided using clustering methods based on the presented design metrics (and possibly others) to generate multiple more specific designs. We also propose to further extend this method and develop a simple guidance system that can be integrated into the manufactured implants. It should indicate the different sections and the way and amount by which the surgeon must deform the plate. An image-guided system that does not require the development of sophisticated shaping tools is an ultimate goal that is brought within reach by the methods and representation presented in this paper. It would allow for a smooth transfer from the design workbench to the operating theater. We also plan to carry out mechanical tests to assess the weighted contribution of bending and torsion to the design criteria and the impact of screw holes deformation.

Acknowledgments. This work was carried out within the frame of the National Center of Competence in Research, Computer-Aided and Image-Guided Medical Interventions (NCCR Co-Me), supported by the funds of the Swiss National Science Foundation (SNSF).

References

1. Perren, S.M.: Evolution and rationale of locked internal fixator technology. Introductory remarks. Injury 32(suppl.2), B3–B9 (2001)
2. Wagner, M.: General principles for the clinical use of the LCP. Injury 34(suppl. 2), B31–B42 (2003)
3. Schmutz, B., Wullschleger, M.E., Kim, H., Noser, H., Schutz, M.A.: Fit assessment of anatomic plates for the distal medial tibia. J. Orthop. Trauma 22, 258–263 (2008)
4. Frankle, M.A., Cordey, J., Frankle, M.D., Baumgart, F., Perren, S.: A retrospective analysis of plate contouring in the tibia using the conventional 4.5 (narrow) dynamic compression plate. J. Orthop. Trauma 8, 59–63 (1994)
5. Goyal, K.S., Skalak, A.S., Marcus, R.E., Vallier, H.A., Cooperman, D.R.: Analysis of anatomic periarticular tibial plate fit on normal adults. Clin. Orthop. Relat. Res. 461, 245–257 (2007)
6. Taljanovic, M.S., Jones, M.D., Ruth, J.T., Benjamin, J.B., Sheppard, J.E., Hunter, T.B.: Fracture fixation. Radiographics 23, 1569–1590 (2003)
7. Kozic, N., Weber, S., Büchler, P., Lutz, C., Reimers, N., Ballester, M.A.G., Reyes, M.: Optimisation of orthopaedic implant design using statistical shape space analysis based on level sets. Med. Image Anal. 14, 265–275 (2010)
8. Seiler, C., Pennex, X., Ritacco, L., Reyes, M.: Femur specific polyaffine model to regularize the log-domain demons registration. In: SPIE Med. Imaging. Florida (2011)
9. Davies, R., Twining, C.J., Taylor, C.J.: Statistical models of shape: optimisation and evaluation. Springer, London (2008)
10. McKay, M.D., Beckman, R.J., Conover, W.J.: A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. Technometrics 21, 239–245 (1979)

Iterative Refinement of Point Correspondences for 3D Statistical Shape Models

Sharmishtaa Seshamani, Gouthami Chintalapani, and Russell Taylor

Department of Computer Science, Johns Hopkins University, Baltimore, MD

Abstract. Statistical atlases of bone anatomy are traditionally constructed with point-based models. These methods establish initial point correspondences across the population of shapes and model variations in the shapes using a variety of statistical tools. A drawbacks of such methods is that initial point correspondences are not updated after their first establishment. This paper proposes an iterative method for refining point correspondences for statistical atlases. The statistical model is used to estimate the direction of "pull" along the surface and consistency checks are used to ensure that illegal shapes are not generated. Our method is much faster that previous methods since it does not rely on computationally expensive deformable registration. It is also generalizable and can be used with any statististical model. We perform experiments on a human pelvis atlas consisting of 110 healthy patients and demonstrate that the method can be used to re-estimate point correspondences which reduce the hausdorff distance from 3.2mm to 2.7mm and the surface error from 1.6mm to 1.4mm for PCA modelling with 20 modes.[1]

1 Introduction

Statistical atlas modelling is a popular tool for analysis of several types of medical images. Applications range from modelling of anatomical variations within shape populations to pathological anomaly detection. Shape atlases are typically built by selection of a representation such as points [1,2,3], curves [4] or level sets [5] followed by modelling using various statistical techniques [6]. In this paper, we address the application of statistical bone atlases. Traditionally, bone (and other rigid body) atlases are modelled using dense point-based methods coupled with linear statistical models. A set of shapes are first represented as either surface meshes or volumetric tetrahedral meshes. Subsequently, vertex correspondences are established using a template mesh. These vertices are then used to represent the shapes as high dimensional vectors which can be statistically analyzed for shape variation and other applications.

One drawback of this approach is that point correspondences are established beforehand and never updated once the model is built. As a consequence, the error and uncertainty of this initialization is propagated through the model. Chintalapani et al [2] propose a solution for updating the model with a bootstrapping method which employs the intensity values of voxels. However, this method

[1] Supported by NIH with Grant 1-R01-EB00683 and JHU internal funds.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 417–425, 2011.

is computationally very expensive since it requires multiple passes through an intensity-based deformable registration algorithm. Secondly, in some cases, the 3D intensity data may not be available. In more recent work, Davies et al [3] propose a minimum description length framework for optimizing a PCA based shape model. Although this method implicitly disallows generation of illegal shapes, it is also computationally very expensive. Secondly, their method is specific to PCA based modelling. In this paper, we propose an iterative bootstrapping method for improving point correspondences without any outside information with the use of any statistical model. A set of initial point correspondences are established (just as in most prior work) to estimate a statistical model, which is then used to re-estimate point correspondences. This iterative procedure can be repeated until convergence. We describe the method for the re-estimation with a PCA based model. However, our method can be applied to any generic statistical model. Leave-out validation and comparison of the vertex error, surface error and volume difference is carried out. Results show that our method is able to improve modelling performance with all these metrics.

2 Methods

We begin with a set of N surface meshes described as: $\mathcal{S} = \{S_i = (T, V_i)|i = 1 \dots N\}$, where $V_i \in \mathcal{R}^{PX3}$ contains the P vertices of the shape and T is the set of triangles. The set $\mathcal{V} = \{V_1 \dots V_N\}$ contains initial point correspondences, ie: the nth vertices of each shape correspond. The objective is to generate a statistical model of this point (vertex) based representation and then improve the initial point correspondences. In the next few sections, we define a few preliminaries before describing the general algorithm.

2.1 Preliminaries

Given the set $\mathcal{V}$, we select one test sample V_i. The rest of the samples form a set $\mathcal{V}'_i$ which is used to generate a PCA model consisting of a mean shape $\bar{M}$ and the modes of variation Y. V_i can now be represented by a set of mode weights λ, computed as: $\lambda = Y^T(V_i - \bar{M})$. The mode weights can be used to reconstruct the test sample as: $V_i^{rec} = \bar{M} + \sum_{i=1}^{n} \lambda_i Y_i$. This type of leave-out modelling and reconstruction can be performed for all $V_i \in \mathcal{V}$. Figure 1(a) shows an example of an original shape and some points on it in red (V_i). The grey points are those that are reconstructed using the model (V_i^{rec}).

Reconstruction with leave-one-out modelling produces an estimate of the points on the left out shape using a model generated by all the other shapes. For each vertex point, the vector relating the original point to the reconstructed points has two components: an out of surface component $\boldsymbol{d_{out}}$ and an along surface component $\boldsymbol{d_{along}}$(Figure 1(b)). In order to ensure point correspondence update only along the surface, we can compute the closest point on the original surface to the reconstructed point (Figure 1(a)). This now gives us the direction that the statistical model "pulls" the original points along the surface. Carrying this out for all vertex points in V_i generates a set of

vertices $V_i^{closest}$. In some cases, the closest point may move very far away from the original point. Update with the closest point could therefore generate large amounts of distortion in the shapes and collapsing of triangles. Instead of performing a closest point update, we can now move the original point in the direction of the closest point along the surface by applying a scale factor μ. $\mu = 1$ generates the closest point itself. However, this may not always correspond to a consistent shape. Hence, we can select a μ which does satisfy this requirement. Once μ is selected, the point can be moved by the vector $\mu \boldsymbol{d_{along}}$ and then projected back onto the surface (closest point to the surface) to generate the updated points $V_i^{updated}$(Figure 1(c)). In the experimental section, we show how the fractional updates preserve consistency and reduce shape distortion for improved point correspondences. Modifying the vertices of a surface mesh can affect the consistency of the reconstruction. In particular, we are interested in avoiding generation of illegal shapes. For mesh based representations, comparison of the normals of the triangles in the original and updated shapes gives a quantitative measure of illegal shapes (flipped triangles) generation. Given a surface $S_i = (V_i, T)$ and the new surface $S_i^{updated} = (V_i^{updated}, T)$, the normals of all triangles of S_i and $S_i^{updated}$ can be computed to generate the sets N_i and $N_i^{updated}$ respectively. The dot products of normals of corresponding triangles can be computed. If a triangle has been flipped, the dot project will be negative.

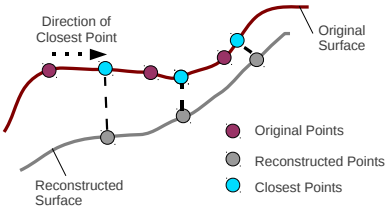

(a) Original surface with a few points marked (Red), Surface reconstructed with PCA with points marked (Grey), Closest Points (Blue) and direction of closest point from original point

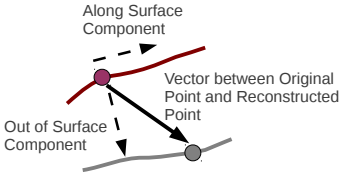

(b) Out of Surface and Along Surface Components of vector between original and reconstructed point

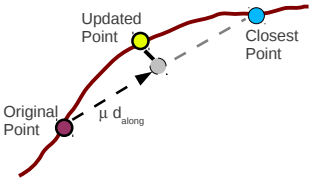

(c) Fractional Update of Points: The point is moved in the direction of the closest point at a scale of μ and then projected onto the surface to generate the updated point (correspondence).

Fig. 1. Explanatory figures for preliminaries

2.2 Algorithm Summary

Algorithm 1 and Figure 2 outline the procedure. Initial set of point correspondences are iteratively updated until convergence. In each iteration, the input set of surfaces is $\mathcal{S}^{updated}$ which was generated by the previous iteration. Closest points are always computed to the original set of surfaces $\mathcal{S}$ in order to maintain the shape of the surfaces.

Algorithm 1. Point Update Algorithm

Input: $\mathcal{S} = \{(V_i, T) | i = 1 \dots N\}$, Stepsize s
for all $S_i = (V_i, T) \in \mathcal{S}$ **do**
  Generate $\mathcal{V'}_i$ and compute PCA model $\Rightarrow (\bar{M}, Y)$.
  Reconstruct V_i with $\bar{M}$ and $Y \Rightarrow V_i^{rec}$
  Compute closest points to S_i. $\Rightarrow V_i^{closest}$
  Compute direction vector d_{along} for every vertex
  Set $\mu = 1$ and compute fractional update $\Rightarrow V_i^{updated}$
  while $V_i^{updated}$ is not consistent **do**
    $\mu = \mu - s$
    Compute fractional update $V_i^{updated}$
  end while
end for
Output: $\mathcal{S}^{updated} = \{(V_i^{updated}, T) | i = 1 \dots N\}$

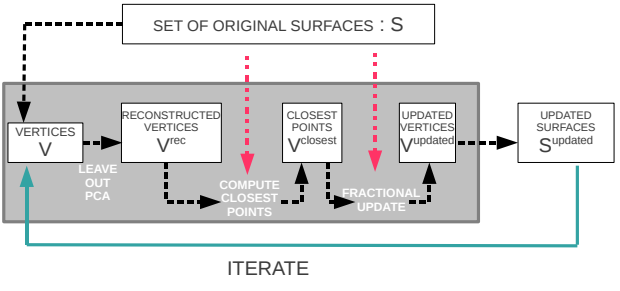

Fig. 2. Algorithm outlining the procedure for iteratively updating points in shape model: A surface is described with a set of vertices and triangles. Vertices are reconstructed using leave out reconstruction. Closest surface points are then computed and a fractional update is applied in the direction of the closest point to generate new vertices. This procedure repeated till convergence.

3 Experimental Results

3.1 Data Acquisition and Validation

We experiment with a pelvis atlas consisting of 110 male patient samples. The data was collected by a physician and anonymized before we performed any processing. Each sample is a segmented 512x512x256 CT volume. Surface meshes (11663 points, 23414 triangles) were extracted for each volume. All meshes were rigidly registered to a template sample to establish initial point correspondences - Iteration 0 set. For validation, we performed **leave-20 out** cross-validation for all the data. We computed the mean vertex error (Distance between original points and reconstructed points), mean surface error (Distance between reconstructed points and original surface) and mean volume error (Sum of surface error and distance from original points to the reconstructed surface (Hausdorff distance)). We also report the standard deviations for all folds of validation.

3.2 Results

We first performed leave-20 out cross-validation using initial point correspondences (iteration 0) by varying the number of modes. The black lines in Figure 3 show the average vertex, surface and volume errors for reconstruction with 1-90 modes. Next, we performed the point update (Algorithm 1). The number of modes selected for reconstruction in step 4 was 30 modes since the black line in Figure 3 flattens out after 30 modes. The reconstructed points (step 4) and closest points (step 5) were used to compute the vertex error, the average along surface distance and the average out of surface distance. The maximum and minimum vertex errors were 5.8mm and 2.3mm. The maximum and minimum along surface distances were 5.1mm and 1.4mm. The maximum and minimum out of surface distances were 3.1mm and 0.6mm. Figure 5 shows an example of point update with this first iteration. The full pelvis on top shows the target (original) shape, which is the mesh generated by initial point correspondences for one of the instances. The figures at the bottom show zoomed versions of the some of the regions. The zoomed view on the left shows regions on the target shape and the zoomed view in the middle shows the same mesh with point correspondences updated with the closest points computed. A lot of distortion can be observed by this step. Secondly, some of triangles are flipped with this point update. In order to deal with both of these problems, we performed the adaptive selection of μ starting with $\mu = 1$ and a step size of $s = 0.1$. This gives us a μ value of 0.3 which is the largest value of μ for generating a consistent dataset. We then performed the leave-20 out cross-validation of these samples just as in Iteration 0. The average vertex, surface and volume errors and the error bars are shown in blue in Figure 3. We then demonstrate the effect of μ on these metrics as follows. We generated datasets for $\mu < 0.3$ (consistent datasets) at step sizes of 0.1. This gave us two more datasets at $\mu = 0.1$ and $\mu = 0.2$. In addition, we also generated one more dataset for $\mu = 1$. This dataset was not consistent (at least 30 out of the 23414 triangles were flipped in the meshes). Leave-20 out cross-validation was performed on these three extra datasets. The average vertex, surface and volume errors and error bars are plotted in Figure 3. The red curves correspond to $\mu = 0.1$. The green curves correspond to $\mu = 0.2$ and the black dashed curves correspond to $\mu = 1$. Thus, we observe that the increase of μ affects these metrics consistently towards $\mu = 1$.

Following this, we performed two more iterations of the leave out reconstruction. Again, we used 30 modes for the reconstruction and μ was adaptively selected at each subsequent iteration. In iteration 2, μ was selected as 0.3 and at iteration 3, μ was selected as 0.4. Leave-20 out cross-validation was then performed for the same 20 samples and the three metrics are plotted in Figure 4. With validation at 20 modes, the average vertex error drops

Table 1. Vertex, Surface & Volume Errors, leave-20-out validation, 20 modes

	Vertex Error(mm)	Surface Error(mm)	Volume Error(mm)
Iter 0	4.0	1.6	3.2
Iter 1	2.9	1.5	2.9
Iter 2	2.4	1.4	2.8
Iter 3	2.1	1.4	2.7

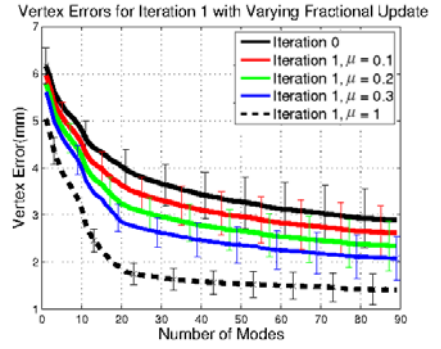

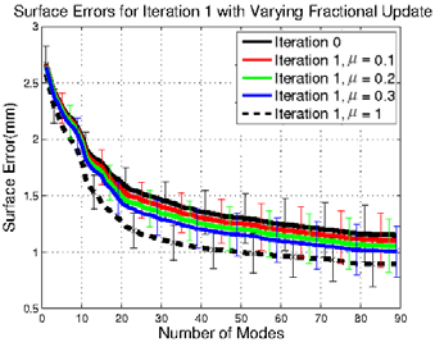

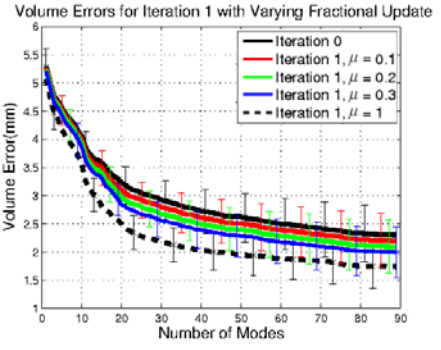

Fig. 3. Graphs showing the Vertex error, Surface error and Volume error with error bars from Leave-20 out cross-validation for Iteration 1 with varying fractional updates. Leave out reconstruction was performed with 30 modes and values of the fractional updates are $\mu = 0.1$, 0.2 and 0.3.

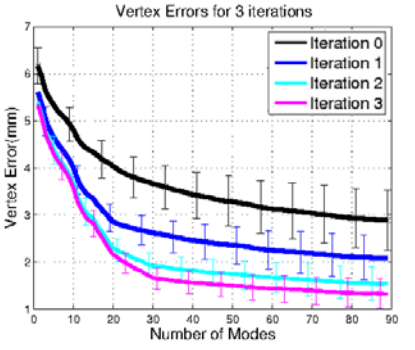

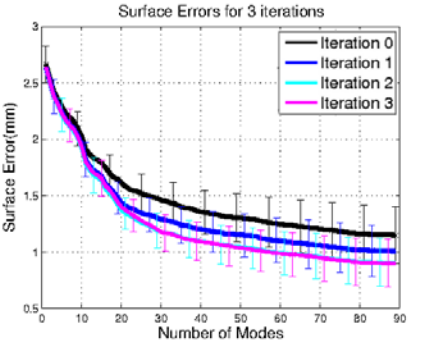

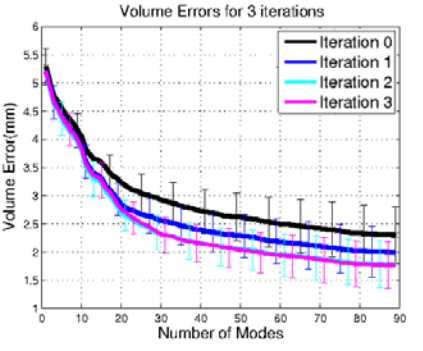

Fig. 4. Graphs showing the Vertex error, Surface error and Volume error with error bars from Leave-20 out cross-validation for three iterations. Leave out reconstruction was performed with 30 modes. μ = was adaptively selected as 0.3, 0.3 and 0.4 in iterations 1,2 and 3 respectively.

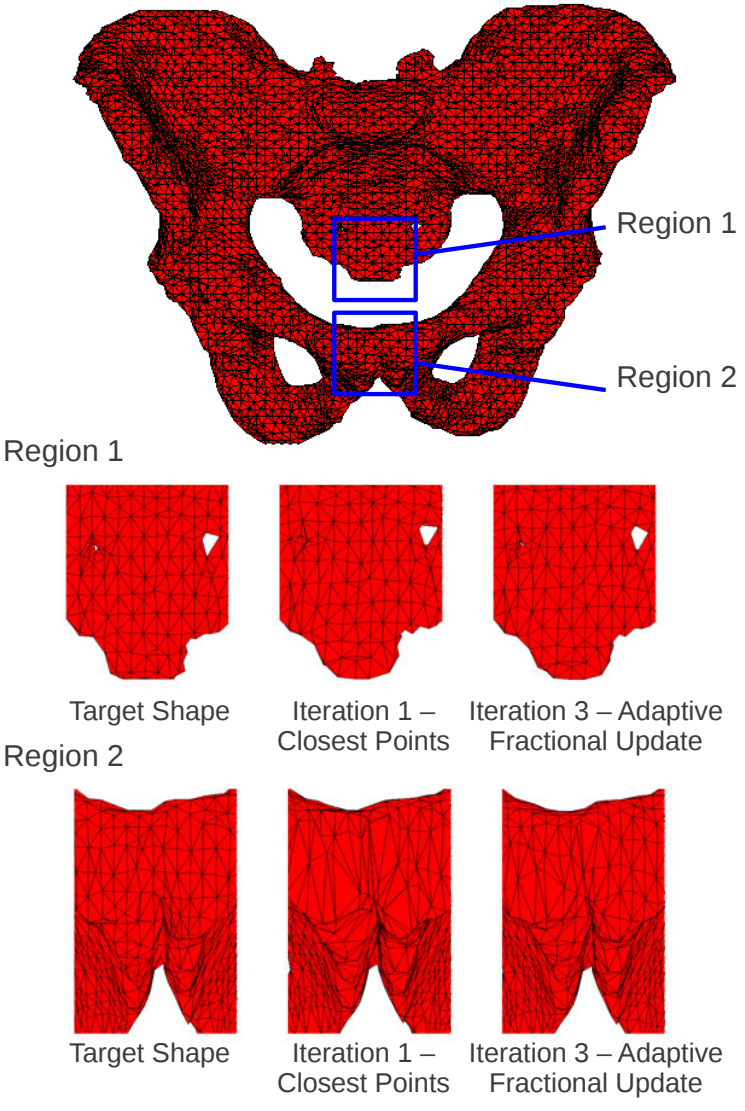

Fig. 5. Top: Original Mesh of a pelvis instance, Bottom: Zoomed view of mesh in two regions. The left most views show the region on the original pelvis, the middle view shows the region with updates using the closest points in iteration 1 (Black dashed curve in Figure 3), the right most views show the region after iteration 3 with adaptive fractional updates (Pink curve in Figure 4). Note that although the metrics are similar, there is a lot more shape distortion in the closest point reconstruction of iteration 1. This shows that the fractional update method is also limiting the shape distortions.

from 4.0mm to 2.1mm, the average surface error drops from 1.6mm to 1.4mm and the average volume error drops from 3.2mm to 2.7. Table 1 shows these values for all iterations. We also note that the difference between these errors in the last two iterations is very small, showing convergence of the procedure. Finally, we compare the reconstruction of the adaptive fractional updates with the reconstruction with the closest points in iteration 1. By inspection of the graphs in Figures 3 and 4, we note that the closest points in iteration 1 generate the black dashed curve in figure 3 and iteration 3 of the adaptive fractional updates generate the pink curve in Figure 4. The metrics are comparable for these two curves. Figure 5 shows a comparison of the original shape and these two types of point updates. We note that the adaptive fractional update not only maintains consistency of the shape but also reduces the amount of shape distortion.

4 Discussion

We have presented a general iterative method for refinement of point 3D point correspondences for statistical atlas modelling. This method is non-specific to the statistical method and does not use outside information to update the point correspondences. Instead, the statistical model's estimate is used to drive the direction of point update. Since our method explicitly computes reconstructions based on the statistical model and then uses the reconstruction to update the points, this method could can be generalized to the use of any statistical model. It is also much faster than previous methods which use intensity based deformable registration or complex optimization procedures. In our experiments, we show that with leave-20 out validation with 20 modes, the vertex error can be reduced from 4mm to 2.1mm, the surface error can be reduced from 1.6mm to 1.4mm and the volume error can be reduced from 3.2mm to 2.7mm. We also note that the fractional update algorithm not only maintains consistency but also reduces the amount of shape distortion. One of the drawbacks of the current procedure is that no intensity information is used at all. Thus, although the shapes are consistent, it is possible that some of the detail may be lost. The work so far has addressed only one specific atlas. In future work, we plan to extend the evaluation of the applicability of our method to other types of statistical atlases.

Acknowledgments. This work was supported by NIH grant RO1 HD05 5741 (ACE, project IBIS) and by NIH grant U54 EB005149 (NA-MIC).

References

1. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active shape models their training and application. Comput. Vis. Image Underst. 61, 38–59 (1995)
2. Chintalapani, G., Ellingsen, L.M., Sadowsky, O., Prince, J.L., Taylor, R.H.: Statistical atlases of bone anatomy: Construction, iterative improvement and validation. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) MICCAI 2007, Part I. LNCS, vol. 4791, pp. 499–506. Springer, Heidelberg (2007)

3. Davies, R.H., Twining, C.J., Cootes, T.F., Taylor, C.J.: Building 3-d statistical shape models by direct optimization. IEEE Trans. Med. Imaging 29(4), 961–981 (2010)
4. Joshi, S.H., Cabeen, R.P., Sun, B., Joshi, A.A., Gutman, B., Zamanyan, A., Chakrapani, S., Dinov, I., Woods, R.P., Toga, A.W.: Cortical sulcal atlas construction using a diffeomorphic mapping approach. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. 6361, pp. 357–366. Springer, Heidelberg (2010)
5. Osher, S.J., Fedkiw, R.P.: Level Set Methods and Dynamic Implicit Surfaces, 1st edn. Springer, Heidelberg (2002)
6. Bishop, C.M.: Pattern Recognition and Machine Learning (Information Science and Statistics). Springer-Verlag New York, Inc., Secaucus (2006)

A New Shape Diffusion Descriptor for Brain Classification

Umberto Castellani[1], Pasquale Mirtuono[1], Vittorio Murino[1,2], Marcella Bellani[3], Gianluca Rambaldelli[3], Michele Tansella[3], and Paolo Brambilla[4,5]

[1] VIPS lab, University of Verona, Italy
[2] Istituto Italiano di Tecnologia (IIT), Italy
[3] Department of Public Health and Community Medicine, Inter-University Center for Behavioural Neurosciences (ICBN), University of Verona, Italy
[4] Department of Experimental Clinical Medical Sciences (DISM), Inter-University Center for Behavioral Neurosciences, University of Udine, Italy
[5] Scientific Institute IRCCS "E. Medea", Udine, Italy

Abstract. In this paper, we exploit spectral shape analysis techniques to detect brain morphological abnormalities. We propose a new shape descriptor able to encode morphometric properties of a brain image or region using diffusion geometry techniques based on the local *Heat Kernel*. Using this approach, it is possible to design a versatile signature, employed in this case to classify between normal subjects and patients affected by schizophrenia. Several diffusion strategies are assessed to verify the robustness of the proposed descriptor under different deformation variations. A dataset consisting of MRI scans from 30 patients and 30 control subjects is utilized to test the proposed approach, which achieves promising classification accuracies, up to 83.33%. This constitutes a drastic improvement in comparison with other shape description techniques.

1 Introduction

Brain morphology techniques using Magnetic Resonance Imaging (MRI) are playing an increasingly important role in understanding pathological structural alterations of the brain [1,2]. A typical approach is to investigate the presence of morphological differences of selected brain structures between neuropsychiatric patients and healthy controls [2]. To this aim, methods for shape analysis can be exploited in order to extract the geometric information which provides the best statistical performance in separating the two populations (i.e., healthy and non-healthy people)[3]. Classic approaches evaluate volumetric variations [2] to explain atrophy or dilation due to such kind of illnesses. Nevertheless, more advanced shape analysis techniques have been proposed aiming at exploiting new aspects of the shape such as spectral [4,5] or local geometric properties [6]. A typical methodology consists of encoding such geometric properties into a *descriptor* which compactly represents the shape. In this fashion, the comparison between shapes can be carried out by measuring the descriptors' similarities in

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 426–433, 2011.

the descriptor space. Thus, the effectiveness of shape descriptors can be evaluated in terms of discriminativeness and robustness against shape variations due to noise or deformations.

In this paper, we propose a new shape descriptor based on advanced *diffusion* geometry techniques. Local geometric properties are encoded by the so-called *Heat Kernel* [7] which exploits heat diffusion characteristics at different scales. The general idea consists of capturing information about the neighborhood of a point on the shape by recording the dissipation of heat over time from that point onto the rest of the shape. In this way, *local* shape characteristics are highlighted through the behavior of heat diffusion over short time periods, and, conversely, *global* shape properties are observed while considering longer periods [7,8]. So doing, simply varying a single parameter (the time), it is possible to characterize the properties of a shape at different scales. Therefore, local heat kernel values observed at each point are accumulated into a histogram for a fixed number of scales leading to the so-called *Global Heat Kernel Signature* (GHKS).
The method is inspired by [7] which proposed the *Heat Kernel signature* (HKS) for a single vertex of a mesh. Here, we extend the HKS for the whole shape for both surface mesh (i.e., external surface) and volumetric representation. The proposed descriptor has several nice properties which are shared with very few other work. GHSK allows for shape comparisons using minimal shape preprocessing, in particular, no registration, mapping, or remeshing is necessary. GHKS is robust to noise since it implicitly employs surface smoothing by neglecting higher frequencies of the shape. Finally, GHKS is able to encode isometric invariance properties of the shape [7] which are crucial to deal with shape deformations.

The proposed descriptor has been tested in the context of the analysis of the schizophrenia illness. A Region-of-Interest (ROI)-based approach [1] is employed by studying the left-thalamus, which is known to be impaired by such disease [2]. Experiments, carried out on a dataset of 30 patients and 30 controls, lead to promising classification results in distinguishing between the two populations also in comparison with other methods.

2 Related Work

Several work has been proposed for detecting alterations of the brain structure by using advanced shape analysis techniques [4,6]. A common approach consists of capturing *global* shape information from the (shape-)spectral domain [5,4]. In [5], geometric properties are encoded by computing spherical harmonic descriptors (SPHARM) on brain surfaces. Although results are interesting, the method is not invariant to surface deformations and therefore it requires shapes registration and data resampling. This pre-processing is avoided in [4], where the so called Shape-DNA signature has been introduced by taking the eigenvalues of the Laplace-Beltrami operator as region descriptor for both the external surface and the volume. Although *global* methods can be satisfying for some classification tasks, they do not provide information about the localization of the morphological anomalies. To this aim, *local* methods have been proposed.

In [6] the so called *feature*-based morphometry (FBM) approach is introduced. Taking inspiration from feature-based techniques proposed in computer vision, FBM identifies a subset of features corresponding to anatomical brain structures that can be used as disease biomarkers. Other approaches are able to combine both *global* and *local* information. More specifically, a recent and important class of methods has been introduced for generic object analysis which employs heat diffusion procedures on 3D shapes [7,8,9,10]. In this class of techniques, global information is provided by the spectral parameters of the Laplace-Beltrami operator employed on 3D data, and local information is defined by the heat diffusion at small scales. In [7], Sun et al. have proposed the so called *Heat Kernel Signature*(HKS): the main idea was to describe the diffusion from a point to itself at several time instants. The HKS provides a natural and efficiently computable multi-scale way to capture information about neighborhoods of a given point. A similar approach has been proposed in [8] by introducing the so called *Auto Diffusion Function* (ADF). The idea and formulation is the same as in [7], but the procedure has been applied for object segmentation and skeleton extraction. In order to obtain a global signature from local measures two main strategies has been proposed [10,9]. In [9] the well known *Bag-of-features* approach is employed starting from the HKS value at each point of the shape. Coversely, in [10] the global shape is captured by computing the distribution of diffusion distances among the points of the shape. In [9], a study on isometry-invariance property of the geometric diffusion process is proposed in order to highlight the differences between *volume*-isometry and *boundary*-isometry. In the former case, the diffusion is computed at voxel level, whereas in the latter the diffusion is computed only on the external surface.
Our approach extends the use of heat kernel on MRI data for classification purposes on medical domain. The method proposed improves [5] since our descriptor is isometry invariant. Moreover, differently than [4], our approach implements a multi-scale analysis to increase the discriminativeness properties of the descriptor. Finally, the main idea of the heat kernel signature[7] to describe the diffusion from a point to itself at different scales has been revised to work on global shape.

3 The Heat Diffusion Process

Given a shape M as a compact Riemannian manifold, the heat diffusion on shape[1] is defined by the *heat* equation:

$$(\Delta_M + \frac{\partial}{\partial t})u(t,x) = 0; \tag{1}$$

where u is the distribution of heat on the surface, Δ_M is the *Laplace-Beltrami* operator which, for compact spaces, has discrete eigendecomposition of the form $\Delta_M = \lambda_i \phi_i$. In this fashion the *heat kernel* has the following eigendecomposition:

$$k_t(x,y) = \sum_{i=0}^{\infty} e^{-\lambda_i t}\phi_i(x)\phi_i(y), \tag{2}$$

[1] In this section, we borrow the notation from [7,9].

where λ_i and ϕ_i are the i^{th} eigenvalue and the i^{th} eigenfunction of the Laplace-Beltrami operator, respectively. The heat kernel $k_t(x, y)$ is the solution of the heat equation with point heat source at x at time $t = 0$, i.e., the heat value at point y after time t. The heat kernel is *isometric invariant*, it is *informative*, *multi-scale*, and *stable* [7]. In order to estimate the Laplace-Beltrami and the heat kernel in discrete domains several strategies can be employed [9]. In the following we describe the cases of surface meshes and volumetric representations.

Heat kernel on surface meshes. In the case of surface mesh only the boundary of the shape is considered. In order to work on a discrete space, we estimate the Laplace-Beltrami operator by employing linear Finite Elements Methods (FEM) [4]. More in detail, given a triangular mesh composed by $v_1, \cdots, v_m$ vertices, with *linear* finite elements the *generalized eigendecomposition* problem [4] becomes:

$$A_{cot}\Phi = -\Lambda B\Phi, \tag{3}$$

where Λ is the diagonal matrix of the Laplace Beltrami eigenvalues λ_i, and Φ is the matrix of corresponding eigenfunctions ϕ_i. The matrices A_{cot} and B are defined as:

$$A_{cot}(i,j) = \begin{cases} \frac{cot\alpha_{i,j} + cot\beta_{i,j}}{2} & if \quad (i,j) \in E, \\ -\sum_{k \in N(i)} A_{cot}(i,k) & if \quad i = j, \\ 0 & otherwise. \end{cases} \tag{4}$$

$$B(i,j) = \begin{cases} \frac{|t_1| + |t_2|}{12} & if \quad (i,j) \in E, \\ \frac{-\sum_{k \in N(i)} |t_k|}{6} & if \quad i = j, \\ 0 & otherwise. \end{cases} \tag{5}$$

where E is the set of edges of the triangular mesh, $\alpha_{i,j}$ and $\beta_{i,j}$ are the two angles opposite to the edge between vertices v_i and v_j in the two triangles sharing the edge (i, j), $|t_i|$ is the area of the triangle t_i, and t_1, t_2 are the triangles that shares the edge (i, j). Indeed, the *heat kernel* can be approximated on a discrete mesh by computing Equation 2 and retaining the k smallest eigenvalues and the corresponding eigenfunctions.

Heat kernel on volumetric representations. In the case of volumetric representations, the interior part of the shape is also considered. The volume is sampled by a regular Cartesian grid composed of voxels, which allows the use of standard Laplacian in R^3 as the Laplace-Beltrami operator. We use finite differences to evaluate the second derivative in each direction of the volume. The heat kernel on volumes is invariant to volume isometries, in which shortest paths between points inside the shape do not change. Note that in real applications exact volume isometries are limited to the set of rigid transformations [9]. However, also non-rigid deformations can faithfully be modelled as approximated volume isometries in practice. Moreover, differently from spectral surface representation,

volumetric approach is able to capture volume atrophy. It is worth noting that, as observed in [7,9], for small t the heat kernel $k_t(x,x)$ of a point x with itself is directly related to the *scalar* curvature $s(x)$ [9]. More formally:

$$k_t(x,x) = (4\pi t)^{-3/2}(1 + \frac{1}{6}s(x)). \tag{6}$$

Note that in the case of surface meshes $s(x)$ can be interpreted as the Gaussian curvature [7]. In practice, Equation 6 states that heat tends to diffuse slower at points with positive curvature, and viceversa. This gives an intuitive explanation about the geometric properties of $k_t(x,x)$ an leads the idea of using it to build a shape descriptor [7].

4 The Proposed Method

The proposed approach is composed of three main phases: i) data gathering, ii) estimation of descriptors, and iii) classification.

Data Gathering. Quantitative data collection and processing in MRI based research implies facing several methodological issues to minimize biases and distortions. The standard approach to deal with these issues is following well established guidelines dictated by international organizations, such as the World Health Organization (WHO), or codified by respected institutions, such as leading universities. In this work we employ a ROI-based approach [1]: only a well defined brain subpart has been considered. Specifically, we focus our analysis on the left-Thalamus whose abnormal activity is already investigated in schizophrenia. Regions have been manually traced by experts, according to well defined medical protocols.

Global Heat Kernel Signature. Once data are collected, a strategy to encode the most informative properties of the shape M can be devised. To this end, a global shape descriptor is proposed, which is inspired by the so-called *Heat Kernel Signature*(HKS) defined as:

$$HKS(x) = [k_{t_0}(x,x), \cdots, k_{t_n}(x,x)]. \tag{7}$$

where x is a point of the shape (i.e., a vertex of a mesh or a voxel) and $(t_0, t_1, \cdots, t_n)$ are n time values. To extend this approach to the whole shape, we introduce the following global shape descriptor:

$$GHKS(M) = [hist(K_{t_0}(M)), \cdots, hist(K_{t_n}(M))], \tag{8}$$

where $K_{t_i}(M) = \{k_{t_i}(x,x), \forall x \in M\}$, and $hist(\cdot)$ is the histogram operator. Note that our approach combines the advantages of [10,9] since it encodes the distribution of local heat kernel values and it works at multiscales. Figure 1 shows a schema of the proposed descriptor. Each point of the shape is colored according to $k_{t_i}(x,x)$. Such values are collected into a histogram for each scale t_i. Finally, histograms are concatenated leading to the global signature.

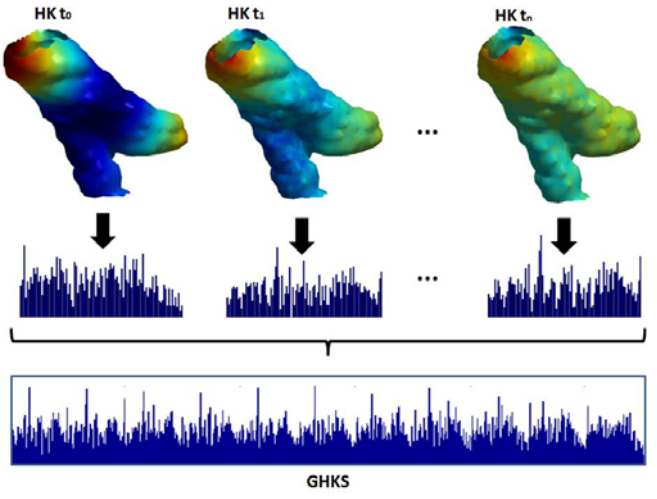

Fig. 1. GHKS: Each point of the shape is colored according to $k_{t_i}(x, x)$. Such values are collected into a histogram for each scale t_i. Finally, histograms are concatenated leading to the global signature.

Support Vector Machine Classification. These descriptors are simply evaluated using a Support Vector Machine (SVM), which is one of the most powerful classifier for object recognition [11]. SVM constructs a maximal margin hyperplane in a high dimensional feature space, by mapping the original features through a kernel function. Here, the input of the SVM are the set of GHKS descriptors extracted for each subject. A learning by example approach is introduced by adopting leave-one-out cross-validation procedure[2].

5 Results

The proposed shape classification method is employed for Schizophrenia detection in Thalamic region. A dataset composed of 30 male patients and 30 male controls has been evaluated. MRI scans were acquired using a 1.5 T Siemens Magnetom Symphony Maestro Class, Syngo MR 2002B. After manual extraction of the ROIs both mesh surfaces and volumetric representations have been recovered. The Laplace-Beltrami operator has been computed as described in Section 3, for both representations, and the *heat kernel* has been computed. In this work, we have used $k = 200$ eigenvalues, and we have scaled the temporal domain logarithmically in $n = 10$ time values, as suggested in [7]. Finally, the GHKS is computed by fixing 100 bins for each histogram. Therefore, for each

[2] A single sample is used as validation data, and the remaining samples as training data. The procedure is repeated such that every sample in the dataset is used once as validation data.

Table 1. Classification rates. The accuracy is computed by Leave-One-Out cross-validation. Three kernels are evaluated and two methods are compared. Both surface and volumetric representations are considered.

Method	Linear-SVM	Polynomial-SVM	RBF-SVM
Surface GHKS	65.00%	66.67%	71.67%
Volumetric GHKS	81.67%	80.00%	**83.33%**
Surface ShapeDNA	50.00%	66.67%	70.00%
Volumetric ShapeDNA	50.00%	71.67%	73.33%

subject the final dimension of the GHKS is $10 \cdot 100 = 1000$. The classification procedure is employed as described in Section 4. Several kernels have been evaluated, namely *linear*, *polynomial* (degree=3), and *radial basis function* (RBF). We compare our descriptor with the so called ShapeDNA descriptor, recently proposed by Reuter et al. [4]. As mentioned above, the ShapeDNA has similar properties of our GHKS descriptor since it encodes the *intrinsic* properties of the shape. Conversely, ShapeDNA does not deal with multiple scales and takes into account of only global information. Table 1 shows the classification performance of the considered approaches. The proposed GHKS descriptor clearly outperforms the ShapeDNA descriptor[3]. Specifically, a drastic improvement is observed when the volumetric approach is employed. In fact, volumetric GHKS reaches the best accuracy (i.e., 83.33%), and it is stable by varying the type of kernel employed. It is worth noting that also in the case of ShapeDNA, better performances are observed with the volumetric procedure. Therefore, from this study we can argue that volumetric approach is more suitable to deal with natural shape variations that raise on brain subparts of different subjects. The computational cost of the proposed GHKS descriptor is not high and effective: the Laplace-Beltrami transform can be employed in around 10 seconds for a mesh of about 3000 vertices. The same eigendecomposition on our volumetric data of $21 \times 37 \times 29$ voxels takes around 25 seconds. Then, the computation of final GHKS takes around a second for both the approaches[4].

6 Conclusions

In this paper, a new shape morphometry approach is introduced to improve the classification between normal subjects and patients affected by schizophrenia. Our GHKS descriptor combines local shape properties into a global signature by exploiting geometric diffusion procedure on MRI data. The approach proposed outperforms previous work, namely ShapeDNA, it is easy to be implemented and efficient. Both volumetric and surface approaches have been evaluated by showing that in our study neuroanatomical variations between different subjects are well modelled by volume isometries. From our experiments, we can highlight

[3] The same number of eigenvalues have been employed.
[4] We used a laptop at $1.66Ghz$. The code is written in Matlab with some parts in C.

the discriminativeness property of the thalamus by confirming the importance of this region to figure out mental disorders, especially in schizophrenia. Future work will address the localization of the disease on both surface and volume by further exploiting the local properties of the heat kernel on MRI data.

Acknowledgements. We acknowledge financial support from the FET programme within the EU-FP7, under the SIMBAD project (contract 213250).

References

1. Giuliani, N.R., Calhouna, V.D., Pearlson, G.D., Francis, A., Buchanan, R.W.: Voxel-based morphometry versus region of interest: a comparison of two methods for analyzing gray matter differences in schizophrenia. Schizophrenia Research 74(2-3), 135–147 (2005)
2. Shenton, M.E., Dickey, C.C., Frumin, M., McCarley, R.W.: A review of MRI findings in schizophrenia. Schizophrenia Research 49(1-2), 1–52 (2001)
3. Fan, Y., Shen, D., Gur, R.C., Gur, R.E., Davatzikos, C.: COMPARE: Classification of morphological patterns using adaptive regional elements. IEEE Transactions on Medical Imaging 26(1), 93–105 (2007)
4. Reuter, M., Wolter, F.-E., Shenton, M., Niethammer, M.: Laplace-Beltrami eigenvalues and topological features on eigenfuntions for statistical shape analysis. In: Computed-Aided Design, vol. 41(10), pp. 739–755 (2009)
5. Gerig, G., Styner, M., Jones, D., Weinberger, D., Lieberman, J.: Shape analysis of brain ventricles using SPHARM. In: Proceedings of the IEEE Workshop on Mathematical Methods in Biomedical Image Analysis, pp. 171–178. IEEE Computer Society, Washington, DC, USA (2001)
6. Toews, M., Wells, W., Collins, D., Arbel, T.: Feature-based morphometry. In: Yang, G.Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (eds.) MICCAI 2009. LNCS, vol. 5762, pp. 109–116. Springer, Heidelberg (2009)
7. Sun, J., Ovsjanikov, M., Guibas, L.: A concise and provably informative multi-scale signature based on heat diffusion. In: Proceedings of the Symposium on Geometry Processing, pp. 1383–1392. Eurographics Association, Berlin (2009)
8. Gebal, K., Baerentzen, J.A., Aanaes, H., Larsen, R.: Shape analysis using the auto diffusion function. In: Proceedings of the Symposium on Geometry Processing, pp. 1405–1413. Eurographics Association, Berlin (2009)
9. Raviv, D., Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Volumetric heat kernel signatures. In: Workshop on 3D Object Retrieval, pp. 39–44. ACM, Firenze (2010)
10. Bronstein, A.M., Bronstein, M.M., Ovsjanikov, M., Guibas, L.J.: Shape recognition with spectral distances. IEEE Trans. Pattern Analysis and Machine Intelligence 33(5), 1065–1071 (2011)
11. Burges, C.: A tutorial on support vector machine for pattern recognition. Data Mining and Knowledge Discovery 2, 121–167 (1998)

Comparison of Shape Regression Methods under Landmark Position Uncertainty

Nora Baka[1,2], Coert Metz[1], Michiel Schaap[1], Boudewijn Lelieveldt[2,4], Wiro Niessen[1,4], and Marleen de Bruijne[1,3]

[1] Erasmus MC, Rotterdam, The Netherlands
n.baka@erasmusmc.nl
[2] Leiden University Medical Center, Leiden, The Netherlands
[3] University of Copenhagen, Denmark
[4] Delft University of Technology, The Netherlands

Abstract. Despite the growing interest in regression based shape estimation, no study has yet systematically compared different regression methods for shape estimation. We aimed to fill this gap by comparing linear regression methods with a special focus on shapes with landmark position uncertainties. We investigate two scenarios: In the first, the uncertainty of the landmark positions was similar in the training and test dataset, whereas in the second the uncertainty of the training and test data were different. Both scenarios were tested on simulated data and on statistical models of the left ventricle estimating the end-systolic shape from end-diastole with landmark uncertainties derived from the segmentation process, and of the femur estimating the 3D shape from one projection with landmark uncertainties derived from the imaging setup. Results show that in the first scenario linear regression methods tend to perform similar. In the second scenario including estimates of the test shape landmark uncertainty in the regression improved results.

1 Introduction

Shape estimation by linear regression has been applied in a great variety of problems, including 3D shape estimation from digitized point cloud [10] and projection images [15], neighboring shape prediction [8,11,14], organ motion prediction based on a few time-points [9,7,1], healthy shape prediction for disease quantification [4], and remaining shape variation prediction from partial field of view [3,2].

The shapes used to construct the regression model are segmentations of images with noise, artifacts, low contrast, etc. Therefore, landmark positions always inhabit a certain amount uncertainty, which may be estimated from the images, imaging setup, segmentation method, etc. Regression methods that are currently applied for shape prediction do not incorporate this additional information.

Studies in chemometrics literature compared the performance of linear regression methods in noisy environments concluding that including data uncertainty may improve prediction in certain situations [12]. These conclusions can not directly be generalized to shape regression, as they apply to single variable outcome

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 434–441, 2011.

cases with low dimensional input and a large training set. Shape regression on the other hand is characterized by high dimensional input and output variables, and a very small training set.

Therefore, we addressed the following questions: (1) Is there a preferred linear regression method for shape regression? (2) In what context is it beneficial to incorporate landmark position uncertainties in the regression?

We answer these questions by comparing several linear regression methods with and without incorporation of uncertainties. We differentiate two scenarios: In the first, uncertainties of the landmark positions are similar in the training and test dataset, whereas in the second uncertainties of the training and test data are different. We evaluate the performances in simulation experiments providing a controlled setup with a Gaussian shape and uncertainty model and known ground truth, and in a real-world dataset for each scenario. First, the end-systolic left ventricle shape is predicted from end-diastole with landmark uncertainties derived from the segmentation process, and second, the 3D femur shape is estimated from one projection with landmark uncertainties derived from the imaging setup.

2 Shape regression

Shape representation. Each training shape consists of landmark points, which are in correspondence across all shapes. After rigid alignment each shape can be represented by its concatenated landmark coordinates as a high dimensional point. The set of training shapes in this high dimensional space is commonly assumed to be Gaussian distributed with lower intrinsic dimensionality.

Regression techniques. In this paper we describe shape regression where both input $\boldsymbol{x}$ and output $\boldsymbol{y}$ are shapes. We focus on linear regression methods, as in low sample size high dimensionality problems non-linear methods tend to over-fit the training data [6]. Ordinary least squares (OLS) is the best unbiased linear estimator. It minimizes the sum of squares prediction error requiring an over-determined system, and often inhabiting large prediction variance[1]. Due to the small sample size and the lower intrinsic dimensionality shape regression is typically under-determined, and OLS is not directly applicable. We compared biased methods as they both reduce prediction variance, and are applicable for under-determined problems.

Standard biased regressions not including landmark uncertainty are principal component regression (PCR) applying OLS on a sub-space calculated by principal component analysis; partial least squares regression (PLS) optimizing the underlying lower dimensional spaces of input and output to maintain the most covariance between them[2]; and ridge regression (RR) applying a norm constraint on the regression coefficients. See [6] for more detailed descriptions.

[1] With multivariate $\boldsymbol{x}$ and $\boldsymbol{y}$ OLS is termed multiple outcome multiple linear regression.
[2] We used the SIMPLS algorithm for PLS.

Maximum likelihood (ML) PCR includes uncertainties of both test and training data by performing OLS on the input subspace A, that is optimal in the ML sense for a given set of samples with known Gaussian uncertainties. A sample $\boldsymbol{x}$ with uncertainty Ψ is represented as $r = (A^T\Psi^{-1}A)^{-1}A^T\Psi^{-1}\boldsymbol{x}$ in this subspace. We defined a minimum uncertainty for all dimensions to maintain stability. Furthermore to cut computational costs, we used an efficient alternating algorithm with independent noise approximation and no intercept optimization [13]. MLPCR has not been used in shape regression before.

To include uncertainty in the test sample only, we used maximum a-posteriori (MAP) ridge regression, proposed in [2]. With known test uncertainty Ψ one obtains the coefficient matrix $B_{\mathrm{MAPR}} = \Sigma_{YX}(\Sigma_{XX} + \Psi + \lambda I)^{-1}$, where Σ_{XX} and Σ_{YX} are the input covariance and cross-covariance matrices, I is the identity matrix, and λ is a regularization constant.

Suitable values for the percentage of variance retained with PCR, number of modes of PLS and MLPCR, and regularization constant for ridge and MAP ridge regressions have to be assessed from the training set. We used cross-validation on the training set to estimate the regularization.

3 Simulation Study

Framework description. To evaluate shape regression with different noise patterns, we simulated data that resembles real world shapes, i.e. (1) the data forms a low-dimensional Gaussian distribution with variance outside this subspace being Gaussian noise; (2) noise of neighboring dimensions is correlated; (3) the variance of modes drops quickly and then smoothly converges to a constant; (4) sample dimensionality greatly exceeds the training set size; and (5) the input only partially predicts the output.

The data was generated with intercepts $\boldsymbol{x}_0$ and $\boldsymbol{y}_0$ as follows

$$\boldsymbol{x}^i = \boldsymbol{x}_0 + \Phi_{in}\boldsymbol{\alpha}^i + \boldsymbol{\delta}^i_{in} \quad , \quad \boldsymbol{y}^i = \boldsymbol{y}_0 + \Phi_{out}(W\boldsymbol{\alpha}^i) + \boldsymbol{\delta}^i_{out} \tag{1}$$

where $\boldsymbol{\alpha}^i$ is the d_{in} dimensional representation of the i-th training input $\boldsymbol{x}^i$, and is sampled per dimension from a Gaussian distribution with decreasing variance. The matrices Φ_{in} and Φ_{out} contain random orthonormal vectors, and W is the random orthogonal linear regression coefficient matrix of size $d_{in} \times d_{out}$. The noise $\boldsymbol{\delta}^i_{in}$ and $\boldsymbol{\delta}^i_{out}$ is sampled from a Gaussian distribution with covariance

$$\Sigma_\delta = \zeta\left[\Sigma_{common} + \Sigma_{sampleSpec}\right] = \zeta\left[\Theta^T\Gamma\Theta + S^TS\right] \quad , \tag{2}$$

where ζ regulates the amount of noise. Noise consists of Σ_{common} producing a shared noise structure, and a sample specific term $\Sigma_{sampleSpec}$. The first is calculated from a random low dimensional space spanned by the columns of Θ with a magnitude defined by the diagonal of Γ, while the second is calculated from a random matrix S consisting of uniform gaussian samples, except of a diagonal band exhibiting larger variance. Such noise models the common uncertainty structure shared by all samples, such as a larger error along the shape surface than perpendicular to it as well as the smoothness of real shapes making noise highly correlated in a neighborhood.

Experiments and results. We conducted experiments with 200 dimensional input and output of intrinsic dimensionality $d_{in} = 5$, $d_{out} = 3$. We chose ζ such that the noise magnitude was 10% of the real data variation. Training set size was 30, unless stated differently. For the first two experiments the exact noise covariance was used as sample uncertainty estimate. We used the root mean squared (RMS) distance of the predicted output and the noise-free test output for evaluation. The reported results are averages of 100 random realizations.

The first experiment simulates applications where training and test shapes are produced with the same segmentation method and imaging protocol. This scenario is most common in real shape regressions. For this simulation training and test input and output is given same amount of noise. Results for different training set sizes shown in Fig. 1 indicate only small differences between regression methods. Results of 2-sample t-tests on the pooled training set sizes are shown in Tab. 1. The same experiment with 50% noise gave similar results.

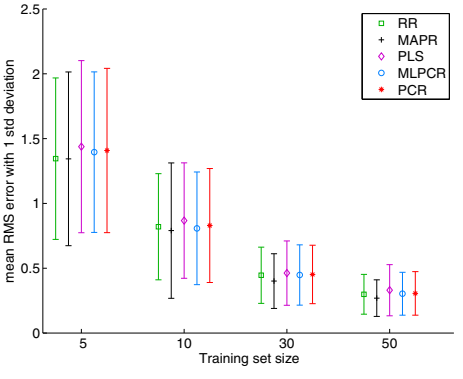

Fig. 1. Error as function of training set size for different regression methods, when noise in training and test data is similar (Simulation experiment 1)

Table 1. Significance table: The upper triangle summarizes results when training and test data are similar (S1 - simulation 1, H - Heart). The lower triangle summarizes experiments with different training and test data (S2 - simulation 2, F - femur 2D-3D). The arrows point to the significantly outperforming method (5% confidence level).

	S1 H	S1 H	S1 H	S1 H
RR	↓ 0	← ←	0 ←	← ←
→ →	MAPR	← ←	← ←	← ←
↑ ↑	↑ ↑	PLS	↓ 0	↓ ↓
0 ↑	↑ ↑	→ ↑	MLPCR	0 ↓
0 ↑	↑ ↑	→ 0	0 →	PCR
S2 F	S2 F	S2 F	S2 F	

The second experiment simulates applications where the training shapes are created from a different modality or with a different method than the test data, e.g. training shapes from CT and test data from X-ray [15]. They exhibit thus different noise structure. We simulated the training set with 10% noise, while test sample noise varied between 5% and 100%. Results are shown in Figure 2 and Table 1.

In the third experiment we assessed the effect of wrongly estimated amount of test uncertainty. We used the setup of the previous experiment with 10% training set noise, and twice as large test noise, and varied the estimated magnitude of the test uncertainty. Note, that the covariance structure of the test uncertainty stayed the same, only the magnitude varied. Results are shown in Figure 3.

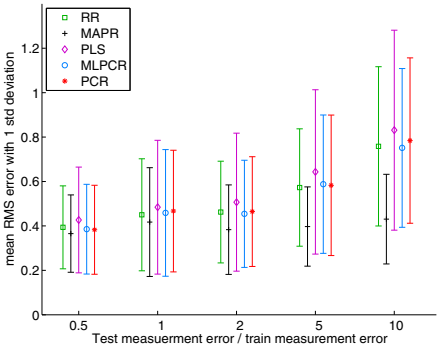

Fig. 2. RMS error for different regression methods as function of the ratio between test and training set noise magnitude (Simulation experiment 2)

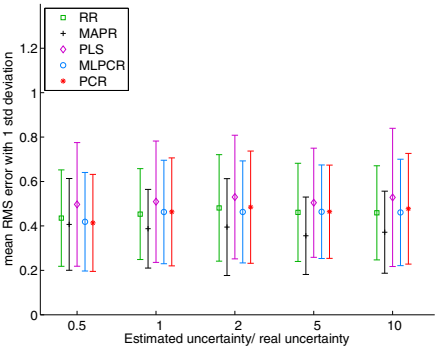

Fig. 3. RMS error for different regression methods as function of the ratio between estimated and actual test noise magnitude (Sim. experiment 3)

4 Real World Datasets

Left ventricle prediction. Estimation of a patient-specific 4D beating heart model is useful for integration of pre-operative 3D CTA in X-ray guided coronary angioplasty procedures [9]. We investigated the accuracy of predicting the end-systolic left ventricle shape from end-diastole as a real world example for the scenario of similar test and train uncertainties.

We extracted 150 end-diastolic and end-systolic left ventricle (LV) shapes from 4D cardiac CTA segmentations, derived by atlas segmentation at end-diastole and subsequent 4D registration as proposed by [9]. The atlases contained 5899 landmark points denoting the LV endo- and epi-cardium. The end-diastolic segmentations were obtained by averaging the landmark points of the eight registered atlases. The spread of these points was used to estimate the landmark uncertainty. Example segmentations with color-coded uncertainty magnitude per landmark are shown in Fig. 4).

The RMS point-to-point (P2P) distance between predicted and segmented end-systolic shape was evaluated in leave-one-out experiments. The regularization parameters in PCR, PLS and RR were optimized via 15-fold cross-validation, and MLPCR was assigned the same number of modes as PCR. Due to the high dimensional input, MAPR was performed in the data subspace, and was optimized per test sample. Results for the different regression methods are shown in Fig. 6 with Table 1 indicating statistical significance.

2D-3D femur reconstruction. Estimation of 3D patient specific bony anatomy from one or more X-Ray images would be beneficial to minimize acquisition costs and radiation dose. Regression based reconstruction after rigid alignment and 2D-3D correspondence generation was proposed by Zheng et al. [15]. We focused on the reconstruction step assuming known projection parameters,

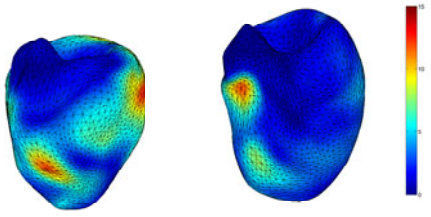

Fig. 4. Two end-diastolic left ventricle shapes with color-coded landmark uncertainty, derived from CTA data by multi-atlas segmentation

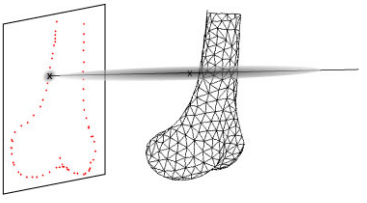

Fig. 5. The 2D-3D uncertainty structure. The 3D position of a projected 2D silhouette point is highly uncertain in the projection direction.

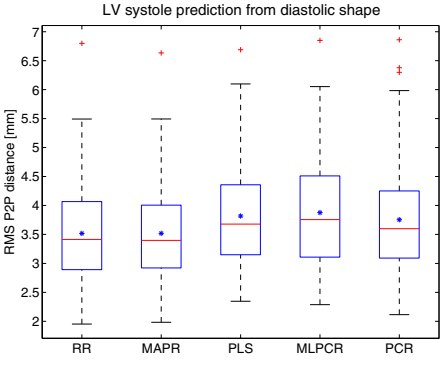

Fig. 6. Box plots showing RMS errors between estimated and reference LV systolic shape. The blue star marks the mean RMS error.

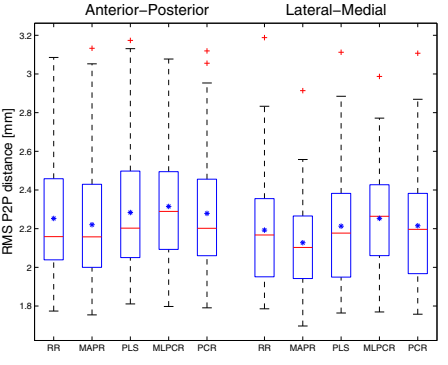

Fig. 7. Box plots showing RMS errors between estimated and reference 3D femur shape for two projection directions

optimal pose, and perfect 2D-3D correspondence. This is an example of the scenario when training and test uncertainty differs.

The regression coefficient matrix was calculated with as input the 3D silhouette landmarks and as output the entire 3D shape. The test input was generated from the 2D projected silhouette landmarks as the 3D point on the projection ray, which is closest to its corresponding landmark of the mean shape. Therefore, the test input was given a large uncertainty along the projection rays, and a small one perpendicular to it, as shown in Fig. 5. The training shapes were rigidly aligned based on the silhouette landmark points only, as proposed in [2]. We performed leave-one-out experiments to estimate the 3D distal femur shape from one projection.

The femur model was created from semi-automatic segmentations of 29 patient CTs and 13 cadavers from both sexes with varying ages. Point correspondence was generated with the GAMES algorithm [5]. The resulting shape model contained 33 modes of variations with 95% retained variance.

We compared the regression methods based on the RMS P2P distance between predicted and test output. Results for the Anterior-Posterior (AP) and the Lateral-Medial (LM) projection directions are shown in Fig. 7 with significance indication in Table 1.

5 Discussion and Conclusions

Despite the growing interest in regression based shape prediction, no study had before systematically compared different regression methods for shape analysis. We aimed to fill this gap by comparing several linear regression methods with a special focus on landmark position uncertainties. We defined two uncertainty scenarios.

In the first scenario test and train uncertainties were similar. Simulation experiment 1 and the LV prediction experiment (Fig. 1, Fig. 6, and Table 1) showed little differences between different linear regression methods. Uncertainty incorporation via MLPCR did not improve results. A similar phenomenon was reported in [12] for MLPCR without intercept optimization. Results might improve without the simplifications of independent noise and no intercept, but the high computational requirements make it currently impractical for shape regression. Including uncertainty of the test sample via MAPR marginally improved predictions in the simulations, but did not improve upon RR on the heart data. While the regularization constant of MAPR depends on the test uncertainty, and has to be re-optimized for every new test sample, the optimization of the standard regression methods is solely based on the training data, can thus be precomputed. We observed that the average cross-validation curve used to chose the shrinkage parameter was smooth for RR, and spiky for PCR and PLS due to the discrete inclusion of directions, and the false correlation magnitudes in case of a small and noisy training set. This might be one reason why RR slightly but consistently outperformed PCR and PLS. Overall, in this scenario we found RR to perform slightly better than to other linear regression methods.

In the second scenario training and test uncertainties differ largely. Simulation experiment 2 and the 2D-3D femur reconstruction experiment (Fig. 2, Fig. 7 and Tab. 1) consistently showed that including knowledge of the test uncertainty in the prediction via MAPR significantly improves results. The practical gain depends on the difference in uncertainty and on how well the assumed normal distribution fits the shape variability and the noise. Real world data might therefore produce a smaller gain with MAPR. Estimation errors of the noise magnitude with the correct covariance structure are, however, tolerable within a large interval (simulation 3, Fig. 3). The performance of RR,PLS,PCR and MLPCR did not change from scenario 1. Therefore, we conclude that in such situations including landmark uncertainty of the test sample may be beneficial.

In our comparison the predicted shape is the end result. Our work could be naturally extended to include the remaining variance of the prediction.

We believe that the presented comparison study gives insights for researchers working on shape prediction in various domains, and helps in the selection of a suitable regression method for future applications.

The Netherlands Organization for Scientific Research (NWO) and SenterNovem, project IGIT4Health is greatly acknowledged for its financial support.

References

1. Ablitt, N.A., Gao, J., Keegan, J., Stegger, L., Firmin, D.N., Yang, G.Z.: Predictive cardiac motion modeling and correction with partial least squares regression. IEEE Trans. Med. Imaging 23(10), 1315–1324 (2004)
2. Baka, N., de Bruijne, M., Niessen, W., Reiber, J.H.C., Lelieveldt, B.: Confidence of model based shape reconstruction from sparse data. In: IEEE ISBI (2010)
3. Blanc, R., Syrkina, E., Székely, G.: Estimating the confidence of statistical model based shape prediction. In: Prince, J.L., Pham, D.L., Myers, K.J. (eds.) IPMI 2009. LNCS, vol. 5636, pp. 602–613. Springer, Heidelberg (2009)
4. de Bruijne, M., Lund, M.T., Tankó, L.B., Pettersen, P.C., Nielsen, M.: Quantitative vertebral morphometry using neighbor-conditional shape models. Med. Image Anal. 11(5), 503–512 (2007)
5. Ferrarini, L., Olofsen, H., Palm, W., van Buchem, M., Reiber, J., Admiraal-Behloul, F.: Games: growing and adaptive meshes for fully automatic shape modeling and analysis. Med. Image Anal. 11, 302–314 (2007)
6. Hastie, T., Tibshirani, R., Friedman, J.: The Elements of Statistical Learning (Data mining, Inference, and Prediction). Springer, Heidelberg (2009)
7. Klinder, T., Lorenz, C., Ostermann, J.: Prediction framework for statistical respiratory motion modeling. In: Jiang, T., Navab, N., Pluim, J., Viergever, M. A. (eds.) MICCAI 2010. LNCS, vol. 6363, pp. 327–334. Springer, Heidelberg (2010)
8. Liu, T., Shen, D., Davatzikos, C.: Predictive Modeling of Anatomic Structures Using Canonical Correlation Analysis. In: IEEE ISBI, pp. 1279–1282 (2004)
9. Metz, C., Baka, N., Kirisli, H., Schaap, M., van Walsum, T., Klein, S., Neefjes, L., Mollet, N., Lelieveldt, B., de Bruijne, M., Niessen, W.: Conditional shape models for cardiac motion estimation. In: Jiang, T., Navab, N., Pluim, J., Viergever, M. (eds.) MICCAI 2010. LNCS, vol. 6361, pp. 452–459. Springer, Heidelberg (2010)
10. Rajamani, K.T., Styner, M.A., Talib, H., Zheng, G., Nolte, L.P., Ballester, M.A.: Statistical deformable bone models for robust 3d surface extrapolation from sparse data. Medical Image Analysis 11(2), 99–109 (2007)
11. Rao, A., Aljabar, P., Rueckert, D.: Hierarchical statistical shape analysis and prediction of sub-cortical brain structures. Med. Image Anal. 12(1), 55–68 (2008)
12. Reis, M.S., Saraiva, P.M.: A comparative study of linear regression methods in noisy environments. J. Chemometrics 18, 526–536 (2004)
13. Wentzell, P.D., Andrews, D.T., Hamilton, D.C., Klaas, F., Kowalski, B.R.: Maximum likelihood principal component analysis. J. Chemometrics 11, 339–366 (1997)
14. Yang, Y.M., Rueckert, D., Bull, A.M.J.: Predicting the Shapes of Bones at a Joint: Application to the Shoulder. Computer Methods in Biomechanics and Biomedical Engineering 11(1), 19–30 (2008)
15. Zheng, G., Gollmer, S., Schumann, S., Dong, X., Feilkas, T., Ballester, M.A.G.: A 2d/3d correspondence building method for reconstruction of a patient-specific 3d bone surface model using point distribution models and calibrated x-ray images. Med. Image Anal. 13(6), 883–899 (2008)

SpringLS: A Deformable Model Representation to Provide Interoperability between Meshes and Level Sets

Blake C. Lucas[1,2], Michael Kazhdan[2], and Russell H. Taylor[2]

[1] Johns Hopkins Applied Physics Laboratory, Laurel, MD, USA
[2] Johns Hopkins University, Baltimore, MD, USA
{blake,misha}@cs.jhu.edu, rht@jhu.edu

Abstract. A new type of deformable model is presented that merges meshes and level sets into one representation to provide interoperability between methods designed for either. The key idea is to use a constellation of triangular surface elements (springls) to define a level set. A Spring Level Set (SpringLS) can be interpreted as a mesh or level set and used in place of them in many instances. There is no loss of shape information in the transformation from triangle mesh or level set into SpringLS. As examples, we present results for joint segmentation/spherical mapping of a human brain cortex and atlas/non-atlas segmentation of a pelvis.

Keywords: mesh, level set, deformable model, segmentation, active contour.

1 Introduction

Deformable models are geometric representations of objects that deform (change shape) due to forces applied at their boundary. Deformable models are applicable to a broad range of problems in image analysis and computer vision including: reconstruction, non-rigid registration, image segmentation, atlasing, and motion tracking. There are two major model representations: meshes [1] and level sets [2, 3]. Deformable model methods usually favor a particular representation (i.e. meshes for 2D/3D registration and level sets for image segmentation). However, large systems that use a mixture of methods are forced to transform one representation into another in order to use the preferred representation for each method. This strategy leads to loss of information and less flexibility in design of the system. For examples of systems that use a mixture of representations, see Tosun et al. [4] and Wand et al. [5].

The Spring Level Set (SpringLS) representation merges meshes and level sets into a single geometric representation that preserves the strengths of both. SpringLS can be interpreted as a mesh or level set, and no shape information is lost in the transformation from triangle mesh or level set into SpringLS. SpringLS is intended for methods that employ a mixture of mesh and level set techniques, but is applicable to almost all deformable model methods. As examples, we apply SpringLS to joint segmentation/spherical mapping of a human brain cortex and atlas/non-atlas segmentation of a pelvis.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 442–450, 2011.

2 Background

Meshes. Meshes were the earliest representation for deformable models [1]. In this framework, the model is deformed by perturbing mesh vertices. The model's boundary is explicitly tracked by remembering the trajectory of each vertex. As a section of the mesh expands or contracts, sharp creases, edges, self-intersections, or triangle flips can develop. Sharp edges and other mesh artifacts violate a common material property that most objects represented by deformable models are smooth and plastic. To reduce artifacts, the mesh must be regularized and re-sampled (remeshed) periodically. Because mesh triangles must be connected to form a watertight model, remeshing is non-trivial [6] and interferes with vertex tracking. Remeshing becomes more of a nuisance if the mesh is not allowed to self-intersect or is allowed to change topology. For these reasons, triangle meshes have become unpopular for applications where the model 1) undergoes large non-rigid deformations that require remeshing; 2) the model is expected to change topology; 3) the model is likely to self-intersect.

Level Sets. The level set method [2, 3] represents a deformable model as a 3D image where the image intensity at each voxel is a distance measurement to the surface of the object. Distance measurements are signed: negative values are inside and positive values are outside the object. A triangle mesh can be extracted by computing the iso-surface corresponding to the zero level set of the image. The level set representation has several advantages over deformable meshes: 1) no need for self-intersection removal; 2) topology change is easy; 3) no need to remesh. These properties have made level sets the popular choice for image segmentation and fluid-like non-rigid deformation.

Level sets are difficult to use for registration and tracking tasks because there's no innate ability to track vertices as in the mesh deformation framework. The surface only exists when an iso-surface is extracted from the level set. Furthermore, the level set is stored as an image that is re-sampled at each time step. Re-sampling an image acts as a low-pass filter that results in feature loss as a function of the number of time steps, even if the motion is rigid (i.e. global registration) or divergence free (i.e. incompressible fluid flow).

Hybrid Representations. Attempts have been made to unify deformable model representations with varying success [7-13]. Of these, the Marker Level Set (MLS) [11] is closest to this work. The MLS method maintains a set of particles located on the level set's zero iso-level. Since particles lie exactly on the zero iso-level, they can be used for tracking the model's boundary. After each level set and particle advection step, the level set is corrected so that particles continue to lie on the level set's zero iso-level. Particles are added to cover the zero iso-level and deleted to prevent over-sampling. The MLS method associates a color with each particle and interpolates the color for new particles based on their neighbors.

The philosophical difference between SpringLS and MLS is that springl surface elements define the model's level set, whereas MLS use particles to correct errors in the level set. MLS requires the deformation method to have an equivalent level set and parametric interpretation in order to deform both representations. Movement of the auxiliary level set with SpringLS is passive and independent of the deformation

method. SpringLS can be applied more broadly to deformation methods for which there is only a parametric interpretation (i.e. Point Distribution Models (PDM) [14]), enabling true interoperability between methods designed for level sets and meshes.

3 Method

A springl is a triangular surface element consisting of a particle and 3 springs connecting the particle to each of the triangle's vertices (see Fig. 1a). Each springl describes a basis function which defines the unsigned clamped distance $d_n(\boldsymbol{X})$ to the triangle. The combination of these distance functions form the unsigned level set $\omega(\boldsymbol{X})$. Support of a springl is represented by a capsule, whose boundary is the iso-surface corresponding to $d_{max} = 0.5$ voxels. Springs, vertices, and particles are coplanar, and the angles between springs are fixed. An auxiliary signed level set $\varphi(\boldsymbol{X})$ is maintained and evolves with the particles. The level set augments the particle representation in several ways: 1) the signed distance function indicates regions that are inside or outside the model; 2) the signed level set indicates when new springls need to be added or removed; 3) the iso-surface extracted from the level set is a watertight triangle mesh representation of the model. The deformable model is stored as three data structures: a triangle mesh representing surface elements, a point cloud of particles, and a 3D image representing the signed level set.

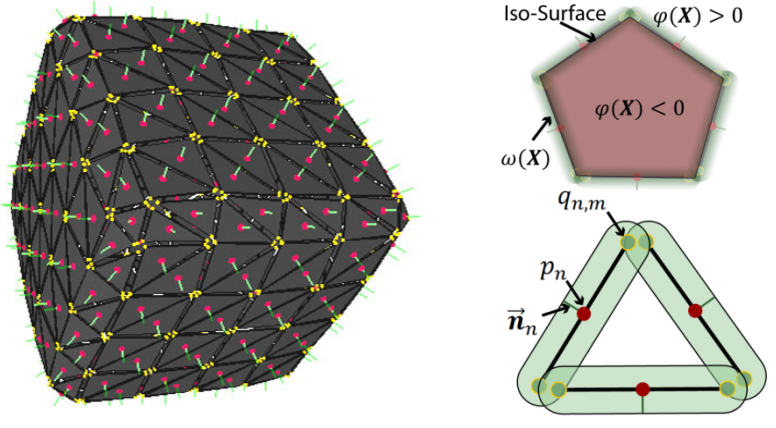

Fig. 1. (a) Model represented by springls showing particles (red), vertices (yellow), normals (green), and surface elements (black). (b) Signed level set, springls, and their capsules (green).

A springl ($\boldsymbol{S}_n$) is represented as five points describing the particle location p_n, correspondence point location a_n, and triangle vertices $q_{n,m}$ (see Fig. 1b). The model is deformed by incrementally advecting springls with first-order lagrangian methods (**`Advect`**). The deformation may be driven by pressure $\rho(\cdot)$ in the normal direction $\vec{\boldsymbol{n}}_n$ of a springl, external velocity field $\vec{\boldsymbol{\sigma}}(\cdot)$, and/or local affine transformation $\mathbf{A}(\cdot)$ derived from an atlas.

$$\frac{\partial p_n}{\partial t}(\mathrm{t}) = \lambda_\rho \vec{\boldsymbol{n}}_n \rho\big(p_n(t)\big) + \lambda_\sigma \vec{\boldsymbol{\sigma}}\big(p_n(t)\big) + \lambda_{atlas}\mathbf{A}(\boldsymbol{S}_n)p_n(t), \text{ and} \tag{1}$$

$$\frac{\partial q_{n,m}}{\partial t}(\mathrm{t}) = \lambda_\rho \vec{\boldsymbol{n}}_n \rho\big(p_n(t)\big) + \lambda_\sigma \vec{\boldsymbol{\sigma}}\big(p_n(t)\big) + \lambda_{atlas}\mathbf{A}(\boldsymbol{S}_n)q_{n,m}(t). \tag{2}$$

After each advection step, particles are fixed and the shape/orientation of each springl is adjusted through a relaxation process. The purpose of relaxation is to orient surface elements so their normals point outward. Motion is dictated by forces that attract neighboring vertices to the edges of nearby triangles. The force ($\overrightarrow{v_{n,m}}$) acting on a vertex due to its k closest points on nearby triangle edges $q_{n,m,k} \epsilon \mathbb{R}^3$ is as follows:

$$\overrightarrow{v_{n,m}} = \lambda_v \tanh\big(\lambda_c \big|\overrightarrow{c_{n,m}}\big|\big)\, \overrightarrow{c_{n,m}} / \big|\overrightarrow{c_{n,m}}\big|, \text{ where} \tag{3}$$

$$\overrightarrow{c_{n,m}} = \textstyle\sum_k \operatorname{atanh}(w_{n,m,k}) \frac{q_{n,m,k} - q_{n,m}}{|q_{n,m,k} - q_{n,m}|}, \text{ and} \tag{4}$$

$$w_{n,m,k} = R^{-1}(|q_{n,m,k} - q_{n,m}| - 2r). \tag{5}$$

The resultant force due to the spring attached to the vertex is

$$f_\kappa = \kappa\big(l - \big|q_{n,m} - p_n\big|\big), \tag{6}$$

where $\lambda_v = 0.05$ is the max force, $\lambda_c = 4$ is the kernel smoothness, $l = 0.1$ voxels is the spring rest length, $\kappa = 0.08$ is the spring constant, $r = 0.05$ voxels is the vertex radius, and $R = 0.6$ voxels is the nearest-neighbor range. The rotational moment due to forces applied on the vertex is

$$\overrightarrow{m_{n,m}} = \overrightarrow{v_{n,m}} \times \overrightarrow{s_{n,m}}, \text{ where} \tag{7}$$

$$\overrightarrow{s_{n,m}} = \big(q_{n,m} - p_n\big)/\big|q_{n,m} - p_n\big|. \tag{8}$$

After summing moments, the resultant moment $\overrightarrow{m_n}$ indicates the amount and axis of rotation, described by the 3x3 matrix $\boldsymbol{M}_n$. The tanh / atanh weighting functions in eq. (3) and eq. (4) dampen rotational motion that can lead to instability in a surface element's orientation. The final update equation is

$$q_{n,m}^{t+1} = p_n + \boldsymbol{M}_n\Big(q_{n,m}^t - p_n + \overrightarrow{s_{n,m}}\big(\overrightarrow{v_{n,m}} \cdot \overrightarrow{s_{n,m}} + f_\kappa\big)\Big). \tag{9}$$

The relaxation process (**`Relax`**) in eq. (9) is repeated for 20 iterations (see *Algorithm 1*). Parameter choices express a tradeoff between minimizing the number of springls needed to cover the zero iso-level while minimizing gap formation. Parameters were selected based on a small number of examples apart from the experiments presented in this paper and have not been changed in these or any subsequent experiments.

Relaxation helps insure that the union of all springl capsules covers the zero iso-level of the signed level set. If springls are unable to cover the zero iso-level through relaxation, gaps must be filled by adding more springls in a subsequent step (**`FillGaps`**) to prevent the model from tearing. When a zero-crossing of the level set is exposed, a springl is added to cover the zero-crossing. Zero-crossings are computed from the centroids of triangles generated from the signed level set's iso-surface. These triangles are reused to fill the hole with a springl in the same shape and position.

The level set is updated at each time step to track the particles (**`Evolve`**). This is done by constructing an unsigned level set $\omega(\boldsymbol{X})$ (eq. (10)) that is the clamped minimum distance to all springls. The signed level set $\varphi(\boldsymbol{X})$ is evolved to minimize the energy function in eq. (11). The free parameter λ controls the model's smoothness. A springl is destroyed if after evolving the signed level set, a springl's particle is more than $(1+\varepsilon)d_{max}$ from the zero iso-level (**`Contract`**). We choose $\varepsilon = 0.25$ in all cases.

$$\omega(\boldsymbol{X}) = \min\{d_{max}, d_1(\boldsymbol{X}) \quad \ldots \quad d_N(\boldsymbol{X})\}, \text{ and} \tag{10}$$

$$E = \int (\omega(\boldsymbol{X}) + \lambda|\nabla\varphi(\boldsymbol{X})|)\delta(\varphi(\boldsymbol{X}))d\boldsymbol{X}. \tag{11}$$

For large parametric deformations where the CFL number exceeds 1, it is usually faster to convert the unsigned level set to a signed level set [15, 16] than to evolve the signed level set with active contour methods. We perform the conversion by growing the background region and then negating the unsigned level set in the foreground region. This can be done robustly with a coarse-to-fine strategy [16] to prevent the background region from leaking through gaps between springls.

Springls are re-sampled every $M = 5$ iterations to regularize the sampling distribution and triangle quality (**`Resample`**). Triangles are split along their longest edge if the length of that edge exceeds a threshold (1.5 voxels). If a triangle's angles fall outside a tolerable range $[20^{\circ}, 160^{\circ}]$, then the springl is removed. Removing poor quality triangles reduces unstable rotation of springls in the relaxation phase.

Springls maintain a mapping from each particle to the centroid of a triangle on the original model. The initial mapping is an identity mapping ($a_n = p_n$). When a springl is split, the mapping is duplicated and when a springl is added, the mapping is chosen to be the average point mappings for neighboring springls. This method produces mappings that lie slightly off the original surface. To prevent correspondence points from drifting from the original surface, correspondence points are moved along the gradient of the distance (φ_{ref}) to the original surface until convergence (eq. (12)).

$$a_n(t+1) = a_n(t) - \lambda\varphi_{ref}\big(a_n(t)\big)\nabla\varphi_{ref}\big(a_n(t)\big). \tag{12}$$

The model deformation process (**`Deform`**) is outlined in *Algorithm 2*.

Algorithm 1. **Relax**
foreach springl n **do**
  foreach vertex m **do**
    foreach neighbor k of vertex m **do**
      Compute $w_{n,m,k}$
      Accumulate $\overrightarrow{c_{n,m}}$
    Compute $\overrightarrow{s_{n,m}}$, $\overrightarrow{v_{n,m}}$, and $\overrightarrow{m_{n,m}}$
    Accumulate $\overrightarrow{m_n}$
  foreach vertex m **do**
    Compute $q_{n,m}^{t+1}$

Algorithm 2. **Deform**
for $k = 1{:}K$ **do**
  Advect
  Relax
  if $k \bmod M == 0$ **then**
    Contract
    Resample
    Relax
    Evolve
    FillGaps
  else Evolve

4 Results

SpringLS was applied to active contour image segmentation [17] of objects driven under pressure forces from image intensities. It was implemented as a mixture of Java and OpenCL for the CPU on a PC with Dual 2.53GHz Intel Xeons and 12GB of RAM. Parameter settings and grid size ($256 \times 256 \times 256$) were fixed for all experiments. Reported segmentation errors are measured as the average minimum distance from mesh vertices on the segmented mesh to target iso-surface. Experiments not initialized with an atlas were repeated with 4 different initializations (Fig. 2a).

Fig. 2b,c shows simultaneous segmentation of a T1 MRI image and spherical mapping of a human brain cortex, which are two important tasks in cortical surface analysis [18]. The MRI image was pre-processed with TOADS to produce WM/GM soft-membership images [19]. Spherical mapping was accomplished by initializing the segmentation with a sphere and then explicitly tracking the surface with springls to find the WM/GM surface in the WM membership image. The model was then evolved outward to find the Pial surface in the WM+GM membership image. The resulting reconstruction has a mapping from each point on either surface to a location on the sphere. Table 1 reports runtime statistics and compares SpringLS to an equivalent level set implementation in terms of surface-to-surface distance measured from the SpringLS iso-surface and DICE coefficient between level set segmentations.

Table 1. Results comparing SpringLS to an equivalent level set method. Algorithm terminated when the DICE coefficent between successive resampling cycles exceeded a threshold.

Experiment	Surface Distance	DICE	Iterations	Springls	Triangles	Time
Pelvis	0.14±0.20 mm	0.9993	290-390	50K-55K	200K-217K	15-19 min
Pelvis \w atlas	0.25±0.32 mm	0.9994	110	39K	149K	4 min
WM/GM	0.18±0.21 mm	0.9985	280-370	116K-118K	366K-372K	19-30 min
Pial	0.16±0.16 mm	0.9989	60	112K-113K	306K	5-6 min

Fig. 3a shows segmentation of a pelvis from a CT image when initialized with a cube and highlights SpringLS's ability to change topology. Initializing the segmentation process with a cube or other object shown in Fig. 2a causes over segmentation of the pelvis to include the femurs and spine as well. To mitigate this

problem, we incorporate an atlas based approach. A PDM statistical atlas of the pelvis was constructed with the method from Seshamani et al. [20]. Analogous statistical atlas methods have been developed for level sets [21], but these representations are not equivalent.

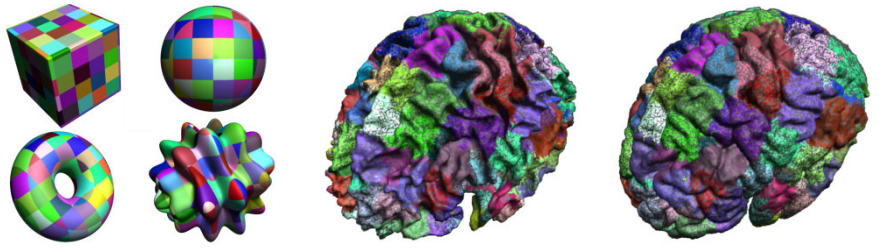

(a) (b) (c)

Fig. 2. Ray-cast renderings of SpringLS showing springls projected onto iso-surface. (a) Initial shapes. (b) WM/GM surface and (c) Pial surface segmentations when initialized with a sphere.

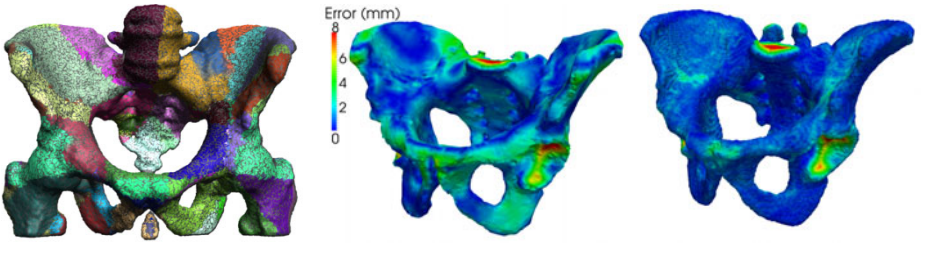

(a) (b) (c)

Fig. 3. (a) Pelvis segmentation when initialized with a cube. Atlas based segmentation showing (b) registered PDM and (c) final segmentation.

The segmentation result in Fig. 3a can be improved by combining level set techniques with a parametric atlas. The atlas was registered (rigid + global scale) to the CT image. The first 10 mode weights were optimized in increasing order to reduce the average distance from the atlas to target iso-surface in CT. Because the initial registration was fairly good (1.56±1.36 mm), optimizing the mode weights modestly improved the segmentation result to 1.40±1.32 mm. The registered mesh was then treated as a constellation of springls and advected towards the target iso-level with an external velocity field produced by Gradient Vector Flow (GVF) [22] and pressure forces, reducing the error to 0.95±1.02 mm. This atlas based method produces a better pelvis segmentation (Fig. 3c) than the non-atlas approach and provides a mapping from each springl back to the atlas, enabling transfer of region labels on the atlas to the segmented pelvis. This technique is also significantly faster than the non-atlas based approach (see Table 1).

5 Conclusion

Spring Level Sets (SpringLS) merge meshes and level sets into a single representation to provide interoperability between methods designed for either. The key idea is to use triangular surface elements to define a level set. Insisting the surface elements be triangle shaped insures no shape information is lost in the transformation from triangle mesh into SpringLS. One may choose not to relax or resample a subset of springls to preserve sharp features or tracking information. Because SpringLS uses disconnected surface elements, the object can change topology, track points, and undergo parametric deformations. The auxiliary level set provides a watertight representation of the model's boundary that cannot self-intersect, and simple rules have been described for adding and destroying surface elements based on the level set representation. We have demonstrated that image segmentation with SpringLS produces results very similar to an equivalent level set implementation, and a registered PDM atlas can be converted into a SpringLS and deformed to produce a better segmentation than without an atlas. SpringLS is open source and distributed as part of the Java Image Science Toolkit (http://www.nitrc.org/projects/jist) to encourage the development of new image analysis systems that are true mixtures of mesh and level set methods.

Acknowledgments. This research was supported by a graduate student fellowship from the Johns Hopkins Applied Physics Laboratory and internal funds from Johns Hopkins University.

References

1. Terzopoulos, D., Platt, J., Barr, A., Fleischer, K.: Elastically deformable models. Computer Graphics 21, 205–214 (1987)
2. Sethian, J.: Level set methods and fast marching methods: evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science. Cambridge Univ. Pr., Cambridge (1999)
3. Osher, S., Fedkiw, R.: Level set methods and dynamic implicit surfaces. Springer, Heidelberg (2003)
4. Tosun, D., Prince, J.L.: A geometry-driven optical flow warping for spatial normalization of cortical surfaces. IEEE Trans. on Medical Imaging 27, 1739–1753 (2008)
5. Wand, M., Jenke, P., Huang, Q., Bokeloh, M., Guibas, L., Schilling, A.: Reconstruction of deforming geometry from time-varying point clouds, pp. 49–58 (2007)
6. Alliez, P., Ucelli, G., Gotsman, C., Attene, M.: Recent advances in remeshing of surfaces. In: Shape Analysis and Structuring, pp. 53–82 (2008)
7. Enright, D., Fedkiw, R., Ferziger, J., Mitchell, I.: A Hybrid Particle Level Set Method for Improved Interface Capturing. Computational Physics 183, 83–116 (2002)
8. Müller, M., Keiser, R., Nealen, A., Pauly, M., Gross, M., Alexa, M.: Point based animation of elastic, plastic and melting objects, pp. 141–151 (2004)
9. Bargteil, A., Goktekin, T., O'brien, J., Strain, J.: A semi-Lagrangian contouring method for fluid simulation. ACM Trans. on Graphics 25, 38 (2006)
10. Müller, M.: Fast and robust tracking of fluid surfaces. In: SIGGRAPH, pp. 237–245 (2009)

11. Mihalef, V., Sussman, M., Metaxas, D.: Textured liquids based on the marker level set. Computer Graphics Forum 26, 457–466 (2007)
12. Pons, J., Hermosillo, G., Keriven, R., Faugeras, O.: Maintaining the point correspondence in the level set framework. Computational Physics 220, 339–354 (2006)
13. Pauly, M., Keiser, R., Kobbelt, L.P., Gross, M.: Shape modeling with point-sampled geometry. ACM Trans. on Graphics (TOG) 22, 641–650 (2003)
14. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active shape models-their training and application. Computer Vision and Image Understanding 61, 38–59 (1995)
15. Mullen, P., De Goes, F., Desbrun, M., Cohen Steiner, D., Alliez, P.: Signing the Unsigned: Robust Surface Reconstruction from Raw Pointsets. Computer Graphics Forum 29, 1733–1741 (2010)
16. Sharf, A., Lewiner, T., Shamir, A., Kobbelt, L., Cohen–Or, D.: Competing fronts for coarse–to–fine surface reconstruction. Computer Graphics Forum 25, 389–398 (2006)
17. Caselles, V., Kimmel, R., Sapiro, G.: Geodesic active contours. International Journal of Computer Vision 22, 61–79 (1997)
18. Chung, M.K., Robbins, S.M., Dalton, K.M., Davidson, R.J., Alexander, A.L., Evans, A.C.: Cortical thickness analysis in autism with heat kernel smoothing. Neuroimage 25, 1256–1265 (2005)
19. Bazin, P.L., Pham, D.L.: Topology-preserving tissue classification of magnetic resonance brain images. IEEE Trans. on Medical Imaging 26, 487–496 (2007)
20. Seshamani, S., Chintalapani, G., Taylor, R.: Alternative Statistical Models for Bone Modelling. In: SPIE Medical Imaging (2011)
21. Tsai, A., Yezzi Jr., A., Wells, W., Tempany, C., Tucker, D., Fan, A., Grimson, W.E., Willsky, A.: A shape-based approach to the segmentation of medical imagery using level sets. IEEE Trans. on Medical Imaging 22, 137–154 (2003)
22. Xu, C., Prince, J.L.: Snakes, shapes, and gradient vector flow. IEEE Trans. on Image Processing 7, 359–369 (2002)

Deformable Segmentation via Sparse Shape Representation

Shaoting Zhang[2] , Yiqiang Zhan[1], Maneesh Dewan[1], Junzhou Huang[2], Dimitris N. Metaxas[2], and Xiang Sean Zhou[1]

[1]Siemens Medical Solutions, Malvern, PA, USA
[2]Department of Computer Science, Rutgers University, Piscataway, NJ, USA

Abstract. Appearance and shape are two key elements exploited in medical image segmentation. However, in some medical image analysis tasks, appearance cues are weak/misleading due to disease/artifacts and often lead to erroneous segmentation. In this paper, a novel deformable model is proposed for robust segmentation in the presence of weak/misleading appearance cues. Owing to the less trustable appearance information, this method focuses on the effective shape modeling with two contributions. First, a shape composition method is designed to incorporate shape prior on-the-fly. Based on two sparsity observations, this method is robust to false appearance information and adaptive to statistically insignificant shape modes. Second, shape priors are modeled and used in a hierarchical fashion. More specifically, by using affinity propagation method, our deformable surface is divided into multiple partitions, on which local shape models are built independently. This scheme facilitates a more compact shape prior modeling and hence a more robust and efficient segmentation. Our deformable model is applied on two very diverse segmentation problems, liver segmentation in PET-CT images and rodent brain segmentation in MR images. Compared to state-of-art methods, our method achieves better performance in both studies.

1 Introduction

In various applications of medical image segmentation, deformable model has achieved tremendous success, which should be contributed to its joint employment of shape and appearance characteristics. While appearance features provide low level clues of organ boundaries, shape imposes high level knowledge to infer and refine deformable model. However, in some medical image analysis, appearance cues are relatively weaker or even misleading (Fig. 1). In those cases, the best "guess" of the organ boundaries can only come from shape priors, which should be effectively modeled from training shapes. However, effective shape modeling is confronting these challenges, 1) shape variation is complex and cannot always be modeled by a parametric probability distribution; 2) a shape instance derived from image appearance cues (input shape) may have gross errors; and 3) local details of the input shape are difficult to preserve if they are not statistically significant in the training data. Traditional deformable model, e.g., Active Shape Model its extensions [1,6], can not tackle them uniformly.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 451–458, 2011.

In this paper, we propose a deformable model aiming to achieve robust segmentation in the presence of weak/misleading appearance cues. In particular, two novel methods are designed for robust and effective shape prior modeling. *First*, instead of assuming any parametric model of shape statistics, we propose to incorporate shape priors on-the-fly through sparse representation. More specifically, we have two sparsity observations: 1) Given a large shape repository of an organ, a shape instance of the same organ can be approximated by the composition of a sparse set of instances in the shape repository; and 2) gross errors from local appearance cues might exist but these errors are sparse in spatial space. Incorporating these two sparsity priors, our deformable model becomes robust to gross errors and can preserve shape details even they are not statistically significant in the training repository. This shape composition method benefits both the model initialization and refinement. *Second*, instead of modeling global shape priors, we propose to decompose the deformable surface to multiple parts and build shape models on them independently. The partition is accomplished by affinity propagation method [4] based on image and geometry features. Since the shape statistics of local structures often has more compact distribution than global structures, this strategy facilitates better shape modeling and increases algorithm runtime efficiency.

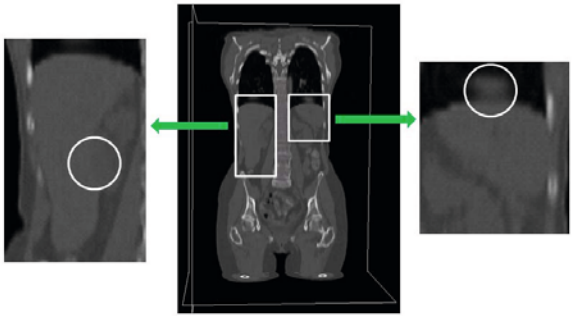

Fig. 1. Middle: whole body low-dose CT data. Left: zoom in of the liver and lung. In the marked region the boundary between the liver and the kidney is hardly observed. The appearance cue is weak because of the low contrast around the boundary. Right: zoom in of the spleen and lung. In the marked region, there is artifact induced by breath. It is part of the lung. Since the image information is misleading here, segmentation methods solely relying on appearance cue may accidentally include this region as spleen.

2 Methodology

Segmentation Framework: To achieve generality, our segmentation framework is designed in the spirit of "data-driven". Fig. 2 shows the workflow of our segmentation system, which consists of offline learning and runtime segmentation stages. In offline learning, 3D volume images along with the manually labeled ground truths are employed to learn the appearance and shape characteristics of the organ under study. More specifically, methods proposed in [10,9] are used to learn landmark detectors and a set of spatially adaptive boundary detectors. Meanwhile, organ surfaces are stored in a shape repository, which will

be exploited to derive shape priors during runtime. Runtime segmentation starts from the initialization of the model surface (represented by a triangular surface) based on automatically detected landmarks and shape priors. The surface then deforms under the guidance of both image appearance cues and shape priors. More specifically, two steps are performed iteratively until convergence. First, the surface model deforms to local places where the learning-based boundary detectors generate higher responses. Next, the locally deformed surface is refined by the shape priors derived from the shape repository.

As discussed before, although learning-based landmark/boundary detectors can tackle reasonable appearance variations [10,9], they might generate wrong responses in the presence of severe imaging artifacts/diseases, and hence mislead the deformable model. In this scenario, shape prior is the only information source to initialize/correct the deformable surface. (Note that shape priors are employed in both landmark-based model initialization and shape refinement in Fig. 2.) Therefore, the effective modeling of shape priors becomes extremely critical to achieve a robust segmentation. Due to page limits, we will focus on the modeling of shape priors in the remainder of this paper.

Shape Prior Modeling via Sparse Composition: Instead of assuming any parametric probabilistic distributions of the shape statistics, our shape prior model is based on two observations: 1) After being aligned to a common canonical space, any shape can be approximated by a sparse linear combination of other shape instances in the same shape category. Approximation residuals might come from inter-subject variations. 2) If the shape to be approximated is derived by appearance cues, residual errors might include gross errors from landmark/boundary detections. But such errors are sparse as well. Accordingly, we aim to incorporate shape priors *on-the-fly* through shape composition, i.e., a shape derived by appearance cues is refined by the approximation of a set of annotated shape instances following the two sparsity observations. It is worth mentioning that sparsity has been adopted in segmentation algorithms in different manners, such as the sparse information models [3], which reconstruct a 3D surface from sparse subcomponents.

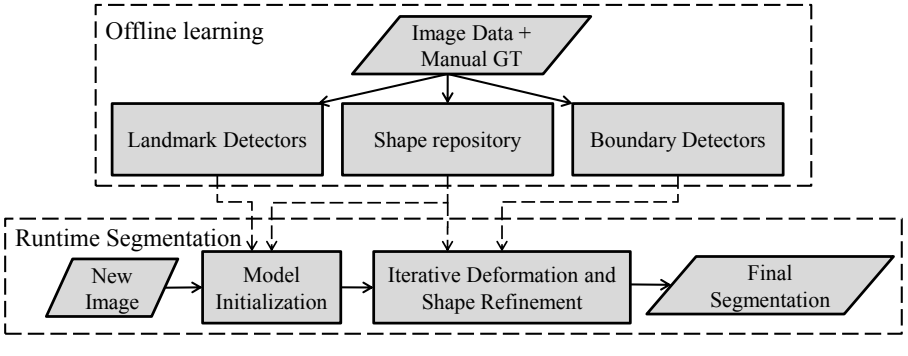

Fig. 2. The workflow of our segmentation framework

In this work, a shape is represented by a triangle mesh which consists of a set of vertices. Denote the input shape as $\mathbf{v}$, where $\mathbf{v} \in \mathbf{R}^{3N}$ is a vector concatenated by 3D coordinates of its N vertices. (In the remainder of this paper, any shape instance is defined as a vector in the same way and has the same dimensionality.) Assume $D = [\mathbf{d}_1, \mathbf{d}_2, ..., \mathbf{d}_K] \in \mathbf{R}^{3N \times K}$ is a large shape repository that includes K accurately annotated shape instances $\mathbf{d}_i$. Note that $\{\mathbf{d}_i, i = 1, 2, 3, ..., K\}$ are pre-aligned using generalized Procrustes analysis [5,1].The approximation of $\mathbf{v}$ by D is then formulated as an optimization problem:

$$\arg\min_{\mathbf{x},\mathbf{e},\beta} \|T(\mathbf{v}, \beta) - D\mathbf{x} - \mathbf{e}\|_2^2, \quad s.t. \|\mathbf{x}\|_0 < k_1, \|\mathbf{e}\|_0 < k_2, \tag{1}$$

where $T(\mathbf{v}, \beta)$ is a global transformation operator with parameter β, which aligns the input shape $\mathbf{v}$ to the common canonical space of D. The key idea of our shape prior modeling lies in the two constraints of (1). In the first constraint, $\mathbf{x} \in \mathbf{R}^K$ denotes the coefficient/weights of linear combination. The $L0$-norm of $\mathbf{x}$ ensures that the number of nonzero elements in $\mathbf{x}$ is less than k_1. In other words, only a sparse set of shape instances can be used to approximate the input shape, which prevents the overfitting to errors from missing/misleading appearance cues. In the second constraint, $\mathbf{e} \in \mathbf{R}^{3N}$ is a vector that models the large residual errors. The sparsity constraint is imposed on $\mathbf{e}$ to incorporate the observation that gross errors might exist but are occasional.

(1) is a NP hard problem owing to the non-convex $L0$ norm. Thanks to the recent proof of the sparse representation theorem [2], $L1$ norm relaxation can be employed to make the problem convex while still preserving the sparsity property. However, to solve (1), we still need to simultaneously optimize multiple variables and deal with the nonlinearty if $T(\mathbf{v}, \beta)$ is modeled as a rigid or similarity transformation. Our solution is to use an Expectation-Maximization (EM) style algorithm (or alternating minimization) to solve (1). It is divided into two sub-problems: 1) estimate the transformation parameter β and 2) efficiently minimize the simplified linear inverse problem with the aligned shape. In the "E" step, β is estimated using Procrustes analysis, which aligns the shape $\mathbf{v}$ to the canonical space as $\mathbf{v}\prime = T(\mathbf{v}, \beta)$. In the "M" step, the following simplified problem is minimized:

$$\arg\min_{\mathbf{x},\mathbf{e}} \|\mathbf{v}\prime - D\mathbf{x} - \mathbf{e}\|_2^2 + \lambda_1\|\mathbf{x}\|_1 + \lambda_2\|\mathbf{e}\|_1, \tag{2}$$

Since (2) now becomes a typical linear inverse problem, it can be solved using existing solvers. The "E" and "M" steps are iteratively performed until $\mathbf{x}$, $\mathbf{e}$ and β converge. $D\mathbf{x}$ is then computed as a refined version of the input shape, which imposes the shape priors on-the-fly.

Multi-resolution Shape Refinement and Local Shape Priors: It has been widely accepted that multiresolution/hierarchical scheme should be employed to improve the efficiency and robustness of deformable segmentation [7]. In a multiresolution scheme, only a small set of sparsely distributed vertices are used as driving vertices to estimate a rough segmentation of the initial stages. As the

iterations increase, more and more vertices join the driving set to gradually reach accurate segmentation. Our sparse shape composition method naturally supports this scheme by estimating a sparse linear combination from an incomplete input. Assume $\mathbf{v}_{sub} = \mathbf{S}\mathbf{v}$ is a subset of all vertices in shape $\mathbf{v}$, where $\mathbf{S}$ is a binary diagonal matrix which indicates if the ith vertex is in the subset ($\mathbf{S}_{ii} = 1$). (1) can then be naturally extended as:

$$\arg\min_{\mathbf{x},\mathbf{e},\beta} \|T(\mathbf{v_{sub}}, \beta) - \mathbf{S}D\mathbf{x} - \mathbf{S}\mathbf{e}\|_2^2, \quad s.t. \|\mathbf{x}\|_0 < k_1, \|\mathbf{e}\|_0 < k_2, \tag{3}$$

(3) can be solved using the same EM optimization. The only difference is that the optimized $\mathbf{x}$ will be finally applied on the full space of D, such that the entire input shape is refined.

One extreme situation of (3) is that $\mathbf{S}$ becomes very sparse and only includes a few vertices (usually with the most distinctive appearance/geometry characteristics). In this situation, (3) indeed becomes the formula of landmark-based surface initialization, which is the first step of our runtime segmentation system. Again, by incorporating shape priors with the assumption of "sparse gross errors", our initialization method becomes robust to erroneous landmark detections due to severe diseases/imaging artifacts.

The merit of (3) is actually beyond the support of multiresolution deformation scheme. In practice, many 3D deformable models include many thousands of points to give an accurate description of organ shapes. The optimization of (2) thus has high computational complexity. In addition, since local shape statistics often lie in a more compact space than global ones, shape priors built on sub-surface are expected to improve the performance. To achieve this goal, we propose a "mesh partitioning" method, which can also be seamlessly incorporated in our sparse shape composition formula. Affinity propagation clustering [4] is employed to divide the model shape into multiple partitions. Since one-to-one correspondences are already constructed among all shapes, affinity propagation only needs to perform once for the model shape. The similarity used in the affinity propagation is defined as the combination of the image similarity and geodesic distances between vertices [9].

In our implementation, each divided partition is further "dilated" for several levels to produce overlaps with neighboring partitions. Finally, partitions are converted to a set of indication matrices $\mathbf{S}_1, \mathbf{S}_2, ..., \mathbf{S}_p$ used in (3). The optimization problem defined on the entire surface is thus decomposed to a set of sub-problems. Each partition is refined independently but the refined partitions are averaged in these overlapping regions to guarantee the smoothness of the entire surface.

The computational complexity of an existing solver (e.g., interior point method) is $\mathcal{O}(N^3)$, where N is the number of vertices of the whole surface. After dividing the whole surface into p partitions with about $\frac{N}{p}$ vertices in each partition. The computational complexity is decreased to only $\frac{1}{p^2}$ of the original one, which highly improves the efficiency.

3 Experiments

Liver Segmentation from Low-dose CT: We evaluate the segmentation performance of our system, using 3D low-dose CT data from PET-CT. In wholebody PET-CT scan, CT images usually have low dose and large slice thickness, which result in low contrast and fuzzy boundaries between organs. Hence, organ segmentation in PET-CT becomes more challenging than traditional CT. In our experiment, the 3D ground truth of low-dose CT is manually segmented by multiple clinical experts. 40 out of 67 CT scans are used to train the landmark detector and also used to construct the shape repository D. The other 27 are left for testing. To obtain the one-to-one correspondence for vertices among all shapes, we choose one shape as a reference and register it to all the others using adaptive-focus deformable model (AFDM) [8]. The shape has around $1,000$ vertices, and 20 are selected as landmarks for model initialization. The proposed method is compared to three popular algorithms: 1) PA: Procrustes Analysis [5] is used to find a similarity transformation to fit a mean shape to detected landmarks. There is no shape refinement during deformation. 2) SMS: It is the Shape Model Search module in ASM [1], which employs the PCA method to learn shape statistics and refine the input shape. 3) SI-NN: k-nearest neighbors method, which uses nearest neighbors to find the closest prototypes in the expert's structure annotations. For a fair comparison, same landmark/boundary detectors and deformation strategy are used in all methods. They only differ in model initialization and model refinement, which involve shape priors.

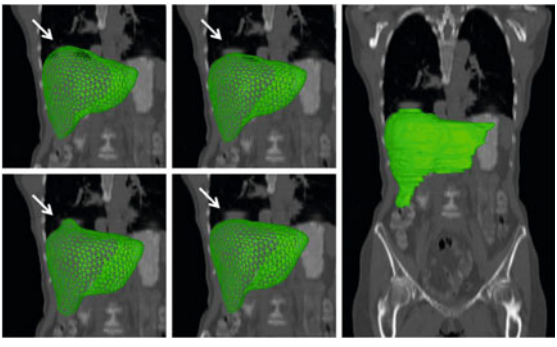

Fig. 3. Initialization results (1st row) and deformation results (2nd row) from the corresponding initialization. Results from PA (1st column) and our method (2nd column) and ground truth (rightmost figure). PA incorrectly includes part of the lung because of the artifacts inducing by breath (see arrows).

Fig. 3 compares the landmark detection based initialization. Since the image contrast of low-dose CT is very low and there are breathing artifacts in the lung region, the landmark detector may easily fail to locate correct positions. Our method is less sensitive to such errors. Its initialization result is already very close to the object boundary. We also compare the deformation results starting from different initializations. A better initialized model also benefits the deformation performance.

To quantitatively evaluate the 3D segmentation accuracy, we report the mean value and standard deviation of the distances between shape surfaces in Tab. 1.

Our proposed framework achieves the best performance. The standard deviations in Tab. 1 show that our method also achieves the best stability among all methods. To evaluate the benefit of mesh partitioning, the surface mesh is divided into 30 regions. The shape refinement step takes several minutes when applied to the whole surface directly. Using mesh partitioning, it significantly improves the efficiency and only takes 2-3 seconds. The whole system takes around 20 seconds (a Python implementation on a PC with 2.4GHz Intel Quad CPU) to segment liver in a 512x512x300 CT volume. Note that the shape refinement module not only improves the robustness of the deformable model, but also decreases the iteration times of deformation since it helps avoid local minima of image information.

Table 1. Quantitative comparisons of the mean values and standard deviations of the distances (voxel) between surfaces. 1st column compares initialization results. Note that SMS and PA uses the same initialization.

Method	**Fig. 3(Init)**	**All data**
PA	2.26 ± 1.72	3.96 ± 3.21
SMS	2.26 ± 1.72	2.16 ± 1.68
SI-NN	4.88 ± 3.61	3.82 ± 3.12
Ours	$\mathbf{1.31 \pm 0.95}$	$\mathbf{1.13 \pm 0.83}$

Rodent Brain Structure Segmentation from MRI: In this study, we use the proposed method to segment rodent brain structures in MR images. In our experiments, 58 data are delineated by clinical experts. 40 are used as training data, and the rest 18 are used as testing. We focus on the 3D segmentation of the cerebellum (Fig. 4). This task is challenging in two aspects. First, there are complex textures and high gradient values inside of the cerebellum region, which adversely affect the deformation module. Second, rodent cerebellum contains two protruding parts, which are easily to be falsely "smoothed out" by traditional shape prior modeling. The visual comparison of the segmentation results is shown in Fig. 4. With regular shape constraint, the protruding parts are shrunk, and

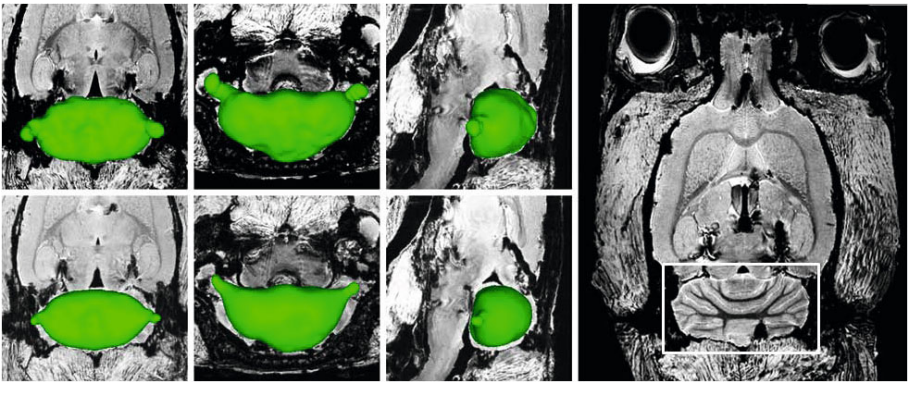

Fig. 4. Segmentation results of the rodent cerebellum. The 1st row is from the proposed method. The 2nd row is from the same framework but without shape prior constraint. The rightmost figure is one slice of the MR image with cerebellum highlighted.

the whole shape is attracted by the high internal gradient. Most regions are under-segmented. The error is 5.86 ± 3.68, in terms of voxel distance. These problems are well tackled by the proposed method. The protruding parts are well preserved, and the global shape are properly constrained. It achieves better segmentation results, with error 1.31 ± 0.91.

4 Conclusions

In this paper, we proposed a deformable model aiming to robustly segment organs in the presence of weak/misleading appearance cues. A sparse shape composition method is proposed to model and incorporate shape priors on-the-fly. It is able to tackle three challenges in a unified framework and naturally supports multi-resolution deformation scheme. Furthermore, we use the affinity propagation method to partition the surface shape local shape priors. Besides a more efficient shape prior modeling, this strategy also dramatically increase run-time efficiency. The majority of the work was carried out when Shaoting Zhang was a research intern at Siemens Medical Solutions, USA. http://www.research.rutgers.edu/∼shaoting/research/siemens2010/project.htm

References

1. Cootes, T., Taylor, C., Cooper, D., Graham, J.: Active shape model - their training and application. Computer Vision and Image Understanding 61, 38–59 (1995)
2. Donoho, D.: For most large undetermined systems of equations, the minimal II-norm near-solution approximates the sparest near-solution. Communications on Pure and Applied Mathematics 59(7), 907–934 (2007)
3. Florin, C., Paragios, N., Funka-Lea, G., Williams, J.: Liver segmentation using sparse 3D prior models with optimal data support. In: Karssemeijer, N., Lelieveldt, B. (eds.) IPMI 2007. LNCS, vol. 4584, pp. 38–49. Springer, Heidelberg (2007)
4. Frey, B., Dueck, D.: Clustering by passing messages between data points. Science 315(5814), 972 (2007)
5. Goodall, C.: Procrustes methods in the statistical analysis of shape. Journal of the Royal Statistical Society 53, 285–339 (1991)
6. Heimann, T., Meinzer, H.P.: Statistical shape models for 3D medical image segmentation: A review. Medical Image Analysis 13(4), 543–563 (2009)
7. Langs, G., Paragios, N., Essafi, S.: Hierarchical 3D diffusion wavelet shape priors. In: IEEE 12th International Conference on Computer Vision, pp. 1717–1724 (2010)
8. Shen, D., Davatzikos, C.: An adaptive-focus deformable model using statistical and geometric information. IEEE Transactions on Pattern Analysis and Machine Intelligence 22(8), 906–913 (2000)
9. Zhan, Y., Dewan, M., Zhou, X.: Cross modality deformable segmentation using hierarchical clustering and learning. In: Yang, G.Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (eds.) MICCAI 2009. LNCS, vol. 5762, pp. 1033–1041. Springer, Heidelberg (2009)
10. Zhan, Y., Zhou, X., Peng, Z., Krishnan, A.: Active scheduling of organ detection and segmentation in whole-body medical images. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part I. LNCS, vol. 5241, pp. 313–321. Springer, Heidelberg (2008)

Pattern Based Morphometry

Bilwaj Gaonkar, Kilian Pohl, and Christos Davatzikos

Section for Biomedical Image Analysis, University of Pennsylvania

Abstract. Voxel based morphometry (VBM) is widely used in the neuroimaging community to infer group differences in brain morphology. VBM is effective in quantifying group differences highly localized in space. However it is not equally effective when group differences might be based on interactions between multiple brain networks. We address this by proposing a new framework called pattern based morphometry (PBM). PBM is a data driven technique. It uses a dictionary learning algorithm to extract global patterns that characterize group differences. We test this approach on simulated and real data obtained from ADNI . In both cases PBM is able to uncover complex global patterns effectively.

Keywords: machine learning, pattern based morphometry, voxel based morphometry.

1 Introduction

VBM[4] is widely used in the neuroimaging community to quantify structural and functional group differences in the brain using 3D and 4D images. VBM involves mapping image data to a standard template space. This is followed by application of voxelwise statistical tests on deformation maps. However, VBM has several weaknesses which are discussed at length in [7] and [5]. For example, VBM fails to account for multivariate group differences, such as interactions between several voxels. Secondly, VBM uses mass univariate testing so that one has to correct for multiple testing. We address these issues by proposing a multivariate morphometric framework called pattern based morphometry(PBM).

PBM is based on K-SVD[3]. K-SVD is a dictionary learning algorithm that has been successfully applied to problems in computer vision [8]. The driving principle behind K-SVD is that it can represent a large set of images as a sparse linear combination of a small set of 'basis images'.

The rest of the paper is structured as follows. Section 2 gives a brief introduction to K-SVD and details how we apply K-SVD to explore group differences. Section 3 describes the results obtained by applying our technique on simulated and real data. We conclude the paper in Section 4 with a discussion of applications and potential avenues for further development of this algorithm.

2 Pattern Extraction Methodology

Our algorithm consists of three steps: 1) Generation of images that represent the difference between two groups(e.g. patients and controls); 2)Application of

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 459–466, 2011.

K-SVD to obtain a dictionary of patterns which can represent these difference images in a sparse way; 3)Ranking and normalization of these patterns. We now describe each step in detail.

2.1 Step 1: Generating Difference Images from Data

Recall that we wish to identify patterns that represent group differences. Our driving assumption here is: Any image generated by subtracting an image in Group 1 from its neighbor in Group 2 can be expressed as a linear combination of a dictionary of image patterns that distinguish the two groups. Our objective is to discover this dictionary of image patterns.

To do this, a large set of difference images is generated by subtracting every image in Group 1 from its neighbors in Group 2. From here on we assume all images to be vectors. Let us denote Group1 by $\mathcal{S} = \{S_1, ..., S_n\}$ and Group 2 by $\mathcal{Z}$. For every element $S_i \in \mathcal{S}$ we first compute the r-nearest neighbors in $\mathcal{Z}$ using a Euclidean metric. Let us denote these by $\{Z_1, ..., Z_r\} \subset \mathcal{Z}$. Subsequently we subtract these r images from S_i to obtain difference vectors D_{ij}

$$D_{ij} \doteq S_i - Z_j \;\; \forall \;\; i \in \{1, ..., n\} \;\; and \;\; \forall \;\; j \in \{1, ..., r\} \tag{1}$$

We collect these difference vectors into a matrix $X = \{D_{11}, ..., D_{nr}\}$ where $X \in \mathbb{R}^{d \times nr}$, d is the number of voxels in the image(usually $d >> nr$).

2.2 Step 2: K-Singular Value Decomposition

For our approach, K-SVD solves a matrix decomposition problem to extract a dictionary of K patterns from X. The patterns in this dictionary can be combined to reconstruct any element of X. We want the image patterns discovered to be as global as possible. This is enforced in K-SVD by using the sparsity constraint which, prohibits the disintegration of large global image patterns into smaller local ones. Formally K-SVD decomposes the matrix X into a basis matrix $B \in \mathbb{R}^{d \times K}$ and a sparse loadings matrix $C \in \mathbb{R}^{K \times nr}$ such that $X \approx BC$. Hence, K-SVD attempts to solve:

$$min_{B,C} ||X - BC||_F^2 \;\; subject \;\; to \;\; \forall i, \;\; ||c_i||_0 \leq T \tag{2}$$

where c_i are columns of the matrix C ,$T \leq K$ is an integer that controls sparsity, $||.||_F$ is the Frobenius matrix norm and $||.||_0$ is the L_0 vector norm. The columns of B are the discriminative morphological patterns extracted from data. The loadings matrix C is used to rank the extracted patterns according to their importance. K-SVD uses an iterative two stage approach to solve (2). In the first stage, the c_i are computed for the current estimate of B. This is done by solving nr separate minimization problems using the orthogonal matching pursuit [13] algorithm. Note that each c_i corresponds to the column x_i of X:

$$\forall i \in \{1, 2, ..., nr\} \quad min_{c_i} \{||x_i - Bc_i||_2^2\} \;\; subject \, to \;\; ||c_i||_0 \leq T \tag{3}$$

where $||.||_2$ denotes the L_2 norm. In the second stage the algorithm updates B one column at a time. For each column of B, denoted by b_k, the corresponding error matrix E_k is computed as:

$$E_k = X - \sum_{j \neq k}^{K} b_j \tilde{c}_j \tag{4}$$

where b_j are the columns of B and $\tilde{c}_k$ are the rows of C. This error matrix quantifies the estimation error that would result from the removal of b_k. Next the basis vector b_k is updated by solving:

$$min_{b_k} ||E_k - b_k \tilde{c}_k||_F^2 \tag{5}$$

This is solved using singular value decomposition (SVD). The b_k obtained by minimization of (5) is used in the computation of b_{k+1} . When the entire matrix B is computed the algorithm returns to stage 1 and iterates. A finite number of iterations are run to generate the basis set B whose columns are the high dimensional morphological patterns representing group differences. For more details on K-SVD we consult the readers to [3].

2.3 Step 3: Ranking and Normalization of the K-SVD Bases

The larger the value of the L_2 norm of the row $\tilde{c}_k$ the more prominently is the pattern represented by b_k expressed in the difference images. For instance if there existed only two patterns in the data then we would expect only two rows of C to contain non zero values. If the k^{th} pattern was more prevalent than the i^{th} one, then $\tilde{c}_k$ would contain more(and possibly larger) non zero numbers than $\tilde{c}_i$. Hence we use the L_2 norm of the row $\tilde{c}_k$ to rank the basis b_k. A higher L_2 norm of c_k assigns a higher rank to the pattern represented by the basis b_k.

3 Results

The results of our experiments on simulated and real datasets are presented in the following section. The simulated data as well as real data was obtained from structural magnetic resonance images(MRI) acquired as a part of the ADNI [1] study. The MR scans are all T1-weighted, acquired sagittally using volumetric 3D MPRAGE with 1.25×1.25 mm in plane spatial resolution and 1.2 mm thick sagittal slices. All images were acquired on a 1.5 T scanner.

3.1 Preprocessing Protocol

Each image was rigidly aligned using the Anterior Commissure Posterior Commissure (AC-PC) points followed by skull removal using the BET algorithm[12]. Skull stripped images were warped to a template using HAMMER[11]. Tissue density maps for grey matter (GM), white matter (WM) and ventricles (VN)

were then generated using the RAVENS approach [6]. These maps were generated from 127 images of Alzheimer's patients and 127 images of normal controls. All RAVENS maps were smoothed by an 8mm full width at half maximum (FWHM) Gaussian kernel before further processing.

3.2 Results on Simulated Datasets

Generation of Simulated Data We first tested PBM on simulated data. We select 63 GM RAVENS maps from normal subjects. Three regions labeled $\{a, b, c\}$ are selected as shown in Fig. 1. Two patterns of atrophy are introduced by reducing the RAVENS maps values in selected locations. Pattern P_1 at locations $\{a, b\}$ and pattern P_2 at locations $\{b, c\}$. In 31 of the 63 images atrophy is introduced, either according to pattern P_1 or according to pattern P_2. These 31 images now represent the patient group in our simulated dataset. The remaining 32 images are not modified and serve as controls.

Analysis of Results. We apply the method described in Section 2 to this simulated dataset with $r = 3$, $T = 1, ..., K$ and for $K = 2, ..., 7$. For every K, the algorithm produces K ranked basis images. Figure 1 shows the first two bases by rank for three different parameter settings.For visualization purposes the intensity histograms of the images corresponding to the patterns represented by each basis are scaled so that all voxel intensities lie between 0 and 1. It should be noted that the results for other parameter settings are almost the same. These results indicate that PBM can discover the patterns P_1 and P_2 introduced in the

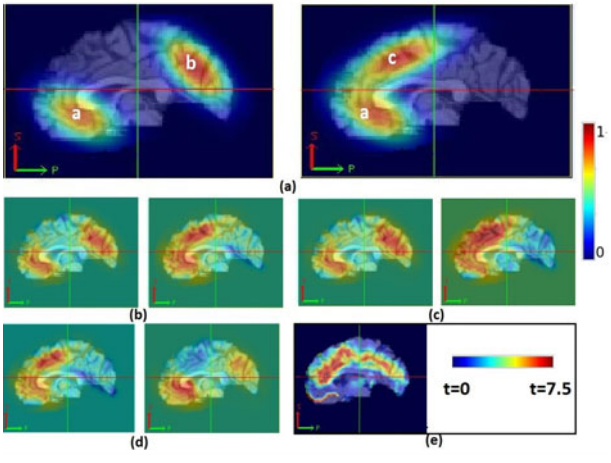

Fig. 1. PBM on simulated data .(a)simulated patterns introduced in the data(b)PBM results, basis ranked 1 and 2 for parameter settings r=3,K=2,T=1 (c)r=3,K=4,T=4 (d) r=3,K=7,T=2 (e)VBM t-statistic map corresponding to p-values not corrected for multiple testing.

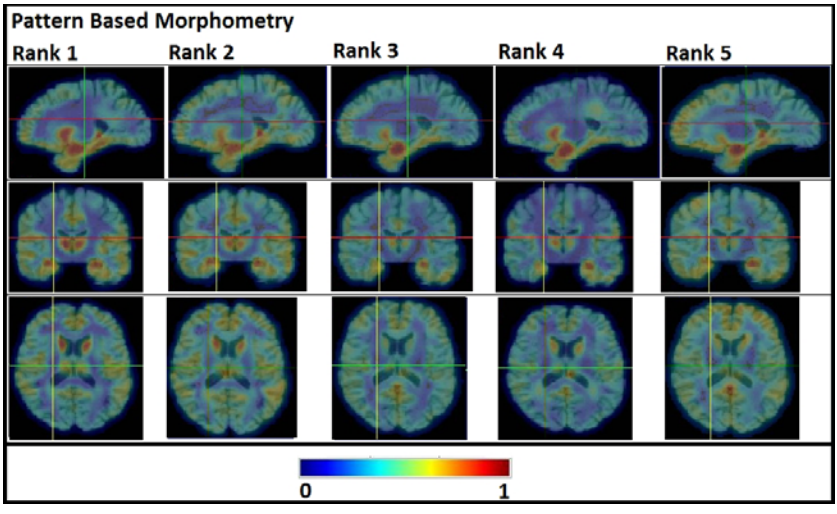

Fig. 2. The 5 top ranked basis/patterns obtained from PBM in WM from ADNI data

data as the first two basis.The figure also shows VBM analysis. It can be clearly seen that the t-statistic from the VBM analysis compares poorly against the PBM analysis. VBM identifies regions that are relatively consistently involved in all subjects, but it doesn't identify the distinct sub-patterns. Neither is it effective when only some of the subjects show involvement of some region (e.g. regions a and c)

3.3 Results on Real Datasets

In this section we apply PBM on the GM,WM and VN RAVENS maps of 127 Alzheimer's patients and 127 normal controls from the ADNI study. We present results for parameter settings of $R = 3$, $K = 5$ and $T = 3$. Note that repeating experiments with several other parameter values does not change the results greatly. The results are presented in Fig.2, Fig. 3 and Fig. 4.

PBM Analysis of GM Maps. Figure 2 shows that the first two patterns detected by PBM in GM are distinct from each other. One involves the putamen more prominently than the other. The insula and thalamic nuclei and the putamen show up in the PBM analysis and have been previously associated with Alzheimer's disease[9]. The other regions that are shown to be strongly associated with disease in GM include the hippocampus, the parahippocampus, the temporal lobe, as well as the frontal lobe and occipital lobes. These regions are known to play a pivotal role in Alzheimer's pathology. These preliminary results indicate that PBM could offer diagnostic value in the analysis of medical image data.

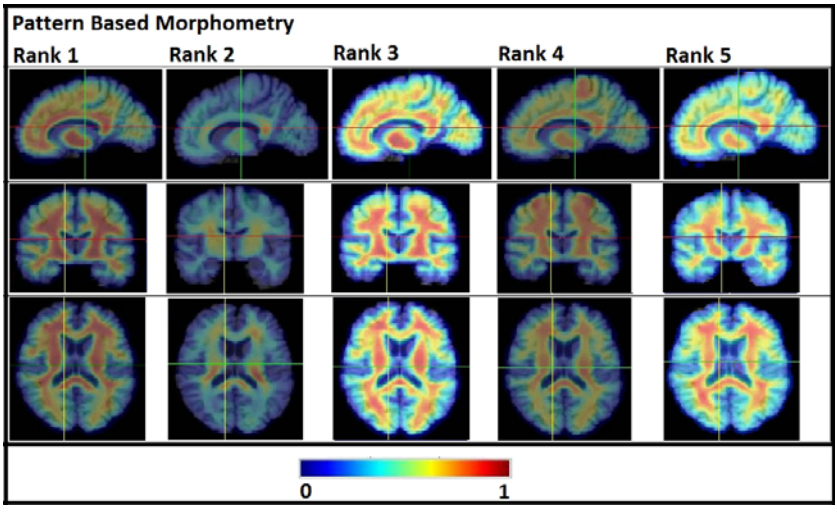

Fig. 3. The 5 top ranked basis/patterns obtained from PBM in VN from ADNI data

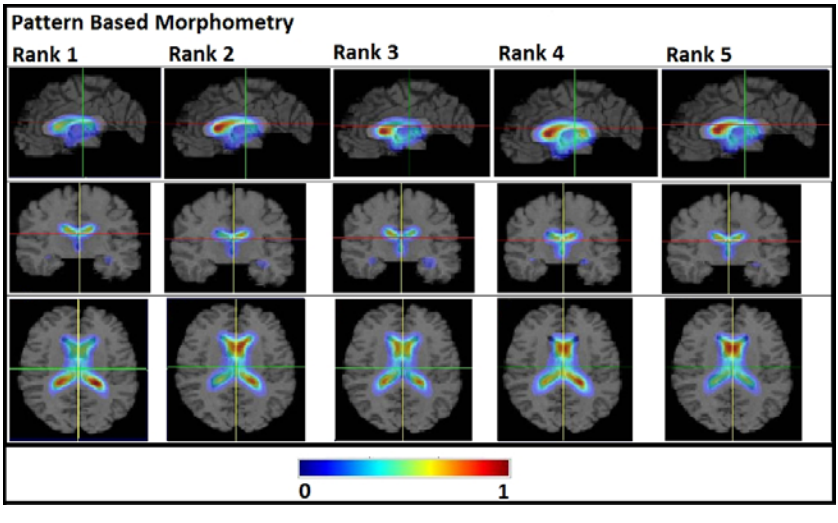

Fig. 4. The 5 top ranked basis/patterns obtained from PBM in GM from ADNI data

PBM Analysis on WM Maps. Figure 3 associates atrophy in most of the major white matter tracts with Alzheimer's pathology. There is recently published evidence[2] to suggest that this is indeed the case. In case of WM the 2^{nd} ranked pattern identified by PBM, highlights periventricular WM atrophy, a hallmark of late onset Alzheimer's disease[10]. Corpus callosum and splenium also show up in PBM analysis as being strongly associated with Alzheimer's pathology.

PBM Analysis of VN Maps. Figure 4 presents the results of a PBM analysis of the ventricles. The top ranked patterns highlight ventricular expansion and periventricular disease oriented pathology. The ventricles are known to be dilated in Alzheimer's disease and the results of PBM are consistent with this fact.

4 Discussion

In this paper we present a novel multivariate approach to morphometry, PBM, which does not suffer from the limitations of VBM described in [7]. PBM can identify subtypes of patterns that don't necessarily involve the same brain regions and facilitate a global analysis of heterogeneous diseases. A limitation of the presented work is that it lacks formal ways of establishing statistical significance. Although the results presented here seem stable with respect to the parameters R, K and T further work is needed in the direction of optimal parameter estimation as well as to exploration of robustness of the method with respect to parameter settings. A second direction of future exploration relates to the Euclidean distance metric used here to evaluate 'nearness' between images. This could be replaced by a different metric in future work possibly yielding better results. Also PBM could be extended to diffusion imaging, fMRI and longitudinal analysis. In summary we have developed and applied a multivariate morphometric framework to quantify group differences between different populations which overcomes some of the limitations of VBM.

Acknowledgment. Data used in preparation of thiis article were obtained from the Alzheimers Disease Neuroimaging Initiative (ADNI) database (adni.loni.ucla.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: http://adni.loni.ucla.edu/wp-content/uploads how_to_apply/ADNI_Authorship_List.pdf

The research was also supported by an ARRA supplement to NIH NCRR (P41 RR13218).

References

1. The alzheimer's disease neuroimaging initiative, http://www.adni-info.org/
2. Agosta, F., Pievani, M., Sala, S., Geroldi, C., Galluzzi, S., Frisoni, G.B., Filippi, M.: White matter damage in alzheimer disease and its relationship to gray matter atrophy. Radiology 258(3), 853–863 (2011), http://dx.doi.org/10.1148/radiol.10101284

3. Aharon, M., Elad, M., Bruckstein, A.: K-svd: An algorithm for designing overcomplete dictionaries for sparse representation. IEEE Transactions on Signal Processing 54(11), 4311–4322 (2006)
4. Ashburner, J., Friston, K.J.: Voxel-based morphometry–the methods. Neuroimage 11(6 Pt.1), 805–821 (2000), http://dx.doi.org/10.1006/nimg.2000.0582
5. Bookstein, F.L.: "voxel-based morphometry" should not be used with imperfectly registered images. Neuroimage 14(6), 1454–1462 (2001), http://dx.doi.org/10.1006/nimg.2001.0770
6. Davatzikos, C., Genc, A., Xu, D., Resnick, S.M.: Voxel-based morphometry using the ravens maps: methods and validation using simulated longitudinal atrophy. Neuroimage 14(6), 1361–1369 (2001), http://dx.doi.org/10.1006/nimg.2001.0937
7. Davatzikos, C.: Why voxel-based morphometric analysis should be used with great caution when characterizing group differences. Neuroimage 23(1), 17–20 (2004), http://dx.doi.org/10.1016/j.neuroimage.2004.05.010
8. Elad, M., Aharon, M.: Image denoising via sparse and redundant representations over learned dictionaries. IEEE Trans. on Image Processing 15(12), 3736–3745 (2006)
9. de Jong, L.W., van der Hiele, K., Veer, I.M., Houwing, J.J., Westendorp, R.G.J., Bollen, E.L.E.M., de Bruin, P.W., Middelkoop, H.A.M., van Buchem, M.A., van der Grond, J.: Strongly reduced volumes of putamen and thalamus in alzheimer's disease: an mri study. Brain 131(Pt. 12), 3277–3285 (2008), http://dx.doi.org/10.1093/brain/awn278
10. Scheltens, P., Barkhof, F., Valk, J., Algra, P.R., van der Hoop, R.G., Nauta, J., Wolters, E.C.: White matter lesions on magnetic resonance imaging in clinically diagnosed alzheimer's disease. evidence for heterogeneity. Brain 115 (Pt. 3), 735–748 (1992)
11. Shen, D., Davatzikos, C.: Hammer: hierarchical attribute matching mechanism for elastic registration. IEEE Trans. Med. Imaging 21(11), 1421–1439 (2002), http://dx.doi.org/10.1109/TMI.2002.803111
12. Smith, S.M.: Fast robust automated brain extraction. Hum. Brain Mapp. 17(3), 143–155 (2002), http://dx.doi.org/10.1002/hbm.10062
13. Tropp, J.A.: Greed is good: algorithmic results for sparse approximation. IEEE Transactions on Information Theory 50(10), 2231–2242 (2004)

Longitudinal Cortical Thickness Estimation Using Khalimsky's Cubic Complex

M. Jorge Cardoso[1], Matthew J. Clarkson[1], Marc Modat[1], and Sebastien Ourselin[1,2]

[1] Centre for Medical Image Computing (CMIC), University College London, UK
[2] Dementia Research Centre (DRC), University College London, UK

Abstract. Longitudinal measurements of cortical thickness is a current hot topic in medical imaging research. Measuring the thickness of the cortex through time is normally hindered by the presence of noise, partial volume (PV) effects and topological defects, but mainly by the lack of a common directionality in the measurement to ensure consistency. In this paper, we propose a 4D pipeline (3D + time) using the Khalimsky cubic complex for the extraction of a topologically correct Laplacian field in an unbiased temporal group-wise space. The thickness at each time point is then obtained by integrating the probabilistic segmentation (transformed to the group-wise space) modulated by the Jacobian determinant of its deformation field through the group-wise Laplacian field. Experiments performed on digital phantoms show that the proposed method improves the time consistency of the thickness measurements with a statistically significant increase in accuracy when compared to two well established 3D techniques and a 3D version of the same method. Furthermore, quantitative analysis on brain MRI data showed that the proposed algorithm is able to retrieve increasingly significant time consistent consistent group differences between the cortical thickness of AD patients and controls.

1 Introduction

The extraction of 4D consistent measurements of thickness from anatomical structures is an important post processing step in neuroimaging. For example, changes in the thickness of the cerebral cortex are of interest in various diseases such as Alzheimer's and Huntington's disease, having the potential to provide a biomarker for diagnosis and neurodegeneration [1]. However, the reliable extraction of 4D consistent and sub-voxel accurate measurements of thickness from probabilistic segmentations is still an unsolved problem.

In order to provide accurate longitudinal measurements, we require that the measurement of thickness is preformed in a consistent direction and location on all time points. Thickness estimation methods, mostly developed for independent 3D time points, can be separated into surface-based and voxel-based techniques. Surface based methods [2] fit a triangulated mesh to the cerebral cortex, making them computationally expensive, especially due to topological constraints. Also, the parametrisation of the surface can be complex and curvature constraints and smoothness parameters can bias the thickness measurements [3]. Voxel-based methods on the other hand extract the value of thickness

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 467–475, 2011.

directly from the voxel grid and are computationally very efficient, but their accuracy is limited by the image resolution and the quality of the segmentation. Nonetheless, voxel based methods have been shown to preform as well as surface based methods [4]. Overall, voxel-based methods can be clustered into 3 subgroups: mathematical morphology [5] based methods use a combination of skeletonisation and region growing techniques in order to calculate the minimal Euclidean distance between points. Partial differential equation (PDE) based methods [6,7,8] solve the Laplace equation between the inner and outer surfaces as if they were charged conductors, resulting in isopotential electric field lines between them. The thickness is then equal to the sum of the lengths of the normals to these isolines. Even though certain topological constraints over the shape of the surfaces are theoretically required, Laplace equation based methods are normally used without enforcing them [8]. Finally, line integral based methods calculate thickness of the structure of interest by finding the direction that minimises the line integrals over its probabilistic segmentation at each voxel position. This method was recently extended to 4D [9] by finding a time consistent directionality between time-points. However, it still lacks topological consistency. In order to solve the problems regarding topology, [10] proposed a method that combines the features of all the above voxel-based methods in a unified, fully automated Khalimsky based thickness estimation algorithm, that is topologically correct and partial-volume aware. This method uses the properties of the Khalimsky grid and an iterative set of element collapse operations to correct the topology of the segmentation. The corrected segmentation is then used to create a multi-stage Laplacian field that encompasses the partial-volume containing areas. The streamlines of this Laplacian field are then integrated using a PDE based method with a spatially varying speed function that is dependent on the probabilistic segmentation.

We propose to extend the above described method [10] in order to encompass 4D consistency. Here, all the time points are registered to an unbiased temporal group-wise space and a group-wise segmentation is then obtained by means of a multivariate EM segmentation algorithm specifically designed for cortical thickness estimation [10]. A multistage Laplacian is then calculated on the temporal group-wise space, and the *per* time point values of thickness are then obtained by integrating over the single time point segmentations along the temporal group-wise derived streamlines.

2 Method

2.1 Proposed Pipeline

The pipeline of most 3D cortical thickness algorithms can be described in 3 steps:

1. The image is segmented into several classes;
2. An implicit or explicit correspondence from one side of the cortex to the opposite side is then found using a multitude of methods;
3. Finally, the thickness is measured according to a specific metric.

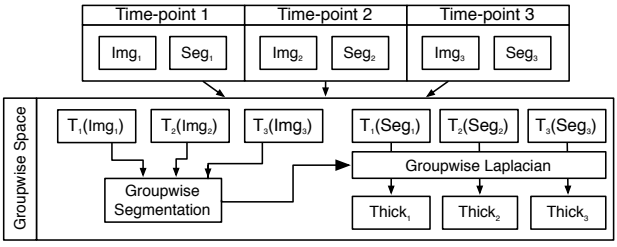

Fig. 1. Top) Individual time points are segmented and registered to the temporal group-wise space; Bottom) A group-wise segmentation is obtained from the transformed images and used to create a group-wise Laplacian field. Thickness is obtained by integrating through the group-wise Laplacian field over the transformed segmentation at each time point.

Any of these steps can be altered in order to introduce temporal consistency. However, altering steps 1 or 3 will directly affect the measurement of thickness, as either the segmentation or the thickness metric itself would be affected by a constraint. This can reduce the statistical significance of difference measurements between groups, leading to increased sample sizes. In order to maintain the sensitivity of the measurement intact but still achieve temporal consistency, we propose to constrain only the direction of the measurement, by changing step 2. The proposed 4D cortical thickness pipeline consists of five steps, as shown on Fig.1. For the sake of simplicity, assume a series of skull stripped brain images acquired at 3 different time points. The five steps can the be described as:

1. The images are segmented independently using a previously published voxel based probabilistic segmentation algorithm [11] specifically designed for cortical thickness. This segmentation will separate the brain into 5 classes: WM, cortical and deep GM and internal and external CSF. The cerebellum is removed within the same pipeline by atlas propagation. Here, the segmentations from the deep GM and internal CSF are added to the WM segmentation in order to create a class containing all the internal structures, simply called WM for the sake of clarity.
2. An unbiased group-wise registration is then created between all the time points. This iterative registration [12] process results in a transformation $\mathbf{T}_t$ for each time point t to the average group-wise space. All the skull stripped images and respective segmentations are transformed to this space.
3. In order to create a group-wise segmentation, a multivariate version of the same algorithm [11] is used. Here, the segmentation model assumes that each label is not only a realisation of one image (time point) but a combined multivariate realisation of all time points, leading to a segmentation with a high level of cortical detail.
4. A topologically correct Laplacian field map is created using the group-wise segmentation as described in section 2.2 and 2.3. This Laplacian map has the directionality information derived from the group-wise segmentation.

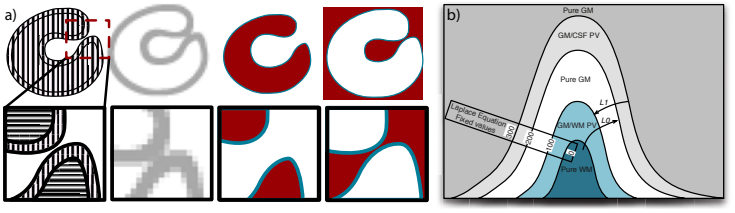

Fig. 2. a) From left to right: An example object; A rasterised version of the object with partial volume effect; the result of step 1 in blue and step 2 in red; the result of the step 3 in blue and step 4 in red; b) Multi-Stage Laplace equation: The iso-lines set to a fixed potential and the Laplace equation is solved. The distance $L0$ and $L1$ is integrated from opposite sides of the object, following the Laplacian field streamlines.

5. This common directionality is then used at each time point to drive a PDE based thickness measure with a speed function proportional to the transformed segmentation modulated by the Jacobian of this transformation, as described in section 2.4

The topologically correct Laplacian field in step 4 constrains the time consistent direction of thickness measurement. This consistency is important as the correspondences between both sides of the cortex can change dramatically on simple 3D models due to sulci and gyri opening and closing.

2.2 Topology Preservation and the Khalimsky's Cubic Complex

Topology-preserving operations are used in many image analysis applications in order to transform an object while leaving its topological characteristics unchanged. Notwithstanding their simplicity, topology-invariant operations in the voxel space have some well-described problems regarding the minimality of the set and the existence of lumps [13]. Abstract complexes, like the Khalimsky space provide a sound topological basis for image analysis. Intuitively, a cubic complex can be seen as a space where every voxel is represented by a structure composed of a set of elements having various dimensions (e.g. cubes, squares, edges, vertices) put together according to some rules. Please refer to [13] for a complete formal description of the cubic complex framework. As shown in [10], this abstract space provides a sound basis for digital topology and topology correction but also an interesting framework for the extraction of thickness measurements.

In order to correct the topology, a series of collapse operations are used to transform a shape into another in a topologically invariant manner [10]. In short, this topology correction step can be described as a 4 step procedure involving Khalimsky based collapse operations, as shown in Fig. 2. The first to steps can be described as a shrink wrap operation in the Khalimsky space in order to obtain a topologically correct WM segmentation. Then, the same operation is performed on the opposite direction in order to correct the topology of the CSF class.

2.3 Multi-Stage Laplace Equation on the Group-Wise Space

In order obtain a 4D consistent measurement of thickness at each time point, a unique association between two sides of the cortex is required. For this purpose, we use the Laplace equation, a second order partial differential equation (PDE), solved between two enclosed boundaries Ω and Ω' in the group-wise space. The classic equation takes the form $\nabla^2\phi = 0$, with the value at ϕ_Ω and $\phi_{\Omega'}$ set up as boundary conditions. Similarly to [10], instead of a single Laplacian field for all the pure voxels as in [6,7,8], a multiple Laplacian field is solved. This obviates the problems regarding the estimation of surface normals for PV integration using ray casting [8]. A set of isolines is generated for each tissue type from the topologically correct group-wise segmentation. The Laplace equation is then solved between these four equipotential lines resulting in a smooth transition field traversing the cortex. The solution of this Laplace equation under an anisotropic voxel grid in the Khalimsky space is presented in [10].

From the resultant Laplacian field, the normals to the direction of the Laplacian isolines in the group-wise space, denoted by N^{GW}, are calculated using finite differences.

2.4 Thickness Measurement at Each Time Point

In order to measure thickness, the length of the streamlines between the inner and outer surface has to be measured at each time point by integrating the vector field N^{GW} on the group-wise space. Because digital topology is not preserved even under diffeomorphic transformations, each time point has to be deformed to the group-wise space using the previously computed group-wise transformation. In order to measure thickness, instead of the basic form partial differential equation proposed by Yezzi [7], where the speed of the advancing front is assumed to be 1, we use a more generalised form of the PDE. Here $\nabla L^t \cdot N^{GW} = f^t$, for an unknown function L^t at time-point t and assuming that N^{GW} and f^t are known. In our case, and differently from [6,7,14,8], the value of f will be spatially varying and equal to the probability of belonging to the cortical GM modulated by the Jacobian determinant of the transformation, $f^t = p^t_{GM}|\mathbf{T}(x)^t|$. This value will act as *time cost* and will make the value of L^t equivalent to the time of arrival in a level-set framework. Modulation by the Jacobian determinant is necessary in order to take the voxel compression into account. Even though collisions of the advancing front might exist, they are not a problem due to the upwind nature of the integration and the existence of the group-wise vector field N^{GW}. Let $L^t_{0_{(x,y,z)}}$ be a function that measures the time of arrival (arc length of the streamline according to the time cost f^t) from the boundary of set $F_{WM_{\text{pure}}}$ on the group-wise space to a point in the object, and $L^t_{1_{(x,y,z)}}$ be the time of arrival from the boundary of set $F_{CSF_{\text{pure}}}$, again in the group-wise space, to the point in the object. The values of $L0$ and $L1$ are calculated using anisotropic finite differences, as described in [10] . The final value of thickness is then defined as $Thick^t = L0^t + L1^t$. In order to reduce the bias of any further statistical analysis, the value of thickness is only calculated at

mid-isopotential line on the group-wise Laplacian field. The ribbon containing the thickness measurements will thus have spherical topology (Euler characteristic of 2) and will be in the same space for all time points.

3 Experiments and Results

The experimental section of this paper is divided into two subsections. First, a digital phantom with time evolving ground truth thickness is used to assess the accuracy and sensitivity of the proposed algorithm compared to three 3D state-of-the-art methods. The proposed method is then applied to brain MRI data in order to assess group separation in terms of cortical thickness between Alzheimer's disease diagnosed patients and controls.

Phantom validation: In order to evaluate if the proposed method can accurately retrieve the underlying thickness of an object, 7 folded 3D digital phantoms with spherical topology and known ground truth thickness were created (Fig. 3), resulting in six high resolution isotropic images with 3 structures equivalent to WM, GM and CSF. The thickness of the object is changing with time from 5.2 to 3.6mm in order to simulate a thickness loss in an object over several time points. Note that the sulci will open after time-point 1. These high resolution phantoms were then down-sampled by 5 in order to simulate PV effect and the thickness of the down-sampled structures was then measured. We compare the proposed 4D method with the 3D version of the same algorithm, the method proposed by Jones et al. [6] and the method proposed by Acosta et al. [8].

Results show that all 3D methods are highly sensitive to temporal structural changes. When compared to the ground truth, the thickness change in time is overestimated, possibly due to sulci opening. The proposed method, on the other hand, uses a 4D consistent directionality derived from all the time-points, resulting in a much more accurate and precise thickness estimation. One should be cautious when reporting cortical thickness loss in time, as 3D methods can severely overestimate it.

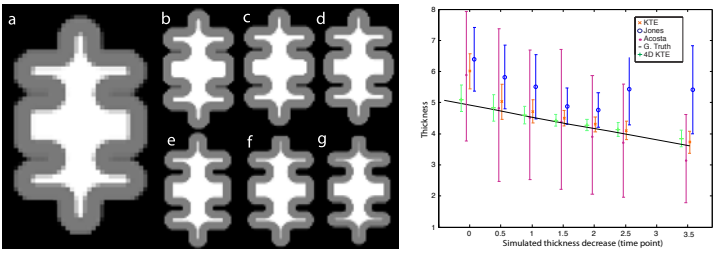

Fig. 3. Left: A 4D simulated high resolution phantom with a time varying thicknesses ranging from 5.2mm to 3.6mm (a-g). Right: Mean and standard deviation of the estimated thickness at all voxel positions from timepoint 1 to 3.5 when compared to the ground truth in black.

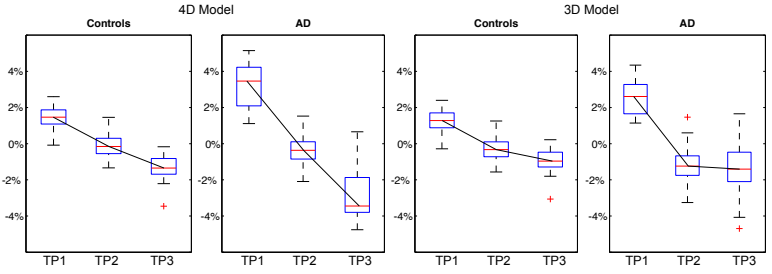

Fig. 4. A plot showing the progression of the average cortical thickness in time, normalised to the average thickness over all time points. From left to right: The normalised thickness for controls and AD patients using the 4D (proposed) method and 3D (time independent) versions of the algorithm.

Brain MRI analysis: To further investigate the temporal consistent of the proposed method, the thickness of the cortical layer was calculated on the ADNI dataset at 3 time points. The main purpose of this study was not to evaluate group separation between different groups but to assess their stability in time. From the full ADNI dataset, a subset of 60 age- and gender-matched subjects (30 AD and 30 controls) were selected. Each subject has T1-weighted 1.5T MRI volumetric images acquired using a 3D MPRAGE sequence (typically $1.20 \times 1.00 \times 1.00mm$) at 0, 12 and 24 months.

Fig. 4 shows the progression of the distribution of the average value of thickness within the cortex at each time point, normalised to the average thickness over all time points. Due to the lack of time consistency, unexpected inversions of the thinning pattern occur on the 3D version of the algorithm. The 4D version of the algorithm shows a marked improvement regarding the stability of the thinning pattern. This leads to a reduction of the standard deviation of the thickness distributions within each group, increasing the statistical power. In order to compare the different groups (AD and controls) on a per area basis, the group-wise space was parcelated into different areas using the an anatomical atlas. The 3D method shows statistically significant differences in thickness ($p < 10^{-3}$) on both the temporal and parietal region. The frontal region is statistically significantly thinner ($p < 10^{-3}$) at both time-points 1 and 3 but it but does not achieve the significance threshold at time-point 2. On the other hand, the proposed 4D method shows statistically significant differences in thickness at the level $p < 10^{-5}$ in the middle and inferior temporal and parietal regions and $p < 10^{-3}$ in the frontal gyrus region in the first time point. From time point 2 onwards, the frontal region becomes significant at $p < 10^{-5}$ and both the superior and occipital regions become statistically significant at $p < 10^{-3}$.

Due to the lack of 4D consistency in the segmentation, cortical lost might be over-estimated. In order to investigate this, future work will explore the

use of a 4D segmentation step for improved consistency. We will also consider the use of the full Jacobian matrix for the modulation step instead of it's determinant.

4 Conclusions

In this paper we present a new method to extract 4D measurements of thickness from cortical segmentations. First, all the time points are registered to an unbiased and temporal consistent group-wise space. Then, a time consistent group-wise point-to-point correspondence is found by means of a multistage Laplacian field derived from a multivariate segmentation in the group-wise space. This common directionality is then used to calculate the thickness at each time point.

Experiments on digital phantoms with known ground truth thickness show that the proposed method is more accurate and precise at retrieving true thickness values than other previously published methods, thereby reducing the overestimation of cortical thinning in the presence of sulci opening. Quantitative analysis on brain MRI data showed that the proposed algorithm is able to retrieve increasingly significant time consistent consistent group differences between the cortical thickness of AD patients and controls.

References

1. Holland, D., Brewer, J.B., Hagler, D.J., Fennema-Notestine, C., Fenema-Notestine, C., Dale, A.M., Initiative, A.D.N.: Subregional neuroanatomical change as a biomarker for Alzheimer's disease. PNAS 106(49) (2009)
2. Fischl, B., Dale, A.M.: Measuring the thickness of the human cerebral cortex from magnetic resonance images. P. Natl. Acad. Sci. USA 97(20), 11050–11055 (2000)
3. Scott, M.L.J., Bromiley, P.A., Thacker, N., Hutchinson, C.E., Jackson, A.: A fast, model-independent method for cerebral cortical thickness estimation using MRI. Medical Image Analysis 13(2), 269–285 (2009)
4. Clarkson, M.J., Cardoso, M.J., Ridgway, G.R., Modat, M., Leung, K.K., Rohrer, J.D., Fox, N.C., Ourselin, S.: A comparison of voxel and surface based cortical thickness estimation methods. NeuroImage (2011)
5. Lohmann, G., Preul, C., Hund-Georgiadis, M.: Morphology-based cortical thickness estimation. In: Taylor, C.J., Noble, J.A. (eds.) IPMI 2003. LNCS, vol. 2732, pp. 89–100. Springer, Heidelberg (2003)
6. Jones, S.E., Buchbinder, B.R., Aharon, I.: Three-dimensional mapping of cortical thickness using Laplace's equation. Human Brain Mapping 11(1), 12–32 (2000)
7. Yezzi, A.J., Prince, J.L.: An Eulerian PDE approach for computing tissue thickness. IEEE Transactions on Medical Imaging 22(10), 1332–1339 (2003)
8. Acosta, O., Bourgeat, P., Zuluaga, M.A., Fripp, J., Salvado, O., Ourselin, S.: Alzheimer's Disease Neuroimaging Initiative: Automated voxel-based 3D cortical thickness measurement in a combined Lagrangian-Eulerian PDE approach using partial volume maps. Medical Image Analysis 13(5), 730–743 (2009)

9. Li, Y., Wang, Y., Xue, Z., Shi, F., Lin, W., Shen, D., The Alzheimer's Disease Neuroimaging Initiative: Consistent 4D cortical thickness measurement for longitudinal neuroimaging study. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. 6362, pp. 133–142. Springer, Heidelberg (2010)
10. Cardoso, M.J., Clarkson, M.J., Modat, M., Ourselin, S.: On the extraction of topologically correct thickness measurements using khalimsky's cubic complex. In: IPMI (2011)
11. Cardoso, M.J., Clarkson, M.J., Ridgway, G.R., Modat, M., Fox, N.C., Ourselin, S.: Alzheimer's Disease Neuroimaging Initiative. LoAd: A locally adaptive cortical segmentation algorithm. NeuroImage 56(3), 1386–1397 (2011)
12. Rohlfing, T., Brandt, R., Menzel, R., Maurer, Jr., C.R.: Evaluation of atlas selection strategies for atlas-based image segmentation with application to confocal microscopy images of bee brains. NeuroImage 21(4), 1428–1442 (2004)
13. Cointepas, Y., Bloch, I., Garnero, L.: A cellular model for multi-objects multidimensional homotopic deformations. Pattern Recognition 34(9), 1785–1798 (2001)
14. Rocha, K.R., Yezzi Jr., J.L., Prince, J.L.: A hybrid Eulerian-Lagrangian approach for thickness, correspondence, and gridding of annular tissues. IEEE TIP, 72–81 (2005)

Spatiotemporal Morphometry of Adjacent Tissue Layers with Application to the Study of Sulcal Formation

Vidya Rajagopalan[1,2], Julia Scott[1,2], Piotr A. Habas[1,2], Kio Kim[1,2], François Rousseau[3], Orit A. Glenn[4], A. James Barkovich[4], and Colin Studholme[1,2]

[1] Biomedical Image Computing Group, {vidyaraj,jascott1,habas,kiokim,studholm}@uw.edu
[2] Department of Pediatrics, University of Washington, Seattle, USA
[3] LSIIT, UMR 7005 CNRS/University of Strasbourg, 67412 Illkirch, France
[4] Department of Radiology and Biomedical Imaging University of California San Francisco, San Francisco, CA 94143, USA

Abstract. The process of brain growth involves the expansion of tissue at different rates at different points within the brain. As the layers within the developing brain evolve they can thicken or increase in area as the brain surface begins to fold. In this work we propose a new spatiotemporal formulation of tensor based volume morphometry that is derived in relation to tissue boundaries. This allows the study of the directional properties of tissue growth by separately characterizing the changes in area and thickness of the adjacent layers. The approach uses temporally weighted, local regression across a population of anatomies with different ages to model changes in components of the growth radial and tangential to the boundary between tissue layers. The formulation is applied to the study of sulcal formation from in-utero MR imaging of human fetal brain anatomy. Results show that the method detects differential growth of tissue layers adjacent to the cortical surface, particularly at sulcal locations, as early as 22 gestational weeks.

1 Introduction

This work is motivated by the study of growth patterns in the developing brain. In particular, there is considerable interest in exploring geometric characteristics of growth that underlie the formation of sulci and gyri. Multiple hypotheses regarding the ontogeny of cortical folding patterns have emerged that consider factors like axonal tension [1], differential cortical expansion [2], mechanical constraints [3] and genetic determination [4] as well as combinations of these factors [e.g. [5], [6]]. A more complete characterization of the local growth patterns associated with cortical folding in the human fetus will provide evidence for the development of theories on gyrogenesis

Tensor based morphometry is a powerful tool that makes use of accurate spatial normalization to study local differences in tissue size across a population,

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 476–483, 2011.

that can be used to study changes in anatomy with age. It has recently been used to study fetal brain growth [7] and brain growth in early childhood [8], [9], [10], [11]. One of the key challenges in TBM analysis is interpreting not just the scalar volume increase with growth, but the directional characteristics that induce the formation of a sulcated brain from a smooth fetal brain. In this paper we describe an approach to examining how tissue growth occurs in adjacent tissue layers as the shape of the boundary between them evolves. We derive a statistical framework to model the components of tissue growth over time in relation to the local coordinate frame of the tissue boundary. We apply this to look at how primary sulci are formed in the developing human fetus in relation to the boundary between the cortical plate and sub-plate, by specifically mapping normal and tangential components of volume increases on either side of the tissue boundary. We make use of temporally weighted, local regression to test whether the rate of area or thickness is increasing or decreasing over time. This allows us to study the pattern of increases in the area of layers in relation to increases in thickness of the layers at specific gestational ages.

2 Methods

For a cross-sectional population of n subject at different ages, we compute a transformation $D_i, \quad i = 1, 2, ...n$ for each subject, which captures the geometric changes required to spatially normalize each subject image to a common space. At each voxel p, the local differences size and shape can be derived from the deformation tensor which is defined as the gradient of the transformation $(D_i^p[x, y, z])$ at that voxel and is given by $(J_i^p) = \left[\frac{\partial D_i^p[x,y,z]}{\partial x \partial y \partial z}\right]$. We form an average brain surface by constructing a triangulated mesh from the cortical surface of the average image. The Jacobian matrices within 4 mm of each surface vertex were averaged (in the log domain [12] and placed onto the surface mesh for surface modeling. For each subject, this then provided a map that summarizes shapes changes on the cortical surface.

2.1 Discriminating Shape Change into Area or Thickness Changes

Shape changes adjacent to anatomical boundaries occur as a series of changes in local area and thickness of the adjacent tissue layers. At each location on the boundary, change in neighboring tissue area and thickness can be quantified by resolving the Jacobian matrix into its radial and tangential components respectively. We form a surface normal map (N) by computing the surface normal at each vertex of the average surface. Vertex-wise approximation of change in neighboring layer thickness T_i^p, is the scalar projection of the Jacobian tensor $((J_i^p))$ of the tissue location projected onto the corresponding surface normal $(\mathbf{n})$,

$$T_i^p = abs(\mathbf{n}^T(J_i^p)\mathbf{n}) \tag{1}$$

In order to compute the change in local surface area, we compute maps of vectors tangential to the surface and perpendicular to the surface normal vector. The

cross-product of **n** and an arbitrary, non-colinear vector **r** results in one possible tangent (**t1**) to the surface. The other surface tangent (**t2**) is the cross-product of (**t1**) and n. Local change in area is the scalar component of the Jacobian tensor onto **t1** and **t2**,

$$A_i^p = \sqrt{(\mathbf{t1}^T(J_i^p)\mathbf{t1}) \times (\mathbf{t2}^T(J_i^p)\mathbf{t2})} \tag{2}$$

For each subject, we compute a map of surface area A_i and thickness T_i change which can then be used to analyze local changes in shape.

2.2 Modeling Local Changes at Specific Time Points

In order to examine the temporal progression of changes in surface shape, we use locally weighted regression (LWR) [13] to examine the relationship between surface area or thickness and a time-variable of interest (e.g. age of subjects). LWR is a non-parametric method that fits a smooth regression model to the data by fitting simple models to localized subsets of the data. For each independent variable of interest (IVOI), a lower degree polynomial is fitted using weighted least squares, giving more weight to points near the IVOI and less weight to points further away. The regression is considered complete when a local model is fit to all independent variables. In order to study age-specific models of surface shape change, we performed LWR with area (or thickness) as the dependent variable and age as the independent variables. Let Y correspond to the vector of dependent variables (at a single vertex), X is the vector of independent variables and and ε are the errors. As shown in Equation 2.2, for each independent variable X_i, a linear model is fit by using weighted least squares where W_i is the weight matrix. A generalized bell-shaped function is usually used to form the weight matrix. (The matrix dimensions of each of the variables are indicated below each variables in Equation 2.2). LWR results in a vector of coefficients $(\beta_1, \ldots, \beta_n)$ corresponding to the total number of independent variables.

$$\underset{(n\times 1)}{W_iY} = \underset{(n\times n)(n\times 1)(1\times 1)}{W_iX\beta_i} + \underset{(n\times 1)}{\varepsilon}$$

Hypothesis testing: Resulting regression coefficient $(\beta_1, \ldots, \beta_n)$ maps are tested for significance using a standard t-test. Statistical significance was computed and these were corrected multiple comparisons using permutation tests [14].

The $\beta_1, \ldots, \beta_n$ values are estimates of increase or decrease in area or thickness at a particular time point and the hypothesis tests estimate statistical significance of these changes.

3 Application – Comparing Growth Patterns Between Cortical Plate (CP) and Subplate (SP) During Early Fetal Development

The following experiments were performed using clinical MR scans of 40 fetal subjects at gestational ages ranging from 20 to 28 weeks. The mothers were

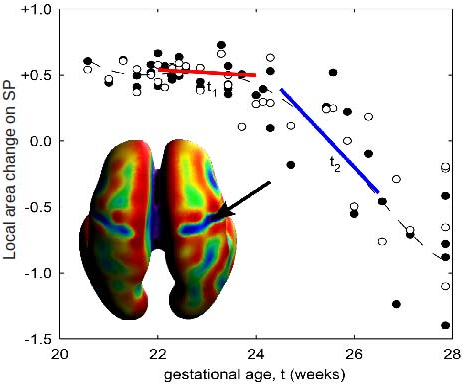

Fig. 1. Example of using LWR to model area changes at a particular location on the subplate (SP). A smooth function is fit to the entire time interval by fitting linear models to smaller subsets of data. For each subset of data, the independent variables closest to the independent variable of interest are weighted higher than those further away.

referred for fetal MRI due to questionable abnormalities on prenatal ultrasound or a prior abnormal pregnancy. All women had normal fetal MRI and all newborns have had normal postnatal neurodevelopment. Fetal imaging was performed in our institution on a 1.5T scanner (GE Healthcare, Milwaukee, WI) without sedation of the mother or the fetus. For each subject multiple stacks of single-shot fast spin-echo (SSFSE) T2-weighted slice images (pixel size 1 mm $\times$ 1 mm, slice thickness $\approx$ 3 mm) are planned in the approximately axial, sagittal and coronal planes with respect to the fetal brain and are acquired with MR sequence parameters (TR = 4500 ms, TE = 91 ms). High resolution 3D volumes were reconstructed from 2D slice MR images using the slice intersection motion correction (SIMC) technique [15]. The reconstructed volumes were automatically segmented into individual tissue types (developing grey matter, developing white matter, the germinal matrix) using an atlas-based approach with probabilistic atlases generated from a spatiotemporal model of the fetal brain [16]. The tissue label maps for each of the 40 scans were co-aligned using an unbiased groupwise registration algorithm. The algorithm simultaneously estimated an average brain shape and a deformable mapping to each of the anatomies being studied. This average shape was estimated in such a way to ensure that the average distance from each point in that space, when mapped to the individuals in the group, is forced to be zero, forming a so-called minimum deformation anatomy. For each subject, the Jacobian matrix maps were computed, from the resulting deformation fields, to quantify the pattern of deformation required to spatially normalize individual anatomies. The inner surface of the cortical plate in the group average tissue segmentation map was formed into a triangulated mesh using a topology preserving marching cubes algorithm [17].

For each subject, we computed two surface Jacobian maps corresponding to two layers of the cerebral mantle namely the cortical plate (CP) and the subplate

(SP). For each layer, we compute a set of maps describing local changes in surface thickness and surface area changes. For each layer and metric, we compute a model of spatial changes from 21 to 28 weeks GA at 1 week intervals.

3.1 Results and Discussion

Figure 2 shows a matrix of the statistically significant variations in area and thickness between the CP and the SP. The T Maps were overlaid on the average surface and displayed using the Rview software[1]. Due to space constraints, we have shown growth patterns for 22, 24 and 26 weeks only. Overall, we see distinctly different patterns of area and thickness changes between the two layers. We also see a distinct temporal pattern of area and thickness change from 21 to 28 weeks. On the SP, the medial side of the fronto-parietal region (as indicated by black arrows) shows a progressive, relative increase in area from 21 to 27 weeks which corresponds to relative decrease in the thickness component of growth during this period. We also notice that the emergence of the post-central sulcus (purple arrows) as increase in area of the SP beginning at 23 weeks which is correlated with a increase in area on the CP plate at the corresponding location after 24 weeks. We observe in the lateral views that the operculization process manifests as changes in area and volume both on the SP and CP. Local area increases in the SP begin at the operculum at 22 weeks. By 24 weeks, the increase in area has shifted to accommodate the emergence of the superior temporal gyrus and by 26 weeks area increases also occurs at the location of the superior temporal sulcus (orange arrow). A similar pattern of area increase is also visible on the CP although the process seems to be delayed by 1 week. During this period, the local thickness component of growth on the SP decreases as expected. Local area increases at the presumptive location of the fronto-parietal gyrus is visible in both the SP and the CP between 21- 24 weeks. The lack of change in thickness of the CP is supported by the findings of Rakic et al. [18], who postulate that the CP enlarges primarily by tangential expansion as a uniform surface and not by increase in thickness.

We see a distinct difference in the directional growth of the CP and SP. In regions of cortical folding, the CP and SP expanded in overlapping areas. Moreover, where SP expanded in the tangential direction, SP also significantly grew in the normal direction. In contrast, the CP did not increase in the normal direction, from which we infer there is no detectable change in thickness at our image resolution. The total expansion of SP underlying developing gryi is in agreement with the histological observation in gyrencephalic animals of large visible SP at gyri and very thin SP at sulci [19], [20]. Coincident expansion of the CP and SP, which is primarily composed of axonal fibers held at transient contacts prior to reaching final targets in the CP [20], may support cortical folding hypotheses that suggest that connectivity is a major driver of cortical folding patterns [1], [21].

[1] http://rview.colin-studholme.net

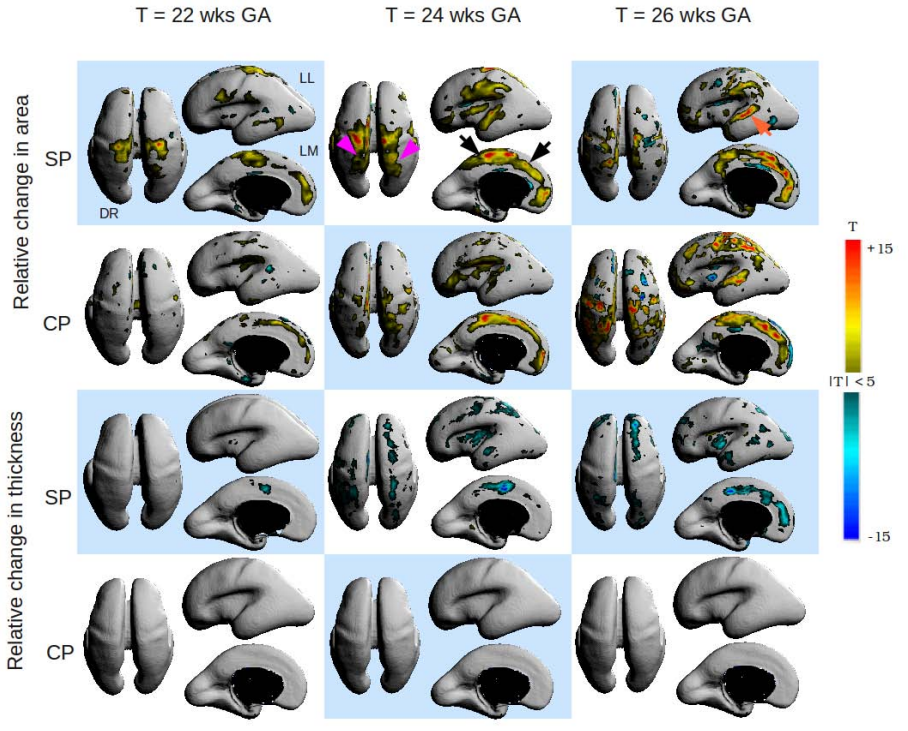

Fig. 2. Spatiotemporal patterns of variational growth between the cortical plate (CP) and subplate (SP) during early fetal development. Warm colors indicate significant, relative increase in area or thickness component of growth with time and cool colors indicate statistically significant reductions in area or thickness component. DR = Dorsal view, LL = left lateral view and LM = left medial view. Changes are bilateral unless noted otherwise.

4 Conclusion

In this paper, we have introduced descriptors that allow us to characterize growth in adjacent tissue layers in terms of changes in area and thickness. Also for the first time, we use temporally weighted, local regression to create spatiotemporal models of tissue growth along an evolving tissue boundary between the adjacent layers. The two methods are incorporated into the TBM framework to model differential patterns of tissue growth between two adjacent anatomical layers. Using the proposed spatiotemporal morphometry framework, we modeled variational growth patterns between the CP and SP which underly the mechanism of sulcation in a fetal brain. The spatiotemporal method allowed us to precisely stage growth patterns of various tissue layers corresponding to the emergence of the primary sulci. Future work will include extending this analysis to other tissue boundaries within the fetal brain and examining the correlation between

changes across the boundary (e.g volume increases in SP vs area increases in CP). The general framework of spatiotemporal morphometry can be adapted to other applications.

Acknowledgments. This research was funded by NIH/NINDS grants: R01 NS 061957 and R01 NS 055064. Imaging for this study was also partially supported by the National Institutes of Health (NIH) Grant No. K23 NS52506-03 and NIH/NCRR UCSF-CTSI Grant No. UL1 RR024131. The work of F. Rousseau was supported by European Research Council under FP7/2007-2013 Grant Agreement 207667.

References

1. Van Essen, D.C.: A tension-based theory of morphogenesis and compact wiring in the central nervous system. Nature 385, 313–318 (1997)
2. Welker, W.: Why does cerebral cortex fissure and fold? a review of determinants of gyri and sulci. Cerebral Cortex 8, 3–136 (1990)
3. Armstrong, E., Schleicher, A., Omran, H., Curtis, M., Zilles, K.: The ontogeny of human gyrification. Cerebral Cortex 5(1), 56–63 (1995)
4. Rakic, P.: Genetic control of cortical convolutions. Science 303(5666), 1983–1984 (2004)
5. Toro, R., Burnod, Y.: A morphogenetic model for the development of cortical convolutions. Cereb. Cortex 15, 1900–1913 (2005)
6. Lefevre, J., Mangin, J.F.: A reaction-diffusion model of human brain development. PLoS Comput. Biol. 6, e1000749 (2010)
7. Rajagopalan, V., Scott, J.A., Habas, P.A., Corbett-Detig, J.M., Kim, K., Rousseau, F., Barkovich, A.J., Glenn, O.A., Studholme, C.: Local tissue growth patterns underlying normal fetal human brain gyrification quantified in utero. Journal of Neuroscience 31(8), 2878–2887 (2011)
8. Dubois, J., Benders, M., Cachia, A., Lazeyras, F., Ha-Vinh Leuchter, R., Sizonenko, S., Borradori-Tolsa, C., Mangin, J., Huppi, P.: Mapping the early cortical folding process in the preterm newborn brain. Cereb. Cortex 18(6), 1444–1454 (2008)
9. Lenroot, R.K., Gogtay, N., Greenstein, D.K., Wells, E.M., Wallace, G.L., Clasen, L.S., Blumenthal, J.D., Lerch, J., Zijdenbos, A.P., Evans, A.C., Thompson, P.M., Giedd, J.N.: Sexual dimorphism of brain developmental trajectories during childhood and adolescence. NeuroImage 36(4), 1065–1073 (2007)
10. Lee, J., Fonov, V., Evans, A.: Mapping brain growth of early childhood using deformation based morphometry. NeuroImage 47(supplement 1), S153–S153 (2009)
11. Aljabar, P., Bhatia, K.K., Murgasova, M., Hajnal, J.V., Boardman, J.P., Srinivasan, L., Rutherford, M.A., Dyet, L.E., Edwards, A.D., Rueckert, D.: Assessment of brain growth in early childhood using deformation-based morphometry. Neuroimage 39(1), 348–358 (2008)
12. Arsigny, V., Commowick, O., Pennec, X., Ayache, N.: A log-euclidean framework for statistics on diffeomorphisms. In: Larsen, R., Nielsen, M., Sporring, J. (eds.) MICCAI 2006. LNCS, vol. 4190, pp. 924–931. Springer, Heidelberg (2006)
13. Cleveland, W.S., Devlin, S.J.: Locally weighted regression: An approach to regression analysis by local fitting. Journal of the American Statistical Association 83, 596–610 (1988)

14. Nichols, T.E., Holmes, A.P.: Nonparametric permutation tests for functional neuroimaging: A primer with examples. Hum. Brain Mapp. 15(1), 1–25 (2002)
15. Kim, K., Habas, P.A., Rousseau, F., Glenn, O.A., Barkovich, A.J., Studholme, C.: Intersection based motion correction of multislice MRI for 3-D in utero fetal brain image formation. IEEE Trans. Med. Imaging 29(1), 146–158 (2010)
16. Habas, P.A., Kim, K., Corbett-Detig, J.M., Rousseau, F., Glenn, O.A., Barkovich, A.J., Studholme, C.: A spatiotemporal atlas of MR intensity, tissue probability and shape of the fetal brain with application to segmentation. Neuroimage 53(2), 460–470 (2010)
17. Lopes, A., Brodlie, K.: Improving the robustness and accuracy of the marching cubes algorithm for isosurfacing. IEEE Trans. Viz. and Comput. Graph. 9(1), 16–29 (2003)
18. Rakic, P., Ayoub, A.E., Breunig, J.J., Dominguez, M.H.: Decision by division: Making cortical maps. Trends Neurosci. 32(5), 291–301 (2009)
19. Smart, I.H., McSherry, G.M.: Gyrus formation in the cerebral cortex in the ferret. I. Description of the external changes. J. Anat, 146, 141–152 (1986)
20. Kostovic, I., Rakic, P.: Developmental history of the transient subplate zone in the visual and somatosensory cortex of the macaque monkey and human brain. J. Comp. Neurol. 297(3), 441–470 (1990)
21. Hilgetag, C.C., Barbas, H.: Role of mechanical factors in the morphology of the primate cerebral cortex. PLoS Comput. Biol. 2(3), e22 (2006)

Fast Shape-Based Nearest-Neighbor Search for Brain MRIs Using Hierarchical Feature Matching

Peihong Zhu, Suyash P. Awate, Samuel Gerber, and Ross Whitaker

Scientific Computing and Imaging Institute, University of Utah, USA

Abstract. This paper presents a fast method for quantifying shape differences/similarities between pairs of magnetic resonance (MR) brain images. Most shape comparisons in the literature require some kind of deformable registration or identification of exact correspondences. The proposed approach relies on an optimal matching of a large collection of features, using a very fast, hierarchical method from the literature, called *spatial pyramid matching* (SPM). This paper shows that edge-based image features in combination with SPM results in a fast similarity measure that captures relevant anatomical information in brain MRI. We present extensive comparisons against known methods for shape-based, k-nearest-neighbor lookup to evaluate the performance of the proposed method. Finally, we show that the method compares favorably with more computation-intensive methods in the construction of local atlases for use in brain MR image segmentation.

1 Introduction

Large collections of medical images are becoming ubiquitous as public resources, and within specific clinical practices. Currently, large studies consist of thousands of images, but in the coming years databases of images of various types will grow to tens of thousands. The availability of such data demands new techniques for image analysis and associated algorithms that are able to efficiently take advantage of these large collections. We begin with a brief discussion of the kinds of algorithms that utilize these large sets of images and how these algorithms demand new technologies for fast image lookup, and end this section with a discussion of how the proposed method addresses this challenge.

One use for a large collection of medical images is to aid in segmentation or tissue classification. *Atlases*, comprising voxel-wise tissue probabilities, incorporate information about spatial location of biological structures. Atlases with hard/soft tissue/object assignments can be used alone, or as "priors", for segmentation and combined with voxel measurements of a specific image to generate label maps in previously unseen images. Atlases are typically constructed by summarizing information concerning image intensities and anatomical shapes from a training set that includes manual segmentations. The information is summarized in (i) an average image (template) and (ii) a tissue probability map. To segment a test image, the template is warped to the test image and the tissue probabilities in the atlas are then transferred to the test image.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 484–491, 2011.

Recent work [1,11,6] shows that there can be significant loss of information in the averaging/summarizing process underlying conventional atlas construction schemes. Alternative strategies rely on *multiatlases* or *nonparametric atlases*. Such schemes consider every image (and associated manual segmentation) in the training set as an atlas. They transfer information from the training set to the test subject by independently warping atlas templates to the latter and, subsequently, warping atlas segmentations to the latter. An important aspect of multi or nonparametric atlases is that one typically selects a small subset of atlases from the training set, comprising the atlas templates that are *most similar* to the test image. Lotjonen et al. [11] use a weighted average of these images, where weights reduce monotonically with *distance*, thus forming a nonparametric estimate of segmentations in the space of images. A small collection of similar images typically registers quite well and results in a *crisper* average, thus providing better quality segmentations compared with those obtained using a summary/average atlas from the entire training set. This has been demonstrated on brain magnetic resonance (MR) images [1] and cardiac MR angiography images [6]. Despite these impressive developments in nonparametric atlases, finding images representing similar anatomy requires registration between the test image and (potentially) every image in the training set. Thus, the problem of efficiently applying these methods to very large training sets remains open.

There are other types of analyses demanding some kind of fast nearest-neighbor (NN) lookup on sets of anatomical images. For instance, Gerber et al. [7] propose a method to learn the underlying manifold structure on sets of brain images, and rely on comparisons (deformable registrations) of each image against every other, resulting in many days of computation time—but ultimately they use only NN relationships to construct the manifold. Wolz et al. [13] similarly use NN relationships to extract biomarkers from a low-dimensional representation under the hypothesis that the low-dimensional coordinates capture information about shape and appearance. Other approaches for learning representations on sets of images exist [2,5], but they share the same need for quickly identifying and quantifying shape-based relationships between similar images.

The problem of fast searching in image spaces has received a lot of attention in computer vision for content-based image retrieval and recognition. In this context, Grauman and Darrell [8] proposed an efficient hierarchical approximation for measuring similarity between histograms of features using a *pyramid match kernel* (PMK). While [8] applies PMK as bags of image features in which spatial location is typically ignored, Lazebnik et al. [10] proposed a variation on the PMK approach, called *spatial pyramid matching* (SPM), which computes approximate geometric correspondences based on histograms of feature locations. They apply SPM primarily to content-based image retrieval, where relatively coarse approximations to feature localization and matching are justified. However, these methods *do* offer bounded-error approximations of L1, bipartite feature matching—raising the possibility that they would have utility in evaluating anatomical shape similarity. This observation is supported, in part, by medical image registration algorithms that rely on hierarchical feature matching [12].

The purpose of this paper is to examine the application of a fast, hierarchical approach for matching image features to the problem of shape-based lookups into large databases of medical images. The result is a novel method for fast, shape-based, approximate search of medical images. The proposed method is several orders of magnitude faster than typical 3D registration methods, thereby making NN searches practical for large databases of medical images. We describe a mechanism by which MR brain images can be reduced to feature maps that, combined with the matching algorithms above, offer acceptable approximations to deformation distance. This paper also validates the accuracy of the fast searches relative to accuracies produced by registration methods. Furthermore, the paper presents results on the application of the proposed method for brain-tissue segmentation, where the results demonstrate segmentation performance as good or better than with registration-based metrics.

2 Methods

The proposed method for measuring the similarity between two brain MR images relies on hierarchical feature matching using SPM [10]. This section reviews the formulation for SPM and describes its application to shape comparison on brain MR images. The use of SPM on image features entails a pipeline of four steps: (i) image preprocessing—intensity and spatial normalization and edge-preserving filtering, (ii) feature extraction—Canny edge detection, (iii) feature labeling—data-driven codebook generation and label assignment, and (iv) SPM comparisons—construction and comparison of a collection of labeled, multilevel feature histograms. Steps (i)–(ii) rely on public domain implementations (i.e. ITK) for affine registration, intensity normalization, filtering, and edge detection that are well known from the literature. We describe steps (iii) and (iv), next.

Feature Extraction and Codebook Construction: The proposed method relies on edges to capture anatomical shape. SPM matches a set of feature maps in a way that approximates L1, bipartite matching, but does not enforce any kind of smoothness on the matching, which is typcially important for medical image registration. In order to avoid mismatching nearby edges associated with different anatomies, we classify edges into different types based on local shape. For edge classification, we use the orientation and curvature of the level sets of the input image at each edge point, which are computed by first- and second-order derivatives of Gaussians. Then each edge is classified by (i) a clustering process on the orientation-curvature features extracted at all edge voxels followed by (ii) a quantization/coding process that assigns a label to every edge voxel based on the cluster center to which it is nearest. We use k-means for clustering and refer to cluster centers as the *codebook* of possible edge patterns $\mathcal{C} = \{c_1, \cdots, c_k\}$ in brain MR images. Figure 1 shows examples of the hard-coded edge maps. To alleviate errors/artifacts in the hard quantization/coding process, we assign "soft" labels as in fuzzy-c-means, where each edge point has a set of membership values of belonging to every cluster.

Shape Similarity using Spatial Pyramid Matching: Each edge type in the codebook is considered as a different feature map, and each feature map is represented as a multilevel spatial histogram, or spatial *pyramid*. For each edge type SPM approximates the distance between nearby edges by the *overlap* of the spatial histograms of the edge images at each level. Features in two different pyramids, are "matched" if they lie in the same bin, at a specific level in the pyramid. The algorithm assigns matches starting from the finest level and proceeds progressively through coarser levels—to a final single bin, where any remaining features are matched. Once labels get matched at a specific level of the pyramid, they are not considered for matches at coarser levels. This pyramid matching process is essentially an approximate, greedy algorithm for bipartite matching between labeled point sets. For each feature match at each level there is an associated weight that assigns a cost (or reward) for that match. To approximate L1 matching, one would use a cost proportional to the bin size at each level in the hierarchy, and sum over each bin at all levels. This is, however, sensitive to mismatches, outliers, and inherent inaccuracies in binning (worse at coarser levels). As an alternative, Grauman and Darrell [8] assign weights that are the reciprocal of the bin size (i.e. rewarding finer-level matches) and show that the resulting similarity measure is a Mercer kernel.

Given two images A and B of a particular edge type $c \in \mathcal{C}$, consider their spatial pyramid representations h_A and h_B, with L levels in order of increasing resolution. Each level l is a spatial histogram with 2^{l-1} bin along each image dimension d, with a total of $M_l = 2^{d(l-1)}$ bins. Let h_A^l and h_B^l be the representations of A and B, respectively, at level l with $h_A^l(i)$ and $h_B^l(i)$ the *soft* count of edges of type c in the ith bin. Then, the number of matches between h_A^l and h_B^l is given by the *histogram intersection*:

$$I(h_A^l, h_B^l) = \sum_{i=1}^{M_l} \min(h_A^l(i), h_B^l(i)). \tag{1}$$

Given the matches that have already occurred at levels $[l+1, L]$ finer than l, the *new matches* occurring at level $l < L$ are: $N_l = I(h_A^l, h_B^l) - I(h_A^{l+1}, h_B^{l+1})$. At the finest level L, all the matches $I(h_A^L, h_B^L)$ are new: $N_L = I(h_A^L, h_B^L)$. If h_A and h_B are the same, then all (new) matches occur at the finest level L.

Similarity between point-sets A and B is measured using the PMK,

$$\kappa(A, B) = \sum_{l=1}^{L} w_l N_l = I(h_A^L, h_B^L) + \sum_{l=0}^{L-1} \frac{1}{2^{L-l}} (I(h_A^l, h_B^l) - I(h_A^{l+1}, h_B^{l+1})), \tag{2}$$

which is a weighted combination of the new matches at each level. The matches at finer levels are given higher weights because matches at finer levels indicate lower separation distance between the matched points. Specifically, the weight w_l, for matches at level l, is $w_l = 1/(2^{L-l})$ that decreases exponentially with level coarseness; this rate is consistent with the exponential increase of the bin size with level coarseness. For the finest level, $w_L = 1$.

To ensure a maximum PMK similarity $\kappa(\cdot, \cdot)$ of 1, $\kappa(\cdot, \cdot)$ is normalized as $\widetilde{\kappa}(A, B) = \kappa(A, B)/\sqrt{\kappa(A, A)\kappa(B, B)}$. Then, **SPM similarity** between two

images is the sum of normalized PMK similarities for each edge type: $S(A,B) = \sum_{c=1}^{k} \widetilde{\kappa}(A_c, B_c)$, where $\widetilde{\kappa}(A_c, B_c)$ is the normalized PMK similarity for the cth code. SPM's algorithmic complexity is linear in the point-set cardinality, number of pyramid levels, and number of codes.

3 Evaluation Methodology

The SPM approach is compared to diffeomorphic registration (LDDMM) [4], which captures geodesic distances of deformation fields, and an elastic registration approach (not strictly a metric) as employed in [7] . This is not an exhaustive list of metrics for registration-based shape comparison, but is representative of the state of the art in terms of compute time and performance. Our hypothesis is that SPM is comparable to the various options for deformation metrics (different algorithms and parameters) in the evaluation of shape differences and k-NNs.

Consider a training set of brain images $\mathcal{B} = \{B_1, \cdots, B_M\}$ and a test set $\mathcal{A} = \{A_1, \cdots, A_N\}$. To compare k-NN performance, we find the k-NNs $B_j \in \mathcal{B}$ of all $A_k \in \mathcal{A}$ using (i) SPM with 2 different codebook sizes, (ii) elastic registration with 2 values of the standard parameter, say λ, which balances the weight between the image-match and deformation smoothness, and (iii) diffeomorphic registration with two λ values. Our evaluation metrics are:

1. *Accuracy* (π): For image A_i, let the k-NNs given by two different methods (among SPM, diffeomorphic registration, and elastic registration) be $\eta_S(A_i,k) \subset \mathcal{B}$ and $\eta_R(A_i,k) \subset \mathcal{B}$. Then, π for S relative to R is: $\pi = (1/N)\sum_{i=1}^{N} |\eta_R(A_i,k) \cap \eta_S(A_i,k)|/|\eta_R(A_i,k)|$.
2. ϵ-ball radius ratio (γ): If images in $\mathcal{B}$ are very similar to each other, there are many neighbors within a relatively small distance. Thus, we propose an ϵ-ball metric, from topology theory, to quantify the extra distance introduced by differences in k-NNs. Given the k-NNs $\eta_R(A_i,k)$, we find the smallest set $\eta_S(A_i,k^*)$ such that $\eta_R(A_i,k) \subset \eta_S(A_i,k^*)$. Note: $k^* \geq k$. Then, $\gamma = (1/N)\sum_{i=1}^{N} [\max_{B\in\eta_S(A_i,k^*)} d_R(A_i,B)]/[\max_{B\in\eta_R(A_i,k)} d_R(A_i,B)]$, where $d_R(A_i,B)$ is the registration-based distance between A_i and B; $\gamma \geq 1$.
3. To compare performance for atlas-based segmentation, we use the *Dice* overlap measure between an atlas-based tissue segmentation A and the ground-truth tissue segmentation G. The Dice overlap is: $(2|A \cap G|)/(|A| + |G|)$.

The evaluation employs: (i) a database $\mathcal{B}$ of 259 skull stripped, T1-weighted brain MR images (normal humans; image size $160 \times 192 \times 176$ voxels) with probabilistic tissue segmentations and (ii) a database $\mathcal{A}$ of 20 BrainWeb [3] images.

4 Results and Discussion

The main motivation of using SPM similarity is to enable *fast* searches for NNs. Thus, we begin by discussing the computational cost and times for matching tasks between two 3D brain images SPM, elastic registration, and diffeomorphic registration. Diffeomorphic registration between two brain MR images with 200^3

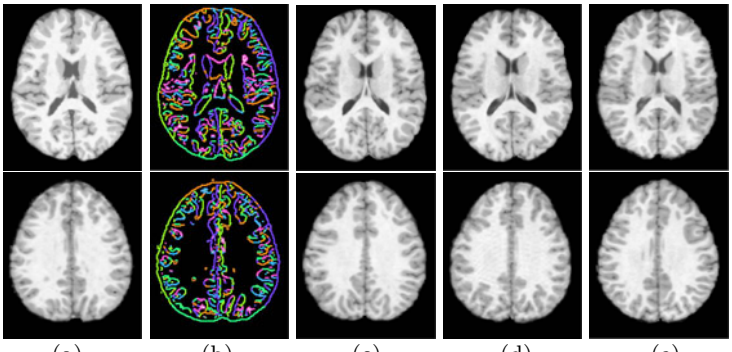

Fig. 1. (a) Axial slices from 2 examples of 3D T1w brain MR images in the 259-brain database (see text). (b) Coded edge maps associated with (a). (c)–(e) 3 nearest neighbors (NNs), for the images in (a), found by the proposed method. The NNs clearly reflect the overall shape of the skull, ventricles, etc. SPM similarities between the top and bottom row images are very small, reflecting the large differences in skull shape.

voxels, using LDDMM [4] with 100 iterations and 5 intermediate timepoints, requires roughly 8×10^{13} flops. An elastic registration requires approximately 50-times fewer, i.e. 2×10^{12}, flops. For the same example, SPM requires less than 10^8 flops. Typical registration methods have a high demand for trilinear interpolations, which can result in poor memory-access patterns. SPM, on the other hand, eliminates interpolations entirely. Codebook construction and edge labeling are computationally light and need to be performed just once, during preprocessing, and require computation of the same order as in preprocessing for registration. Run times on conventional, serial CPUs confirm, approximately, these calculations. A careful C++ CPU implementation of LDDMM, registration between two brain images requires around 4-6 hours, while a similar implementation of elastic registration requires roughly 10 minutes. A Matlab implementation of SPM takes roughly 40 seconds. The experiments in this paper used a GPU implementation [9] of LDDMM (roughly 10 minutes required). In real systems to be deployed one would implement SPM on a parallel architecture such as a GPU (a 100× speedup over Matlab is expected). Such experiments becomes architecture specific and are beyond the scope of this paper.

Selecting 3500 random pairs, we compute SPM similarities and registration-based distances. Linear regression between: (i) registration-based distances provided by different methods gives slopes around 0.83, R^2 around 0.75; (ii) SPM similarity and registration-based distance gives slopes around -0.73, R^2 around 0.6. This indicates that SPM similarity correlates very well with registration-based distances.

For accuracy we consider a subset of 100 from the 259 brain MR images. We compute SPM similarities and registration-based distances for all pairs in the dataset. Subsequently, we compute π for k-NN searches with $k = 10$ for every image in the dataset and average the values over the dataset (the standard deviation is around 0.13 for all results). Figure 2 indicates that (i) π for SPM, relative

k-NN Accuracy

	Diff_1	Diff_2	Elas_1	Elas_2	SPM_1	SPM_6	SPM_{18}
Diff_1	1	0.39	0.22	0.35	0.25	0.32	0.32
Diff_2		1	0.51	0.69	0.45	0.53	0.53
Elas_1			1	0.45	0.36	0.36	0.36
Elas_2				1	0.42	0.52	0.53
SPM_1					1	0.56	0.52
SPM_6						1	0.86
SPM_{18}							1

Average ϵ-Ball Radius Ratio

	Diff_1	Diff_2	Elas_1	Elas_2	SPM_1	SPM_6	SPM_{18}
Diff_1	1	1.24	1.30	1.25	1.38	1.33	1.32
Diff_2	1.26	1	1.20	1.16	1.33	1.29	1.26
Elas_1	1.29	1.19	1	1.23	1.29	1.27	1.27
Elast_2	1.13	1.07	1.10	1	1.12	1.09	1.09

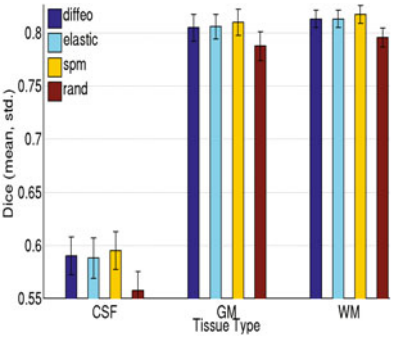

Fig. 2. [**Top-Left**] k-NN accuracy π for different parameter values (smoothness of warp) of diffeomorphic and elastic registration as well as different numbers of edge codes (subscripts) for SPM. [**Bottom-Left**] Average ϵ-ball radius ratio γ for different methods compared to reference distances listed in the first column. The first two rows use diffeomorphic registration to give the reference $d_R(\cdot,\cdot)$, while the last two use elastic registration. [**Right**] Dice overlaps for atlas-based segmentation of 20 BrainWeb images [3]. Atlases are constructed using 10 nearest neighbors selected by diffeomorphic registration, elastic registration, SPM, and random selection.

to each registration method, is very close to π measured for one registration method relative to another; and (ii) using codebooks, rather than simply edges, improves SPM's performance. Figure 2 also shows that SPM's performance is also robust respect to the size of the codebook. The second table in Figure 2 shows ϵ-ball radius ratios, indicating that SPM and registration-based methods misidentify kNNs in a way that introduces similar errors in image distances.

Next, the utility of SPM is evaluated for *atlas-driven segmentation* by comparing the quality of local kNN based atlases. Notice that SPM does *not provide a coordinate transformation*, thus after finding the k-NNs with each method, we use the LDDMM registration to warp the k-NN images (and segmentations) to the test image and average their tissue probabilities to construct the atlas. In order to isolate k-NN search performance, tissue values are assigned based only on the local atlas and no intensity based segmentation is used. We used 20 test images from the BrainWeb dataset [3] and build the k-NN atlases from the 259-brain image database. To evaluate the efficacy of each atlas we use the Dice overlap of segmentation with respect to ground truth. We average these Dice values over all 20 test brains. Figure 2 shows that Dice overlaps using SPM outperform the others for this dataset. Dice overlaps for elastic and diffeomorphic registration are very close to each other. Random selection of images reduces performance by several percent, which is consistent with findings in [6,1].

Conclusion: This paper demonstrates the effectiveness, in terms of both compute times and accuracy, of an SPM-based k-NN search for brain MR images. Some engineering decisions, such as limiting features to edges and coding based on orientation and curvature, are effective, but warrant further investigation. For instance, one might learn features directly from image patches, especially if

the images in question are not dominated by high-contrast edges. The question of generally applicability is also quite interesting. While the results for brain images are quite promising, successful demonstrations on other types of data, such as cardiac, have to be completed. However, the basic properties of the proposed framework, including a great deal of flexibility in possible features and robustness to feature outliers, promises to be more general.

Acknowledgements. This work is supported by the NIH/NCRR Center for Integrative Biomedical Computing - 2P41 RR0112553-12, and the NIH/NCBC National Alliance for Medical Image Computing - U54-EB005149.

References

1. Aljabar, P., Heckemann, R., Hammers, A., Hajnal, J., Rueckert, D.: Multi-atlas based segmentation of brain images: Atlas selection and its effect on accuracy. NeuroImage 46(3), 726–738 (2009)
2. Aljabar, P., Wolz, R., Srinivasan, L., Counsell, S., Boardman, J., Murgasova, M., Doria, V., Rutherford, M., Edwards, A., Hajnal, J., Rueckert, D.: Combining morphological information in a manifold learning framework: Application to neonatal MRI. In: Jiang, T., Navab, N., Pluim, J.P., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. 6363, pp. 1–8. Springer, Heidelberg (2010)
3. Aubert-Broche, B., Collins, D., Evans, A.: A new improved version of the realistic digital brain phantom. Neuro Image 32(1), 138–145 (2006)
4. Beg, F., Miller, M., Trouve, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. Int. J. Comp. Vis. 61(2) (2005)
5. Chen, T., Rangarajan, A., Eisenschenk, S.J., Vemuri, B.C.: Construction of neuroanatomical shape complex atlas from 3D brain MRI. In: Jiang, T., Navab, N., Pluim, J.P., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. 6363, pp. 65–72. Springer, Heidelberg (2010)
6. Depa, M., Sabuncu, M.R., Holmvang, G., Nezafat, R., Schmidt, E.J., Golland, P.: Robust atlas-based segmentation of highly variable anatomy: Left atrium segmentation. In: Camara, O. (ed.) MICCAI Workshop Stat. Atlases Comp. Models Heart, pp. 1–8 (2010)
7. Gerber, S., Tasdizen, T., Fletcher, P., Joshi, S., Whitaker, R.: ADNI: Manifold modeling for brain population analysis. Med. Imag. Analysis 14(5), 643–653 (2010)
8. Grauman, K., Darrell, T.: The pyramid match kernel: Efficient learning with sets of image features. J. Mach. Learn. Res. 2, 725–760 (2007)
9. Ha, L., Kruger, J., Fletcher, T., Joshi, S., Silva, C.T.: Fast parallel unbiased diffeomorphic atlas construction on multi-graphics processing units. In: Euro. Symp. Parallel Graph. Vis., pp. 65–72 (2009)
10. Lazebnik, S., Schmid, C., Ponce, J.: Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In: IEEE Conf. Computer Vision and Pattern Recognition, vol. 2, pp. 2169–2178 (2006)
11. Lotjonen, J., Wolz, R., Koikkalainen, J., Thurfjell, L., Waldemar, G., Soininen, H., Rueckert, D.: ADNI: Fast and robust multi-atlas segmentation of brain magnetic resonance images. NeuroImage 49(3), 2352–2365 (2010)
12. Shen, D., Davatzikos, C.: HAMMER: Hierarchical attribute matching mechanism for elastic registration. IEEE Trans. Med. Imag. 21(11), 1421–1439 (2002)
13. Wolz, R., Aljabar, P., Hajnal, J.V., Rueckert, D.: Manifold learning for biomarker discovery in MR imaging. In: Wang, F., Yan, P., Suzuki, K., Shen, D. (eds.) Conf. Mach. Learn. Med. Imag., pp. 116–123 (2010)

3D Active Shape Model Segmentation with Nonlinear Shape Priors

Matthias Kirschner, Meike Becker, and Stefan Wesarg

Graphisch-Interaktive Systeme, Technische Universität Darmstadt,
Fraunhoferstraße 5, 64283 Darmstadt, Germany
{matthias.kirschner,meike.becker,stefan.wesarg}@gris.tu-darmstadt.de

Abstract. The Active Shape Model (ASM) is a segmentation algorithm which uses a Statistical Shape Model (SSM) to constrain segmentations to 'plausible' shapes. This makes it possible to robustly segment organs with low contrast to adjacent structures. The standard SSM assumes that shapes are Gaussian distributed, which implies that unseen shapes can be expressed by linear combinations of the training shapes. Although this assumption does not always hold true, and several nonlinear SSMs have been proposed in the literature, virtually all applications in medical imaging use the linear SSM. In this work, we investigate 3D ASM segmentation with a nonlinear SSM based on Kernel PCA. We show that a recently published energy minimization approach for constraining shapes with a linear shape model extends to the nonlinear case, and overcomes shortcomings of previously published approaches. Our approach for nonlinear ASM segmentation is applied to vertebra segmentation and evaluated against the linear model.

1 Introduction

Accurate segmentation of organs in medical images is challenging, because adjacent structures are often mapped to the same range of intensity values, which makes it hard to detect their boundaries. In these cases, prior knowledge of the shape of an organ can be used to avoid that the segmentation leaks into the neighboring structures. One of the most popular segmentation algorithms with a shape prior is the Active Shape Model (ASM) [1], which uses a linear, landmark-based Statistical Shape Model (SSM).

The linear SSM is learned by a Principal Component Analysis (PCA) of the training shapes, which implies the assumption that the shapes are Gaussian distributed. While this model has been successfully applied to the segmentation of various structures in medical imaging, there are cases in which the assumption does not hold true. An example for such a case is shown in Figure 1 (left), which shows a training data set consisting of 14 lumbar (L1-L3) and 9 thoracic (Th10-Th12) vertebrae, projected to the first two principal components. One can clearly see that each vertebra type (thoracic, lumbar) builds a cluster. The mean shape of this training data set shows characteristics of both lumbar and thoracic shapes

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 492–499, 2011.

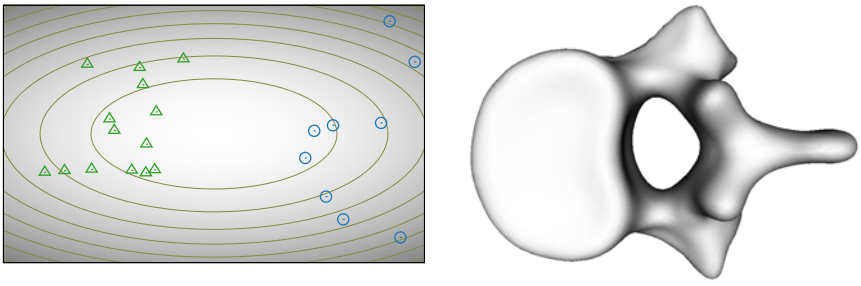

Fig. 1. Left: A set of 14 lumbar vertebrae (green triangles) and nine thoracic vertebrae (blue circles), projected to the first two principal components. The shading encodes values of the log-likelihood function of the learned Gaussian. The brighter the shading, the higher is the probability. The curves show isocontours of the function. Right: The mean shape, which shows characteristics of both lumbar and thoracic vertebrae.

(Figure 1, right). While it is the most likely shape with respect to the linear SSM, we would not expect a vertebra of this shape in the human body.

To describe such datasets more accurately, nonlinear, multimodal models are required. Several nonlinear extensions of SSMs for ASM segmentation have been proposed in the literature, for example based on Kernel Principal Component Analysis (KPCA) [2,3] or Gaussian Mixture Models [4]. However, to the best of our knowledge, these approaches have never been applied to segmentation in 3D, which typically suffers from the huge gap between the high dimensionality of the shapes and the small number of training examples. In fact, Heimann and Meinzer observe in their recent review [5] that nonlinear, landmark-based SSMs have hardly attracted the attention of the community so far. This stands in contrast to level set segmentation with shape priors, where nonlinear techniques such as KPCA [6] or Parzen Density estimation [7] are becoming increasingly popular for learning priors from signed distance functions.

In this paper, we take a first step towards bridging the gap between theoretical models on one side and real world applications on the other. We summarize our main contributions as follows:

1. We discuss methodological shortcomings of previously published approaches for constraining shapes in nonlinear ASMs, especially in the case of high dimensional training data sets with few training examples.
2. We present a unified formulation for constraining shapes in ASMs, which contains linear shape models as special cases.
3. We present a particular nonlinear ASM based on KPCA (KPCA-ASM).
4. We apply the KPCA-ASM to the real world problem of segmenting vertebrae in 3D CT scans. Our evaluation provides insights into the properties of the nonlinear ASM, and shows that it can increase the segmentation accuracy in certain applications.

2 Statistical Shape Models and the Active Shape Model

In this section, we briefly review Statistical Shape Models (SSM) and the segmentation algorithm Active Shape Model (ASM), which were both introduced by Cootes et al. [1]. In SSMs, each of the S training shapes is represented by N corresponding landmarks $\boldsymbol{x}_i^{(j)} = (x_{ij}, y_{ij}, z_{ij})$, which are concatenated to the shape vector $\boldsymbol{x}_i = (x_{i1}, y_{i1}, z_{i1}, \ldots, x_{iN}, y_{iN}, z_{iN}) \in \mathbb{R}^{3N}$. Procrustes alignment is used to compute a common coordinate system for the shapes. The linear SSM uses PCA to capture the statistics of the training shapes. One computes the eigenvectors $\boldsymbol{p}_1, \ldots, \boldsymbol{p}_{3N}$ and corresponding eigenvalues $\lambda_1 \geq \ldots \geq \lambda_{3N}$ of the covariance matrix $\boldsymbol{C} = \frac{1}{S-1} \sum_{i=1}^{S} (\boldsymbol{x}_i - \bar{\boldsymbol{x}})(\boldsymbol{x}_i - \bar{\boldsymbol{x}})^T$ of the shapes. We discard eigenvectors with index $i > t$, were $t = \min\{t' | \sum_{i=1}^{t'} \lambda_i / \sum_{i=1}^{3N} \lambda_i > 0.98\}$. The $3N \times t$-matrix of retained eigenvectors is denoted by $\boldsymbol{P} = (\boldsymbol{p}_1 | \ldots | \boldsymbol{p}_t)$, and the subspace of $\mathbb{R}^{3N}$ spanned by these eigenvectors is denoted by $\mathcal{U}$.

The ASM is an iterative algorithm, which is initialized by placing the mean shape into the image. Each iteration consists of two steps: Landmarks are displaced to optimal image features in their vicinity, which have been selected by an appearance model. Then the deformed shape is constrained with the SSM.

3 Shape Constraints for Linear and Nonlinear ASMs

In the original ASM algorithm, a shape $\boldsymbol{x}'$ is constrained to a plausible shape $\boldsymbol{x}$ (a) by projecting $\boldsymbol{x}'$ to $\mathcal{U}$ using the formula

$$\boldsymbol{b} = \boldsymbol{P}^T (\boldsymbol{x}' - \bar{\boldsymbol{x}}), \tag{1}$$

(b) by imposing bounds on the subspace vector $\boldsymbol{b}$, e. g. by enforcing that $b_i \in [-3\sqrt{\lambda_i}, 3\sqrt{\lambda_i}]$, and (c) by generating the plausible shape by $\boldsymbol{x} = \bar{\boldsymbol{x}} + \boldsymbol{P}\boldsymbol{b}$. This approach cannot be trivially extended to nonlinear shape models. For example, Romdhani et al. [2] place upper bounds on the KPCA components like in the linear model, but it has been shown that this approach is not valid in general [3].

Instead, Cootes et al. [4] and Twining and Taylor [3] propose to constrain shapes with nonlinear shape models by minimizing an energy $E_{\text{shape}}(\boldsymbol{x})$ until $E_{\text{shape}}(\boldsymbol{x}) \leq \theta$, where θ defines a plausibility threshold. In contrast to the least squares projection in Eq. 1, this approach does not guarantee a result which is close to the deformed shape $\boldsymbol{x}'$. Furthermore, it restricts the space of allowed shapes to those shapes which are close to the isocontour $E_{\text{shape}}(\boldsymbol{x}) = \theta$, and it remains open how to choose θ. Finally, the 'proximity to data' measure for KPCA [3] does not penalize the possible loss of information that may occur when projecting to a feature space. This loss can be considerably high, especially in case of high dimensional shapes and small training data sets.

A unified approach for constraining shapes: The first two of the aforementioned problems can be solved by adding an image energy $E_{\text{image}}(\boldsymbol{x})$ that penalizes deviations from the boundary detected by the appearance model. A deformed shape $\boldsymbol{x}'$ is constrained by minimizing $E(\boldsymbol{x}) = \alpha \cdot E_{\text{image}}(\boldsymbol{x}) + E_{\text{shape}}(\boldsymbol{x})$

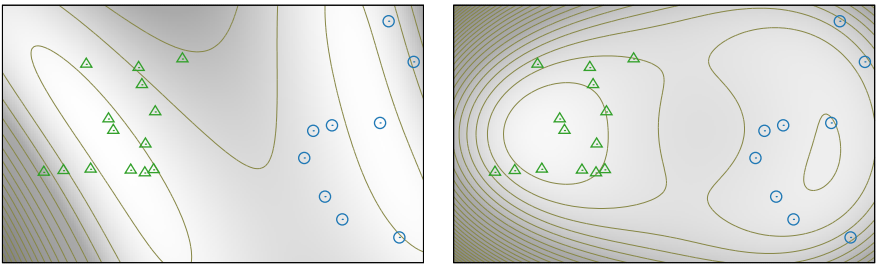

Fig. 2. Visualization of two KPCA-based shape energies (Left: σ_3; Right: σ_9). In contrast to the Gaussian energy (Figure 1), these energies are multimodal.

until a (possibly local) minimum is reached. The parameter α balances image and shape energy, and must be chosen appropriately. In order to account for the loss of information that occurs when projecting to a feature space, we use the general approach of Moghaddam and Pentland [8]: The shape energy is split into two terms, namely the *distance in feature space* (DIFS) and the *distance from feature space* (DFFS). The DIFS computes the probability in feature space, whereas the DFFS penalizes the costs of the projection to the feature space. In the linear case, the feature space is $\mathcal{U}$, and the energy becomes

$$E_{\text{shape-pca}}(\boldsymbol{x}) = \frac{1}{2}\sum_{i=1}^{t}\frac{\boldsymbol{b}_i^2}{\lambda_i} + \frac{1}{2\rho}\|\boldsymbol{r}\|^2, \tag{2}$$

where $\|\boldsymbol{r}\|^2 = \|\boldsymbol{x} - \bar{\boldsymbol{x}}\|^2 - \|\boldsymbol{b}\|^2$ is the costs of the projection of $\boldsymbol{x}$ to $\mathcal{U}$, and $\rho = \sum_{i=t+1}^{3N}\frac{\lambda_i}{3N-t}$. We have recently shown [9] that an ASM based on this energy has better segmentation accuracy than the standard approach which uses Eq. 1 as it is less restrictive: By using the DFFS, shapes can leave the subspace $\mathcal{U}$. It permits a limited amount of variability that has not been previously observed in the training data, such that the model adapts better to unseen shapes.

KPCA shape energy: KPCA [10] is a nonlinear extension of PCA: The basic principle is to map the input data to a feature space F and subsequently do a linear PCA of the mapped data. The mapping is implicitly defined by a kernel function $k(\boldsymbol{x}, \boldsymbol{y})$. We use a Gaussian kernel function $k(\boldsymbol{x}, \boldsymbol{y}) = \exp -\frac{\|\boldsymbol{x}-\boldsymbol{y}\|^2}{2\sigma^2}$, where the parameter σ controls the width of the kernel. The centered kernel is defined by $\tilde{k}(\boldsymbol{x}, \boldsymbol{y}) = k(\boldsymbol{x}, \boldsymbol{y}) - \frac{1}{S}\sum_{i=1}^{S}(k(\boldsymbol{x}, \boldsymbol{x}_i) + k(\boldsymbol{y}, \boldsymbol{x}_i)) + \frac{1}{S^2}\sum_{i=1}^{S}\sum_{j=1}^{S}k(\boldsymbol{x}_i, \boldsymbol{x}_j)$.

In order to do KPCA of the training shapes, we first compute the Gram matrix $\boldsymbol{K} \in \mathbb{R}^{S\times S}$ with the entries $\boldsymbol{K}_{ij} = \tilde{k}(\boldsymbol{x}_i, \boldsymbol{x}_j)$, and then determine the eigendecomposition of $\boldsymbol{K}$, which yields its eigenvectors $a_1, \ldots, a_S$ and eigenvalues $\tilde{\nu}_1, \ldots, \tilde{\nu}_S$. We introduce a regularization parameter $\tilde{\epsilon} = 0.001$ and discard all eigenvectors with $\tilde{\nu}_i \leq \tilde{\epsilon}$. The remaining r eigenvectors are re-normalized such that $\tilde{\nu}_i^{-\frac{1}{2}} a_i^T a = 1$. Moreover, we define $\nu_i = \frac{1}{S-1}\tilde{\nu}_i$ and $\epsilon = \frac{1}{S-1}\tilde{\epsilon}$.

The r KPCA components of $\boldsymbol{x}$ can be computed by $\beta_i = \sum_{j=1}^{S} a_{ij}\tilde{k}(\boldsymbol{x}_i, \boldsymbol{x})$. The shape energy for a KPCA shape model is defined by

$$E_{\text{shape-kpca}}(\boldsymbol{x}) = \sum_{i=1}^{t} \frac{\beta_i^2}{\nu_i} + \frac{1}{\epsilon}(\tilde{k}(\boldsymbol{x}, \boldsymbol{x}) - \sum_{i=1}^{t} \beta_i^2). \tag{3}$$

Note the similarity between Eq. 3 and Eq. 2. In fact, the above definitions are designed so that in the case of $\epsilon = \rho$ and $k_{\text{lin}}(\boldsymbol{x}, \boldsymbol{y}) = x^T y$, both energies are identical up to a constant factor of 2.

We want to mention that Eq. 3 has been previously used by Cremers et al. [11] in a Mumford-Shah based level set segmentation system. However, to the best of our knowledge, it has never been used in an ASM. Instead of evolving a curve, we find a local minimum of the combined image and shape energy in each ASM iteration. We use the L-BFGS method [12] to minimize the energy, which is designed for high-dimensional problems.

Figure 2 visualizes two instances of the function $E_{\text{shape-kpca}}(\boldsymbol{x})$, learned from the vertebral shape data set. We used here the Gaussian kernel with widths σ_3 and σ_9, where $\sigma_K = K \cdot \frac{1}{S}\sum_{i=1}^{S} \min_{j \in \{1,\ldots,S\}} \|\boldsymbol{x}_i - \boldsymbol{x}_j\|$, a choice inspired by Cremers et al. [11]. In both cases, the KPCA shape energies are multimodal, in contrast to the linear shape energy (Figure 1), which has a single peak in a region which is not occupied by any training examples.

Enforcing smoothness: A high image energy and a high ρ or ϵ permits larger deviations from the training shapes. While additional flexibility of the SSM is certainly desired, because it allows for more accurate segmentations, no force controls the nature of this deformation. This means that individual landmarks can be drawn towards outliers of the appearance model, and the segmentations become jagged. We therefore explicitly enforce smooth segmentations by smoothing the shape after energy minimization with a Laplacian filter: Let $\boldsymbol{q}$ be the position of a landmark, and let $\boldsymbol{m}$ be the mean of its neighbors. We move $\boldsymbol{q}$ to $\boldsymbol{m}$ if $\|\boldsymbol{q} - \boldsymbol{m}\| > 1.2 \cdot \delta_{\boldsymbol{q}}$. For each landmark, the constant $\delta_{\boldsymbol{q}}$ is set to the maximal distance between $\boldsymbol{q}$ and $\boldsymbol{m}$ in the training data. In each ASM iteration, we iterate five times through all landmarks in order to smooth the shape.

4 Appearance Model and Image Energy

Klinder et al. [13] achieve good segmentation results for vertebrae with an appearance model that considers the image intensity, gradient magnitude and gradient direction. While their model uses various ad-hoc parameters, we combine these features in a more principled, purely statistically motivated way. Each image feature vector sampled at a boundary voxel v consists of a five dimensional feature vector $\boldsymbol{f} = (I(v), \|G(v)\|, G_1(v), G_2(v), G_3(v))$. Here, $I(v)$ is the image intensity at voxel v, $G(v) = (G_1(v), G_2(v), G_3(v))$ is the gradient and $\|G(v)\|$ its magnitude. In our implementation, $G(v)$ is computed in normal direction $\boldsymbol{n}$ of the current landmark. We compute an orthonormal basis $\{\boldsymbol{n}, \boldsymbol{n}', \boldsymbol{n}''\}$, and use

these directions to compute $G(v)$ with finite differences, instead of using the image axes. We model the distribution of $\boldsymbol{f}$ by the product of two independent probability distributions: A Gaussian distribution over $\boldsymbol{f}_1 = (I(v), \|G(v)\|)$ with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$, and a Fisher distribution [14] over $\boldsymbol{f}_2 = G(v)$ with mean $\boldsymbol{\eta}$ and concentration parameter κ. This motivates our fitness function $g(\boldsymbol{f}) = -\frac{1}{2}(\boldsymbol{f}_1 - \boldsymbol{\mu}^T)\boldsymbol{\Sigma}^{-1}(\boldsymbol{f}_1 - \boldsymbol{\mu}) + \kappa \cdot \boldsymbol{f}_2^T\boldsymbol{\eta} - \kappa$, which is the log-likelihood of the assumed distribution up to a constant term. We have $g(\boldsymbol{f}) \leq 0$ for all $\boldsymbol{f}$. The parameters $\boldsymbol{\Sigma}$, $\boldsymbol{\mu}$, κ and $\boldsymbol{\eta}$ are learned from the training data. We define the image energy by $E_{\text{image}}(\boldsymbol{x}) = \sum_{j=1}^{N}(1 - g(\boldsymbol{f}^{(j)}))^{-1}\|\boldsymbol{x}^{(j)} - \boldsymbol{x}'^{(j)}\|^2$, where $\boldsymbol{x}'$ is the deformed shape and $\boldsymbol{f}^{(j)}$ the feature at landmark j.

5 Experiments

In our evaluation, we used seven CT scans of the abdomen from a publicly available database[1]. We manually segmented the vertebrae in these scans, and used the data from the first five scans as training data and the remaining as test data. All in all, the training data set consists of nine thoracic (Th10: 1; Th11: 3; Th12: 5) and 14 lumbar (L1: 5; L2: 5; L3: 4) vertebrae. The test data set consists of five thoracic (Th10: 1; Th11: 2; Th12: 2) and six lumbar (L1: 2; L2: 2; L3: 2) vertebrae. In the experiments, we segmented the test data with three different types of ASMs: The standard linear ASM, the linear ASM with energy minimization [9], and the KPCA-ASM with different kernel widths. Dependent on the chosen shape model, the shape energy returns different numerical values, which made it necessary to balance α for each method individually in order to achieve good segmentation results. This was done manually by trying different values for α and bracketing the optimal choice. Initial pose parameters of the SSM were determined manually for each test data set. The same manual initialization was used for each tested approach. Each ASM was executed for 50 iterations. We evaluated the segmentation quality with the measures Volumetric Overlap Error (VOE), Average Surface Distance (ASD), and Hausdorff Distance (HD).

6 Results

The quantitative results of our experiments are shown in Table 1. It can be seen that the KPCA-ASMs outperform the linear ones. Compared to the standard linear ASM, the VOE decreases by over 20 %, the ASD by more than 25 %. Depending on σ, the HD decreases by up to 13 %. The increase of accuracy is especially large for thoracic vertebrae, where the average HD decreases by more than 20 % from 8.7 mm (standard ASM) to 6.8 mm (KPCA-ASM σ_6). A qualitative result for a thoracic vertebra is shown in Fig. 3. The standard ASM required 13 seconds, and the KPCA-ASM between 40 and 65 seconds.

[1] 3D-IRCADb-01: http://www.ircad.fr/softwares/3Dircadb/3Dircadb1/

Table 1. Segmentation results obtained with different ASM algorithms. The table shows average results and standard deviations (SD).

	VOE [%]		ASD [mm]		HD [mm]	
Manual Initialization	60.66	SD: 7.01	3.83	SD: 0.82	16.40	SD: 2.97
Linear ASM (standard)	22.33	SD: 6.55	0.83	SD: 0.33	7.99	SD: 3.41
Linear ASM ($\alpha = 2000$)	20.38	SD: 7.62	0.73	SD: 0.38	7.27	SD: 3.64
KPCA-ASM (σ_3, $\alpha = 1750$)	**17.40**	SD: 6.66	0.61	SD: 0.32	7.95	SD: 3.52
KPCA-ASM (σ_6, $\alpha = 100$)	17.55	SD: 3.80	**0.60**	SD: 0.22	**6.88**	SD: 2.45
KPCA-ASM (σ_9, $\alpha = 50$)	17.45	SD: 4.14	**0.60**	SD: 0.22	6.95	SD: 2.72

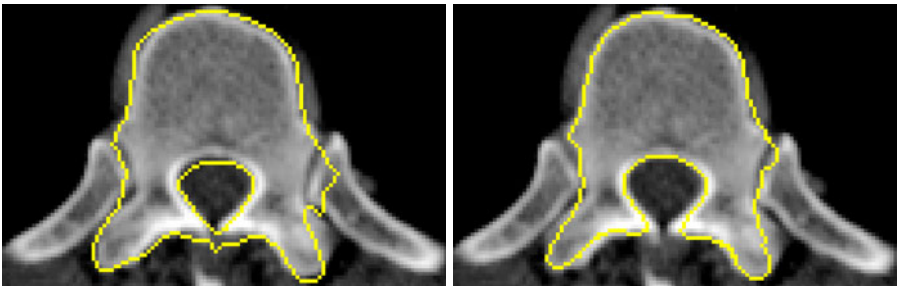

Fig. 3. Comparison of segmentation results for a thoracic vertebra (Th11): The linear ASM cannot always separate ribs and vertebra (left). The KPCA-ASM has more specific shape constraints, and separates vertebra and ribs well (right).

7 Discussion

In this paper, we presented a unified approach for constraining linear and nonlinear shapes in ASMs. Based on this model, we devised a nonlinear KPCA-ASM, which we used for segmenting vertebrae in CT scans. Our evaluation clearly shows that a KPCA-ASM can increase the segmentation accuracy compared to linear ASMs in applications where the training data has a nonlinear distribution. In our opinion, the reason for this is that the KPCA-ASM is more specific than the linear ASM, and thus detects outliers of the appearance model more effectively. While the nonlinearity of our training data is obvious, because it can be separated into lumbar and thoracic vertebrae, the nonlinear structure of the data may be more subtle in other cases. However, our nonlinear SSM was trained in a completely unsupervised way, without prior knowledge about clusters. This is why we expect that a KPCA-ASM will also perform well in other scenarios with nonlinearly distributed training data.

Compared to the results recently reported by Klinder et al. [13], we achieve a significantly better ASD for segmenting individual vertebrae, although it must be mentioned that we have used manual initialization here. Note that Klinder et al. first try to identify the vertebra (e. g. L1), and then segment it with models dedicated to this vertebra. Our quantitative results indicate that, given a

reasonably well initialization, the KPCA-ASM is able to automatically determine the correct vertebra class during segmentation.

A drawback of the KPCA-ASM is that several parameters must be chosen, such as the kernel width σ, the regularization parameter $\tilde{\epsilon}$ and the image energy force α. However, our evaluation shows that the effort for tuning pays off.

To the best of our knowledge, this is the first time that a nonlinear ASM has been successfully applied to the segmentation of organs in tomographic 3D scans. Given the positive evaluation results of our study, we expect that nonlinear ASMs will play an important role in medical image segmentation in the future.

References

1. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active shape models - their training and application. Comput. Vis. and Image Underst. 61(1), 38–59 (1995)
2. Romdhani, S., Gong, S., Psarrou, A.: A multi-view nonlinear active shape model using kernel PCA. In: BMVC 1999, pp. 483–492 (1999)
3. Twining, C.J., Taylor, C.J.: Kernel principal component analysis and the construction of non-linear active shape models. In: BMVC 2001 (2001)
4. Cootes, T.F., Taylor, C.J.: A mixture model for representing shape variation. In: Image and Vision Computing, pp. 110–119 (1997)
5. Heimann, T., Meinzer, H.P.: Statistical shape models for 3D medical image segmentation: A review. Medical Image Analysis 13(4), 543–563 (2009)
6. Dambreville, S., Rathi, Y., Tannenbaum, A.: A framework for image segmentation using shape models and kernel space shape priors. IEEE Trans. Pattern. Anal. Mach. Intell. 30(8), 1385–1399 (2008)
7. Wimmer, A., Soza, G., Hornegger, J.: A generic probabilistic active shape model for organ segmentation. In: Yang, G.-Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (eds.) MICCAI 2009. LNCS, vol. 5762, pp. 26–33. Springer, Heidelberg (2009)
8. Moghaddam, B., Pentland, A.: Probabilistic visual learning for object representation. IEEE Trans Pattern Anal. Mach. Intell. 19, 696–710 (1997)
9. Kirschner, M., Wesarg, S.: Active shape models unleashed. In: Proc. SPIE Medical Imaging 2011: Image Processing (2011)
10. Schölkopf, B., Smola, A.J., Müller, K.R.: Nonlinear component analysis as a kernel eigenvalue problem. Neural Computation 10(5), 1299–1319 (1998)
11. Cremers, D., Kohlberger, T., Schnörr, C.: Shape statistics in kernel space for variational image segmentation. Pattern Recognition 36(9), 1929–1943 (2003)
12. Liu, D.C., Nocedal, J.: On the limited memory BFGS method for large scale optimization. Mathematical Programming B 45, 503–528 (1989)
13. Klinder, T., Ostermann, J., Ehm, M., Franz, A., Kneser, R., Lorenz, C.: Automated model-based vertebra detection, identification, and segmentation in CT images. Medical Image Analysis 13(3), 471–482 (2009)
14. Fisher, N.: Robust estimation of the concentration parameter of Fisher's distribution on the sphere. J. Roy. Statistical Society. Series C 31(2), 152–154 (1982)

Automatic Construction of Statistical Shape Models for Vertebrae

Meike Becker, Matthias Kirschner, Simon Fuhrmann, and Stefan Wesarg

GRIS, TU Darmstadt, Fraunhoferstraße 5, 64287 Darmstadt, Germany
meike.becker@gris.tu-darmstadt.de

Abstract. For segmenting complex structures like vertebrae, a priori knowledge by means of statistical shape models (SSMs) is often incorporated. One of the main challenges using SSMs is the solution of the correspondence problem. In this work we present a generic automated approach for solving the correspondence problem for vertebrae. We determine two closed loops on a reference shape and propagate them consistently to the remaining shapes of the training set. Then every shape is cut along these loops and parameterized to a rectangle. There, we optimize a novel combined energy to establish the correspondences and to reduce the unavoidable area and angle distortion. Finally, we present an adaptive resampling method to achieve a good shape representation. A qualitative and quantitative evaluation shows that using our method we can generate SSMs of higher quality than the ICP approach.

1 Introduction

Segmentation of vertebrae is necessary for several clinical applications such as the treatment of severe herniated vertebral disks or the insertion of pedicle screws. Segmenting vertebrae is a challenging task [12]. Besides local image artifacts or noise, different vertebrae lie close together and are difficult to distinguish from each other and from the adjacent rib cage. Therefore, a general approach is to use a-priori information by means of statistical shape models (SSMs) to make the segmentation more robust. An SSM contains information about the mean of the training set and the possible variations from the mean. For a detailed discussion of SSMs we refer to the review of Heimann et al. [10].

One of the main challenges using SSMs is to establish correspondences in the training data set. Each training shape has to be represented by the same number of 3D points (also called landmarks) and landmarks representing the same anatomical feature should have the same index. A common way to solve this problem is to parameterize every surface to a common base domain and to establish the correspondences on this parameter space. Current work in this context focuses on genus 0 objects such as liver or kidney [2,5,9,14]. Visually spoken, the genus describes the number of holes of a surface. Only Lamecker et al. [13] presented a method independent of the genus, where they employ manually defined patches for each shape and parameterize every patch to a disk. Using this approach, discontinuities may occur along the cuts.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 500–507, 2011.

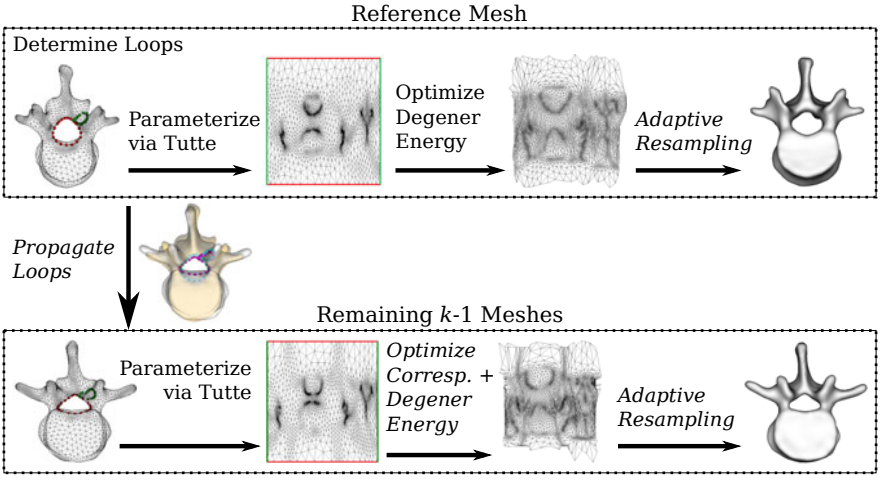

Fig. 1. Overview of our algorithm. Our contributions are printed in italics.

In this work we present an *automated* approach designed for closed surfaces of genus 1 like vertebrae (see Figure 1). We do not rely on a simple Iterative-Closest-Point (ICP) based approach [1] only, since this algorithm does not take into account the triangle structure of the mesh. Hence, flipped triangles in the aligned shapes may occur (see Figure 2).

Related work: Concerning the establishment of correspondences on the parameter space, approaches similar to ours are those of Brett and Taylor [2] and Meier and Fisher [14]. Both work on a parameter space but do not reduce any distortion. Another approach is the work of Davies et al. [5], which is based on information theory. Here, the so-called Minimum Description Length (MDL) function is minimized over the parameter space. We do not use this objective function, since it is more costly than our approach.

A lot of research has been done in the area of mesh parameterization. For a detailed overview we refer to the Siggraph Course Notes of Hormann et al. [11]. When mapping a 3D shape to a plane rectangle, distortion of angle and area is unavoidable in most cases. Tewari et al. [17] presented an algorithm for meshing genus-1 point clouds via parameterization. Since they do not sample on the parameter space but reconstruct the shape by inverting the parameterization, they do not have to consider distortion. Tailored to the parameterization of genus-1 meshes is the work of Steiner and Fischer [16]. While they consider angular distortion, no area distortion is taken into account. However, for a high-quality resampling on the parameter space, it is crucial to minimize both angular and area distortion [6,11]. Only few algorithms exist which minimize both kinds of distortion simultaneously. We choose Degener et al.'s approach [6] since they use a differentiable energy and can obtain a parameterization which is optimal for uniform sampling.

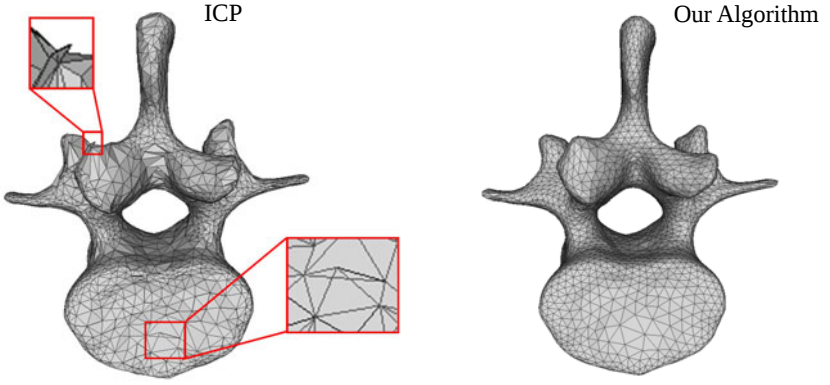

Fig. 2. Comparison of vertebrae after correspondence establishment (with the same number of triangles). Using ICP (*left*), triangle flips occur (red boxes), while our algorithm prevents them (*right*).

Contributions: Our work has the following contributions:

1. We automatically construct a statistical shape model for vertebrae. To the best of our knowledge, this is the first approach without any manual interaction apart from the simple ICP approach.
2. We present an algorithm for the consistent propagation of two reference loops to the remaining training set.
3. We introduce a correspondence term based on point-to-point distances and combine it with Degener's energy [6] in order to establish both accurate correspondences and a good reconstruction quality after uniform sampling.
4. We present a new algorithm for curvature adaptive resampling of the parameter space to further improve the shape representation.

2 Methods

The input of our algorithm is a set $S = \{s_i : i = 1, ..., k\}$ of training surfaces, the so-called shapes, extracted from expert segmented image data. Each shape is approximated by a triangle mesh $\mathcal{M} = (\mathcal{P}, \mathcal{T})$, where $\mathcal{P}$ consists of the approximating points $P_j, j = 1, ..., n$ and $\mathcal{T}$ contains the triangles describing the connectivity of the points. For a given index j of a node we write P_j for the 3D world space coordinates and p_j for the 2D parameter space coordinates. In the following we describe how to establish the correspondences for this training set.

Reference Mesh: We select a randomly chosen reference shape s_{Ref}. From topology we know that we have to cut a surface of genus 1 along two cut loops in order to make it homeomorphic to a rectangle in the plane. For an introduction to topology we refer to the book of Munkres [15]. Let $G_{\text{Ref}} = (V_{\text{Ref}}, K_{\text{Ref}})$ be the corresponding graph to s_{Ref} with vertices V_{Ref} and edges K_{Ref}. On this

graph, we determine the shortest set of loops with a common basepoint using the automatic algorithm of Erickson and Whittlesey [7]. We do this for every point of the triangle mesh and choose the shortest set of loops over all basepoints. Here, we use the Euclidean distance to weight the edges. The found loops span the fundamental group [15] of the vertebra given the basepoint b_{Ref}.

After the two cut loops have been found on the reference mesh, we parameterize it to a rectangle using Tutte's graph embedding method [18] with uniform weights. For mapping the boundary, we map each loop to an edge of the rectangle and create a copy which we map to the opposite edge (see Figure 1). In doing so, opposite edges are identified, so that the corresponding quotient space [15] of the rectangle is homeomorphic to a torus.

The distortion of angle and area on the parameter space leads to a bad shape representation when sampling uniformly on the parameter space. Degener et al. [6] minimize in every node an energy which reduces both angle and area distortion. It is defined as $E_j^{\text{Deg}} = \sum_{T \in 1\text{-ring}(j)} E_{\text{angle}}(T) E_{\text{area}}(T)$ for every node $j = 1, ..., n$, where the set 1-ring(j) contains all triangles T such that $j \in T$,

$$E_{\text{angle}}(T) = \frac{a^2 \cdot \cot\alpha + b^2 \cdot \cot\beta + c^2 \cdot \cot\gamma}{2\,\text{area}(T)} \text{ and} \tag{1}$$

$$E_{\text{area}}(T) = \frac{\text{area}(T_{3D})}{\text{area}(T)} + \frac{\text{area}(T)}{\text{area}(T_{3D})}. \tag{2}$$

The variables a, b and c represent the edge lengths of the triangle T in the parameter space and α, β and γ describe the opposite interior angles of the corresponding triangle T_{3D} in the world space. We use this differentiable energy to reduce the area and angle distortion on the parameter space. We additionally reduce the distortion at the boundary by employing a free-boundary optimization that allows the boundary to move. Points leaving the rectangle on one side, are inserted again on the opposite, identified side.

Remaining training set: Since we want to cut approximately along the same anatomical paths for each training shape, i. e. find consistent loops, we propagate the two reference loops to the remaining training shapes as follows: We first align every remaining training shape $s_i, i = 2, ..., k$, to the reference mesh using the Iterative-Closest-Point (ICP) method, integrating normals in the similarity criterion. We then propagate the reference basepoint b_{Ref} by choosing the nearest neighbor on s_i. Again, we apply Erickson and Whittlesey's algorithm to the propagated basepoint b_i, but we introduce a new cost function for the edges based on the distance to the reference loops. Let L^{Ref} contain the points of the two reference loops. Then the cost function $c_i : K_i \longrightarrow \mathbb{R}_+$ for the graph $G_i = (V_i, K_i)$ of the shape s_i is defined as follows:

$$c_i(j_1, j_2) = 0.5[\min_{\zeta \in L^{\text{Ref}}}(||P_{j_1}^i - P_\zeta^{\text{Ref}}||) + \min_{\xi \in L^{\text{Ref}}}(||P_{j_2}^i - P_\xi^{\text{Ref}}||)]. \tag{3}$$

In this way we ensure that the propagated loop L^i is close to the reference loop L^{Ref}. Furthermore, Erickson and Whittlesey's algorithm guarantees that

we obtain a shape homeomorphic to a rectangle by cutting along the propagated loops. Hence, the propagation of the loops is assured to be successful.

Subsequently, the current shape s_i is parameterized in the same way as the reference mesh but with the difference that we also establish correspondences: Remember that every shape s_i is aligned to the reference mesh via ICP. This global alignment for establishing correspondences can lead to flipped triangles as the ICP algorithm ignores the triangulation. Therefore, we minimize an energy over the parameter space. This energy is defined as $E_j = E_j^{\text{Cor}} + \lambda E_j^{\text{Deg}}$, where E_j^{Cor} describes the correspondence term, E_j^{Deg} is responsible for reducing the distortion and $\lambda \in \mathbb{R}_0^+$ is a free parameter which has to be chosen appropriately. Let $N_{\text{ICP}}^q(j)$ contain the indices of the q nearest neighbors (in our case $q = 8$) on the reference shape of the node j on the current shape s_i. For the correspondence term we use a soft-nearest neighbor approach and define our energy as

$$E_j^{\text{Cor}} = \sum_{k \in N_{\text{ICP}}^q(j)} w_{jk} ||p_j^i - p_k^{\text{Ref}}||, \tag{4}$$

where the weights $w_{jk} = \exp(-||P_j^i - P_k^{\text{Ref}}||)$ describe the distance between the point j on s_i and the point k on s_{Ref} in the 3D world space. To achieve a good shape representation, we add to this correspondence term the distortion energy E_j^{Deg} which we already used for s_{Ref}. Like Degener et al. [6], we minimize the combined energy E_j for $j = 1, ..., n$ using the Polak Ribière method. Triangle flips are avoided by restricting every node to the kernel of its 1-ring.

In this way we parameterize every shape consistently to the parameter space, where we establish correspondences and reduce the resulting distortion.

Adaptive Resampling: Once the parameterizations are in correspondence, we can define a common sampling grid in order to reconstruct corresponding landmarks for each shape from its parameterization. During reconstruction, a compromise between two conflicting goals must be established: The more landmarks we use, the greater becomes the gap between the number of training shapes and their dimension. Conversely, we need enough landmarks to obtain an accurate model of the shape which contains every relevant feature. Obtaining a sparse, yet accurate description of the vertebral shape is particularly challenging. While the vertebral body has a relatively simple geometry, the vertebral processes are thin regions with high curvature. Therefore, it is reasonable to use a curvature adaptive strategy for landmark sampling. Adaptive sampling schemes for SSM construction have been proposed by Heimann et al. [9] and Cates et al. [3]. Heimann et al. use adaptive sampling exclusively to compensate for area distortion in the parameterizations. The particle system method of Cates et al. adaptively oversamples features with high curvature, but does not inherently produce a consistent triangulation for all shapes.

We first construct a dense, uniform sampling grid that is used to sample the parameter space. The large number of sampling points and the low distortion of the parameter spaces ensure that all details of the input shapes are preserved. We average all densely reconstructed shapes to a mean mesh $\overline{\mathcal{M}}$, and use a

state-of-the-art remeshing algorithm [8] in order to obtain a sparse, curvature adaptive representation $\overline{\mathcal{M}}_{\text{Adap}}$ of $\overline{\mathcal{M}}$. Because the initial, uniform sampling grid defines a parameterization of $\overline{\mathcal{M}}$, and all points of $\overline{\mathcal{M}}_{\text{Adap}}$ lie on the surface of $\overline{\mathcal{M}}$, we can compute the parameter space coordinates of the points of $\overline{\mathcal{M}}_{\text{Adap}}$ by simple linear interpolation. These parameter space coordinates define the final, adaptive sampling grid. It is used to sample sparse landmark representations of all input shapes, from which we compute the SSM.

3 Experiments and Results

For our experiments, we used a training set of 14 manually segmented lumbar vertebrae (L1–L3) from CT scans of five different patients. We established correspondences for this training set with our algorithm using different λ. For comparison, we also used the ICP algorithm including normals in the similarity criterion as done by Brett and Taylor [2]. Then we constructed the corresponding SSMs using the standard approach by Cootes et al. [4]. The different λ for our algorithm indicate how much weight we give relatively to the Degener term for the different SSMs. As the Degener term is numerically much higher than the correspondence term, we note that the factor λ does not reveal anything about the absolute ratio of the two energies.

In Figure 3 we show that the local artifacts caused by ICP can be avoided with our algorithm, which leads to SSMs with higher quality. Furthermore, we evaluated the SSMs using the measures of generalization and specificity as proposed by Davies et al. [5] (see Figure 4). For calculating those measures, we sampled 1000 normally distributed model instances and compared them to the

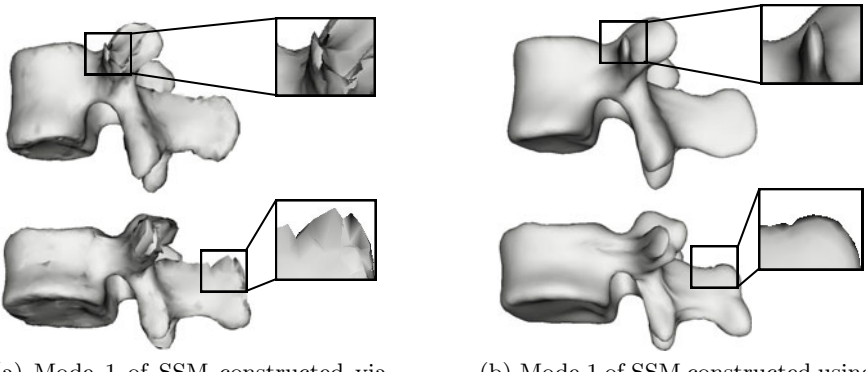

(a) Mode 1 of SSM constructed via ICP using normals

(b) Mode 1 of SSM constructed using our algorithm with $\lambda = 0.00025$

Fig. 3. Comparison of SSMs. The first mode of the respective SSM is shown for 2σ (*first row*) and -2σ (*second row*), where σ describes the standard deviation. Local artifacts occur in the SSM constructed via ICP using normals. These artifacts are eliminated with our algorithm, which leads to an SSM of higher quality.

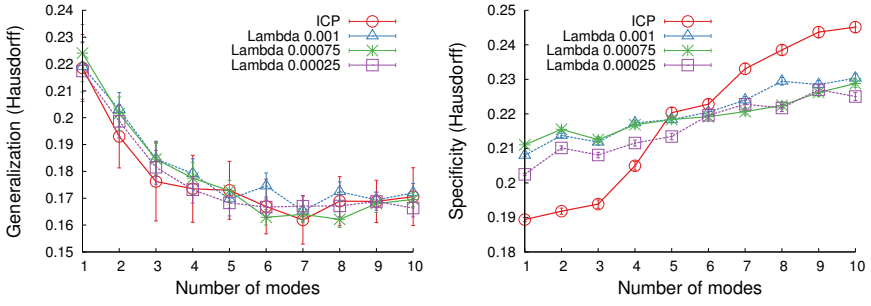

Fig. 4. Evaluation of SSMs using the generalization and specificity measure [5], vertical bars indicating the standard error of the mean. The x-axis gives the number of modes used in the experiment. The generalization values (*left*) are very similar, while our algorithm has better specificity values (*right*) from mode 5 on.

training set using the Hausdorff metric [15]. We did not use the compactness measure since the ICP algorithm may lead to a loss of detail, but loss of detail leads to a better compactness. Hence, we cannot get reliable information about the quality of the SSM using the compactness measure.

We observe that from mode 5 on the specificity measure of our algorithm is better than that of the ICP, while for the generalization, results lie too close together to determine any rank order. To be more specific, all values of our algorithm lie within the interval of the standard error of the ICP.

4 Discussion

We presented an approach for automatically solving the correspondence problem for vertebral shapes. To the best of our knowledge it is the first automated approach to construct SSMs for vertebrae aside from the simple ICP approach. We ensure that for every shape of the training set we cut approximately along the same anatomical paths. When every shape is parameterized to the rectangle, we optimize the correspondence by minimizing a novel combined energy function based on point-to-point distances which additionally penalizes distortion.

For a good reconstruction of the shape, we developed an enhanced method for adaptive resampling which is independent of the underlying parameter space. In contrast to the method of Heimann et al. [9], we do not need to use different maps to sample different regions of the parameter space, even if the parameter space is the unit sphere. The remeshing algorithm does not only distribute landmarks adaptively according to curvature, but also favors equiangular triangles. Thus, it implicitly compensates for small distortion in the parameterizations.

The benefit of our approach over ICP is that it is designed to avoid triangle flips and thus ensures a topological consistent shape model. Experiments show that our SSMs generalize as well as the ICP model, while their specificity is even better. Another advantage of our method is its generality in the sense that it can be used for any surface of genus 1 and furthermore, any (differentiable) objective

function can be used for the correspondence establishment on the parameter space. Hence, it would be interesting to test other correspondence energies such as the MDL function [5] to see if the correspondences can be further improved.

References

1. Besl, P.J., Mckay, H.D.: A method for registration of 3-D shapes. IEEE Trans. Pattern Anal. Mach. Intell. 14(2), 239–256 (1992)
2. Brett, A., Taylor, C.: Construction of 3D shape models of femoral articular cartilage using harmonic maps. In: Delp, S., DiGoia, A., Jaramaz, B. (eds.) MICCAI 2000. LNCS, vol. 1935, pp. 1205–1214. Springer, Heidelberg (2000)
3. Cates, J.E., Fletcher, P.T., Styner, M.A., Shenton, M.E., Whitaker, R.T.: Shape modeling and analysis with entropy-based particle systems. In: Karssemeijer, N., Lelieveldt, B. (eds.) IPMI 2007. LNCS, vol. 4584, pp. 333–345. Springer, Heidelberg (2007)
4. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active shape models their training and application. Comput. Vis. Image Underst. 61, 38–59 (1995)
5. Davies, R.H., Twining, C.J., Taylor, C.J.: Statistical Models of Shape: Optimisation and Evaluation. Springer, London (2008)
6. Degener, P., Meseth, J., Klein, R.: An adaptable surface parameterization method. In: 12th International Meshing Roundtable 2003, Santa Fe, pp. 201–213 (2003)
7. Erickson, J., Whittlesey, K.: Greedy optimal homotopy and homology generators. In: 16th ACM-SIAM Symposium on Discrete Algorithms, pp. 1038–1046. SIAM, Vancouver (2005)
8. Fuhrmann, S., Ackermann, J., Kalbe, T., Goesele, M.: Direct resampling for isotropic surface remeshing. In: Vision, Modeling, and Visualization 2010, pp. 9–16. Eurographics Association, Siegen (2010)
9. Heimann, T., Wolf, I., Meinzer, H.P.: Automatic generation of 3D statistical shape models with optimal landmark distribution. Methods Inf. Med. 46(3), 275–281 (2007)
10. Heimann, T., Wolf, I., Meinzer, H.P.: Statistical shape models for 3D medical image segmentation: A review. Medical Image Analysis 13(4), 543–563 (2009)
11. Hormann, K., Lévy, B., Sheffer, A.: Mesh parameterization: Theory and practice. In: SIGGRAPH 2007 Course Notes, vol. 2. ACM Press, San Diego (2007)
12. Klinder, T., Ostermann, J., Ehm, M., Franz, A., Kneser, R., Lorenz, C.: Automated model-based vertebra detection, identification, and segmentation in CT images. Medical Image Analysis 13(3), 471–482 (2009)
13. Lamecker, H., Lange, T., Seebass, M.: A statistical shape model for the liver. In: Dohi, T., Kikinis, R. (eds.) MICCAI 2002. LNCS, vol. 2489, pp. 421–427. Springer, Heidelberg (2002)
14. Meier, D., Fisher, E.: Parameter space warping: shape-based correspondence between morphologically different objects. IEEE Trans. Med. Imaging 21(1), 31–47 (2002)
15. Munkres, J.R.: Topology. Prentice Hall, Upper Saddle River (2000)
16. Steiner, D., Fischer, A.: Planar parameterization for closed manifold genus-g meshes using any type of positive weights. J. Comput. Inf. Sci. Eng. 5, 118–125 (2005)
17. Tewari, G., Gotsman, C., Gortler, S.J.: Meshing genus-1 point clouds using discrete one-forms. Computers and Graphics 30(6), 917–926 (2006)
18. Tutte, W.T.: How to draw a graph. London Math. Soc. 13, 743–768 (1963)

Graph Based Spatial Position Mapping of Low-Grade Gliomas

Sarah Parisot[1,2,3], Hugues Duffau[4], Stéphane Chemouny[3], and Nikos Paragios[1,2,⋆]

[1] Laboratoire MAS, Ecole Centrale Paris, Chatenay Malabry, France
[2] Equipe GALEN, INRIA Saclay - Ile de France, Orsay, France
[3] Intrasense SAS, Montpellier, France
[4] Département de Neurochirurgie, Hopital Gui de Chauliac, CHU Montpellier, France

Abstract. Low-grade gliomas (WHO grade II) are diffusively infiltrative brain tumors arising from glial cells. Spatial classification that is usually based on cerebral lobes lacks accuracy and is far from being able to provide some pattern or statistical interpretation of their appearance. In this paper, we propose a novel approach to understand and infer position of low-grade gliomas using a graphical model. The problem is formulated as a graph topology optimization problem. Graph nodes correspond to extracted tumors and graph connections to the spatial and content dependencies among them. The task of spatial position mapping is then expressed as an unsupervised clustering problem, where cluster centers correspond to centers with position appearance prior, and cluster samples to nodes with strong statistical dependencies on their position with respect to the cluster center. Promising results using leave-one-out cross-validation outperform conventional dimensionality reduction methods and seem to coincide with conclusions drawn in physiological studies regarding the expected tumor spatial distributions and interactions.

1 Introduction

Low-grade gliomas (WHO grade II) are diffusively infiltrative brain tumors that are generally revealed by seizures in young patients with a normal social and professional life [1]. Although surgery of low-grade gliomas has been a controversial subject for decades, it is now considered as the best therapeutic option as reported in the recent European Guidelines [2]. However, surgery can cause functional deficit, especially if the tumors are located near or within a functional area. For that matter, the tumor's location in the brain is of great importance.

Inferring spatial dependencies regarding tumor appearance in the brain is a research direction that has gained little attention [3,4]. The most natural approach to such a problem is through statistical modeling with respect to the position of low-grade gliomas. This could be achieved using dimensionality reduction techniques like principal [5] or independent [6] component analysis. This

⋆ This work was supported by ANRT (grant 147/2010), Intrasense and the European Research Council Starting Grant Diocles (ERC-STG-259112).

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 508–515, 2011.

can be easily achieved at the voxel level, but a huge number of observations is needed and the linearity assumption regarding the correlation between the spatial positions of tumors is imposed. Non-linear methods embed the observation space into a low-dimensional space and then seek for correlations in this space. Examples refer to isomap [7] or laplacian graphs [8]. Within the considered clinical setting, neither the expected rank of the reduced space is known nor the number of samples is sufficient to approximate the manifold.

Network connectivity analysis is an alternative to the above mentioned methods, a quite popular idea in functional imaging when targeting task-specific brain understanding through mutual activations of brain regions [9]. In such a context, the brain is parceled according to a certain criteria and considered to be a fully connected network, each node is associated with a multi-dimensional variable explaining the activations as a function of time/task. The aim is then to group parcels with important statistical correlations in terms of behavior [10].

In this paper, we are inspired from these methods but we amend them to deal with a static setting. While the brain parcels are known in functional imaging, in our case these parcels do depend on the observations and correspond to manually annotated low-grade gliomas. Our method first registers all samples to the same reference pose while explicitly taking into account the tumor position during the registration. The registered data are considered to define a statistical measure of coherence between different tumors based on their spatial position as well as their geometric form. This measure is used to determine a graph where messages are exchanged between nodes. The strength of the message depends on the statistical measure of coherence between tumors, while the overall capacity of a node depends on the total volume of messages being passed to it. The problem of understanding spatial dependencies between tumors appearance is then casted as an unsupervised clustering problem[11] where both the number of clusters, the cluster centers and the nodes assignments are to be determined.

The reminder of this paper is organized as follows: in section 2 we detail the preprocessing step and introduce the spatial position representation network. The optimization of the network is discussed in section 3 while experimental validation and comparisons with linear methods are presented in section 4. Discussion concludes the paper.

2 Network Representation of Low-Grade Gliomas Spatial Dependencies

Our data set consisted of 95 3D MRI images of 95 different patients with low-grade gliomas: 90 FLAIR T2 weighted, and 5 T2 weighted. We had a majority of male patients, age between 21 and 65 and tumor size between 3.5 and 123 cm^3. The image size ranged from 256x256x24 to 512x512x33, and the pixel resolution from 0.4x0.4 to 0.9x0.9 mm^2 in the (x,y) plane and 5.3 to 6.4 mm in the z plane. Each image had been manually annotated by experts to indicate the position of the tumor. The atlas for registration consisted of a high resolution FLAIR image (size 256x256x24, resolution 0.9x0.9x5.45 mm^3) of a tumor free brain.

Let us consider without loss of generality n acquisitions $V_i()$ of brain volumes as well as the corresponding segmentations $S_i()$ binary maps indicating whether a pixel belongs or not to the tumor. Both, intensity volumes and binary maps are rigidly registered to the same pose. The first step towards spatial position mapping of low-grade gliomas consists in removing the brain local anatomical variability through deformable registration. This is achieved using the discrete optimization method presented in [12] that is amended to deal with the lack of visual correspondences in the brain tumor areas by not taking into account tumor voxels during registration. Using the optimal deformable registration parameters, let us consider now $S(\mathbf{x}) = S(T(\mathbf{x}))$ being the deformed segmentation map.

We then measure the proximity between two tumors by adopting as suggested in [13] the Mahalanobis distance. Let us consider the spatial coordinates of all pixels belonging to tumors S_i and S_j. Let $\bar{\mathbf{x}}_\mathbf{i}$ and $\bar{\mathbf{x}}_\mathbf{j}$ be the center of mass of S_i and S_j. The Mahalanobis distance between the tumors is computed as:

$$d_M(S_i, S_j) = \sqrt{(\bar{\mathbf{x}}_\mathbf{i} - \bar{\mathbf{x}}_\mathbf{j})^T \Sigma^{-1} (\bar{\mathbf{x}}_\mathbf{i} - \bar{\mathbf{x}}_\mathbf{j})} \qquad \Sigma = \frac{(n_i - 1)\Sigma_i + (n_j - 1)\Sigma_j}{n_i + n_j - 2} \quad (1)$$

where Σ_i and Σ_j are the covariance matrices of pixels coordinates of S_i and S_i while n_i and n_j are the number of pixels in tumors S_i and S_j respectively. The complete network corresponding to our data-set is shown in [Fig. (1)] where the nodes correspond to the tumors and the strength and color of the edges correspond to the proximity of the observed tumors.

3 Network Optimization

Let us consider the set of distances $d(S_i, S_j)$ between all pairs of low-grade gliomas. Let us assume that the observations can be expressed through a compact sub-graph of k central nodes. Both the number and the position of the central nodes are to be determined while the remaining observations should be expressed from one of these nodes. In order to determine the quality of a central node, we adopt two measurements, a global and a local one. The global measurement assumes that a central node should have significant overlap with the remaining measurements. This can be quantified by the overall distance of the node from the rest of the data. The assignment of an observation to one of the central nodes (local measurement) should be based on their distance. The optimal network connectivity should minimize the local and the global criteria both in terms of central node selection and tumor assignments. This is an unsupervised clustering problem,

$$\min_k \min_{c_1,\ldots,c_k} \min_{l_1,\ldots,l_n} \left(\sum_{i=1}^{n} d(c_i, S_j) + \alpha \sum_{j=1}^{k} \delta(c_i - l_j) d(S_{l_j}, S_j) \right) \quad (2)$$

where l_j is the assignment of observation j and α is a constant coefficient balancing the contributions of the two terms. To recover the lowest potential of the

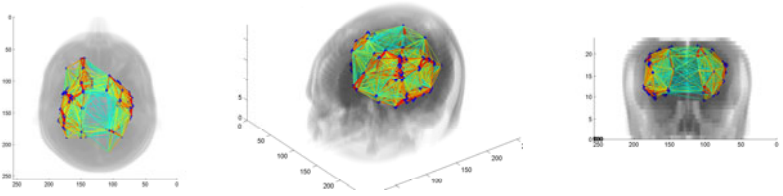

(a) (b) (c)

Fig. 1. Network representation of the data-set. Arcs corresponding to a distance greater than 5 are not displayed.

above functional, we will rely on a recently proposed clustering algorithm [11], which does not require any initialization and provides near-optimal clustering results both in terms of number of central nodes as well as in terms of remaining nodes assignments. The optimization results will heavily depend on the relative importance given between the two terms (*alpha* value). In order to determine the optimal number of central nodes, one should consider the quality of clusters being associated with them. We consider the following criteria to determine the compactness of clusters and the quality of the overall representation:

Dunn Index [14]: For a partition of K clusters, it is computed as

$$D = \min_{i\in[1:K]}\{\min_{j\in[1:K],j\neq i}\{\frac{d(c_i,c_j)}{d_{max}}\}\} \tag{3}$$

where d_{max} corresponds to the maximal distance of a sample to the center of the cluster it belongs to, and $d(c_i, c_j)$ is the distance between the centers c_i and c_j of clusters i and j. Intuitively, a good clustering will be characterized by a high Dunn index (see [Fig. (2b)]): compact and well separated clusters yield a low d_{max} and a high distance inter clusters.

Davies-Bouldin Index [15]: this index also identifies compact and well separated clusters. It computes the maximum similarity between clusters:

$$DB = \frac{1}{K}\sum_{i=1}^{K}\max_{j\in[1:K],j\neq i}\frac{\sigma_i+\sigma_j}{d(c_i,c_j)} \tag{4}$$

where σ_i is the average distance of all points in cluster C_i to its center, K is the number of clusters and $d(c_i, c_j)$ is the distance between the centers of clusters i and j. A small DB value indicates little similarities between clusters and therefore, a better clustering ([Fig. (2c)]).

Global Silhouette Index [16]: For a given cluster $C_k = (S_1, ...S_n)$, the silhouette index assigns to each of its members a quality measure $s(S_i)$ (i=1,...,n) that is a confidence indicator on the membership of S_i in C_k. It is defined as

$$GS = \frac{1}{K}\sum_{j=1}^{K}\frac{1}{n_j}\sum_{i=1}^{n_j}s(S_i) \text{ with } s(S_i) = \frac{b(S_i)-a(S_i))}{\max(a(S_i),b(S_i))} \tag{5}$$

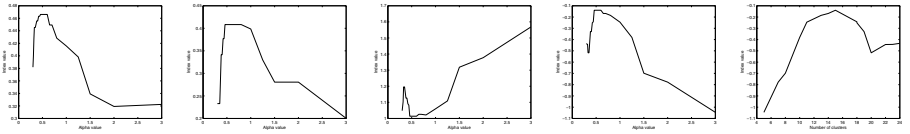

(a) (b) (c) (d) (e)

Fig. 2. Values of the 3 criteria in function of *alpha* (a,b,c,d) and of the number of clusters (e). (a) Global Silhouette index, (b) Dunn index, (c) Davies-Bouldin index, (d,e) Combination of the three indexes.

where n_j is the number of elements in cluster C_j, $a(S_i)$ is the average distance between S_i and all of the remaining elements in C_k, and $b(S_i)$ is the minimum average distance between S_i and all of the elements clustered in $C_j, (j = 1, ..., K; j \neq k)$. $s(S_i)$ takes values between -1 and 1. A value close to 1 indicates that S_i has been assigned to the appropriate cluster and a value close to -1 infers that S_i has been misclassified. A value close to zero suggests that S_i lies equally far away from 2 clusters. Since the largest silhouette index is indicative of the best clustering, we select the clustering that yields the maximum global silhouette index ([Fig. (2a)].

We have considered increasing relative importance between the global and the local term as shown in [Fig. (2)]. All the considered criteria reached their best value for the same 15 nodes graph-structure. The optimal network connectivity is shown in [Fig. (3)]. One can observe a notion of symmetry between the left and the right hemisphere. The distribution of the individual distance values before clustering is shown in [Fig. (4a)] while [Fig. ((4b)], and ((4c)] show the distribution per cluster of distances and individual silhouette indexes.

4 Experimental Validation

In order to evaluate the performance of the method we have considered a leave-one-out cross validation strategy. We have used $n-1$ measurements to learn the topology of the graph and the remaining observation to predict its position in the network. We have performed this test $n = 95$ times.

We denote 3 different results that satisfy the 3 optimality criteria. In most cases, the number of central nodes has reached the same value as for the whole data-set, that is $K = 15$. In one case, the optimal number of clusters was 16. The removed sample is a large tumor (about 110 cm^3) that is the center of an important cluster (12 nodes) which splits in 2 without its center. The removal of 7 samples yielded 14 clusters networks. Those samples were assigned to small clusters (2 or 3 nodes) that were merged with a neighboring cluster when the sample is removed.

We have considered 2 criteria to determine the ability of the retained centers to express their respective populations. First the cluster members should be

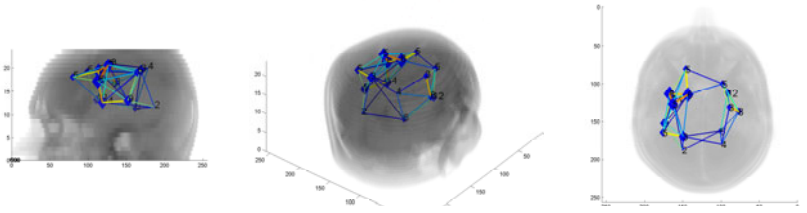

(a) (b) (c)

Fig. 3. 15 nodes network corresponding to the optimal criteria, connections corresponding to distances bigger than 4 are not displayed

overlapping with the center and their cluster's center c_k should be the closest, i.e. $d(S_i, c_k) = \min_{j\in[1:K]} d(S_i, c_j)$ and $d(S_i, c_k) \leq 2$. Secondly, we use the individual silhouette index $s(S_i)$ to evaluate the membership of a sample to a cluster. If the value is high, there is no doubt about the membership of the sample. If it is close to 0, we compare its value to the average silhouette index of the corresponding cluster : $s(S_i) \geq 0.3$ or $s(S_i) \geq \frac{1}{n_k} \sum_{j=1,j\neq i}^{n_k} s(S_j) - std(s_k)$, where $std(s_k)$ is the standard deviation of individual silhouette index values for cluster C_k, and n_k the number of elements in C_k. Such a configuration was able to properly assign 80% of the whole training example.

In order to evaluate the network prediction strength, we have considered the cluster that optimally represents the removed observation, and measured the quality of the cluster once this new sample has been added to it. Ideally, the quality of the clusters should remain the same if the network is able to express the variability of new samples. This was the case in 73% of the 95 cases. Failure can occur if the new sample is an outlier, equally close to 2 existing clusters or was the center of a cluster. We also estimate the effect of the removal of a sample on the clustered graph's structure by evaluating the quality of correspondences between the clustered graph obtained from the whole data-set G_0, and the 15 nodes graphs obtained from cross validation experiments $G_k, (k = 1, .., 87)$. To this end, we seek a matching between the nodes of the graphs by using the algorithm proposed by [17]. We find complete match in 83% of the cases and only one node didn't correspond for the remaining graphs.

Finally, comparison with principal component analysis (PCA) was considered. PCA computes a new orthogonal coordinates system that regroups the maximum variance in a minimum vectors. It is a simple way to find correlation between data: the fewer vectors necessary to represent the data, the bigger the correlation. We performed PCA on the data-set at the voxel level. Each [256,256,24] binary segmentation map was converted to a [256x256x24,1] vector, so that each voxel was a variable and each image was a sample. While our experiments using network representation suggest that 15 clusters represent 80% of the data, 15 PCA vectors regroup only 65% of it. Another drawback of PCA was that the obtained vectors did not correspond to specific preferential locations.

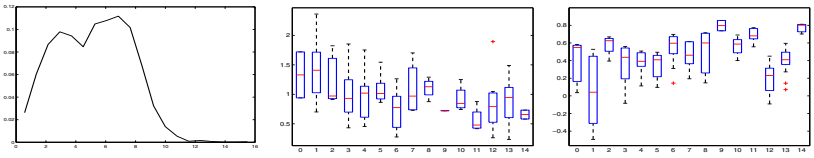

(a) (b) (c)

Fig. 4. Distance distribution between all the tumors (a),Box-plots of the distances (b) and individual Silhouette (c) values for each cluster

5 Discussion

Very few studies have dealt with spatial dependencies of gliomas appearances in the brain. In this paper, we proposed graph theory to represent our data-set spatial dependencies and clustering to regroup statistically dependent tumors. Cross-validation results on an important volume of clinical data support the idea that the complete graph could be reduced to a handful nodes, which indicates statistical preferential locations for low-grade gliomas in the brain that our clustered network enables to identify. There are very few tumors in or near the occipital and prefrontal lobes (7%). We find (symmetrically) the higher amount of tumors around the insula or the temporal lobe (27% in the right hemisphere and 33% in the left hemisphere). The remaining tumors are in the frontal and parietal lobes, mostly close to the motor areas. Those results are consistent with previous observations [4]. Furthermore, preliminary results indicate that the graph's structure remains unchanged while at the same time the method outperforms standard dimensionality reduction techniques.

Several open questions remained unanswered from this study. Despite the sizable validation set, increasing the number of cases considered in the study could further enhance the claims of the paper. On a more theoretical view-point, the impact of the registration process is critical since the selection of the reference pose introduce a strong bias on the results. The use of population registration methods [18] that simultaneously deform all the data while taking into account tumoral regions is a promising alternative. The estimation [19] and propagation of registration uncertainties is another mean of eliminating the bias while at the same time producing a qualitative interpretation of the results. This can also produce better means of measuring spatial coherence between tumors through high-dimensional embedding and Gaussian processes distance definition. Creating a statistical representation of tumor appearances can be of great interest and a valuable clinical tool for computer aided diagnosis. The network obtained through leave-one-out cross validation, through n-to-m graph matching and network topology optimization could lead to a unique representation for the spatial mapping of low-grade gliomas in the brain. Automatic segmentation is a task that could automatically benefit from the obtained graphical model since it

consists of a powerful position prior queue. Another possible outlook would be to integrate functional data and observe the consequences of a tumor on the functional brain and its plasticity.

References

1. Duffau, D.: Surgery of low-grade gliomas: towards a 'functional neurooncology'. Current Opinion in Oncology 21(6), 543–549 (2009)
2. Sofetti, R., Baumert, B., Bello, L., et al.: Guidelines on management of low-grade gliomas: report of an EFNS-EANO task force. Eur. J. Neurol. 17, 1124–1133 (2010)
3. Larjavaara, S., Mantyla, R., Salminen, T., et al.: Incidence of gliomas by anatomic location. Neuro-oncology 9, 319–325 (2007)
4. Duffau, H., Capelle, L.: Preferential brain locations of low-grade gliomas. Cancer 100, 2622–2626 (2004)
5. Shlens, J.: A tutorial on Principal Component Analysis. Systems Neurobiology Laboratory, Salk Institute for Biological Studies (2005)
6. Hyvarinen, A., Oja, E.: Independent component analysis: algorithms and applications. Neural Netw. 13, 411–430 (2000)
7. Tenenbaum, J.B., Silva, V., Langford, J.C.: A Global Geometric Framework for Nonlinear Dimensionality Reduction. Science 290, 2319–2323 (2000)
8. Belkin, M., Niyogi, P.: Laplacian eigenmaps and spectral techniques for embedding and clustering. In: NIPS (2001)
9. Zheng, X., Rajapakse, J.C.: Learning functional structure from fMR images. Neuroimage 31, 1601–1613 (2006)
10. Bullmore, E., Sporns, O.: Complex brain networks: graph theoretical analysis of structural and functional systems. Nat. Rev. Neurosci. 10, 186–198 (2009)
11. Komodakis, N., Paragios, N., Tziritas, G.: Clustering via LP-based Stabilities. In: NIPS, pp. 865–872 (2008)
12. Glocker, B., Komodakis, N., Tziritas, G., Navab, N., Paragios, N.: Dense image registration through MRFs and efficient linear programming. Medical Image Analysis 12, 731–741 (2008)
13. Pokrajac, D., Megalooikonomou, V., Lazarevic, A., Kontos, K., Obradovic, Z.: Applying spatial distribution analysis techniques to classification of 3D medical images. Artificial Intelligence in Medicine 33, 261–280 (2005)
14. Dunn, J.C.: Well separated clusters and optimal fuzzy partitions. J. Cybern. 4, 95–104 (1974)
15. Davies, D.L., Bouldin, D.W.: A cluster separation measure. IEEE Trans. on PAMI 1(4), 224–227 (1979)
16. Rousseeuw, P.J.: Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. J. Comput. Appl. Math. 20, 53–65 (1987)
17. Torresani, L., Kolmogorov, V., Rother, C.: Feature correspondence via graph matching: Models and global optimization. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) ECCV 2008, Part II. LNCS, vol. 5303, pp. 596–609. Springer, Heidelberg (2008)
18. Sotiras, A., Komodakis, N., Glocker, B., Deux, J.-F., Paragios, N.: Graphical models and deformable diffeomorphic population registration using global and local metrics. In: Yang, G.-Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (eds.) MICCAI 2009. LNCS, vol. 5761, pp. 672–679. Springer, Heidelberg (2009)
19. Glocker, B., Paragios, N., Komodakis, N., Tziritas, G., Navab, N.: Optical Flow Estimation with Uncertainties through Dynamic MRFs. In: CVPR (2008)

Automated Registration of Whole-Body Follow-Up MicroCT Data of Mice

Martin Baiker[1], Marius Staring[1], Clemens W.G.M. Löwik[2], Johan H.C. Reiber[1], and Boudewijn P.F. Lelieveldt[1,3]

[1] Div. of Image Processing, Leiden University Medical Center, The Netherlands
b.p.f.lelieveldt@lumc.nl
http://www.lkeb.nl
[2] Dept. of Endocrinology, Leiden University Medical Center, The Netherlands
[3] ICT group, Dept. of Mediamatics, Delft University of Technology, The Netherlands

Abstract. In vivo MicroCT imaging of disease models at multiple time points is of great importance for preclinical oncological research, to monitor disease progression. However, the great postural variability between animals in the imaging device complicates data comparison.

In this paper we propose a method for automated registration of whole-body MicroCT follow-up datasets of mice. First, we register the skeleton, the lungs and the skin of an articulated animal atlas (Segars et al. 2004) to MicroCT datasets, yielding point correspondence of these structures over all time points. This correspondence is then used to regularize an intensity-based B-spline registration. This two step approach combines the robustness of model-based registration with the high accuracy of intensity-based registration.

We demonstrate our approach using challenging whole-body in vivo follow-up MicroCT data and obtain subvoxel accuracy for the skeleton and the skin, based on the Euclidean surface distance. The method is computationally efficient and enables high resolution whole-body registration in ≈17 minutes with unoptimized code, mostly executed single-threaded.

1 Background

The possibility to scan the entire body of small animals with dedicated hardware in vivo offers great benefits for preclinical research, because it allows to follow e.g. pathology development over time within the same subject. This excludes intersubject variability and has ethical and economical benefits.

A problem that arises with imaging entire bodies is the potentially large postural variability of animals that are imaged at different time points (Fig. 1). This significantly complicates data examination, because researchers have to 'align' structures of interest visually and navigate through large whole-body datasets. For some applications, dedicated animal holders can be used to reduce the postural variability. However, such holders may influence the study, e.g. by obstructing light in optical imaging based studies.

To deal with the problem of high postural variability, in [1] we presented a robust method for registration between the skeleton, the lungs and the skin of

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 516–523, 2011.

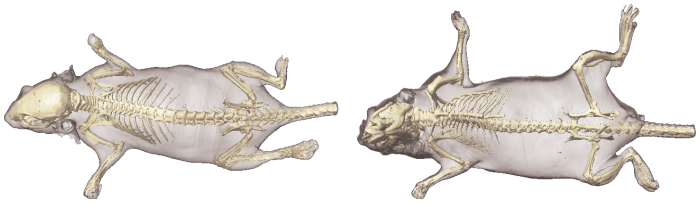

Fig. 1. Demonstration of the large variability in animal posture between different scans. Shown are subjects scanned in prone (left) and supine (right) position.

a mouse atlas (MOBY [2]) and whole-body MicroCT data of mice. We subsequently used the point correspondences on these structures to map the remainder of the body using Thin Plate Spline (TPS) interpolation. However, in areas with few correspondences, the accuracy of the mapping may be limited.

In this paper we aim at improving the accuracy of the TPS mapping by integrating intensity information during the registration. We present an accurate, time efficient and highly robust method for registration of follow-up MicroCT datasets that contain articulated objects. This we achieve by regularizing an intensity-based registration criterion with the Euclidean distance metric, based on pointsets of anatomical correspondences. We evaluate the method using non-contrast-enhanced MicroCT data of eight animals, imaged at two time points.

2 Previous Work

Several strategies are described in the literature that focus on registration of images with multiple structures of interest with varying structural properties. Staring et al. [3] describe an approach that adds a local rigidity penalty term to the registration function to penalize rigid object deformations. Somayajula et al. [4] present an intensity-based registration of whole-body MicroCT follow-up datasets of mice using a scale-space approach. A method that relies on skeleton segmentations from MicroCT is described in Li et al. [5]. The skeletons are aligned using nonrigid robust point matching, followed by intensity-based nonrigid registration based on radial basis functions. Suh et al. [6] register the skeleton using extended demons with subsequent intensity-based registration using normal demons. These approaches exploit the high CT contrast to avoid unrealistic bone deformation without [4] and with [3,5,6] using the skeleton explicitly. All methods may suffer from local minima when bones are in close proximity, but especially in case of large postural variability.

A possibility to increase the robustness of whole-body registration is to model and register individual parts of an animal. Approaches range from registration of individual Volumes Of Interest and subsequent interpolation (block-matching), that do not take relationships between VOIs into account, to methods that register structures of interest simultaneously or hierarchically. Articulated registration methods are based on realistic modeling of joints and were applied for example to mouse hind limbs [7].

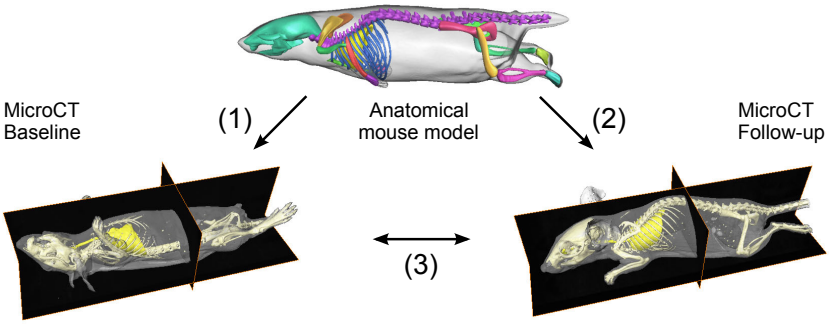

Fig. 2. First, an anatomical animal atlas (skeleton, lungs, skin) is registered to a baseline (fixed image, 1) and one or multiple follow-up (moving image, 2) MicroCT datasets. The point correspondence between the atlas and the datasets allows to establish point correspondences between the datasets as well, which can subsequently be used to regularize intensity-based registrations (3).

3 Method: Whole-Body Mouse Registration

In the following, we shortly describe an atlas-based framework for articulated registration presented in earlier work [1] and then the proposed extension for intensity-based registration. An overview of the framework is shown in Fig. 2. The fixed and moving images are denoted with I_F and I_M respectively, and the transformation relating the two by $\boldsymbol{T}_{\boldsymbol{\mu}}$, with parameters $\boldsymbol{\mu}$.

3.1 Articulated Whole-Body Registration

The mouse atlas used in this work is the publicly available MOBY atlas [2] that we modified by manually segmenting individual bones and organs, identifying joint locations and adding anatomically realistic joint models. The registration of this atlas to MicroCT was presented in previous work [1] and will be described briefly. Using a hierarchical anatomical model of the skeleton, each atlas bone is registered individually to an *unlabeled* skeleton surface representation, using the Iterative Closest Point (ICP) algorithm [8]. In each step, the Degrees of Freedom (DoFs) of the transformation function are defined by the joint type, by which the current bone is connected to the bone that is higher in the hierarchy. To account for differences in bone size, anisotropic scaling is added to the motion parameters of each bone. Thus, the DoFs vary between seven for a hinge joint (translation, non-isotropic scaling, one rotation) and nine for a ball joint. The surfaces of the lungs and the skin are subsequently registered, initialized by the skeleton registration result. The final result is a dense set of corresponding points on the skin, the skeleton and the lungs. Establishing such a point correspondence between the atlas and a target for data of several timepoints, allows to subsequently establish point correspondence between the timepoints as well (see Fig. 2). Corresponding pointsets of two different timepoints are in the following denoted as $\mathcal{Z}_F$ and $\mathcal{Z}_M$.

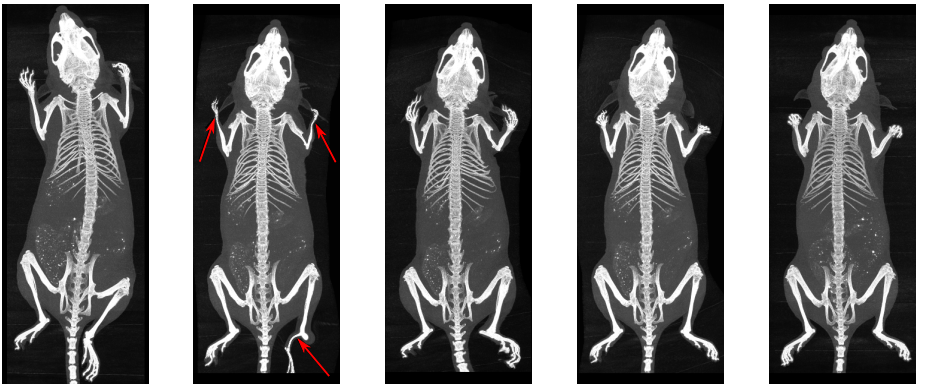

(a) Moving I_M (b) IFFD (c) TPS (d) IFFD Reg (e) Fixed I_F

Fig. 3. Dorsal-ventral maximum intensity projections of MicroCT volumes before and after registration with different methods. Note that (a) was acquired in prone, whereas (e) was acquired in supine position. The arrows indicate erroneous limbs after registration based on intensity information only. This is the animal shown in Fig. 5

3.2 Regularized Intensity-Based Registration

The articulated skeleton registration is surface-based and mostly neglects intensity information in the data. To combine the robustness of the articulated registration with the accuracy of intensity-based methods, we propose to regularize an intensity-based registration with the point correspondence from the articulated registration. Registration is formulated as an optimization problem:

$$\arg\min_{\boldsymbol{\mu}} \mathcal{C} = \arg\min_{\boldsymbol{\mu}} \mathcal{S}_{\text{sim}}(\boldsymbol{T}_{\boldsymbol{\mu}}; I_F, I_M) + \alpha \mathcal{S}_{\text{CP}}(\boldsymbol{T}_{\boldsymbol{\mu}}; \mathcal{Z}_F, \mathcal{Z}_M), \tag{1}$$

where the cost function $\mathcal{C}$ is optimized with respect to the transformation parameters $\boldsymbol{\mu}$. $\mathcal{S}_{\text{sim}}$ measures the image intensity similarity. We chose Normalized Cross Correlation (NCC), because all datasets are acquired with the same modality. We thus assume a linear relationship between the intensity values of I_F and I_M. $\mathcal{S}_{\text{CP}}$ is a metric incorporating the similarity of the corresponding pointsets $\mathcal{Z}_F$ and $\mathcal{Z}_M$ and is defined as the mean Euclidean distance between them:

$$\mathcal{S}_{\text{CP}} = \frac{1}{P} \sum_{\boldsymbol{x}_F^i \in \mathcal{Z}_F} \left\| \boldsymbol{x}_M^i - \boldsymbol{T}_{\boldsymbol{\mu}}(\boldsymbol{x}_F^i) \right\|, \tag{2}$$

where P is the number of corresponding points, and $\boldsymbol{x}_F^i, \boldsymbol{x}_M^i$ corresponding points from the fixed and moving image pointsets, respectively. The two terms of Eq. 1 are weighted by the parameter α. The optimization problem is solved using a parameter-free Adaptive Stochastic Gradient Descent (ASGD) optimization routine [9], in a multiresolution fashion, using Gaussian pyramids. For each resolution, the optimal value of α is set manually, depending on how much the

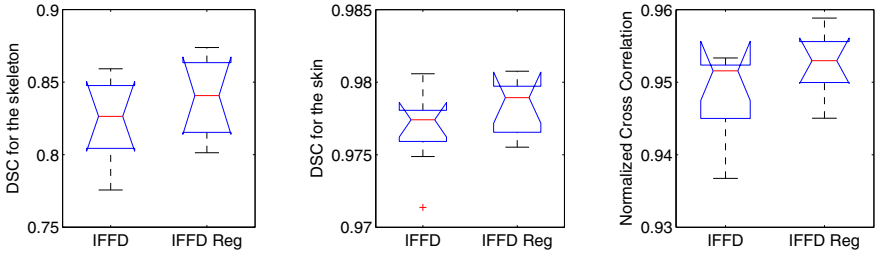

Fig. 4. Boxplots of the DSC for the skeleton and the skin and the NCC, for IFFD and IFFD Reg. Notch overlap indicates no significant difference ($p \geq 0.05$) between medians. Note that after initialization, the medians are: DSC skeleton 0.15, DSC skin 0.81, NCC 0.65 and using TPS interpolation 0.42, 0.91 and 0.81.

image intensity and the point distance measure should contribute to $\mathcal{C}$. In the first resolutions, $\mathcal{S}_{\text{CP}}$ should have a relatively large impact on $\mathcal{C}$, to remove large postural differences. Thus, α is set to a relatively large value because otherwise the optimization may get stuck in local minima. Assuming that afterwards I_F and I_M are coarsely aligned, the influence of $\mathcal{S}_{\text{CP}}$ can be gradually decreased and removed from $\mathcal{C}$ in the last resolution ($\alpha = 0$).

The intensity-based registration was initialized by a similarity registration (motion and isotropic scaling), followed by nonrigid registration with the transformation $\boldsymbol{T}_{\boldsymbol{\mu}}$ parameterized by B-splines [10]. They were employed in a multigrid setting, gradually refining the B-spline control point grid over the resolutions.

4 Experimental Setup

Eight female mice (Balb/c nu/nu, Charles River, L'Arbresle, France), 6 weeks old at baseline, were scanned twice, three weeks apart, once in prone and once in supine position and with arbitrary limb position. MicroCT (SkyScan 1076, Kontich, Belgium) parameters were: 1.5° steps, 180°, 50keV x-ray voltage, 200μA anode current, Al filter 0.5mm and exposure time 100ms. The datasets were reconstructed with built-in software (beam-hardening and ring artifact correction both 10) and a dynamic range of -1000 to 4000 Hounsfield units. No cardiac nor respiratory gating was used. The data was subsampled to 144^3 μm^3 voxelsize ($\approx 250 \times 200 \times 650$ voxels), smoothed with a Gaussian filter ($\sigma = 1$) and segmented using the Color Structure Code technique [13] with $T = 24$ for the skeleton and the skin and $T = 6$ for the lungs. Triangular surface meshes were extracted from the segmentations using Marching Cubes (more details in [1]).

Following the procedure in Section 3.1, we derived ≈2000 correspondences on the skeleton, the lungs and the skin. For the intensity-based registration, we used 5 resolutions (500 iterations) for the similarity registration and 6 resolutions (2000 iterations) for the B-Spline registration. α was kept constant at 0.05 in resolutions 1-4, decreased to 0.005 and 0 in resolutions 5 and 6 respectively (parameter files available at `http://elastix.isi.uu.nl/wiki.php`). Invertibility

Table 1. Skeleton and skin surface distance and landmark localization accuracy (in voxels). Surface distances are based on eight animals and the landmark distances on a subset of three animals. (*) Results are based on a different, yet comparable dataset.

Skeleton distance	**Mean**	**Median**	**Max**	**Min**
Init	9.70 ± 11.68	5.59	81.16	3e-6
TPS	2.01 ± 2.72	1.32	36.91	3e-6
IFFD	1.19 ± 5.15	0.34	71.99	5e-7
IFFD Reg	0.49 ± 0.80	0.33	17.83	3e-7
Li et. al [5] (*)	0.61 ± 0.19	N/A	N/A	N/A
Skin distance	**Mean**	**Median**	**Max**	**Min**
Init	9.56 ± 10.30	6.46	76.62	9e-6
TPS	3.79 ± 3.63	2.71	36.70	1e-6
IFFD	1.37 ± 4.58	0.50	68.64	4e-7
IFFD Reg	0.83 ± 1.16	0.49	16.41	9e-8
Landmark distance	**Mean**	**Median**	**Max**	**Min**
Init	65.24 ± 32.81	64.52	131.62	4.91
TPS	6.25 ± 3.75	5.52	25.63	2.17
IFFD	3.75 ± 7.46	1.90	51.87	0.37
IFFD Reg	1.97 ± 1.72	1.57	11.51	0.37
Li et. al [11] (*)	3.46 ± 1.88	3.64	5.96	1.04

and smoothness of all final transformations was confirmed using the determinant of the Jacobian of the deformation fields, which was > 0 within all animals.

For evaluation, the following metrics were chosen: Normalized Cross Correlation (NCC) to assess the intensity similarity and the Dice Similarity Coefficient (DSC) to assess skeleton and skin segmentation accuracy. The DSC is defined as $2(V_1 \cap V_2)/(V_1 + V_2)$ and measures structural overlap. It is well suited for elongated and thin structures, which occur in our data (Fig. 1). We also determined the Euclidean Point to Surface Distance (EPSD) between the skeletons and skins of registered datasets. We excluded the tail, since it is irrelevant for most studies. Color-coded EPSD mapping to the surfaces allows to detect local registration inaccuracies. Finally, we assessed how well specific bone structures are registered, by measuring the Euclidean Point to Point Distance (EPPD) between 19 anatomical landmarks, manually indicated before and after registration, on distal body parts like the limbs, on the spine and on the ribs. Results are given after initialization, TPS interpolation, intensity-based registration without (IFFD) and with using regularization (IFFD Reg). For comparison with published work, we present results of Li et al. [5], because their datasets are comparable to ours.

Correspondence determination was done with Matlab 2010b (The Mathworks, Natick, USA) and the intensity-based registration using the ITK-based and publicly available `elastix` software [12] on an Intel Xeon E5620 8 cores (2.4GHz) and 24GB RAM. The time requirements were ≈5 mins. for IFFD and ≈17 mins. for IFFD Reg. (including ≈5 mins. to determine correspondence).

5 Results and Discussion

Qualitative results of the registration are shown in Fig. 3, quantitative results for the DSC and the NCC are presented in Fig. 4 and the surface distances and

landmark localization accuracy before and after registration are given in Tab. 1 and Fig. 5. The very large difference between the metrics after initialization and after IFFG are an indication for the large postural differences between the animals. Comparing TPS and IFFD, the average error is smaller for IFFD, but the maximum is much larger. The reason is the large initial postural differences between animals. TPS can deal with that and therefore, all body parts are registered equally well. IFFD is very accurate, when body parts lie within the registration capture range, but fails completely otherwise. Generally, the more distal to the body, the higher the error becomes. Fig. 5 and Fig. 3 support this because the error increases significantly at the limbs. The results of IFFD Reg reveal that our approach can handle large variability in the data without losing accuracy. The DSC plot (Fig. 4) shows excellent overlap for the skeleton and the skin. We obtain subvoxel accuracy for bone and the skin in the surface distance measure (Tab. 1). The maximum distances mainly stem from the very distal ends of the limbs and the ribs for the skeleton, and folds for the skin (Fig. 5). In addition, IFFD Reg yields higher intensity similarity than IFFD (Fig. 4). For all presented metrics, IFFD Reg outperforms both, TPS and IFFD, proving that relying on point correspondence or intensity only is not sufficient for highly accurate registration, in case of large postural differences.

Compared to published data by Li et al. [5,11], we have similar results for the skeleton distance and better results for the landmark localization. Their method pays special attention to registration of the ribs, thus it might yield more accurate results for these structures. However, they evaluate using ex vivo data, excluding rib movement artifacts. If accurate rib registration is required, an additional stiffness penalty could easily be added to our registration criterion [3]. In addition, we want to stress that the method in Li et al. requires 260 minutes for registration and our method takes ≈17 minutes. We realize that those experiments were performed on outdated hardware (Pentium PC, 2GHz, 1GB RAM), but most of our code was executed single-threaded and in addition, our image domain was approximately twice as big. It would be interesting to compare our method to the promising approach of Suh et al. [6] as well, which seems to be more time efficient and more accurate, compared to Li et al.

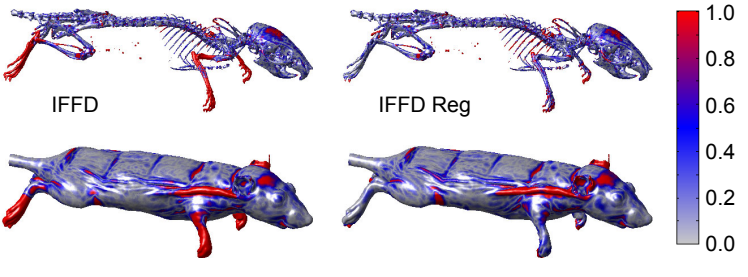

Fig. 5. The skeleton and the skin of an animal at baseline with color-coded Euclidean distance to the nearest surface point on the mapped skeleton and skin after registration using IFFD and IFFD Reg respectively. Values (in voxels) are based on one animal.

Finally we want to point out, that the registration of an atlas yields a segmentation of the skeleton as a by-product.

6 Conclusion

We presented a highly robust and accurate approach for registration of articulated objects with application to whole-body MicroCT data of mice. This we obtained by regularizing an intensity-based registration criterion with a distance metric, derived from point correspondence among datasets. We performed registration of in vivo whole-body MicroCT data with high resolution in ≈17 minutes and obtained subvoxel accuracy for the skeleton and the skin. Compared to competing methods, our approach is very time efficient.

References

1. Baiker, M., Milles, J., Dijkstra, J., Henning, T.D., Weber, A.W., Que, I., Kaijzel, E.L., Löwik, C.W.G.M., Reiber, J.H.C., Lelieveldt, B.P.F.: Atlas-based whole-body segmentation of mice from low-contrast Micro-CT data. Med. Image Anal. 14(6), 723–737 (2010)
2. Segars, W.P., Tsui, B.M.W., Frey, E.C., Johnson, G.A., Berr, S.S.: Development of a 4D digital mouse phantom for molecular imaging research. Mol. Im. Biol. 6(3), 149–159 (2004)
3. Staring, M., Klein, S., Pluim, J.P.W.: A Rigidity Penalty Term for Nonrigid Registration. Med. Ph. 34(11), 4098–4108 (2007)
4. Somayajula, S., Joshi, A.A., Leahy, R.M.: Mutual Information Based Non-rigid Mouse Registration Using A Scale-Space Approach. In: Proc. IEEE Int. Symp. on Biomedical Imaging, pp. 1147–1150. IEEE Press, New York (2008)
5. Li, X., Yankeelov, T.E., Peterson, T.E., Gore, J.C., Dawant, B.M.: Automatic nonrigid registration of wholebody CT mice images. Med. Ph. 35(4), 1507–1520 (2008)
6. Suh, J.W., Scheinost, D., Dione, D.P., Dobrucki, L.W., Sinusas, A.J., Papademetris, X.: A non-rigid registration method for serial lower extremity hybrid SPECT/CT imaging. Med. Image Anal. 15(1), 96–111 (2011)
7. Papademetris, X., Dione, D.P., Dobrucki, L.W., Staib, L.H., Sinusas, A.J.: Articulated rigid registration for serial lower-limb mouse imaging. In: Duncan, J.S., Gerig, G. (eds.) MICCAI 2005. LNCS, vol. 3750, pp. 919–926. Springer, Heidelberg (2005)
8. Besl, P.J., McKay, N.D.: A method for registration of 3D shapes. IEEE Trans. Pattern Anal. Machine Intelligence 14(2), 239–256 (1992)
9. Klein, S., Pluim, J.P.W., Staring, M., Viergever, M.A.: Adaptive stochastic gradient descent optimisation for image registration. J. Com. Vis. 81(3), 227–239 (2009)
10. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L.G., Leach, M.O., Hawkes, D.J.: Nonrigid registration using free-form deformations: application to breast MR images. IEEE TMI 18(8), 712–721 (1999)
11. Li, X., Yankeelov, T.E., Peterson, T.E., Gore, J.C., Dawant, B.M.: Constrained non-rigid registration for whole body image registration: method and validation. In: Cleary, K.R., Miga, M.I. (eds.) SPIE Med. Im., vol. 6512, pp. 11–18 (2007)
12. Klein, S., Staring, M., Murphy, K., Viergever, M.A., Pluim, J.: Elastix: a toolbox for intensity-based medical image registration. IEEE TMI 29(1), 196–205 (2010)
13. Sturm, P., Priese, L., Wang, H.: A CSC Based Classification Method for CT Bone Images. In: Proc. 3DPVT, pp. 1080–1084. IEEE Press, New York (2006)

Evaluating Volumetric Brain Registration Performance Using Structural Connectivity Information

Aleksandar Petrović[1,2] and Lilla Zöllei[3]

[1] University of Oxford, FMRIB Centre, John Radcliffe Hospital, Headington, Oxford, OX3 9DU, UK
[2] Siemens AG, Healthcare Sector, 91052 Erlangen, Germany
[3] Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Boston, MA 02176, USA
petrovic@fmrib.ox.ac.uk, lzollei@nmr.mgh.harvard.edu

Abstract. In this paper, we propose a pipeline for evaluating the performance of brain image registration methods. Our aim is to compare how well the algorithms align subtle functional/anatomical boundaries that are not easily detectable in T1- or T2-weighted magnetic resonance images (MRI). In order to achieve this, we use structural connectivity information derived from diffusion-weighted MRI data. We demonstrate the approach by looking into how two competing registration algorithms perform at aligning fine-grained parcellations of subcortical structures. The results show that the proposed evaluation framework can offer new insights into the performance of registration algorithms in brain regions with highly varied structural connectivity profiles.

Keywords: Volumetric registration, MRI, probabilistic tractography, structural connectivity, human thalamus.

1 Introduction

There have been numerous studies comparing the performance of volumetric registration algorithms by measuring the correspondence between manually outlined labels of the reference and the warped images (e.g. see [1]). These types of evaluations shed light on which algorithm performs better and in which areas. However, they often rely on information that was also used to drive the registration algorithms (e.g. intensity contrast in T1-weighted images). In order to avoid such circularity, ideally, an independent testing mechanism should be used. In other words, one should avoid the circularity of the information flow as well as 'using registration to assess registration'. This paper tackles the former issues and tries to minimize the influence of the latter.

Our paper proposes a new framework for evaluating performance differences of volumetric registration algorithms. It relies on a modality different from the one used for registration (in our case, T1-weighted images). Diffusion Weighted Imaging (DWI) is a suitable choice for this purpose as it contains information complementary to that found in other structural modalities. In this paper we assess the inter-subject alignment of anatomical boundaries within subcortical structures derived through the

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 524–531, 2011.

analysis of white matter connectivity. Note, the aim of this paper is to present an evaluation methodology and not to provide a comprehensive comparison of a range of registration methods.

DWI has been used to map the architecture of cortical white matter (WM) and also to determine WM consistency [2,3,4]. Furthermore, WM connectivity patterns highly correlate with the functional segregation of the brain [5]. Thus, evaluating registration performance in the areas of distinct connectivity patterns might yield additional insights into the inter-subject alignment of functional homologues.

We were inspired by the work of [6] which shows that thalamic connectivity analysis can result in anatomically and functionally plausible and highly consistent parcellations of the anatomy, revealing structures (i.e. different thalamic nuclei) that cannot be seen on T1- or T2-weighted images alone. Therefore, if we were to: 1) parcellate all the thalami in a group of subjects, e.g. as proposed in [7], and 2) consistently label anatomically corresponding clusters across subjects, it would be possible to measure inter-subject cluster overlap. Importantly, the registration algorithms assessed by this measure do not use information used for parcellation. The thalamus and the putamen were used as examples to demonstrate our findings.

2 Methods

In order to demonstrate our hypothesis, this paper evaluates the differences between the Combined Volumetric and Surface-based registration (CVS) algorithm [8, 9] and FNIRT (FMRIB's non-linear image registration tool) [10] by examining the post-registration overlap of connectivity-based segmentation labels across a group of control subjects. CVS and FNIRT belong to two complementary and widely used software packages and it is of great interest to the neuroimaging community to know which registration method of these two is superior in what case.

First, for each subject, image distortion and bias field corrections were performed prior to the DW-to-T1 affine registration (FSL/FLIRT [13]). In that way, pairs of DW and T1 images were brought to the same coordinate system prior to further analysis. Secondly, in order to perform parcellation of subcortical structures using structural connectivity, every voxel in the region of interest (ROI, e.g. the thalamus) has to be associated with a feature vector encoding measures of connectivity to other brain areas. These feature vectors become inputs into the clustering algorithm. Commonly, each entry in the feature vector is the probabilistic tractography result (a measure of uncertainly of the WM tract originating in the seed voxel and passing through the target) to one of the neocortical targets, e.g. a major brain lobe [6]. In order to boost the sensitivity in discriminating between areas of varying connectivity, we decided to identify a greater number of cortical targets using the spherical registration and parcellation algorithm of FreeSurfer [12]. Targets for probabilistic tractography (PROBTRACKX/FSL [13]) were formed in the following steps:

1) Using Freesurfer, parcellate surfaces of both cortical hemispheres into 80 equilateral triangles (calculated as the first subdivision of an icosahedron) that are consistently labeled across all subjects. Triangular features were chosen as they do not follow any particular anatomical divisions and do not therefore bias the results;

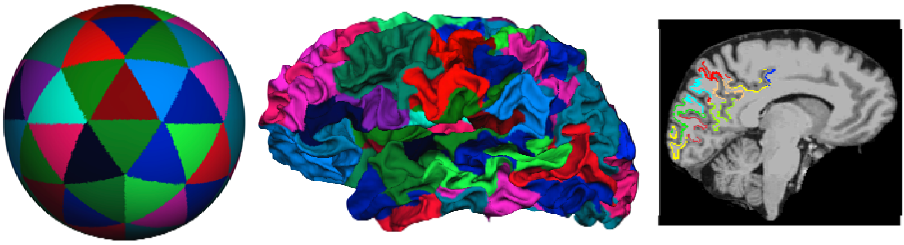

Fig. 1. Triangular targets were drawn on an inflated and registered sphere (left) thus imposing anatomical consistency in triangle positions across subjects. Figures (middle and right) show projections of triangles to the white matter surface – different colours correspond to different triangle labels, i.e. different cortical targets.

2) Propagate surface parcellations into the volume. A 1 mm thick sheath on the grey/white matter surface was constructed as depicted in 1) (see Fig. 1);
3) Concatenate volumetric masks/triangular labels from both hemispheres to form a list of cortical targets (targets for probabilistic tractography).

Therefore, the final output of probabilistic tractography for every ROI is an MxN connectivity matrix where M is the number of voxels in the ROI and N is the number of targets, which in our current implementation is set to 160 (80 targets for each hemisphere). Typically, connectivity matrices are very sparse having 90-95% of all values zero. Clustering of the thalamic and putamen connectivity matrices is done for all 41 subjects using Hierarchical Dirichlet Process Mixture Models (HDPM) with spatial constraints [7]. Both structures were segmented by FreeSurfer [11]. Each of the connectivity matrices was preprocessed and grouped into 9 clusters. Preprocessing consisted of PCA data dimensionality reduction to size Mx20, typically retaining more than 95% of explained variance. We decided not to use infinite mixture models in our analysis as they regularly produced more than 40 clusters, many of which contained just a few voxels. By varying the number of clusters for a small subset of subjects, we found that 9 clusters is a good compromise between capturing anatomically plausible divisions and achieving reasonable inter-subject consistency. Clustering consistency was also found to be robust with respect to the number of cortical targets. The parameter regulating intra-subject spatial smoothness (called 'beta', as implemented in [7]) was set to 1, which is the default value suggested in [7]. We should note that HDPM-based clustering does not make use of inter-subject registrations and the correspondence among cluster labels is established solely according to the data matrix fed into the HDPM pipeline.

After the clustering step, all subjects were non-linearly registered to a template (which was part of the hierarchical clustering cohort, a randomly chosen subject) using two methods: CVS and FNIRT (see Sections 2.1 and 2.2). We decided not to use a standard space (such as the MNI152) for our analysis, as we test for the alignment of very subtle and subject-specific parcellations and are thus in need for a template image with similar detail in it. We believe that if the T1-weighted intensities within our ROIs are insufficient to discriminate between internal subregions, then the

alignment of other cortical areas would constrain the deformations, hopefully maximizing the alignment of functionally distinct areas.

Finally all clustering results were warped using the two registration methods. We then computed the Jaccard overlap measure to test the inter-subject alignment accuracy among the HDPM clustering of the connectivity descriptors in the ROIs.

Fig. 2. Thalamus and Putamen clustered using proposed methodology (HDPM clustering of connectivity matrices). Different colours/numbers represent different clusters (9 in total) of distinct WM connectivity patterns, e.g. thalamic nuclei (image left).

2.1 CVS

Combined Volumetric and Surface-based Registration is a brain image registration method that maximizes the alignment of both cortical and subcortical structures. It consists of three image processing steps. First, a surface registration algorithm finds correspondences between the input surfaces from two brain scans [14] and these correspondences are transformed into a sparse displacement field in Euclidean space. This morph is then diffused into a dense displacement field in the volume using a nonlinear elastic model. Finally, a nonlinear volumetric registration refines the alignment, bringing subcortical structures, which are not near the surfaces, into accurate alignment. This technique has been shown to produce state-of-the-art alignment of cortical folding patterns, architectonics and subcortical structures [8].

2.2 FNIRT

The nonlinear FNIRT registration tool [10] uses a B-splines representation of the registration warp field and optimizes the sum of squared differences as its objective function. It is not typically run in a subject-to-subject setting as it was initially optimized for subject-to-template registrations. For our project, together with the original developers of the code, we established a particular set of parameters that achieve high quality subject-to-subject correspondence.

2.3 Data Description

The experiments were run on data provided through collaboration with Dr. Randy Gollub and the Medical Investigation of Neurodevelopmental Disorders (MIND) Institute. Forty-one data sets were selected all of which have been acquired by our

collaborators using an identical MRI sequence on a Siemens scanner. The structural data is of 256x256x256 size with 1 mm^3 voxel resolution and TR =12ms, TE=4.76ms, TI=4.76ms, flip angle=20°. The diffusion scans use single shot EPI, and a twice-refocused SE pulse sequence, optimized to minimize eddy current-induced distortions (TR/TE=7400/89 ms, b=700 s/mm^2, 256x256 mm FOV, 128x128 matrix, 2 mm (0 mm gap) slice thickness, 10 T2 + 60 DWI. Sixty-four slices were acquired in the AC-PC plane. The 60 diffusion-encoding gradient directions were determined using the electrostatic shell method, and result in a high signal-to-noise diffusion volume.

3 Results

Comparison between the FNIRT and CVS registration methods was performed for thalamic and putamen clusterings. Table 1 summarizes the results of compound (calculated across all labels in parallel) Jaccard coefficient measures. Using this measure every subject's clustering was compared to the template clustering. In summary, FNIRT performs consistently better than CVS in the left thalamus.

Table 1. Mean Jaccard coefficient between the clustering of the template and all the other thalami/putamen (40 comparisons per table entry – see Fig. 3). Statistically significant difference at $p<0.01$ level is indicated in bold (unpaired T-test, DOF correction).

THALAMUS	FNIRT	CVS	PUTAMEN	FNIRT	CVS
Left	**0.19**	0.14	Left	0.20	0.18
Right	0.16	0.13	Right	0.21	0.20

We also calculated how well both methods perform when each ROI is considered a single label, without any clustering. In both the left and right hemispheres, FNIRT performed better ($p<0.001$) over CVS. If this effect is regressed out of the findings from Table 1, previous conclusions do not change (FNIRT performs better in the left thalamus with $p<0.05$). Regressing out an effect, in this case, means calculating the compound Jaccard coefficient just for the intersection of the two ROIs.

We also calculated the pairwise compound Jaccard coefficient (compound Jaccard coefficient calculated between each pair of subjects) for the whole dataset excluding the template. These results are shown in Table 2. Similarly to the previous analysis, we computed this metric with the ROIs treated as a single label. In both the left and right hemispheres, FNIRT performed better ($p<10^{-5}$). When the effect of the global alignment is regressed out of the findings from Table 2, FNIRT still outperformed CVS for both the left and right putamen and for the left thalamus ($p<0.01$).

Table 2. Mean pairwise compound Jaccard coefficient (see [8] for computational details) among the clusterings of all ROIs excluding the template (40x40 comparisons per table entry – see Fig. 3). Statistically significant difference at $p<0.01$ level is given in bold (unpaired T-test, DOF correction).

THALAMUS	FNIRT	CVS	PUTAMEN	FNIRT	CVS
Left	**0.28**	0.23	Left	**0.26**	0.22
Right	**0.27**	0.20	Right	**0.18**	0.15

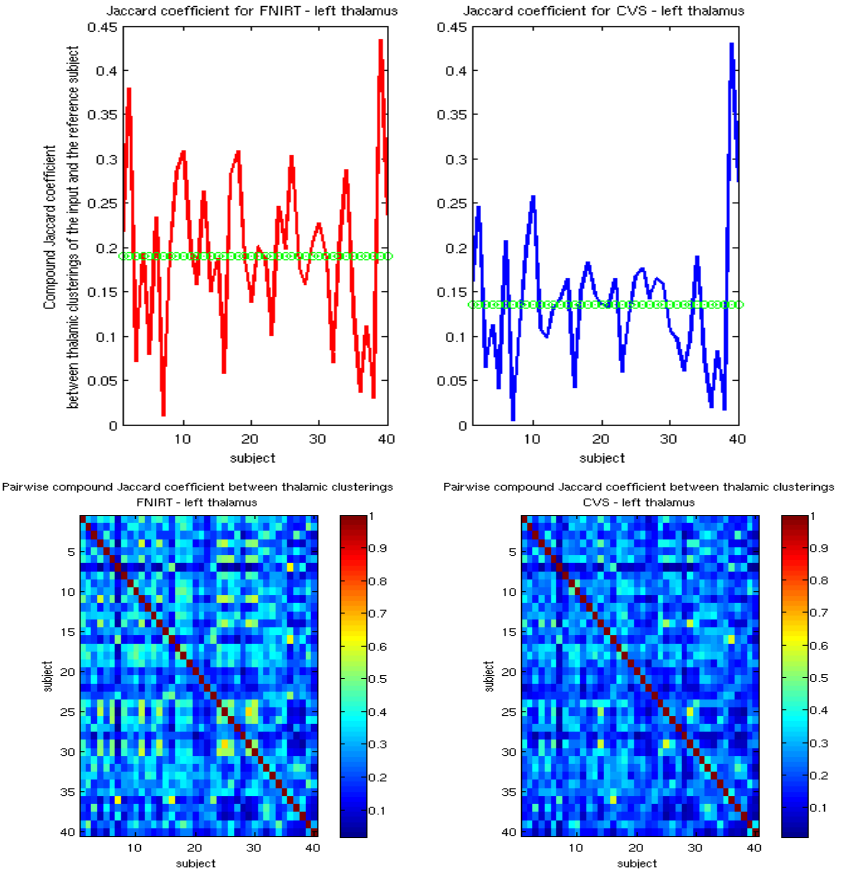

Fig. 3. (**FIRST ROW**) Compound Jaccard coefficient for the left thalamus for FNIRT and CVS algorithms. Mean Jaccard coefficient value is shown with a green line (see Table 1). (**SECOND ROW**) Pairwise compound Jaccard coefficient for the clusterings of the left thalamus for FNIRT and CVS algorithms (see Table 2). Subjects are numbered from 1 to 40.

4 Discussion

A new method for evaluating registration performance is proposed. It tests for registration differences within areas of uniform T1/T2 contrast, but of highly variable connectivity (such as the thalamus and the putamen).

The results show that it is possible to make highly local distinctions between competing registration methods possibly giving a new insight into how and where they can be improved. However, our pipeline so far includes automatically segmented subcortical labels, which is suboptimal to having expert segmentations. The questions such as what are the best ways to cluster connectivity matrices, impose inter-subject cluster correspondences and how these affect final findings also need further investigation. Furthermore, we would like the registration methods (that are

compared) to be independent of the evaluation pipeline. In this particular demonstration, CVS relies on FreeSufer cortical surface registrations which might introduce a potential confound (possibly in favor of CVS).

Although our goal is to compare performance of volumetric registration methods by looking into the overlap of anatomical boundaries found through WM connectivity analysis, we acknowledge that the accuracy of registration itself can be application specific [15]. E.g., this particular evaluation pipeline can yield an "optimal" registration method to be used when analyzing the group results of a functional paradigm with activations in deep brain structures. A "better" registration method could, in that case, result in statistically stronger conclusions and/or better spatial localization of the measured effect.

Finally, we hope to motivate further research into the integration of connectivity information (both structural and functional) into registration algorithms. E.g., the structural connectivity matrix utilized in this paper (and comparable across subjects) can be directly integrated into the registration cost functions of CVS or FNIRT [16].

Acknowledgments. This work was supported by the Oxford University Full Clarendon Award and NIH grant 1K99HD061485-01A1. Authors would like to thank Dr. Mark Jenkinson and Prof. Bruce Fischl for useful comments and advice.

References

1. Klein, A., Ghosh, S.S., Avants, B., Yeo, B.T.T., Fischl, B., Ardekani, B., Gee, J.C., Mann, J.J., Parsey, R.V.: Evaluation of volume-based and surface-based brain image registration methods. Neuroimage 51, 214–220 (2010)
2. Smith, S.M., Jenkinson, M., Johansen-Berg, H., Rueckert, D., Nichols, T.E., Mackay, C.E., Watkins, K.E., Ciccarelli, O., Cader, M.Z., Matthews, P.M., Behrens, T.E.: Tract-based spatial statistics: voxelwise analysis of multi-subject diffusion data. Neuroimage 31, 1487–1505 (2006)
3. Behrens, T.E., Berg, H.J., Jbabdi, S., Rushworth, M.F., Woolrich, M.W.: Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? Neuroimage 34, 144–155 (2007)
4. Jbabdi, S., Woolrich, M.W., Andersson, J.L., Behrens, T.E.: A Bayesian framework for global tractography. Neuroimage 37, 116–129 (2007)
5. Klein, J.C., Behrens, T.E., Robson, M.D., Mackay, C.E., Higham, D.J., Johansen-Berg, H.: Connectivity-based parcellation of human cortex using diffusion MRI: Establishing reproducibility, validity and observer independence in BA 44/45 and SMA/pre-SMA. Neuroimage 34, 204–211 (2007)
6. Johansen-Berg, H., Behrens, T.E., Sillery, E., Ciccarelli, O., Thompson, A.J., Smith, S.M., Matthews, P.M.: Functional-anatomical validation and individual variation of diffusion tractography-based segmentation of the human thalamus. Cereb Cortex 15, 31–39 (2005)
7. Jbabdi, S., Woolrich, M.W., Behrens, T.E.: Multiple-subjects connectivity-based parcellation using hierarchical Dirichlet process mixture models. Neuroimage 44, 373–384 (2009)
8. Postelnicu, G., Zöllei, L., Fischl, B.: Combined volumetric and surface registration. IEEE Trans. Med. Imaging 28, 508–522 (2009)

9. Zöllei, L., Stevens, A., Huber, K., Kakunoori, S., Fischl, B.: Improved tractography alignment using combined volumetric and surface registration. Neuroimage 51, 206–213 (2010)
10. Andersson, J., Jenkinson, M., Smith, S.: Non-linear registration, aka spatial normalisation. Oxford Centre for Functional Magnetic Resonance Imaging of the Brain, Department of Clinical Neurology, Oxford University, Oxford, UK (2007)
11. Fischl, B., Salat, D.H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., van der Kouwe, A., Killiany, R., Kennedy, D., Klaveness, S., Montillo, A., Makris, N., Rosen, B., Dale, A.M.: Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. Neuron 33, 341–355 (2002)
12. Fischl, B., Liu, A., Dale, A.M.: Automated manifold surgery: constructing geometrically accurate and topologically correct models of the human cerebral cortex. IEEE Trans. Med. Imaging 20, 70–80 (2001)
13. Woolrich, M.W., Jbabdi, S., Patenaude, B., Chappell, M., Makni, S., Behrens, T., Beckmann, C., Jenkinson, M., Smith, S.M.: Bayesian analysis of neuroimaging data in FSL. Neuroimage 45, S173–S186 (2009)
14. Fischl, B., Sereno, M.I., Tootell, R.B., Dale, A.M.: High-resolution intersubject averaging and a coordinate system for the cortical surface. Hum. Brain Mapp. 8, 272–284 (1999)
15. Yeo, B.T., Sabuncu, M., Golland, P., Fischl, B.: Task-optimal registration cost functions. Med. Image Comput. Comput. Assist. Interv. 12, 598–606 (2009)
16. Petrović, A.: Connectivity Driven Registration of Magnetic Resonance Images of the Human Brain. DPhil Thesis University of Oxford (2011)

Joint Segmentation and Deformable Registration of Brain Scans Guided by a Tumor Growth Model

Ali Gooya*, Kilian M. Pohl, Michel Bilello, George Biros, and Christos Davatzikos

Section for Biomedical Image Analysis,
Suite 380, 3600 Market St.,19104 Philadelphia, US
{ali.gooya,kilian.pohl,michel.bilello,christos.davatzikos}@uphs.upenn.edu
http://www.rad.upenn.edu/sbia

Abstract. This paper presents an approach for joint segmentation and deformable registration of brain scans of glioma patients to a normal atlas. The proposed method is based on the Expectation Maximization (EM) algorithm that incorporates a glioma growth model for atlas seeding, a process which modifies the normal atlas into one with a tumor and edema. The modified atlas is registered into the patient space and utilized for the posterior probability estimation of various tissue labels. EM iteratively refines the estimates of the registration parameters, the posterior probabilities of tissue labels and the tumor growth model parameters. We have applied this approach to 10 glioma scans acquired with four Magnetic Resonance (MR) modalities (T1, T1-CE, T2 and FLAIR) and validated the result by comparing them to manual segmentations by clinical experts. The resulting segmentations look promising and quantitatively match well with the expert provided ground truth.

Keywords: joint segmentation-registration, EM, diffusion-reaction model.

1 Introduction

Statistical atlases constructed from MR scans are powerful tools for aiding the analysis and understanding of brain tumor development. The atlases are used for tasks such as learning the relative location of tumors with respect to healthy tissue [4] or guiding the automatic segmentation of brain tumor scans [13]. An important component in the construction and application of the atlases is the registration of brain tumor MR scans to the a common coordinate system. This coordinate system is often represented by a MR scan of a healthy subject due to the subject specific nature of brain tumors. Although a plethora of methods for image registration exists [2], this registration task is generally considered very challenging as there is no correspondence for the pathology in the healthy scans.

* Corresponding author.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 532–540, 2011.

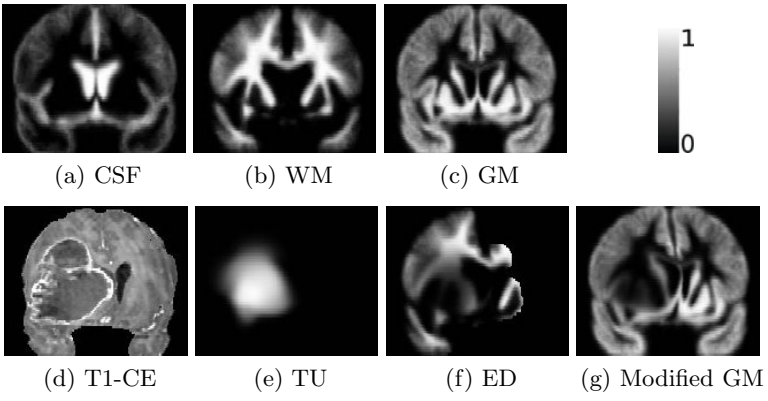

(a) CSF (b) WM (c) GM

(d) T1-CE (e) TU (f) ED (g) Modified GM

Fig. 1. (a)-(c) Healthy white matter (WM), gray matter (GM) and cerebro spinal fluid (CSF) probability maps, (d) sample glioma scan, (e) corresponding estimated tumor (TU), (f) edema (ED), and (g) gray matter probability maps before registration to the patient.

In addition, the healthy tissue in the brain scan is often severely deformed by the mass-effect of the tumor so that its shape is very different from that in the healthy brain scan. In this paper, we address this issue by developing a new approach for brain tumor registration that explicitly models the mass effect of the pathology.

A popular approach for registering brain images to an atlas coordinate system is to mask out the pathology and then perform the registration to the healthy tissue [11,13,14]. Alternative methods use atlas seeding, in which a tumor is seeded in the healthy (atlas) scan [10]. Both of these frameworks essential ignore edema as they fail to explicitly model the diffusion of the tumor cells into the neighboring healthy tissue. This limitation was addressed by [1], whose registration incorporated a biophysical (diffusion-reaction) tumor growth model simulating the mass-effect and diffusion caused by the pathology. However, the approach requires accurate segmentations of the pathology which are difficult to produce automatically. We now derive a joint atlas registration and segmentation framework to circumvent this problem.

Similar to joint segmentation and registration methods targeted towards healthy brains [8,6], our method iteratively estimates the posterior probabilities of tissue classes and registers the atlas via the EM algorithm. However, our approach also computes the diffusion-reaction parameters of the tumor growth model and the coordinates for atlas seeding. These additional parameters greatly increase the complexity of the optimization problem compared to [8,6] which required us to carefully adopt the EM model to this specific application. The remainder of this paper is organized as follows: In Section 2 we review the diffusion-reaction model and describe the construction of the atlas. The atlas is then used in Section 3 to guide our joint registration and segmentation

framework. In Section 4 we present our quantitative evaluation and results of the application to sample patients, and conclude the paper in Section 5.

2 Atlas Generation

In this section, we use a glioma growth model to embed a tumor in an originally healthy atlas. We define the atlas as a set of probability maps that specify the spatial distribution of brain tissues (see Fig.1(a)-(c)). This modified atlas then guides the EM algorithm in registering and segmenting of subjects with glioma.

As in [1], a model of the glioma is implanted into the healthy brain scan by artificially seeding a tumor in the healthy atlas and growing it using the biophysical model proposed in [3]. Let Ω_A and $[0, T]$ denote the space of atlas and a time interval for growing the tumor. The evolution of tumor probability $\pi_{TU}(\mathbf{x}, t)$ in $\Omega_A \times [0, T]$ is determined by the following diffusion-reaction model:

$$\begin{aligned} &\frac{\partial \pi_{TU}}{\partial t} - \nabla(D\nabla\pi_{TU}) + \nabla(\pi_{TU}\mathbf{v}) = \rho\pi_{TU}(1-\pi_{TU}), \\ &\nabla[\lambda\mathbf{I}\nabla\cdot\mathbf{u} + \mu(\nabla\mathbf{u} + \nabla\mathbf{u}^T)] = p\nabla\pi_{TU} \end{aligned} \tag{1}$$

where ∇ is the differential operator, $\mathbf{u}$ is the mass effect displacement field caused by the presence of the tumor, $\mathbf{v} = \partial\mathbf{u}/\partial t$ is the relevant velocity field, p is a scalar which determines the strength of the tumor mass effect, D is a spatially variable function capturing diffusion coefficient within white (D_{WM}) and gray matters (D_{GM}), and ρ is proliferation coefficient. We fix $\rho = 0.025$ and $D_{GM} = 1e^{-10}$ (the default values of [3]) as the method is relatively insensitive to these parameters. Now, if we denote with $\mathbf{x}_0$ the initial seed location of the tumor and with d the voxel size of the Ω_A, then our tumor growth model is completely defined by the parameters $\mathbf{q} \equiv \{\mathbf{x}_0, p, D_{WM}, T\}$ given the initial conditions for the $\mathbf{u}(\mathbf{x}|\mathbf{q}, t = 0) = 0$ and $\pi_{TU}(\mathbf{x}|\mathbf{q}, t = 0) = exp(-(\mathbf{x} - \mathbf{x_0})^2/d^2)$.

Once we solve the above equation for $\mathbf{u}$ and π_{TU}, we combine those results at $t = T$ with the original atlas of healthy brains $\pi^o_X(\cdot)$ to infer tissue probability maps $\pi_X(\cdot|\mathbf{q})$ for tissue X. To simplify notation, we omit $t = T$ from $\mathbf{u}$ and π_{TU} and simply denote with $\pi_{TU}(\mathbf{x}|\mathbf{q})$ the spatial probability map of glioma being present at location $x \in \Omega_A$ at time T and $\mathbf{u}(x|\mathbf{q})$ the corresponding mass effect at that location and time. We then construct $\pi_X(\cdot|\mathbf{q})$ for GM and CSF by deforming the corresponding spatial probabilities $\pi^o_X(\cdot)$ of the healthy population via the mass-effect $\mathbf{u}$ and weighing them with (1 - π_{TU}):

$$\begin{aligned} \pi_{GM}(\mathbf{x}|\mathbf{q}) &\equiv \pi^o_{GM}(\mathbf{u}(\mathbf{x})) \cdot (1 - \pi_{TU}(\mathbf{x})) \\ \pi_{CSF}(\mathbf{x}|\mathbf{q}) &\equiv \pi^o_{CSF}(\mathbf{u}(\mathbf{x})) \cdot (1 - \pi_{TU}(\mathbf{x})) \end{aligned} \tag{2}$$

We construct the spatial probability map of edema π_{ED} based on the assumption that edema is in close proximity of the tumor, which we model via the Heaviside function $H(\pi_{TU}(\mathbf{x}))$ ($H(a) = 0$ for $a \leq 0$ and $H(a) = 1$ otherwise),

and be confined to the mass deformed white matter of the healthy brain, which we model with $\pi^o_{WM}(\mathbf{u}(\mathbf{x}))$. Thus:

$$\pi_{ED}(\mathbf{x}|\mathbf{q}) \equiv \pi^o_{WM}(\mathbf{u}(\mathbf{x})) \cdot (1 - \pi_{TU}(\mathbf{x})) \cdot H(\pi_{TU}(\mathbf{x})). \tag{3}$$

We define the subject specific spatial probability map of the white matter as:

$$\pi_{WM}(\mathbf{x}|\mathbf{q}) \equiv 1 - [\pi_{TU}(\mathbf{x}|\mathbf{q}) + \pi_{ED}(\mathbf{x}|\mathbf{q}) + \pi_{CSF}(\mathbf{x}|\mathbf{q}) + \pi_{GM}(\mathbf{x}|\mathbf{q})] \tag{4}$$

A sample set of the generated probability maps $\pi_{TU}(\cdot|\mathbf{q}), \pi_{ED}(\cdot|\mathbf{q})$ and $\pi_{GM}(\cdot|\mathbf{q})$ is shown in Fig.1, illustrating the impact of the mass effect and tumor invasion in originally healthy atlas.

Finally, we note that we use the same probability map for enhancing tumor and necrosis as our tumor growth model does not distinguish between these two tissue types. In addition, we simplify notation by denoting the probability maps generated in this section with $\pi_k(\cdot|\mathbf{q})$, $1 \leq k \leq K = 6$.

3 Joint Segmentation-Registration

We now describe the framework for joint segmentation-registration which is guided by the atlas defined in the previous section. We assume that a set of J co-registered, inhomogeneity-corrected, and skull stripped MR images is given in the reference (fixed) domain Ω_F so that for any sample voxel $\mathbf{x} \in \Omega_F$, $\mathbf{y}(\mathbf{x}) \equiv [y_1(\mathbf{x}), \cdots, y_J(\mathbf{x})]^T$ is an independent observation vector that corresponds to the J image intensities. We then define observation set as: $\mathbf{Y} = \{\mathbf{y}(\mathbf{x})|\mathbf{x} \in \Omega_F\}$. The goal of this section is to derive an algorithm for estimating the intensity distributions of each structure $\mathbf{\Phi}$, the atlas coefficient $\mathbf{q}$, and the deformation between the atlas and the reference domain $\mathbf{h}$.

We further specify $\mathbf{\Phi}$ by assuming the conditional probability distribution function (pdf) of each $\mathbf{y}(\mathbf{x})$ is a weighted mixture of K Gaussians: $f(\mathbf{Y}|\mathbf{\Phi}, \mathbf{h}, \mathbf{q}, \mathbf{x}) \equiv \sum_{k=1}^{K} \pi_k(\mathbf{h}(\mathbf{x})|\mathbf{q}) f_k(\mathbf{y}(\mathbf{x})|\mathbf{\Phi})$ where $f_k \sim N(\mathbf{m}_k, \mathbf{\Sigma}_k)$ is a multivariate Gaussian distribution with mean $\mathbf{m}_k$ and the covariance matrix $\mathbf{\Sigma}_k$, and $\mathbf{\Phi} \equiv \{\mathbf{m}_1, \cdots, \mathbf{m}_K, \mathbf{\Sigma}_1, \cdots, \mathbf{\Sigma}_K\}$. The mixture weights are determined by $\pi_k(\mathbf{h}(\mathbf{x})|\mathbf{q})$ which are originally defined in the atlas space Ω_A (see Section 2) and registered to the patient space through $\mathbf{h} : \Omega_F \rightarrow \Omega_A$, a vector field mapping the atlas into the patient space. Based upon these assumptions, we write the conditional likelihood of $\mathbf{Y}$ as: $f(\mathbf{Y}|\mathbf{\Phi}, \mathbf{h}, \mathbf{q}) = \prod_{\mathbf{x} \in \Omega_F} f(\mathbf{Y}|\mathbf{\Phi}, \mathbf{h}, \mathbf{q}, \mathbf{x})$.

Our problem of joint segmentation registration and atlas parameter estimation can be defined as the solution of the following:

$$(\mathbf{\Phi}_o, \mathbf{h}_o, \mathbf{q}_o) \equiv \underset{\mathbf{\Phi}, \mathbf{q}, \mathbf{h}}{\text{argmax}} \quad log f(\mathbf{Y}|\mathbf{\Phi}, \mathbf{h}, \mathbf{q}). \tag{5}$$

One way of computing the solution to this problem is via Expectation-Maximization algorithm [8]. EM is an iterative algorithm which in stead of directly solving (5), maximizes a lower bound on $log f(\mathbf{Y}|\mathbf{\Phi}, \mathbf{h}, \mathbf{q})$. At every

iteration, given the current estimate of the unknown parameters $\mathbf{\Phi}'$, $\mathbf{h}'$ and $\mathbf{q}'$, the lower bound of the log-likelihood in (5) is written as:

$$Q(\mathbf{\Phi}, \mathbf{h}, \mathbf{q}|\mathbf{\Phi}', \mathbf{h}', \mathbf{q}') \equiv \sum_{\mathbf{x}\in\Omega_F} \sum_{k=1}^{K} p_k(\mathbf{x}) log(\pi_k(\mathbf{h}(\mathbf{x})|\mathbf{q}) f_k(\mathbf{y}(\mathbf{x})|\mathbf{\Phi})), \tag{6}$$

where $p_k(\mathbf{x})$ stands for the posterior probability of class k at voxel $\mathbf{x}$ (see equ.(7)) The structure of the proposed EM algorithm consists of iterations between the E-Step and M-Step, during which the posteriors and parameters $\{\mathbf{\Phi}, \mathbf{h}, \mathbf{q}\}$ are respectively updated. Further detail is as follows:

E-Step: In this step, label estimation is achieved by updating the computed posterior probabilities given the current estimate of the parameter:

$$p_k(\mathbf{x}) = \frac{f_k(\mathbf{y}(\mathbf{x})|\mathbf{\Phi}')\pi_k(\mathbf{h}'(\mathbf{x})|\mathbf{q}')}{\sum_{l=1}^{K} f_l(\mathbf{y}(\mathbf{x})|\mathbf{\Phi}')\pi_l(\mathbf{h}'(\mathbf{x})|\mathbf{q}')}. \tag{7}$$

M-Step: The update of the distribution parameters, (*i.e.* means and covariance matrices) in $\mathbf{\Phi}$ have closed form solutions which can be found in the literature [9] and are not mentioned here. We optimize $\mathbf{h}$ by the following variational framework which computes the differential of (6) with respect to an infinitely small test function $\mathbf{v}$:

$$0 = Q(\mathbf{\Phi}, \mathbf{h}'+\mathbf{v}, \mathbf{q}'|\mathbf{\Phi}', \mathbf{h}', \mathbf{q}') - Q(\mathbf{\Phi}', \mathbf{h}', \mathbf{q}'|\mathbf{\Phi}', \mathbf{h}', \mathbf{q}') = \sum_{\mathbf{x}\in\Omega_F} \mathbf{v}^t(\mathbf{x}).\{\mathbf{r}(\mathbf{x}) + \mathbf{W}(\mathbf{x}).\mathbf{v}(\mathbf{x})\} \tag{8}$$

In this equation, the gradient vector $\mathbf{r}(\mathbf{x})$ and the matrix $\mathbf{W}(\mathbf{x})$ are defined as:

$$\mathbf{r}(\mathbf{x}) = 2\sum_{k=1}^{K} \frac{p_k(\mathbf{x})}{\pi_k(\mathbf{h}'(\mathbf{x})|\mathbf{q}')} \nabla\pi_k(\mathbf{h}'(\mathbf{x})|\mathbf{q}'), \tag{9}$$

$$\mathbf{W}(\mathbf{x}) = \sum_{k=1}^{K} p_k(\mathbf{x}) [\frac{\mathbf{H}(\pi_k(\mathbf{h}'(\mathbf{x})|\mathbf{q}'))}{\pi_k(\mathbf{h}'(\mathbf{x})|\mathbf{q}')} - \frac{\nabla\pi_k(\mathbf{h}'(\mathbf{x})|\mathbf{q}')(\nabla\pi_k(\mathbf{h}'(\mathbf{x})|\mathbf{q}'))^t}{(\pi_k(\mathbf{h}'(\mathbf{x})|\mathbf{q}'))^2}] \tag{10}$$

where $\mathbf{H}$ is the Hessian matrix. The detailed derivations are omitted due to space limitations. Equation (8) leads to $\mathbf{r}(\mathbf{x}) + \mathbf{W}(\mathbf{x}).\mathbf{v}(\mathbf{x}) = 0$, hence $\mathbf{v}$ can be found as: $\mathbf{v}(\mathbf{x}) = \mathbf{W}^{-1}(\mathbf{x}) \cdot \mathbf{r}(\mathbf{x})$. For further numerical stability, we add an identity matrix component to $\mathbf{W}$, therefore the update of $\mathbf{h}(\mathbf{x})$ can be written as:

$$\mathbf{h}'(\mathbf{x}) \leftarrow \mathbf{h}'(\mathbf{x}) - [\mathbf{W}(\mathbf{x}) + c\mathbf{I}]^{-1}\mathbf{r}(\mathbf{x}), \tag{11}$$

where $\mathbf{I}$ is the identity matrix and c is a constant. In this paper, we found $c = 0.1$ to produce a robust and reasonable deformation field. Notice the update equation is computed independently at every voxel, which in general results in a non-smooth deformation field. In order to apply the smoothness constraint, similar to Thirions' demons framework [7] we diffuse the estimated deformation vectors by a Gaussian convolution filter with a band width of 2.

Table 1. Dice overlap ratios (%) of the segmented tumor and edema with the expert provided ground truths for *fully* and *partly* optimized tumor models. For S1-S5 subjects total volumes of edema and tumor were manually segmented whereas for subjects S6-S10 every third slice was delineated by our specialist.

Opt.	Lab.	S.1	S.2	S.3	S.4	S.5	S.6	S.7	S.8	S.9	S.10	Avg.
Fully	TU	89.4	89.7	87.2	83.0	91.8	70.0	81.6	84.0	82.7	90.8	**85.4**
	ED	85.7	80.0	60.4	59.5	83.3	81.6	81.7	87.6	75.6	66.0	**76.0**
Partly	TU	83.9	89.0	87.9	80.3	86.1	51.8	78.7	71.8	80.2	79.5	**71.7**
	ED	83.7	79.2	59.0	46.2	83.1	63.0	78.2	84.8	74.1	64.8	**71.6**

To update the atlas parameters $\mathbf{q}$, since no analytical expression for the derivatives of $Q(\mathbf{\Phi}', \mathbf{h}', \mathbf{q} | \mathbf{\Phi}', \mathbf{h}', \mathbf{q}')$ w.r.t $\mathbf{q}$ exists, we follow a numerical scheme. We maximize (6) using a derivative free pattern search library [5]. Subsequently, each process returns its corresponding Q value to the library and the procedure is iterated until a maximum is found. Since this operation is computationally expensive it is performed only after having an adequate convergence on estimated deformation field otherwise we keep it fixed.

4 Results

We applied our proposed joint segmentation-registration method to 10 glioma patients. Our preprocessing pipeline starts with skull stripping of all modalities (FLAIR,T2,T1, and T1CE) and MR field inhomogeneity correction [15]. These images are co-registered to the atlas using an affine registration based on mutual information [12]. We solved (1) on a lattice of $64 \times 64 \times 64$ nodes for efficiency reasons. We numerically compared our EM based segmentation results to the expert provided references for edema and tumor labels using Dice volume overlap ratio. For the S1-S5 cases (see the first five columns in Table 1) total volumes of pathology were delineated and for the S6-S10 cases every third slice was segmented by our specialist. We also computed the dice scores with respect to every third slice in S1-S5. The average difference between these scores and those obtained based on entire volume was less than 0.75%.

Sample results of seven patients in Fig. 2 show a high visual correspondence with patients anatomies. Moreover, it is interesting to observe that the registered atlas probability maps closely match the patient segmented labels, which indicates good registration as well.

This observation is further verified by our numerical evaluations in Table 1 which shows that the segmentations using our method have reasonable matches with the reference volumes. Also, the Dice scores favorably compare to the values reported in [13], though the data sets are different. Regarding the observed discrepancy between the expert provided segmentations and our results, we believe that it is due to the fact that the proposed method takes a voxel-based

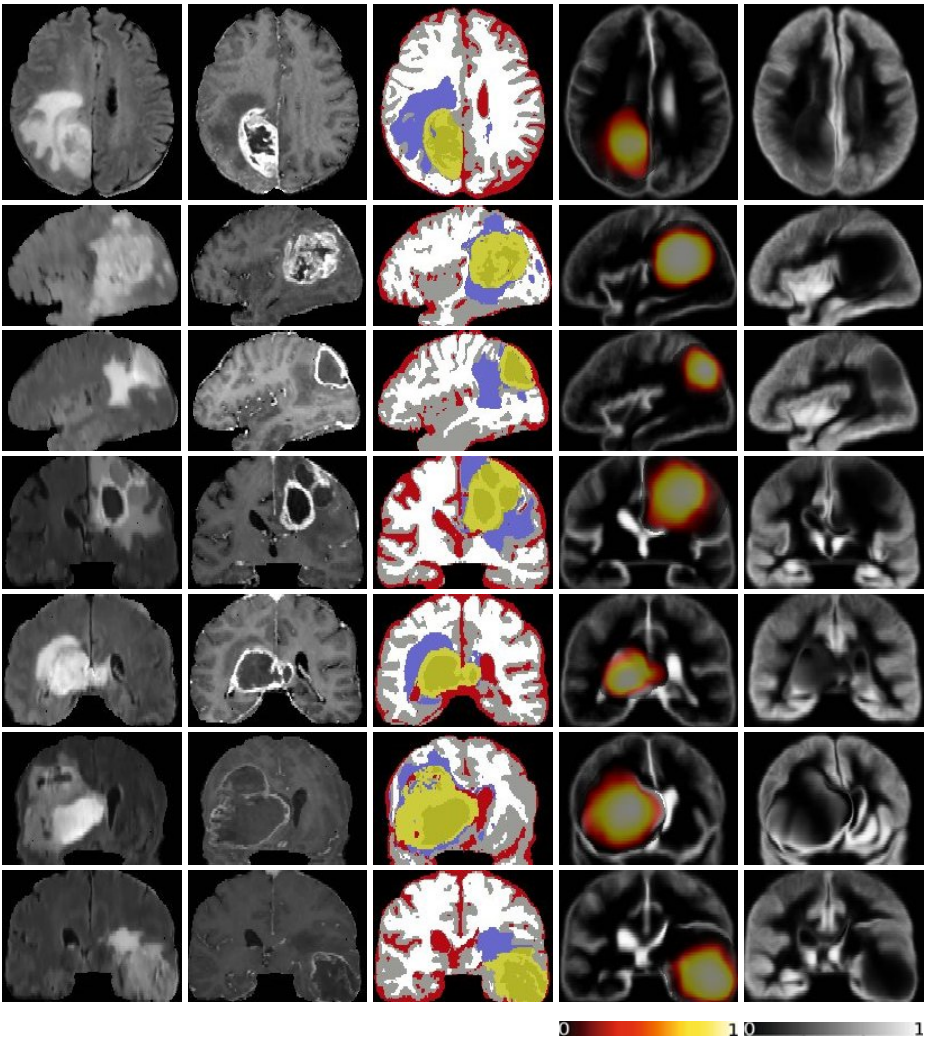

(a) FLAIR (b) T1CE (c) Labels (d) CSF and TU (e) GM

Fig. 2. Segmentation and registration results for seven sample patients. Each row corresponds to a single patient and represents the results in the slice with largest tumor section. (a)-(b) FLAIR,T1-CE images,(c) segmentation results indicating enhancing tumor, necrosis, edema, CSF, gray and white matters in light and dark yellows, purple, red, gray and white colors respectively, (d) overlay of the tumor and CSF probability maps registered to the patient scans, (e) probability map of GM registered to the patient scans.

classification approach, while a human rater considers other complex feature such as the shape and appearance.

Moreover, in order to investigate the sensitivity of the results w.r.t. optimality of parameters of model, two different sets of experiments were performed. In the first case, all model parameters in $\mathbf{q}$ were optimized (denoted by *Fully*), and in the second case only D_{WM} was optimized and the expert chosen parameters such as seed location and tumor growth length were not refined (denoted by *Partly*). As shown in Table 1, the Dice coefficients in fully optimized mode are in general higher and imply better overlaps compared to partly optimized model.

5 Conclusion

We developed a joint segmentation registration tool for glioma images. Our proposed method utilizes multi-channel MR images as the patient feature images, and an originally healthy atlas as the spatial probability maps for various tissue labels. We utilized a tumor growth model to modify the probability maps of the original atlas. The model impacts the atlas original probability maps by both deforming and masking them the due to tumor mass-effect and diffusion. We employed an EM algorithm to iteratively refine the estimates of posterior probabilities of various tissue labels, registration field and tumor growth parameters. Validation using 10 data sets reveals that the method can handle large mass effects and tumor sizes with various tissue types such as necrosis, edema and tumor infiltration. Quantitative evaluations of segmentations of our method were based on Dice overlap ratios with expert provided reference volumes, and in general are higher than values reported in the state-of-the-art literature.

Acknowledgment. The research was supported by an ARRA supplement to NIH NCRR (P41 RR13218).

References

1. Gooya, A., Biros, G., Davatzikos, C.: Deformable registration of glioma images using em algorithm and diffusion reaction modeling. IEEE TMI 30, 375–390 (2011)
2. Zitova, B., Flusser, J.: Image registration methods: a survey. Im. Vis. Comp. (2003)
3. Hogea, C., Davatzikos, C., Biros, G.: An image-driven parameter estimation problem for a reaction-diffusion glioma growth model with mass effects. J. Math. Biol. 56, 793–825 (2008)
4. Duffao, H., Capalle, L.: Preferential brain locations of low grade gliomas. Cancer 15, 2622–2626 (2004)
5. Hough, P.D., Kolda, T.G., Torczon, V.J.: Asynchronous parallel pattern search for nonlinear optimization. SIAM Journal on Scientific Computing 23(1), 134–156 (2001)
6. Ashburner, J., Friston, K.J.: Unified segmentation. NeuroImage 26, 839–851 (2005)
7. Thirion, J.P.: Image matching as diffusion process: An analogy with maxwell's demons. Med. Ima. Anal. 2, 243–260 (1998)
8. Pohl, K.M., Fisher, J., Kikinis, R., Wells, W.M.: A bayesian model for joint segmentation and registration. NeuroImage 31, 228–239 (2006)

9. Leemput, K.V., Maes, F., Vandermeulen, D., Colchester, A., Suetens, P.: Automated segmentation of multiple sclerosis lesions by model outlier detection. IEEE TMI 20, 677–688 (2001)
10. Cuadra, M.B., Craene, M.D., Duay, V., Macq, B., Pollo, C., Thiran, J.: Dense deformation field estimation for atlas-based segmentation of pathological MR brain images. Comput. Methods Programs Biomed. 84, 66–75 (2006)
11. Brett, M., Leff, A.P., Rorden, C., Ashburner, J.: Spatial normalization of brain images with focal lesions using cost function masking. NI 14 (2001)
12. Jenkinson, M., Bannister, P.R., Brady, J.M., Smith, S.M.: Improved optimization for the robust and accurate linear registration and motion correction of brain images. Neuroimage 17, 825–841 (2002)
13. Prastawa, M., Bullitt, E., Moon, N., Leemput, K.V., Gerig, G.: Automatic brain tumor segmentation by subject specific modification of atlas priors. Aca. Rad. 10 (2003)
14. Kyriacou, S.K., Davatzikos, C., Zinreich, S.J., Bryan, R.N.: Nonlinear elastic registration of brain images with tumor pathology using a biomechanical model. IEEE Trans. Med. Imag. 18, 580–592 (1999)
15. Sled, J., Zijdenbos, A., Evans, A.: A nonparametric method for automatic correction of intensity nonuniformity in mri data. IEEE TMI. 17, 81–97 (1998)

Non-local Shape Descriptor: A New Similarity Metric for Deformable Multi-modal Registration

Mattias P. Heinrich[1,2,⋆], Mark Jenkinson[2], Manav Bhushan[1,2], Tahreema Matin[3], Fergus V. Gleeson[3], J. Michael Brady[4], and Julia A. Schnabel[1]

[1] Institute of Biomedical Engineering, University of Oxford, UK
[2] Oxford University Centre for Functional MRI of the Brain, UK
[3] Department of Radiology Churchill Hospital, Oxford, UK
[4] Department of Radiation Oncology and Biology, University of Oxford, UK
mattias.heinrich@eng.ox.ac.uk, julia.schnabel@eng.ox.ac.uk
http://users.ox.ac.uk/~shil3388

Abstract. Deformable registration of images obtained from different modalities remains a challenging task in medical image analysis. This paper addresses this problem and proposes a new similarity metric for multi-modal registration, the non-local shape descriptor. It aims to extract the shape of anatomical features in a non-local region. By utilizing the dense evaluation of shape descriptors, this new measure bridges the gap between intensity-based and geometric feature-based similarity criteria. Our new metric allows for accurate and reliable registration of clinical multi-modal datasets and is robust against the most considerable differences between modalities, such as non-functional intensity relations, different amounts of noise and non-uniform bias fields. The measure has been implemented in a non-rigid diffusion-regularized registration framework. It has been applied to synthetic test images and challenging clinical MRI and CT chest scans. Experimental results demonstrate its advantages over the most commonly used similarity metric - mutual information, and show improved alignment of anatomical landmarks.

1 Introduction

Advances in medical image registration techniques have resulted in a number of robust and accurate methods for deformable registration of scans of the same modality [1]. However, the registration of images from different modalities remains challenging. Alignment of multi-modal images helps to relate relevant information from different scans and to find the corresponding anatomical location of the response of functional imaging in structural scans. Intensity relations between those scans are not functional and can vary locally.

⋆ We would like to thank EPSRC and Cancer Research UK for funding this work within the Oxford Cancer Imaging Centre. JAS also acknowledges funding from EPSRC EP/H050892/1.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 541–548, 2011.

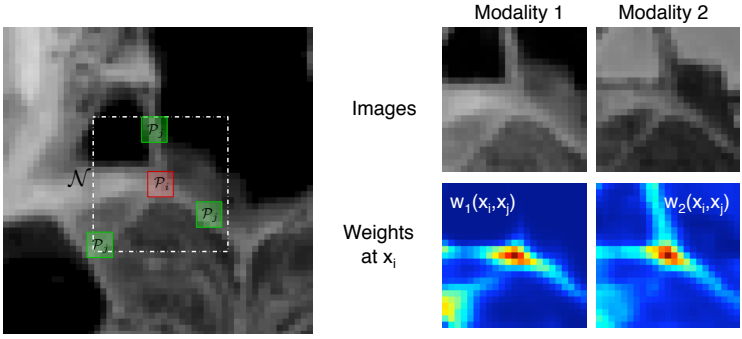

Fig. 1. Estimation of the non-local shape descriptor for the same feature at location $\mathbf{x}_i$ in two different modalities, red and blue colour channels of cryosection (see text for further details)

Mutual information (MI) is derived from information theory and measures the statistical dependency of two random variables. It was first introduced to medical image registration for the rigid alignment of multi-modal scans [2] [3], and later used successfully in a variety of applications, including deformable registration. It is based on the assumption that a lower entropy of the joint intensity distribution corresponds to a better alignment. However, in several practical applications, additional constraints must be made or extensions added. Several weaknesses of MI for non-rigid registration have been identified [4]. For example, it is affected by non-uniform intensity distributions like bias fields. MI is intrinsically a global measure and therefore local deformations can lead to local minima in the solution as shown in [5]. To overcome these difficulties, we introduce a novel similarity metric for multi-modal image registration.

2 Non-local Shape Descriptor

We propose the non-local shape descriptor (NLSD), which defines a response related to the shape of image features at each location in both images to be registered. The shape descriptor is well adapted to medical image registration purposes, because it aims to extract anatomically meaningful geometric shapes.

The proposed similarity term is derived from a very efficient denoising technique, non-local means [6]. For the purpose of denoising it is necessary to find structural similarity in an extended non-local region of an image feature. The values of the most similar patches in the non-local search window contribute to a weighted average of the denoised central voxel. In this paper we will use the non-local weights to extract a geometric descriptor, which forms the basis of the proposed multi-modal similarity metric. A related descriptor, the self-similarity descriptor, has been presented for the application of object detection in [7].

We search for similar patches in a limited non-local region $\mathcal{N}$ around the current voxel of interest $\mathbf{x}_i$. Within $\mathcal{N}$ all patches $\mathcal{P}_j$ are compared to the patch centred on $\mathbf{x}_i$. This concept is illustrated in Fig. 1, showing a magnification

of an image feature in two different modalities, in this case the blue and red colour channel of a cryosection. The dashed white line delimits the non-local search region, while the green squares outline exemplary patches $\mathcal{P}_j$, and the red square the central patch $\mathcal{P}_i$.

A weight $w(\mathbf{x}_i, \mathbf{x}_j)$ is assigned to each location $\mathbf{x}_j$ in $\mathcal{N}$ according to an exponentially decaying distance function based on the Euclidean distance between two patches $\mathcal{P}_i$ and $\mathcal{P}_j$.

$$w(\mathbf{x}_i, \mathbf{x}_j) = \exp\left(-\frac{\sum_{\Delta\mathbf{x}} \|I(\mathbf{x}_i + \Delta\mathbf{x}) - I(\mathbf{x}_j + \Delta\mathbf{x})\|^2}{\sqrt{2}\sigma^2}\right) \tag{1}$$

where $\Delta\mathbf{x}$ is defined over the range of voxels within a patch centred at $\mathbf{0}$. We thereby try to find patches within the non-local search region that are similar to the central patch. The value for σ^2 is the local variance of the noise and can be directly estimated from the 3D image data (see [8] for details).

The similarity metric at a given position $\mathbf{x}$ is defined as the normalised cross correlation (NCC) of the respective weights w_1 and w_2 for the images I_1 and I_2:

$$\mathrm{NLSD}(\mathbf{x}) = \frac{\sum_k \left((w_1(\mathbf{x}, \mathbf{x}_k) - \overline{w_1}) \cdot (w_2(\mathbf{x}, \mathbf{x}_k) - \overline{w_2})\right)}{\sqrt{\sum_k (w_1(\mathbf{x}, \mathbf{x}_k) - \overline{w_1})^2}\sqrt{\sum_k (w_2(\mathbf{x}, \mathbf{x}_k) - \overline{w_2})^2}}, k \in \mathcal{N}_\mathbf{x} \tag{2}$$

where $\overline{w}$ is the mean of all weights within the non-local region $\mathcal{N}_\mathbf{x}$. NCC is robust against noise, but to accommodate for missing correspondences a mutual-saliency weighting [9] could potentially be beneficial.

We have implemented the calculation of this new similarity term as a convolution filter to evaluate the SSD of two patches. The pointwise product of both images is obtained and subsequently convolved with a uniform averaging filter. For the calculation of the weights within the non-local region $\mathcal{N}$ the second image is shifted by $\mathbf{x}_k - \mathbf{x}_0 (\forall k \in \mathcal{N})$ and the averaging filter is applied again. This implementation speeds up the calculation of the similarity metric substantially and avoids the need for preselection of potentially good weights as proposed in [8]. The size of the non-local region should be as large as possible, but for practical applications a size of 15x15 for 2D experiments and 7x7x7 for 3D images together with a patch size of 3x3 and 3x3x3, respectively, has been found to be sufficient to obtain good responses for the shape descriptor.

3 Registration Framework

Within the non-rigid registration, we aim to minimize the following cost function w.r.t. the deformation field $\mathbf{u} = (u, v, w)^T$, consisting of a non-linear similarity term $\mathcal{S}$ (dependent on $\mathbf{u}$) and a diffusion regularization term:

$$\underset{\mathbf{u}}{\mathrm{argmin}} = \int_\Omega \mathcal{S}\left(I_1(\mathbf{x}), I_2(\mathbf{x} + \mathbf{u})\right)^2 + \alpha \,\mathrm{tr}\left(\nabla\mathbf{u}(\mathbf{x})^T \nabla\mathbf{u}(\mathbf{x})\right)^2 d\mathbf{x} \tag{3}$$

Simple gradient descent methods show slow convergence, especially in homogenous regions [10]. Since the objective function to be minimized is of the form

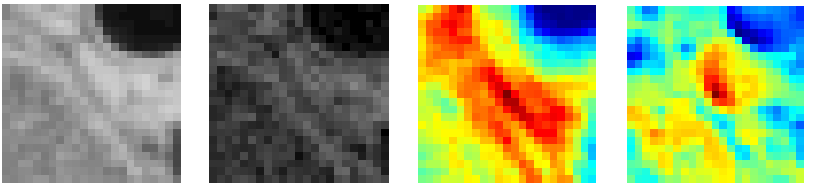

(a) Modality 1 **(b)** Modality 2 **(c)** LNMI **(d)** NLSD

Fig. 2. Feature location in both images ((a) red and (b) blue channel of colour cryosection). Response in search window of (c) LNMI and (d) NLSD. Our proposed method shows a more discriminative peak in the centre.

of $\sum_i f_i^2$, we can apply the Gauss-Newton optimization method, where f is minimized iteratively with the update rule: $(\mathbf{J^T J})\mathbf{u_{gn}} = -\mathbf{J^T} f$, where $\mathbf{J}$ is the derivative of f w.r.t. $\mathbf{u}$. This can be adapted to our regularized cost function. We simplify the notation to $\mathcal{S} = \mathcal{S}\left(I_1(\mathbf{x}), I_2(\mathbf{x})\right)$ and $\nabla\mathcal{S} = (\frac{\delta\mathcal{S}}{\delta u}, \frac{\delta\mathcal{S}}{\delta v}, \frac{\delta\mathcal{S}}{\delta w})^T$ and $\Delta\mathbf{u} = \nabla\left(\nabla(\mathbf{u}(\mathbf{x})\right)$. The regularization term is linear w.r.t. $\mathbf{u}$ as the differential operator is linear. The resulting update step given an initial or previous deformation field $\mathbf{u_{prev}}$ is given by:

$$\left(\nabla\mathcal{S}^T\nabla\mathcal{S} + \alpha\Delta\right)\mathbf{u_{gn}} = -(\nabla\mathcal{S}^T\mathcal{S} + \alpha\Delta\mathbf{u_{prev}}) \tag{4}$$

Equation 4 is solved using an iterative solver. The final deformation field is calculated by the addition of the update steps $\mathbf{u_{gn}}$. The parameter α balances the similarity term with the regularizer. We set $\alpha = 1$ in all experiments.

We have also implemented MI, the classic choice of a multi-modal similarity criterion, within the same deformable registration framework. For the variational optimisation we need to evaluate the similarity function at each location, therefore a local derivation of mutual information is used, as described in [11]. An overview of the possibilities of variational implementations other statistical similarity terms, as well as a discussion of a locally weighted computation of the global measures, is given in [12].

Given the joint probability $p_{12}(\mathbf{i})$ of the co-occurrence of an intensity pair $\mathbf{i} = (i_1, i_2)^T$ in two images I_1 and I_2 and the two marginal intensity probabilities $p_1(i_1)$ and $p_2(i_2)$, local normalised mutual information (LNMI) at location $\mathbf{x}$ is defined as (using the global entropy of I_1 for normalization):

$$\mathrm{LNMI}(\mathbf{x}) = \log\left(\frac{p_{12}(I_1(\mathbf{x}), I_2(\mathbf{x}))}{p_1(I_1(\mathbf{x})) \cdot p_2(I_2(\mathbf{x}))}\right) \frac{1}{\int_{\mathbf{x}} p_1(I_1(\mathbf{x})) \log(p_1(I_1(\mathbf{x})))d\mathbf{x}} \tag{5}$$

The joint and marginal histograms are recalculated at each iteration and smoothed with a Parzen window kernel of size 5x5 with a standard deviation of 0.5. We use 128 histogram bins.

4 Saliency and Robustness of Correspondences

We examine the ability of our new similarity metric to distinguish between anatomical features in different modalities under the influence of local

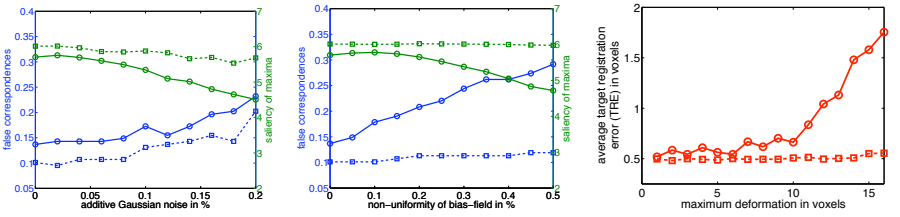

(a) Influence of increasing noise.

(b) Influence of increasing bias field.

(c) TRE with increasing strength of deformations.

Fig. 3. (a,b) Saliency and robustness of both similarity metrics are compared. (c) TRE of registrations of synthetic deformations (see text for details). LNMI is displayed with solid lines and circles, NLSD with dashed lines and squares.

deformations, additive noise and non-uniform bias fields. A ground truth is provided by using two images of different colour channels taken from a cryosection of the *Visible Human Project*, which are intrinsically aligned. A number of landmarks were selected at the same location in both images using the Harris corner detector. The similarity metric was then calculated between a point in the first image and all locations in the second image within a window of 23x23 around that location. Figure 2 shows one selected point and the search window in both images. The proposed metric can better distinguish the local maximum in the centre.

We run these comparisons over all feature locations and quantify the results using two criteria. First, the distance of the maxima of the similarity function is compared with the ground truth location. Deviations of more than 3 pixels are counted as false correspondences. We define the fractional amount of false matches as robustness. Second, the saliency or discrimination between the maximum and its surrounding values is quantified by convolving the similarity response with a Mexican hat function ($\sigma = 1$), so that high positive values are characteristic for a high saliency in the similarity function (see Fig. 3 (a,b)).

5 Results

We performed registrations, using a multi-resolution scheme, for the synthetic test images on which a simulated deformation was applied. The deformations of varying strengths are obtained using a uniform B-spline grid and random control point displacements. In Fig. 3 (c) the average target registration error (TRE) between ground truth deformations and registration are compared for LNMI and NLSD for increasing magnitudes of deformations. For larger deformations, NLSD shows higher accuracy, and the TRE remains almost unaffected, while for LNMI the TRE strongly deteriorates. This demonstrates the disadvantages of mutual information, when the initial estimate of the joint intensity distribution is not close enough to the real distribution, due to large deformations, and the similarity term is susceptible to local optima.

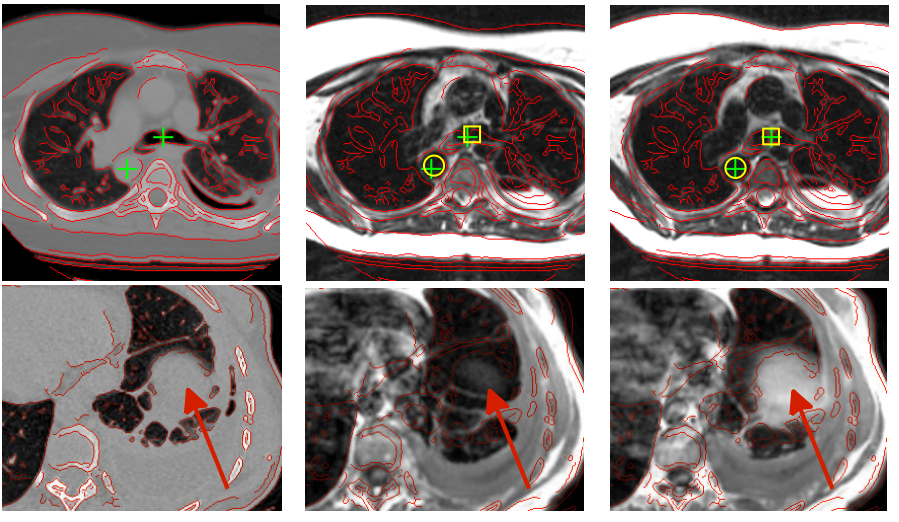

(a) Detail view of slices of CT target volume

(b) MRI, aligned using LNMI

(c) MRI, aligned using proposed metric (NLSD)

Fig. 4. Axial slices through CT and MR image of the lungs of two patients with empyema. Contours of CT are shown for visual guidance. Landmarks in top row: descending aorta (◯) and carina (□) after non-rigid registration compared to the gold standard (+) demonstrate a better alignment for NLSD. The example below shows a substantial improvement for the dome of the diaphragm (arrow).

We then applied our proposed technique to a clinical dataset of 11 patients, which were scanned with both CT and MRI. All patients suffered from empyema, a lung disease where the pleura gets infected and excess fluid fills up the pleural space. This causes the lung to collapse and the extra fluid turns into an abscess. Both modalities are useful for detecting this pathology, but because the patients are scanned in two different sessions and at different levels of breath-hold, there are non-rigid deformations which make it difficult to relate the scans for the clinician. A particular challenge for the registration are large slice thicknesses of up to 8 mm used for the MRI acquisition. The background is removed using a threshold and a morphological filter. For the registration, first a rigid body transformation is estimated using a blockmatching algorithm [13]. In the second step, the proposed non-rigid registration is performed, using a multiresolution scheme with 3 levels. Similarity terms and their derivatives are recalculated before each iteration of the Gauss-Newton optimisation method. The iterations are stopped when the mean of the cost function does not further decrease.

We compare the results of our proposed metric, non-local shape descriptors (NLSD), against local normalised mutual information (LNMI). The running time for one 3D registration of images with a size of 253x253x132 voxels is 43 minutes for NLSD and 24 minutes for LNMI on a single core. The range of values for the determinant of the Jacobian of the deformation fields are [0.25, 1.81] for NLSD

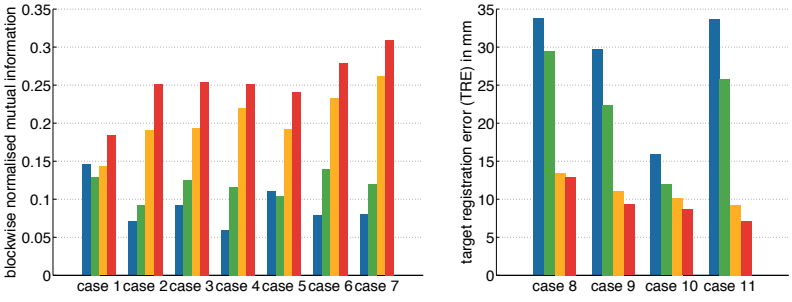

(a) Blockwise mutual information **(b)** Landmark registration error

Fig. 5. Quantitative evaluation of registration outcome: before registration, rigid-body alignment, non-rigid registration using LNMI, and using the proposed metric NLSD.

and [0.26, 1.81] for LNMI, thus no physically implausible folding occurred and all transformations are invertible.

The registration outcome for two cases is displayed in Fig. 4. The target CT volume is shown along with the aligned MR images with deformable registration using LNMI and our proposed metric. Both the contours of the CT and manual anatomical landmarks reveal the advantages and improved accuracy of the presented method. To compare our findings quantitatively, we first calculated the intensity-based similarity before and after registration using the presented metrics. Although an improvement of a similarity function does not necessarily ensure anatomical correspondence, it can highlight differences between methods. We use mutual information calculated within cubic blocks of 30^3 voxels to reduce the influence of the non-uniform bias field in the MRI scans. Figure 5 (a) shows an improvement of this measure using NLSD, for all seven cases, over MI. Additionally a clinical expert manually selected landmarks for the four remaining cases. Between 12 and 15 corresponding landmarks were selected in the four image pairs, containing both normal anatomical locations and disease specific places. It must be noted that some of the landmarks are very challenging to locate, both due to low scan quality (motion artifacts) and changes of the pathology in the diseased areas between scans. On average, the target registration error (TRE) could be further reduced by about 2 mm using our new metric compared to LNMI (see Fig. 5 (b)).

6 Conclusion

In this work a new similarity metric for deformable multi-modal registration is proposed. The non-local shape descriptor (NLSD) aims to extract the most descriptive geometric features in medical images, while being nearly independent of non-functional intensity relations, non-uniform intensity fields and additive noise. This new metric can robustly find correspondences in different modalities and strongly discriminate between salient points. The technique is implemented

in a variational, diffusion-regularized registration framework and compared to the most commonly used alternative - mutual information. We demonstrate that NLSD achieves much improved and more accurate registration results, especially in the case of large deformations. We validate our findings for the application of deformable registration of clinical MR and CT scans of diseased patients. Anatomical landmarks chosen by an expert clinician show improved alignment using our metric. A more thorough evaluation, including more landmarks and a comparison within different transformation models, will be addressed in the future.

References

1. Murphy, K., van Ginneken, B., Reinhardt, J., Kabus, S., Ding, K., Deng, X., Pluim, J.P.W.: Evaluation of Methods for Pulmonary Image Registration: The EMPIRE10 Study. In: Medical Image Analysis for the Clinic: A Grand Challenge (2010)
2. Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., Suetens, P.: Multimodality Image Registration by Maximization of Mutual Information. IEEE Trans. Med. Imaging 16(2), 187–198 (1997)
3. Viola, P., Wells III, W.M.: Alignment by Maximization of Mutual Information. Int. J. Comput. Vision 24(2), 137–154 (1997)
4. Haber, E., Modersitzki, J.: Intensity Gradient Based Registration and Fusion of Multi-modal Images. Methods Inf. Med. 46(3), 292–299 (2007)
5. Loeckx, D., Slagmolen, P., Maes, F., Vandermeulen, D., Suetens, P.: Nonrigid Image Registration Using Conditional Mutual Information. IEEE Trans. Med. Imaging 29(1), 19–29 (2010)
6. Buades, A., Coll, B., Morel, J.M.: A Non-Local Algorithm for Image Denoising. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 60–65. IEEE Computer Society, Los Alamitos (2005)
7. Shechtman, E., Irani, M.: Matching Local Self-Similarities across Images and Videos. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 1–8. IEEE Computer Society, Los Alamitos (2007)
8. Coupé, P., Prima, S., Hellier, P., Kervrann, C., Barillot, C.: An Optimized Blockwise Nonlocal Means Denoising Filter for 3-D Magnetic Resonance Images. IEEE Trans. Med. Imaging 27(4), 425–441 (2008)
9. Ou, Y., Davatzikos, C.: DRAMMS: Deformable Registration via Attribute Matching and Mutual-Saliency Weighting. In: Prince, J., Pham, D., Myers, K. (eds.) IPMI 2009. LNCS, vol. 5636, pp. 50–62. Springer, Heidelberg (2009)
10. Zikic, D., Baust, M., Kamen, A., Navab, N.: Generalization of Deformable Registration in Riemannian Sobolev Spaces. In: Jiang, T., Navab, N., Pluim, J., Viergever, M. (eds.) MICCAI 2010. LNCS, vol. 6362, pp. 586–593. Springer, Heidelberg (2010)
11. Rogelj, P., Kovačič, S., Gee, J.C.: Point similarity measures for non-rigid registration of multi-modal data. Comput. Vis. Image Und. 92(1), 112–140 (2003)
12. Hermosillo, G., Chefd'hotel, C., Faugeras, O.: Variational Methods for Multimodal Image Matching. Int. J. Comput. Vision 50(3), 329–343 (2002)
13. Ourselin, S., Roche, A., Prima, S., Ayache, N.: Block Matching: A General Framework to Improve Robustness of Rigid Registration of Medical Images. In: Delp, S., DiGoia, A., Jaramaz, B. (eds.) MICCAI 2000. LNCS, vol. 1935, pp. 557–566. Springer, Heidelberg (2000)

Preconditioned Stochastic Gradient Descent Optimisation for Monomodal Image Registration

Stefan Klein[1,⋆], Marius Staring[2], Patrik Andersson[3], and Josien P.W. Pluim[3]

[1] Biomedical Imaging Group Rotterdam, Erasmus MC, Rotterdam, The Netherlands
[2] LKEB, Leiden University Medical Center, Leiden, The Netherlands
[3] Image Sciences Institute, University Medical Center Utrecht, The Netherlands
s.klein@erasmusmc.nl

Abstract. We present a stochastic optimisation method for intensity-based monomodal image registration. The method is based on a Robbins-Monro stochastic gradient descent method with adaptive step size estimation, and adds a preconditioning matrix. The derivation of the preconditioner is based on the observation that, after registration, the deformed moving image should approximately equal the fixed image. This prior knowledge allows us to approximate the Hessian at the minimum of the registration cost function, without knowing the coordinate transformation that corresponds to this minimum. The method is validated on 3D fMRI time-series and 3D CT chest follow-up scans. The experimental results show that the preconditioning strategy improves the rate of convergence.

Keywords: image registration, optimisation, stochastic gradient descent, preconditioner.

1 Introduction

Image registration is an important technique in medical imaging applications. It refers to the process of spatially aligning images. Extensive surveys on registration methods are presented in [2, 6] for example.

In this article, we focus on intensity-based image registration with a parameterised coordinate transformation. Let $F(\boldsymbol{x}) : \Omega_F \mapsto \mathbb{R}$ and $M(\boldsymbol{x}) : \Omega_M \mapsto \mathbb{R}$ denote the *fixed* and *moving* image, respectively, with $\Omega_F, \Omega_M \subset \mathbb{R}^D$, and D the dimension of the images. Define the parameterised coordinate transformation $\boldsymbol{T}(\boldsymbol{x};\boldsymbol{\mu}) : \Omega_F \times \mathbb{R}^Q \mapsto \Omega_M$ with the parameter vector $\boldsymbol{\mu}$ of dimension Q. Examples of parameterisations include rigid, affine, and B-spline models. The aim of image registration is to find transformation parameters $\hat{\boldsymbol{\mu}}$ such that the deformed moving image $M(\boldsymbol{T}(\boldsymbol{x};\hat{\boldsymbol{\mu}}))$ resembles the fixed image $F(\boldsymbol{x})$. This can be formulated as a minimisation problem:

$$\hat{\boldsymbol{\mu}} = \arg\min_{\boldsymbol{\mu}} \mathcal{C}(\boldsymbol{\mu}), \tag{1}$$

⋆ Corresponding author.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 549–556, 2011.

where $\mathcal{C}(\boldsymbol{\mu})$ represents the cost function. Examples of intensity-based cost functions are mean squared difference (MSD) and mutual information (MI).

In [3–5] it was demonstrated that a Robbins-Monro (RM) type stochastic gradient descent [9] method efficiently solves the minimisation problem (1). The basic RM method uses the following iterative scheme:

$$\boldsymbol{\mu}_{k+1} = \boldsymbol{\mu}_k - \gamma(k)\tilde{\boldsymbol{g}}_k, \quad k = 0, 1, \ldots, K-1, \tag{2}$$

where $\tilde{\boldsymbol{g}}_k$ denotes an approximation of the true derivative $\boldsymbol{g}_k \equiv \partial\mathcal{C}/\partial\boldsymbol{\mu}(\boldsymbol{\mu}_k)$, $\gamma(k)$ is a scalar gain factor that determines the step size, and K is the number of iterations. The approximated derivative $\tilde{\boldsymbol{g}}_k$ is obtained by computing $\boldsymbol{g}_k$ using only a small subset of voxels $\boldsymbol{x} \in \Omega_F$, randomly selected in every iteration k [5]. The step size $\gamma(k)$ is defined as a slowly decaying function of k:

$$\gamma(k) = a/(k + A)^{\alpha}, \tag{3}$$

with user-specified constants $a > 0$, $A \geq 1$, and $0 < \alpha \leq 1$. It is important to set proper values for these constants. The optimal settings depend on the choice of $\mathcal{C}$, the transformation model, and the image content. To address this issue, an adaptive stochastic gradient descent (ASGD) method was proposed in [3]:

$$\boldsymbol{\mu}_{k+1} = \boldsymbol{\mu}_k - \gamma(t_k)\tilde{\boldsymbol{g}}_k, \quad t_{k+1} = [t_k + \text{sigmoid}(-\tilde{\boldsymbol{g}}_k'\tilde{\boldsymbol{g}}_{k-1})]^+, \tag{4}$$

where $[x]^+$ stands for $\max(x, 0)$, the prime denotes the transpose operation, t_0 and t_1 are user-defined constants, and γ as above. The "time" variable t_k realises an adaptive behaviour, in which the step sizes are increased when consecutive gradients $\tilde{\boldsymbol{g}}_k$ point in a similar direction, and vice versa. Based on the theoretical convergence conditions, reasonable values for a, A and α were estimated.

Both RM and ASGD are gradient descent type methods, which typically expose a low rate of convergence on badly scaled cost functions, characterised by a high ($\gg 1$) condition number of the Hessian $\boldsymbol{H} \equiv \partial^2\mathcal{C}/\partial\boldsymbol{\mu}\partial\boldsymbol{\mu}$ at $\hat{\boldsymbol{\mu}}$ [8]. In this paper, we propose a preconditioning strategy for RM and ASGD, specifically designed for monomodal image registration. The preconditioning is demonstrated to accelerate convergence in both rigid and nonrigid registration problems.

2 Method

2.1 Preconditioned Stochastic Gradient Descent

The use of a preconditioning matrix is a well-known technique to accelerate optimisation methods [8]. Based on the standard RM method, we define the following preconditioned stochastic gradient descent (PSGD) method:

$$\boldsymbol{\mu}_{k+1} = \boldsymbol{\mu}_k - \gamma(k)\boldsymbol{P}\tilde{\boldsymbol{g}}_k, \quad k = 0, 1, \ldots, K-1, \tag{5}$$

where the preconditioner $\boldsymbol{P}$ is a positive definite $Q \times Q$ matrix. It serves to scale the derivative $\tilde{\boldsymbol{g}}_k$, and should be chosen such that larger steps are taken in

directions where the cost function is flat, and smaller steps in directions where the cost function has a high curvature. The theoretically optimal choice $\boldsymbol{P} = \boldsymbol{H}(\hat{\boldsymbol{\mu}})^{-1}$ makes (5) similar to the Newton-Raphson method. Unfortunately, since $\hat{\boldsymbol{\mu}}$ is unknown before registration, this choice of $\boldsymbol{P}$ is impossible to compute. It is however possible to compute an approximation, as follows.

2.2 A Preconditioner for Monomodal Image Registration

In this subsection, a preconditioning matrix for monomodal image registration problems is derived. For explanation, MSD is used as a cost function, but the derivation is similar for other cost functions. The MSD cost function is given by:

$$\mathcal{C}(\boldsymbol{\mu}) = \frac{1}{V} \sum_{\boldsymbol{x} \in \Omega_F} \left(F(\boldsymbol{x}) - M(\boldsymbol{T}(\boldsymbol{x}; \boldsymbol{\mu})) \right)^2, \tag{6}$$

with V the number of $\boldsymbol{x} \in \Omega_F$. For the derivative $\boldsymbol{g}(\boldsymbol{\mu})$ and the Hessian $\boldsymbol{H}(\boldsymbol{\mu})$ we have:

$$\boldsymbol{g}(\boldsymbol{\mu}) \equiv \frac{\partial \mathcal{C}}{\partial \boldsymbol{\mu}} = -\frac{2}{V} \sum_{\boldsymbol{x} \in \Omega_F} (F - M \circ \boldsymbol{T}) \frac{\partial \boldsymbol{T}'}{\partial \boldsymbol{\mu}} \frac{\partial M}{\partial \boldsymbol{x}}, \tag{7}$$

$$\begin{aligned} \boldsymbol{H}(\boldsymbol{\mu}) \equiv \frac{\partial^2 \mathcal{C}}{\partial \boldsymbol{\mu} \partial \boldsymbol{\mu}} = \frac{2}{V} \sum_{\boldsymbol{x} \in \Omega_F} \Bigg[& \frac{\partial \boldsymbol{T}'}{\partial \boldsymbol{\mu}} \frac{\partial M}{\partial \boldsymbol{x}} \frac{\partial M}{\partial \boldsymbol{x}}' \frac{\partial \boldsymbol{T}}{\partial \boldsymbol{\mu}} \\ & - (F - M \circ \boldsymbol{T}) \left(\frac{\partial^2 \boldsymbol{T}'}{\partial \boldsymbol{\mu} \partial \boldsymbol{\mu}} \frac{\partial M}{\partial \boldsymbol{x}} + \frac{\partial \boldsymbol{T}'}{\partial \boldsymbol{\mu}} \frac{\partial^2 M}{\partial \boldsymbol{x} \partial \boldsymbol{x}} \frac{\partial \boldsymbol{T}}{\partial \boldsymbol{\mu}} \right) \Bigg], \end{aligned} \tag{8}$$

where the compact notation $M \circ \boldsymbol{T} \equiv M(\boldsymbol{T}(\boldsymbol{x}; \boldsymbol{\mu}))$ was introduced, and all function arguments were omitted. Our aim is to find an approximation $\widetilde{\boldsymbol{H}}$ to $\boldsymbol{H}(\hat{\boldsymbol{\mu}})$, whose inverse can be used as a preconditioning matrix $\boldsymbol{P}$.

When F and M are images of the same modality, we can exploit the fact that $M \circ \boldsymbol{T}$ will be approximately equal to F after successful registration: $F(\boldsymbol{x}) \approx M(\boldsymbol{T}(\boldsymbol{x}; \hat{\boldsymbol{\mu}}))$. With this approximation, the following two identities are derived:

$$F - M \circ \boldsymbol{T} = 0, \quad \frac{\partial M}{\partial \boldsymbol{x}} = \left[\frac{\partial \boldsymbol{T}'}{\partial \boldsymbol{x}} \right]^{-1} \frac{\partial F}{\partial \boldsymbol{x}}. \tag{9}$$

Substituting (9) in (8) yields the following approximation of the Hessian at $\hat{\boldsymbol{\mu}}$:

$$\widetilde{\boldsymbol{H}} = \frac{2}{V} \sum_{\boldsymbol{x} \in \Omega_F} \frac{\partial \boldsymbol{T}'}{\partial \boldsymbol{\mu}} \left[\frac{\partial \boldsymbol{T}'}{\partial \boldsymbol{x}} \right]^{-1} \frac{\partial F}{\partial \boldsymbol{x}} \frac{\partial F'}{\partial \boldsymbol{x}} \left[\frac{\partial \boldsymbol{T}}{\partial \boldsymbol{x}} \right]^{-1} \frac{\partial \boldsymbol{T}}{\partial \boldsymbol{\mu}}. \tag{10}$$

Since $\hat{\boldsymbol{\mu}}$ is unknown, we approximate the terms $\partial \boldsymbol{T}/\partial \boldsymbol{\mu}(\boldsymbol{x}; \hat{\boldsymbol{\mu}})$ and $\partial \boldsymbol{T}/\partial \boldsymbol{x}(\boldsymbol{x}; \hat{\boldsymbol{\mu}})$ by $\partial \boldsymbol{T}/\partial \boldsymbol{\mu}(\boldsymbol{x}; \boldsymbol{\mu}_0)$ and $\partial \boldsymbol{T}/\partial \boldsymbol{x}(\boldsymbol{x}; \boldsymbol{\mu}_0)$, respectively. By setting $\partial \boldsymbol{T}/\partial \boldsymbol{x} \approx \boldsymbol{I}$ (assuming small deformations) the following expression is finally obtained:

$$\widetilde{\boldsymbol{H}} = \frac{2}{V} \sum_{\boldsymbol{x} \in \Omega_F} \frac{\partial \boldsymbol{T}'}{\partial \boldsymbol{\mu}}(\boldsymbol{x}; \boldsymbol{\mu}_0) \frac{\partial F}{\partial \boldsymbol{x}}(\boldsymbol{x}) \frac{\partial F'}{\partial \boldsymbol{x}}(\boldsymbol{x}) \frac{\partial \boldsymbol{T}}{\partial \boldsymbol{\mu}}(\boldsymbol{x}; \boldsymbol{\mu}_0). \tag{11}$$

We then define the preconditioning matrix by $\boldsymbol{P} \equiv [\widetilde{\boldsymbol{H}} + \beta\lambda\boldsymbol{I}]^{-1}$, with $0 \leq \beta \leq 1$ a user-defined factor, and $\lambda > 0$ the maximum eigenvalue of $\widetilde{\boldsymbol{H}}$, which can be estimated using an iterative block-Lanczos method. By adding the identity matrix, the condition number of $\boldsymbol{P}$ is limited to $(\beta+1)/\beta$, as a safeguard in case of ill-conditioned $\widetilde{\boldsymbol{H}}$, which may arise in nonrigid registration problems. In rigid registration problems, $\beta \downarrow 0$ is a valid choice. Instead of explicitly computing the matrix inverse, a Cholesky decomposition $\widetilde{\boldsymbol{H}} + \beta\lambda\boldsymbol{I} = \boldsymbol{L}\boldsymbol{L}'$ is used, allowing fast application of $\boldsymbol{P} \equiv \boldsymbol{L}^{'-1}\boldsymbol{L}^{-1}$ in each iteration. Note that $\widetilde{\boldsymbol{H}}$ is independent of $\boldsymbol{\mu}_k$, so $\widetilde{\boldsymbol{H}}$ and the Cholesky decomposition only need to be computed once.

2.3 Preconditioned Adaptive Stochastic Gradient Descent

A preconditioned version of ASGD, called PASGD, is derived in this subsection, which combines the adaptive step size mechanism of (4) with the preconditioner.

First we show that the PSGD method (5) is in fact a standard RM algorithm performed in a different parameter space. Let us introduce a new parameter vector $\boldsymbol{\nu} = \boldsymbol{L}'\boldsymbol{\mu}$. The original minimisation problem (1) is equivalent to:

$$\hat{\boldsymbol{\nu}} = \arg\min_{\boldsymbol{\nu}} \mathcal{D}(\boldsymbol{\nu}), \quad \text{with } \mathcal{D}(\boldsymbol{\nu}) \equiv \mathcal{C}(\boldsymbol{L}^{'-1}\boldsymbol{\nu}). \tag{12}$$

Define $\boldsymbol{h}_k \equiv \partial\mathcal{D}/\partial\boldsymbol{\nu}(\boldsymbol{\nu}_k) = \boldsymbol{L}^{-1}\partial\mathcal{C}/\partial\boldsymbol{\mu}(\boldsymbol{\mu}_k) = \boldsymbol{L}^{-1}\boldsymbol{g}_k$. The basic RM scheme (2) in terms of $\boldsymbol{\nu}$ reads $\boldsymbol{\nu}_{k+1} = \boldsymbol{\nu}_k - \gamma(k)\tilde{\boldsymbol{h}}_k$. Substituting $\tilde{\boldsymbol{h}}_k = \boldsymbol{L}^{-1}\tilde{\boldsymbol{g}}_k$ and $\boldsymbol{\nu} = \boldsymbol{L}'\boldsymbol{\mu}$ yields (13), and multiplying both sides of the equation by $\boldsymbol{L}^{'-1}$ yields (14):

$$\boldsymbol{L}'\boldsymbol{\mu}_{k+1} = \boldsymbol{L}'\boldsymbol{\mu}_k - \gamma(k)\boldsymbol{L}^{-1}\tilde{\boldsymbol{g}}_k, \tag{13}$$

$$\boldsymbol{\mu}_{k+1} = \boldsymbol{\mu}_k - \gamma(k)\boldsymbol{L}^{'-1}\boldsymbol{L}^{-1}\tilde{\boldsymbol{g}}_k, \tag{14}$$

in which we can recognise the preconditioner $\boldsymbol{P} \equiv \boldsymbol{L}^{'-1}\boldsymbol{L}^{-1}$.

Doing the same for the ASGD scheme (4) results in the PASGD method:

$$\boldsymbol{\mu}_{k+1} = \boldsymbol{\mu}_k - \gamma(t_k)\boldsymbol{P}\tilde{\boldsymbol{g}}_k, \quad t_{k+1} = [t_k + \text{sigmoid}(-\tilde{\boldsymbol{g}}_k'\boldsymbol{P}\tilde{\boldsymbol{g}}_{k-1})]^+. \tag{15}$$

Since the PASGD scheme is essentially an ordinary ASGD method in the domain of $\boldsymbol{\nu}$, all convergence conditions given in [3] remain valid, but should be interpreted in the space of $\boldsymbol{\nu}$. To estimate a, A, and α, a similar procedure as presented in [3] can therefore be used. The main idea is explained here, but the exact derivation is omitted for brevity. As suggested in [3], we set $\alpha = 1$ and $A = 20$. For a, the following expression was proposed in [3]:

$$a \equiv a_{\text{MAX}}\eta \equiv a_{\text{MAX}}\text{E}||\boldsymbol{g}||^2 / \left(\text{E}||\boldsymbol{g}||^2 + \text{E}||\boldsymbol{g} - \tilde{\boldsymbol{g}}||^2\right), \tag{16}$$

where $a_{\text{MAX}} \equiv 2A/\lambda$, with λ defined as the maximum eigenvalue of the Hessian of the cost function, and E denotes the expectation. The η factor intuitively takes into account that the step size should be reduced with increasing approximation error of $\tilde{\boldsymbol{g}}$. Whereas in [3] λ was unknown and had to be estimated from a user-defined maximum allowed voxel displacement per iteration, in this work

we can simply use $\lambda = 1$, as it can be derived that $\frac{\partial^2 \mathcal{D}}{\partial \boldsymbol{\nu} \partial \boldsymbol{\nu}} \approx \boldsymbol{I}$. Applying the reparametrisation to the definition of η changes the $\mathrm{E}||\boldsymbol{g}||^2$ terms to $\mathrm{E}\boldsymbol{g}'\boldsymbol{P}\boldsymbol{g}$ (and similar for $\boldsymbol{g} - \tilde{\boldsymbol{g}}$). Like in [3], the expectations in the definition of η are replaced by their empirical estimates. Since evaluating $\boldsymbol{g}$ (the exact cost function derivative) is an expensive operation, we apply the approximation $||\boldsymbol{g} - \tilde{\boldsymbol{g}}||^2 \approx ||\boldsymbol{g}||^2 + ||\tilde{\boldsymbol{g}}||^2$ to decouple the $\boldsymbol{g}$ and $\tilde{\boldsymbol{g}}$ terms. The term $\mathrm{E}\boldsymbol{g}'\boldsymbol{P}\boldsymbol{g}$ is estimated by a single measurement of $\boldsymbol{g}'\boldsymbol{P}\boldsymbol{g}$ (evaluated at $\boldsymbol{\mu}_0$), and $\mathrm{E}\tilde{\boldsymbol{g}}'\boldsymbol{P}\tilde{\boldsymbol{g}}$ is estimated by averaging a few (N) measurements at random positions $\boldsymbol{\mu}_n$ generated according to $\boldsymbol{\mu}_n - \boldsymbol{\mu}_0 \sim L^{'-1}\mathcal{N}(\boldsymbol{0}, \sigma_1^2 \boldsymbol{I})$, with $\sigma_1^2 = \mathrm{E}\boldsymbol{g}\boldsymbol{P}\boldsymbol{g}/Q$.

3 Experiments and Results

The proposed PASGD method was compared to the standard RM method, ASGD, and to a deterministic LBFGS quasi-Newton (QN) method [8]. All algorithms were integrated in `elastix` [4], an open source software package for image registration. The Cholesky decomposition was implemented using the CHOLMOD library [1]. Two applications were considered: rigid registration of 3D functional MR images (fMRI) and nonrigid registration of 3D CT chest scans.

3.1 Rigid Registration of fMRI Series

Eight fMRI time-series were acquired in the context of research on brain-computer interfaces (BCI). Seven series were recorded with a 2D EPI sequence; one with 3D PRESTO. Each time-series consisted of $\tau \approx$ 200-400 scans. The image size was $64 \times 64 \times$ [20-40], with $4 \times 4 \times 4$ mm voxels. In the BCI experiments, real-time rigid registration of each scan to the first scan is required [7]. For our experiment, scans at time points $t = 0, 1, 100, 200, (300,)$ and τ were selected. All scans with $t > 0$ were registered to the scan at $t = 0$, which resulted in a total of 37 registrations. Since the head's motion was small in most cases, the experiments were repeated with an extra initial offset to make the registration problem more challenging. The applied translations and rotations were drawn from a uniform distribution between ± 8 mm and $\pm 6°$, respectively.

The parameter vector $\boldsymbol{\mu}$ was formed by $\boldsymbol{t}$ and $\boldsymbol{S}\boldsymbol{\theta}$, where $\boldsymbol{t}$ is the translation vector, $\boldsymbol{\theta}$ represents the Euler angles, and $\boldsymbol{S}$ is a diagonal matrix with elements:

$$s_{ii} = \left(\int_{\Omega_F} \left|\left| \frac{\partial \boldsymbol{T}}{\partial \theta_i}(\boldsymbol{x}; \boldsymbol{\mu}_0) \right|\right|^2 d\boldsymbol{x} \Big/ \int_{\Omega_F} d\boldsymbol{x} \right)^{-\frac{1}{2}}. \tag{17}$$

Matrix $\boldsymbol{S}$ scales the rotation parameters by the average voxel displacement caused by a small perturbation of the rotation angle. This brings the values of the elements of $\boldsymbol{\mu}$ approximately in the same range, thus avoiding a very badly scaled cost function. For PASGD, the rescaling step was omitted ($\boldsymbol{S} = \boldsymbol{I}$), since the preconditioning matrix already should take care of this. To compute $\widetilde{\boldsymbol{H}}$, $V = 50\,000$ samples were used. To compute $\tilde{\boldsymbol{g}}_k$, $V = 2000$ random samples were used (except for QN, which used all voxels). For $\boldsymbol{P}$, $\beta = 10^{-7}$ was used.

Table 1. Results of the experiments with fMRI series

	No additional offset		With additional offset	
	$\|\|\Delta\boldsymbol{T}\|\|$ [mm] avg ± sd	nr. of it. avg ± sd	$\|\|\Delta\boldsymbol{T}\|\|$ [mm] avg ± sd	nr. of it. avg ± sd
QN	0.00 ± 0.00	3 ± 2	0.01 ± 0.02	8 ± 2
RM $a = 0.0025$	0.14 ± 0.19	48 ± 68	1.15 ± 1.48	216 ± 63
RM $a = 0.005$	0.14 ± 0.14	25 ± 39	0.30 ± 0.40	156 ± 82
RM $a = 0.01$	0.13 ± 0.13	17 ± 30	0.14 ± 0.14	80 ± 54
RM $a = 0.02$	0.14 ± 0.13	40 ± 52	0.14 ± 0.13	63 ± 41
RM $a = 0.04$	0.41 ± 0.61	134 ± 83	0.38 ± 0.55	134 ± 84
ASGD	0.14 ± 0.14	15 ± 30	0.15 ± 0.15	102 ± 53
PASGD	0.14 ± 0.13	11 ± 33	0.13 ± 0.13	19 ± 36

The number of iterations was set to $K = 250$. The RM method was tested for $a \in \{0.0025, 0.005, 0.01, 0.02, 0.04\}$, with $A = 50$ and $\alpha = 0.6$.

The result of QN without additional offset was treated as gold standard. The differences between the gold standard's transformation $\boldsymbol{T}(\boldsymbol{x}; \hat{\boldsymbol{\mu}}_{\text{gold}})$ and the other methods' transformations were computed to verify that all methods converged to the same solution. Table 1 presents for each method the average and standard deviation of $||\Delta\boldsymbol{T}|| \equiv ||\boldsymbol{T}(\boldsymbol{x}; \hat{\boldsymbol{\mu}}_{\text{gold}}) - \boldsymbol{T}(\boldsymbol{x}; \hat{\boldsymbol{\mu}})||$ over all $\boldsymbol{x}$ in all images. Both with and without additional offset, all differences were smaller than a voxel.

To compare the convergence rates, we chose ASGD as a baseline, and measured the performance improvement with respect to that method. For each method, we counted the number of iterations before $\mathcal{C}(\boldsymbol{\mu}_k) \leq 1.01 \cdot \mathcal{C}(\hat{\boldsymbol{\mu}}_{\text{ASGD}})$ occurred for the first time for at least 5 consecutive iterations. Note that $\mathcal{C}(\boldsymbol{\mu}_k)$ was calculated based on *all* voxels of the fixed image (not only the $V = 2000$ random samples that were used to compute $\tilde{\boldsymbol{g}}_k$). The results are summarised in Table 1 (nr. of it.) by the average and standard deviation over all images. The QN method required the least number of iterations, as expected, since it uses all voxels in each iteration (which makes it more expensive per iteration). RM performed best with $a \approx 0.01$, which gave similar results as ASGD. The PASGD method outperformed RM and ASGD.

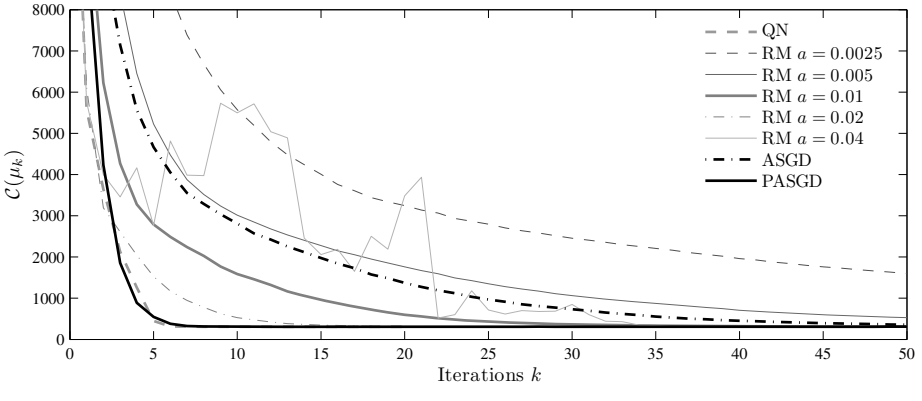

Fig. 1. Cost function plot for one of the fMRI experiments with additional offset

In Fig. 1, for one example image pair, the cost function $\mathcal{C}(\boldsymbol{\mu}_k)$ is plotted as a function of k for all methods. All 37 graphs were visually inspected and the pattern was fairly consistent. In a few cases the RM methods with $a \in \{0.02, 0.04\}$ suffered from instabilities (heavily fluctuating cost function values), indicating that the step sizes were too large.

3.2 Nonrigid Registration of CT Chest Scans

CT chest scans of five patients were obtained from the Department of Radiology, UMC Utrecht. For each patient a baseline and a follow-up scan, taken 3-9 months later, were available. Each scan was manually cropped around the right lung, and downsampled by a factor of two, which gave images of about $120 \times 160 \times 220$ voxels, with voxel size approximately $1.4 \times 1.4 \times 1.4$ mm. As a region of interest for registration, a dilated (kernel radius 10) lung segmentation was used.

Each follow-up scan was registered to the baseline scan using a B-spline transformation model [10]. Initial experiments showed that a regularisation term needs to be added to the cost function, to avoid foldings. The sum of second order spatial derivatives of the deformation field was used as a regularisation term, with a weighting factor of $5 \cdot 10^7$. The Hessian at $\boldsymbol{\mu}_0$ of this term was also included in the preconditioner. A three-level multiresolution strategy was employed. The distance between the B-spline control points was halved in each resolution level, such that at the final level the control points were spaced 20 mm in each direction. The

Table 2. Results of the experiments with CT chest scans. R3 is the finest resolution.

	$\|\|\Delta \boldsymbol{T}\|\|$ [mm] avg ± sd	R1 nr. of it. avg ± sd	R2 nr. of it. avg ± sd	R3 nr. of it. avg ± sd
QN	0.00 ± 0.00	11 ± 0	5 ± 3	3 ± 2
ASGD	1.26 ± 2.77	242 ± 3	214 ± 9	155 ± 28
PASGD $\beta = 1$	0.72 ± 1.47	52 ± 26	22 ± 11	25 ± 18
PASGD $\beta = 0.1$	0.36 ± 0.51	28 ± 36	5 ± 3	13 ± 10
PASGD $\beta = 0.01$	0.36 ± 0.38	39 ± 55	5 ± 3	13 ± 11

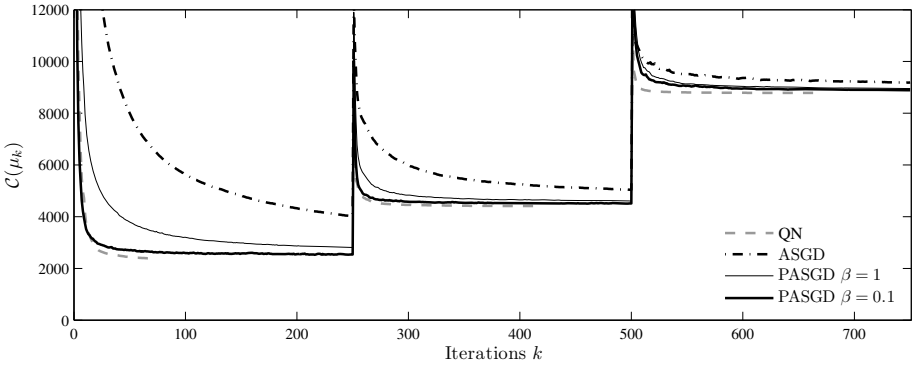

Fig. 2. Convergence results for one of the CT image pairs

images were smoothed using a Gaussian kernel with standard deviation of 2, 1, and 0.5 times the voxel size, for each resolution level respectively. The matrix $\widetilde{\boldsymbol{H}}$ was computed using $V = 100\,000$. PASGD was tested with $\beta \in \{1, 0.1, 0.01\}$. The tests with RM were omitted in this section.

For evaluation the same approach was followed as in the fMRI experiments. Table 2 summarises the evaluation results. The convergence results (nr. of it.) were calculated for each resolution separately (R1-R3). The numerical results in Table 2 indicate that PASGD achieved faster convergence than ASGD. The influence of β was moderate. Figure 2 plots the cost function series for one of the image pairs. PASGD with $\beta = 0.01$ was omitted for clarity, since it was very similar to $\beta = 0.1$.

4 Conclusion

The experiments with fMRI and CT data show that the proposed preconditioning technique has a beneficial effect on the rate of convergence, both in rigid and nonrigid registration problems. The PASGD method is, just as RM and ASGD, designed to work with stochastic estimates of the cost function derivatives, which leads to low computational costs per iteration [5]. The PASGD method couples this with a good rate of convergence by using second order information of the cost function.

References

1. Chen, Y., Davis, T., Hager, W., Rajamanickam, S.: Algorithm 887: CHOLMOD, supernodal sparse cholesky factorization and update/downdate. ACM Trans. Math. Softw. 35(3), 1–14 (2008)
2. Hill, D.L.G., Batchelor, P.G., Holden, M., Hawkes, D.J.: Medical image registration. Phys. Med. Biol. 46(3), R1–R45 (2001)
3. Klein, S., Pluim, J.P.W., Staring, M., Viergever, M.A.: Adaptive stochastic gradient descent optimisation for image registration. Int. J. Comput. Vis. 81(3), 227–239 (2009)
4. Klein, S., Staring, M., Murphy, K., Viergever, M.A., Pluim, J.P.W.: `elastix`: a toolbox for intensity-based medical image registration. IEEE Trans. Med. Imag. 29(1), 196–205 (2010)
5. Klein, S., Staring, M., Pluim, J.P.W.: Evaluation of optimization methods for nonrigid medical image registration using mutual information and B-splines. IEEE Trans. Image Process. 16(12), 2879–2890 (2007)
6. Maintz, J.B.A., Viergever, M.A.: A survey of medical image registration. Med. Image Anal. 2(1), 1–36 (1998)
7. Mathiak, K., Posse, S.: Evaluation of motion and realignment for functional magnetic resonance imaging in real time. Magn. Reson. Med. 45, 167–171 (2001)
8. Nocedal, J., Wright, S.J.: Numerical optimization. Springer, New York (1999)
9. Robbins, H., Monro, S.: A stochastic approximation method. Ann. Math. Stat. 22(3), 400–407 (1951)
10. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L.G., Leach, M.O., Hawkes, D.J.: Nonrigid registration using free-form deformations: Application to breast MR images. IEEE Trans. Med. Imag. 18(8), 712–721 (1999)

Random Walks for Deformable Image Registration

Dana Cobzas and Abhishek Sen

Computing Science, University of Alberta, Canada

Abstract. We introduce a novel discrete optimization method for non-rigid image registration based on the random walker algorithm. We discretize the space of deformations and formulate registration using a Gaussian MRF where continuous labels correspond to the probability of a point having a certain discrete deformation. The interaction (regularization) term of the corresponding MRF energy is convex and image dependent, thus being able to accommodate different types of tissue elasticity. This formulation results in a fast algorithm that can easily accommodate a large number of displacement labels, has provable robustness to noise and a close to global solution. We experimentally demonstrate the validity of our formulation on synthetic and real medical data.

1 Introduction

Image registration is a fundamental problem in medical imaging, central for many clinically relevant applications like statistical studies on a population of patients, analysis of disease progress and multi-modality fusion for better diagnosis and treatment. The registration problem can be formalized as finding the optimal transformation that aligns a source with a target image, based on a similarity score. Depending on the type of transformation, registration methods can be classified into global (rigid, affine) and local (non-linear, non-rigid). Global registration methods involve few parameters to be optimized and are thus well-posed, being constrained in the parameter space. Non-rigid registration methods estimate a dense deformation field that defines, for every location, a vector that locally aligns the two images. This is an inherently ill-posed problem due to the high dimensionality of the parameter space and therefore relies on regularization.

Several ways of imposing regularization have been proposed in the literature [1]. The popular free form deformation model (FFD) [2] restricts the parameter space to a set of control points that define a smooth interpolation field for the rest of the image. Alternative methods explicitly add a regularization term (e.g. fluid, elastic) in the registration energy, that is either optimized together with the data matching term [3,4], or applied as a separate smoothing process (demon's [5]). One other way of imposing regularization is to restrict the space of deformations to a Sobolev space [6]. Some effort has been made to adapt the regularization of deformations to local image content [7,8]. This is particularly important considering that different tissue deform differently and parts of the image might contain an abnormality that does not match the atlas.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 557–565, 2011.

Nonlinear registration is traditionally formulated in a continuous domain and the optimal transformation is recovered using gradient descent. This estimation is often slow and suboptimal due to non-convexity of the energy functional that is optimized. Recently, few works have formulated deformable registration as a discrete labeling problem [9,10]. The space of deformations is discretized and the registration energy is formulated as a Markov random field (MRF) optimization. If the interaction energy is submodular, a graph cut method guarantees a good quality solution [11,10]. For more complex interaction terms, Glocker et al. [9] proposed a linear programming method that uses the primal dual principle.

In this paper we proposed an alternative discrete formulation of the registration problem based on the random walker algorithm [12,13]. Our formulation is equivalent to a Gaussian MRF with an interaction (regularization) term that is convex and image dependent, thus being able to accommodate different types of tissue elasticity. Following the formulation from [13], we incorporate data similarity terms as 'priors' for the displacement labels. Unlike the graph cuts [10] or the primal dual method [9] that would only guarantee a good quality solution for the discrete registration, the random walker method finds a unique global minimum. Probabilities of a particular displacement label are calculated by solving a combinatorial Laplace equation. The random walker formulation of the registration problem results in a fast algorithm that can easily accommodate a large number of displacement labels and has provable robustness to noise [12].

2 Methods

2.1 An Energy Formulation of Deformable Registration

Let I and J respectively be the reference (target) and the floating (source) d-dimensional images $I, J : \Omega \to \mathbb{R}, \Omega \subset \mathbb{R}^d$. Image registration seek an optimal transformation $T : \Omega \to \Omega$ that aligns the two images based on a similarity measure. In deformable registration, T is usually expressed in terms of a displacement field u as $T = \mathrm{Id} + u$, with the identity operator Id. u is found as the minimum of an energy functional:

$$u^* = \mathrm{argmin}_u E_D(I, J \circ T) + \alpha E_R(u) \tag{1}$$

where E_D is a data term that measures the similarity between the two images and E_R is the regularization energy term. Expanding the two energy terms, and denoting the similarity measure with Φ and an image-dependent (adaptable) regularization function with Ψ we get:

$$u^* = \mathrm{argmin}_u \int_{\mathbf{x}\in\Omega} \Phi\left(I(\mathbf{x}), J(\mathbf{x}+u(\mathbf{x}))\right) d\mathbf{x} + \alpha \int_{\mathbf{x}\in\Omega} \Psi(\nabla J(\mathbf{x}), \nabla u(\mathbf{x})) d\mathbf{x} \tag{2}$$

2.2 Discrete Formulation for the Random Walker Algorithm

Regarding I and J as discrete representations for the target and source image, we next formulate registration as a discrete optimization. We consider a discrete

set of labels $\mathcal{L} = \{u^1, u^2, ..., u^K\}$ corresponding to a quantized version of the deformation space $u^i \in \mathcal{D} = \{\mathbf{d}^1, \mathbf{d}^2, ..., \mathbf{d}^K\}$. The registration problem becomes a labeling problem that seeks to assign an optimal label for each image location. A common model for representing such problems are MRFs. The pixel locations of the image are mapped on a graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, where $\mathcal{N}$ represents the set of nodes (image locations) and $\mathcal{E}$ represents a neighboring system of the image grid (typically 4 or 8 in 2D). The labeling problem is then solved by minimizing:

$$E(\mathbf{u}) = \sum_{i \in \mathcal{N}} \Phi_i(u_i) + \alpha \sum_{(i,j) \in \mathcal{E}} \Psi_{ij}(u_i, u_j) \tag{3}$$

where $u_i \in \mathcal{L}$ denotes the displacement label for location i, $\Phi_i(.)$ is the unary potentials representing the data term and $\Psi_{ij}(.,.)$ are the pairwise potentials representing the interaction (smoothing) term. Due to the independence assumption of the unary data potentials Φ_i, we require a point-wise similarity score. This constraint was relaxed in [9] by approximating a local score (e.g. mutual information) in neighborhoods defined by control points. There are few formulations of the traditional MRF for solving the discrete registration problem. When the smoothing term Ψ_{ij} is a metric, the MRF energy can efficiently be optimized using graph cuts [10,8]. For more complex interaction terms, Glocker et al. [9] use a linear programming method (based on the primal-dual principle).

We make few modifications to the traditional MRF from Equation 3 to be able to map the registration problem to a random walker with priors [13]. First, we relax the labeling system to continuous variables u_i^k that represent the probability of node i having the label u^k. Next, we consider a Gaussian MRF, where the interaction term has the form $\Psi_{ij}(u_i^k, u_j^k) = w_{ij}(u_i^k - u_j^k)^2$, with w_{ij} being an image dependent edge weight (e.g. $w_{ij} = \exp(-\beta(J_i - J_j)^2)$ where J_i represents the image intensity for location i). Last, for defining the data term $\Phi_i(u_i^k)$, we consider a set of real-valued nodewise priors λ_i^k that represents the probability density that the displacement vector at location i has the value $\mathbf{d}^k$, $\lambda_i^k = \exp(-\gamma(I_i - J_{i+\mathbf{d}^k})^2)$. By $J_{i+\mathbf{d}^k}$ we denoted the intensity of image J at location i displaced with $\mathbf{d}^k$.

With these three modifications, following [13], we can define the registration energy corresponding to the label u^k as the continuous-valued Gaussian MRF:

$$E^k(u^k) = \sum_{i \in \mathcal{N}} \left(\sum_{l=1, l \neq k}^{K} \lambda_i^l (u_i^k)^2 + \lambda_i^k (1 - u_i^k)^2 \right) + \alpha \sum_{(i,j) \in \mathcal{E}} w_{ij}(u_i^k - u_j^k)^2 \tag{4}$$

While space does not allow a rigorous interpretation of the above equation (details in [13]), intuitively, we see that when λ_i^k is large (meaning that the displacement $\mathbf{d}^k$ matches the similarity score at location i) the data energy term encourages high probability values for label u^k and small probability values for all other labels $u^l, l \neq k$. Compacting notations, we denote L, the combinatorial Laplacian matrix of the graph:

Algorithm 1. Random walker nonlinear registration

1. generate multi-resolution images $I_1(=I), I_2, I_3$ and J_3, with a factor $r(=2)$
2. **for** i=3:1 **do**
3. define a set of discrete labels $\{\mathbf{d}_1, \ldots, \mathbf{d}_K\}$
4. setup the image graph, compute Laplacian L and priors Λ
5. solve for deformations labels $\mathbf{u}^k$ for every k (Eq. 7)
6. assign $u_i = \mathbf{d}^k$ where $k = \text{argmax}(u_i^1, \ldots, u_i^K)$
7. **if** $i > 1$ **then**
8. compute source at next level $J_{i-1} = \text{warp}(\text{upsample}(J_i), \text{interp}(ru_i))$
9. **end if**
10. **end for**
11. registered image = $\text{warp}(J_1, u_1)$

$$L_{ij} = \begin{cases} d_i = \sum_k w_{ik} & \text{degree of } \ i \text{ if } i = j \\ -w_{ij} & \text{if } (i,j) \in \mathcal{E} \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

and by $\Lambda^k = \text{diag}(\lambda^k)$, the matrix having the values of λ^k on the diagonal. Equation 4 can be written as:

$$E^k(\mathbf{u}^k) = \sum_{l=1, l\neq k}^{K} \mathbf{u}^{kT} \Lambda^l \mathbf{u}^k + (1-\mathbf{u}^k)^T \Lambda^k (1-\mathbf{u}^k) + \alpha \mathbf{u}^{kT} L \mathbf{u}^k \tag{6}$$

where $\mathbf{u}^k$ collects all nodes probabilities for label k in a vector. The minimum of this energy is obtained when $\mathbf{u}^k$ is the solution of this equation:

$$\left(\alpha L + \sum_{l=1}^{K} \Lambda^l\right) \mathbf{u}^k = \lambda^k \tag{7}$$

which is a combinatorial Laplace equation of a graph augmented with a node for each label u^ksuch that the weights on the new edges (k,i) have value λ_i^k. The combined matrix on the left side of the equation is guaranteed to be positive definite and therefore the equation has a unique global solution that gives the nodes probabilities for the displacement labels $\mathbf{u}^k$.

2.3 Multi-resolution Framework and Implementation Details

The random walker algorithm is computationally expensive as well as memory expensive. The number of equation systems to be solved is the same as the number of displacement labels, and each of these equations has the number of variables equal to the number of pixels in the image. As the number of displacement labels can be quite large, especially in 3D, not only the solution of the linear systems is time-consuming, but also their assembly. We obtained an efficient approximation of the solution using a multi-resolution framework. Due to

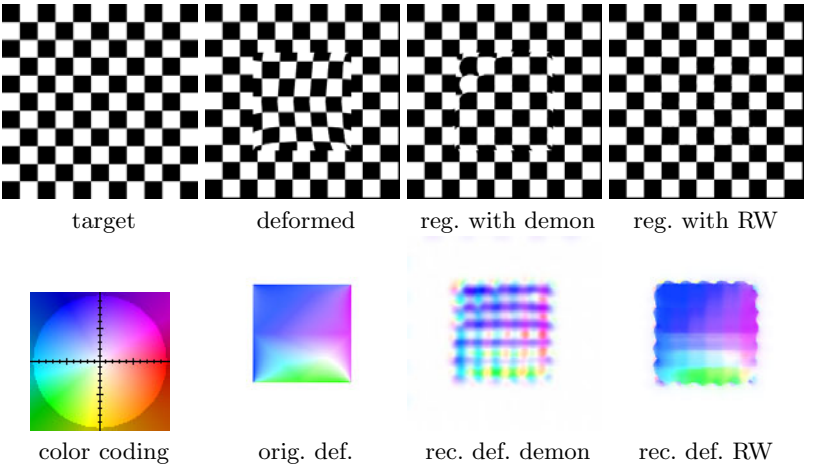

Fig. 1. Comparative results of recovered deformations for **checkerboard** image. From left to right, the top row shows the target image, the deformed image and the SSD error of the registered images with demon's and RW methods. Bottom row shows the angle and magnitude color coding convention, the ground truth deformations and the recovered deformations using demon's and RW methods.

this multiresolution approach, even though the RW solution is optimal at each resolution level, the composite solution is no longer guaranteed to be optimal.

The multi-resolution images were obtained by downsampling the original images based on nearest neighbor interpolation. We defined deformations in an incremental way propagating deformations obtained at a lower resolution to the next higher resolution level. This approximation is carried out by an interpolation based on Delaunay triangulation after scaling the low resolution field by the multi-resolution scale factor. At each resolution we compute the remaining deformations by solving the sparse linear equation system 7 for each discrete label (in practice only $K-1$ times as we impose sum of unity for the probabilities $\sum_k u_i^k = 1$). At each level the magnitude and number of expected deformations encoded as discrete labels $\{\mathbf{d}^1, \mathbf{d}^2, ..., \mathbf{d}^K\}$ decreases. The final displacements u_i are taken as the ones with maximum probability among all labels $u_i = \mathbf{d}^k$ where $k = \text{argmax}(u_i^1, \ldots, u_i^K)$. The algorithm is summarized in Algorithm 1.

3 Experiments

We present results of our experiments on real and synthetic data. We compared the performance of the proposed RW registration with a traditional demon's implementation [14]. The set of parameters are optimized to achieve best SSD scores (ex. $\alpha = 1$, Gaussian weights w_{ij} with $\beta = 0.005$, priors λ_i^k with $\gamma = 10^{-5}$ for real data). For all experiments we used 3 levels of resolution generated by a scale factor of 0.5 and 0.25. The number of displacement labels at each level

Table 1. Comparative numerical results for the three datasets

	Deformation field err				SSD err		Dice coef	
	ang.		magn.					
	D	RW	D	RW	D	RW	D	RW
checkerboard	1.58 ± 1.20	0.45 ± 0.97	1.12 ± 1.54	1.65 ± 2.89	10.88	5.01	-	-
brain MRI	-	-	-	-	8.82	3.38	WM: 0.82	0.84
							GM: 0.79	0.81
							CSF: 0.83	0.84
abdominal CT	-	-	-	-	15.46	9.64	muscle: 0.62	0.79

is dependent on the initial image size (ex. low-high resolution : $60, 40, 30$ for an initial 256×256 image, corresponding to a displacement range of about $[-15, 15]$, $[-10, 10]$, $[-7, 7]$ pixels, respectively). The experiments were run on MATLAB using Intel Core 2 Duo Processor of 2.10 GHz with 4 GB RAM. The algorithm took about 200 sec. to complete on a 256×256 image. Most time was taken by MATLAB's sparse linear equation solver.

3.1 Quality of Recovered Deformations

For testing the accuracy of the recovered deformation field we synthetically deformed a **checkerboard** image with a known deformation field. We tested how deformations are recovered by the RW registration method and the demon's algorithm. We measured the angular and magnitude errors of the recovered deformation fields as well as the SSD error between target and registered images. Qualitative results are presented in Figure 1 and numerical scores in Table 1. We color coded deformations using the same convention as for optic flow as shown in bottom-left of Figure 1. We notice that the recovered deformation fields using RW registration has less artifacts, and the recovered deformations are closer in orientation to the original ones. The magnitude of recovered deformations using RW is slightly larger than the ones recovered using demons, probably due to the fact that regularization is imposed at the energy level for which we obtain a global solution as opposed to demon's iterative approach.

3.2 Results on Real Medical Data

For the experiments with real data, we again compared our RW method and demon's algorithm. To quantify results we measured SSD error between registered images and the target image. Also, both datasets had some ground truth segmentations (WM/GM/CSF for brain MRI data and muscle in the abdominal CT data). We calculated the dice coefficients between ground truth segmentations in the target image, and segmentations from the source image warped on the space of the target image using the recovered deformation field. Note that a larger, closer to 1 value for a dice coefficient indicates a better segmentation.

The first experiment uses the **brain MRI** dataset from Internet Brain Segmentation Repository[1]. We performed registration between two patients, both

[1] http://www.cma.mgh.harvard.edu/ibsr/data.html

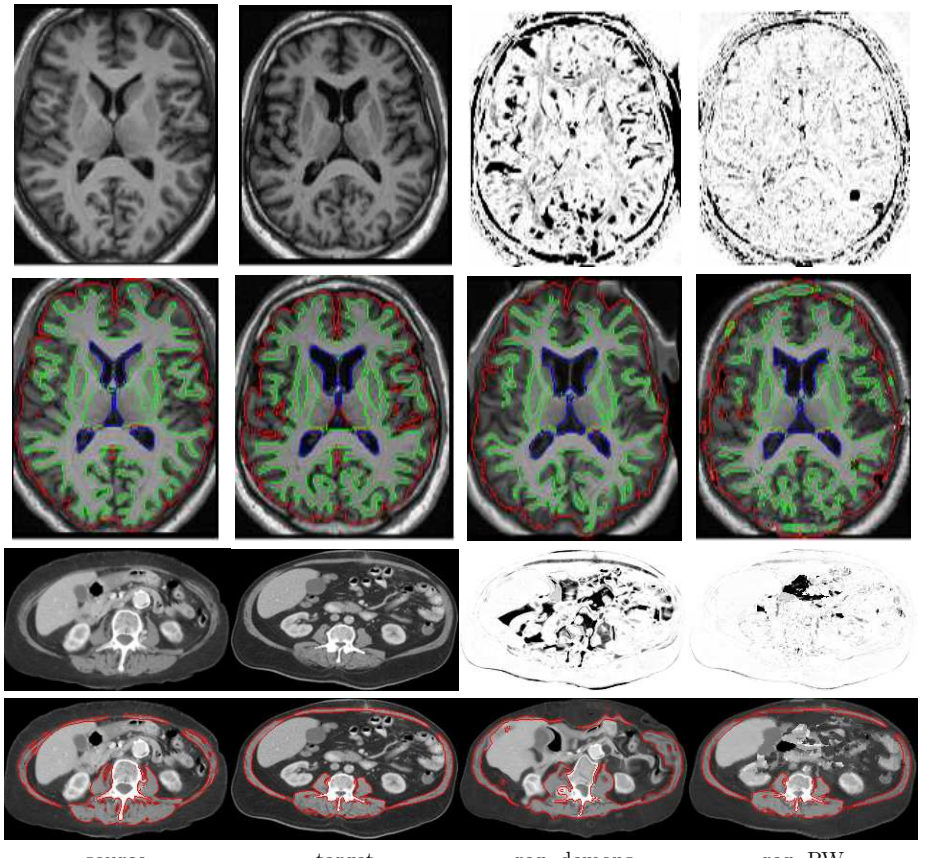

Fig. 2. Results for real data. Left to right: (row 1) **brain MRI** data - source, target, SSD error between target and registered image (bright for small, dark for large errors); (row 2) segmentations on source and target images,warped segmentations on registered images; (row 3,4) same for **abdominal CT** data.

from the database. The dataset contains WM,GM and CSF labels. As shown in the first two rows of Figure 2 and in Table 1, the RW method has better performance both visually and quantified compared to the demon's algorithm.

In a second experiment we registered two **abdominal CT** images from a local cancer institute. Notice the large difference between the two datasets. Demon's method was not able to recover the large deformations but the RW's registration that recovers a global minimum was much better. As the muscle exist in both images, and is therefore the part that is expected to match, we measured the dice coefficient on the muscle segmentation. Ground truth segmentation was provided by a medical student. Figure 2 and Table 1 present the results.

4 Discussion

We have presented a discrete method for non-rigid image registration based on the random walker method. The new formulation has several advantages: at each resolution level, we globally minimize a convex energy, with a regularization term that is image dependent thus being able to accommodate different elasticity depending on the tissue type.

As future work, we are looking into a more efficient implementation of our method. One option is to use a lower-dimensional deformation model like the FFD model that computes displacements only at control points. This approach would also allow approximating non-local data potentials (e.g. mutual information) in neighborhoods around the control points (similar to [9]). This technically violates the independence assumption, but practically the loss of optimality at a particular resolution may be compensated by the richer non-local measure.

References

1. Zikic, D., Kamen, A., Navab, N.: Unifying characterization of deformable registration methods based on the inherent parametrization. In: Fischer, B., Dawant, B.M., Lorenz, C. (eds.) WBIR 2010. LNCS, vol. 6204, pp. 161–172. Springer, Heidelberg (2010)
2. Huang, X., Metaxas, D.T.: Metamorphs: Deformable shape and texture models. In: CVPR (2004)
3. Pennec, X., Cachier, P., Ayache, N.: Understanding the "demon's algorithm": 3d non-rigid registration by gradient descent. In: Taylor, C., Colchester, A. (eds.) MICCAI 1999. LNCS, vol. 1679, pp. 597–605. Springer, Heidelberg (1999)
4. Vercauteren, T., Pennec, X., Malis, E., Perchant, A., Ayache, N.: Insight into efficient image registration techniques and the demons algorithm. In: Karssemeijer, N., Lelieveldt, B. (eds.) IPMI 2007. LNCS, vol. 4584, pp. 495–506. Springer, Heidelberg (2007)
5. Thirion, J.P.: Image matching as a diffusion process: an analogy with maxwell's demons. Medical Image Analysis 2(3), 243–260 (1998)
6. Zikic, D., Baust, M., Kamen, A., Navab, N.: Generalization of deformable registration in riemannian sobolev spaces. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. MICCAI, pp. 586–593. Springer, Heidelberg (2010)
7. Stefanescu, R., Pennec, X., Ayache, N.: Grid powered nonlinear image registration with locally adaptive regularization. In: Medical Image Analysis: MICCAI 2003 Special Issue, vol. 8(3), pp. 325–342 (2004)
8. Tang, L., Hamarneh, G., Abugharbieh, R.: Reliability-driven, spatially-adaptive regularization for deformable registration. In: Fischer, B., Dawant, B.M., Lorenz, C. (eds.) WBIR 2010. LNCS, vol. 6204, pp. 173–185. Springer, Heidelberg (2010)
9. Glocker, B., Komodakis, N., Tziritas, G., Navab, N., Paragios, N.: Dense image registration through mrfs and efficient linear programming. Medical Image Analysis 12(6), 731–741 (2008)
10. Tang, T.W., Chung, A.C.: Non-rigid image registration using graph-cuts. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) MICCAI 2007, Part I. LNCS, vol. 4791, pp. 916–924. Springer, Heidelberg (2007)
11. Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. IEEE Trans. PAMI 23(11), 1222–1239 (2001)

12. Grady, L.: Random walks for image segmentation. IEEE Trans. PAMI 28(11), 1768–1783 (2006)
13. Grady, L.: Multilabel random walker image segmentation using prior models. In: CVPR, pp. 763–770 (2005)
14. Kroon, D.J.: Deamon's implementation, http://www.mathworks.com/matlabcentral/fileexchange/

Laplacian Eigenmaps Manifold Learning for Landmark Localization in Brain MR Images*

Ricardo Guerrero, Robin Wolz, and Daniel Rueckert

Biomedical Image Analysis Group, Imperial College London
{reg09,r.wolz,d.rueckert}@imperial.ac.uk

Abstract. The identification of anatomical landmarks in medical images is an important task in registration and morphometry. Manual labeling is time consuming and prone to observer errors. We propose a manifold learning procedure, based on Laplacian Eigenmaps, that learns an embedding from patches drawn from multiple brain MR images. The position of the patches in the manifold can be used to predict the location of the landmarks via regression. New images are embedded in the manifold and the resulting coordinates are used to predict the landmark position in the new image. The output of multiple regressors is fused in a weighted fashion to boost the accuracy and robustness. We demonstrate this framework in 3D brain MR images from the ADNI database. We show an accuracy of ∼0.5mm, an increase of at least two fold when compared to traditional approaches such as registration or sliding windows.

Keywords: Manifold Learning, Laplacian Eigenmaps, Landmarks.

1 Introduction

The detection of landmarks is a crucial step in many medical imaging applications, including registration, shape modeling and morphometry. Approaches to landmark detection can be roughly classified into three main categories: geometric-, classification- and regression-based techniques.

Geometric-based techniques identify significant points, lines, surfaces or volumes based on features from differential geometry. In [1], 3D differential operators are used to detect salient feature points in brain Magnetic Resonance (MR) images. Using features of two- and three-point combinations, e.g. pairwise and radial distances, angles, etc., a geometrical probability approach is used in [2] to analyze the structure of 3D chromatin in interphase cell nucleii. Rosten et al. [3] used machine learning techniques to develop a high-speed corner detector, that is used to track corner features in real-time video. One of the main disadvantages with the use of geometric-based features is that geometrical saliency does not necessarily imply anatomical relevance. Another popular approach common in the computer vision community includes the use of SIFT and similar types of feature descriptors [4]. Unfortunately, some of these feature detectors cannot be easily extended to 3D.

* This project was partially funded by CONACyT, SEP and the Rabin Enzra trust.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 566–573, 2011.

In the classification-based approach to landmark detection the general idea is to use a classifier to test an image, either exhaustively [5] or in a reduced space [6], whether a patch contains the landmark or not. In [5] the authors used a variation of AdaBoost in combination with Haar-like features to detect human faces. To locate the faces in real time in 2D images a cascade of classifiers is constructed which is applied to a window sliding over the entire image. Instead of sliding a window exhaustively across the image, Lampert et al. [6] proposed a branch and bound algorithm to reduce the search space. In [7] a probabilistic boosting tree is used to learn a discriminative model based on contextual features, with a marginal space learning strategy, for landmarks detection in cardiac MR images.

In Zhou et al. [8], a boosted regression approach is formulated for medical images. Amongst other applications, the authors test their approach for the detection and localization of anatomical structures. However, as explained in [9], the approach has two major drawbacks: First, it assumes that the output variables have a multivariate Gaussian distribution, which for real data is seldom the case. The second drawback is that the weak learner in the boosted regression is generally too weak. Therefore, the training requires the combination of a very large number of weak regressors to converge, making the training time computationally unfeasible, even for 2D images. In [9] the authors address these shortcomings, by adding more representational power to the weak learner.

In this work we propose a manifold learning approach that is capable of learning a low-dimensional embedding of image patches. The assumption is that the local anatomy around a particular landmark is well-represented in this embedding. We can then learn a regressor that predicts the displacements between the patch and the landmark. Patches from unseen images are mapped to the learned manifold using an out-of-sample approach and the regressor is then used to obtain an estimate of the landmark position. Finally, a consensus from the predictions made by several patches (from the same image) is reached, using a weighted average of all the estimates. The approach has been trained on a large dataset of 100 brain MR images from cognitively normal subjects (CN), patients with mild cognitive impairment (MCI) and Alzheimer's disease (AD) from the Alzheimer's Disease Neuroimaging Initiative (ADNI[1]) study database. A different set of 100 images from ADNI is used for testing the proposed approach.

2 Method

2.1 Laplacian Eigenmaps Manifold Learning

Given a set of N_I MR brain images, we extract N_P equally sized patches from each image in a region of interest (ROI) around a landmark. Each patch consisting of D voxels is stored as an intensity vector $\mathbf{x}_n = \{\mathrm{x}_1, ..., \mathrm{x}_D\}$ in $\mathbb{R}^D$, and $\mathbf{X} = \{\mathbf{x}_1, ..., \mathbf{x}_N\}$, where $N = N_I \cdot N_P$. Our aim is to learn the underlying manifold in $\mathbb{R}^d$ ($d \ll D$) that represents the relationship between patches in the vicinity of a given landmark. We intend to learn a manifold that can be used

[1] www.loni.ucla.edu/ADNI

to predict the displacement $\Delta_n = \{\delta_x, \delta_y, \delta_z\}$ between the center of any patch and a landmark. Manifold learning offers a powerful approach to find a representation of images that facilitates the application of machine learning techniques such as regression. Since the patches lie on or near to a non-linear manifold, the Euclidean distance between patches in the original space is not necesarilly meaningful and cannot be used for regression. After uncovering the manifold structure in the data, the Euclidean distance in the embedded space provides a more meaningful approximation of the geodesic distance in the original space.

Laplacian Eigenmaps can be used to find a low-dimensional representation of the data $f : \mathbf{X} \rightarrow \mathbf{Y}$, $\mathbf{y}_i = f(\mathbf{x}_i)$ while preserving the local geometric properties of the manifold [10]. Laplacian Eigenmaps uses a local neighbourhood graph to approximate geodesic distances among data points. In our work we use the Euclidean norm as a distance (similarity) metric to identify the k-neighbourhood around each point. From these distances a sparse neighbourhood graph G is constructed. Furthermore, a weight matrix $\mathbf{W}$ assigns a value to each edge in G, and is computed using a Gaussian heat kernel $K(\mathbf{x}_i, \mathbf{x}_j)$ with standard deviation σ. Laplacian Eigenmaps aims to place points $\mathbf{x}_i$ and $\mathbf{x}_j$ close together in the low-dimensional space, if their weight $w_{i,j}$ is high, i.e. if they are close in the high-dimensional space. This is done by minimizing the cost function given by

$$\phi(\mathbf{Y}) = \text{argmin} \sum_{i,j} \|\mathbf{y}_i - \mathbf{y}_j\|^2 w_{i,j} \ , \tag{1}$$

under the constraint that $\mathbf{y}^T \mathbf{D} \mathbf{y} = 1$ which removes an arbitrary scaling factor in the embedding and prevents the trivial solution where all $\mathbf{y}_i$ are zero. The minimization of Eq. (1) can be formulated as an eigenproblem [11], through the computation of the degree matrix $\mathbf{M}$ of $\mathbf{W}$, and the Laplacian $\mathbf{L}$, where $m_{i,i} = \sum_j w_{i,j}$ and $\mathbf{L} = \mathbf{M} - \mathbf{W}$. Hence, the low-dimensional manifold $\mathbf{Y}$ that represents all the data points can be obtained via solving a generalized eigenproblem, $\mathbf{L}\nu = \lambda \mathbf{M}\nu$, where ν and λ are the eigenvectors and eigenvalues, and in turn the d smallest (non-zero) eigenvectors ν represent the new coordinate system.

2.2 Approximate Nearest Neighbors

Since we are learning a manifold comprised from a relatively large number of examples (219,700, see Sec. 3), the similarity matrix $\mathbf{W}$ that needs to be calculated in Laplacian Eigenmaps is very large (∼50 billion elements), and even though it is strictly k-sparse, calculating exact nearest neighbours would mean that a non-sparse matrix would need to be calculated first, in order to find the k nearest neighbours and then sparsify $\mathbf{W}$; making the calculation of all the exact pairwise distances computationally unfeasible. We therefore, instead calculate approximate nearest neighbors (ANN) using a hierarchical k-means tree, which is constructed by splitting the data points into k_m distinct regions using k-means clustering, then applying the same method recursively. The recursion stops once the number of data points in each region is below k_m [12], as implemented in the FLANN library [13].

2.3 Out of Sample Extension

For the application considered in this work, it is necessary to map new patches into the manifold in order to use the embedded coordinates to make a prediction via regression. For linear dimensionality reduction techniques like PCA this is straightforward, as they provide a projection matrix for exact transformation between the original space and the embedded space. Unfortunately, this is not the case for non-linear methods. Therefore, approximation techniques must be used. We address this problem by using an out of sample technique that employs the Nyström approximation [14], which approximates the eigenvectors of a large matrix based on the eigendecomposition of a submatrix of the large matrix. Laplacian Eigenmaps are based on an initial kernel K, as explained in Sec. 2.1. An equivalent, training set dependent normalized kernel, is:

$$\tilde{K}(\mathrm{x}'_i, \mathrm{x}_j) = \frac{1}{N} \frac{K(\mathrm{x}'_i, \mathrm{x}_j)}{\sqrt{E_{\mathbf{X}'}[K(\mathrm{x}'_i, \mathbf{X})] E_{\mathbf{X}}[K(\mathrm{x}_j, \mathbf{X})]}} , \tag{2}$$

where x_j and x'_i are points from the training $\mathbf{X}$ and test $\mathbf{X}'$ datasets, respectively, the expectations are taken over the empirical data and N is the number of training samples (see [15] for full analysis).

2.4 Spatial Prior Probabilities

Assuming that the brain is in some approximately known orientation and possition, a landmark's spatial location is bounded, to a certain extent, to a particular volume within the brain. Once the images are affinely registered, the possible locations of each landmark are bounded within this space. Thus, we can restrict the search for each landmark to those locations which have a non-zero probability (ROI) for the location of the landmark. We model the spatial prior probabilities of each landmark, based on the position of the landmark in the training set, using kernel (or parzen window) density estimation. This can be formulated as

$$P(\mathrm{x}) = \frac{1}{n} \sum_{i=1}^{n} \frac{1}{h_n^d} K_p \left(\frac{\mathrm{x} - \mathrm{x}_i}{h_n} \right) \tag{3}$$

where x are the 3D voxel coordinates (x, y, z), x_i are all the landmark's coordenates, $K_p(\cdot)$ is the window function or the kernel in a d-dimensional space such that $\int_{\Re^d} K_p(\mathrm{x}) d\mathrm{x} = 1$, n is the number of observations and $h_n > 0$ is the bandwidth parameter that corresponds to the width of the kernel. The kernel function $K_p(\cdot)$ is modelled as a Gaussian function.

2.5 Landmark Prediction

Using the low-dimensional coordinates $\mathbf{Y} = \{\mathbf{y}_1, ..., \mathbf{y}_N\}$ of each patch (that came from the training set) and their corresponding displacements Δ_n (between the centre of the patch in image space and the position of the landmark), we

fit a linear regressor, using $\mathbf{Y}$ as independent variables and Δ_n as dependant variables, to obtain an estimated displacement $\Delta'_n = \{\delta'_x, \delta'_y, \delta'_z\}$.

A test dataset $\mathbf{X}'$ is built from patches belonging to a test image's ROI. These patches are embedded in the landmark specific manifold, to obtain their low-dimentional representation $\mathbf{Y}'$. Using the learned regressor coefficients β, an estimate of the displacement from patch n to the landmark is obtained. For each test image we randomly sample 100 image patches and average the prediction results for all the patches. Since patches that predict small displacements have a higher accuracy, a weighted average was calculated, where predictions' weights are based on the magnitude of Δ'_n. A Gaussian kernel $K(\Delta'_n)$ with $\sigma = 0.65$, and mean zero was used for the weighting.

3 Data and Results

The images that were used to evaluate the proposed method were obtained from the ADNI database [16]. In the ADNI study, brain MR images were acquired at regular intervals after an initial baseline scan from approximately 200 cognitively normal older subjects (CN), 400 subjects with mild cognitive impairment (MCI), and 200 subjects with early Alzheimers disease (AD). In this work, we used a subset of 1.5T T1-weighted baseline images of 100 randomly chosen subjects for training and another 100 randomly chosen subjects for testing. In both (training and testing) datasets there are 24 AD, 48 MCI and 28 healthy subjects, to faithfully represent the full ADNI dataset. All brain images were skull stripped, affinely aligned to the MNI space and normalized using linear intensity rescaling.

The high-dimensional training set $\mathbf{X} = \{\mathbf{x}_1, ..., \mathbf{x}_N\}$, is obtained by collecting 3D image cubic patches of 21^3 voxels around a regular grid that is centred at the landmark, from 100 different brain MR images. The grid has a spacing of 3 voxels and a displacement of ± 18 voxels in each dimension. This volume is chosen so that it includes the non-zero probability volume obtained from the PDF estimation. For each image, we exhaustively sample from this grid: For each image in the training set we sample 13^3 (2197) patches. Doing this for the 100 images in the trainings set and rearranging them so that each patch is represented as column vector (with 21^3 values) yields a 219,700 by 9161 (N, D) matrix that contains all the patches, from all the training images, around the landmark in question. For both the training and testing datasets a total 20 landmarks (table 1) were manually selected by an expert observer using 3 orthogonal views.

From the training set $\mathbf{X}$ we learn the underlying low-dimensional manifold, using Laplacian eigenmaps (Sec. 2.1). The parameters k (nearest neighbors in the neighborhood graph), d (the output dimensionality of the data) and σ from the Gaussian heat kernel, were empirically set to 50, 200 and 1, respectively. The parameter k was chosen to yield a fully connected neighbourhood graph, ensuring that all distances from the landmark were equally represented. The coefficient of determination, R^2, for the linear regression was used as an indicative to determine the final dimensionality d of $\mathbf{Y}$, with values of around 0.9 obtained for 200 dimensions. Finally, tuning the parameter σ shows little improvement.

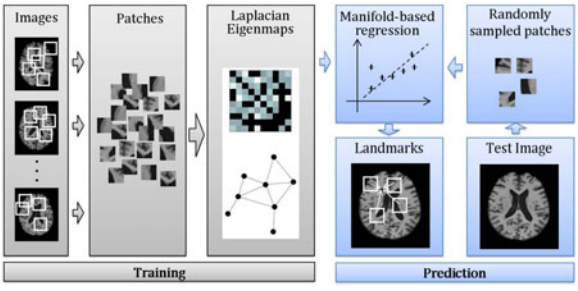

Fig. 1. Diagram of method's training and testing steps

Similarly, for each new test image, we took 100 patches at random locations, within a non-zero probability in the PDF and not necessarily belonging to the grid used in the training. We then embed the new points (patches) into the learned low-dimensional manifold, using the technique described in Sec. 2.3, and use the regression coefficients to obtain a prediction from each point. A final landmark prediction for each image is obtained, using a weighted average, as described in Sec. 2.5. A diagram of the whole process is shown in Fig. 1.

Table 1 shows results of the proposed method (landmark specific manifold, LM) and two other possible approaches: a 3D adaptation of the sliding window algorithm (SW) [5], and non-rigid registration (REG). The 3D sliding window, is trained using 400 positive and 4000 negative examples (taken from the vicinity of the positives), and it is build as a monolithic 100 feature (3D Haar features) classifier. In the registration approach, a non-rigid B-spline registration algorithm (as proposed in [17]), with a final control point spacing of 5mm, was used to propagate the landmarks from the MNI template back to the baseline images, using the transformations obtained from the registration process. A five by two cross validation of the method was carried out in order to asses the results and ensure repreducibility (the average of the five tests is showned, with a varriability among tests of $\sim$0.19mm). Our methods shows its poorest performance on the lower aspect of the cerebellum (point 16), this is mainly attributed to the fact that this particular landmark is very close to the edge of the image and the surface of the brain, meaning that relatively few patches in its vicinity contain useful information. A statistical comparison, between the three classes of patients (AD, MCI and CN), of the landmark positions in the training dataset and the landmark prediction error using our method offered no intra class distinction. It was also observed that landmarks with a higher location variability tend to have a higher prediction error then those with a lower variability.

4 Discussion and Future Work

We have proposed a method that uses Laplacian Eigenmaps to learn a low-dimensional manifold that represents local anatomy around a specific landmark

Table 1. Accuracy of the proposed method on the ADNI database, for the Splenium of corpus callosum (outer aspect, inferior tip and inner aspect (1,2,3)), Genu of corpus callosum (outer and inner aspect (4,5)), Superior and inferior aspect of pons (6,7), Superior and inferior aspect cerebellum (8,16), Fourth ventricle (9), Putamen posterior and anterior (10,11)(left), (12,13)(right), Anterior and posterior commissure (14,15), Anterior tip of lateral ventricle (left and right) (17,18), Inferior tip of lateral ventricle (left and right) (19,20).

	Mean			Std. dev.		
L	SW	ML	REG	SW	ML	REG
1	1.75	0.75	3.95	1.04	0.31	1.43
2	1.46	0.52	2.1	0.75	0.23	0.89
3	1.81	0.51	2.31	0.99	0.25	1.31
4	1.58	0.55	1.73	1.09	0.24	1.01
5	1.28	0.52	1.47	0.67	0.25	0.64
6	1.22	0.52	2.79	0.63	0.26	1.26
7	1.86	0.59	1.7	0.92	0.25	0.85
8	2.27	0.7	2.99	1.71	0.33	1.64
9	1.09	0.48	5.57	0.65	0.22	2.7
10	2.21	0.54	4.36	1.22	0.21	1.81

	Mean			Std. dev.		
L	SW	ML	REG	SW	ML	REG
11	1.86	0.53	2.48	1.13	0.25	1.29
12	2.2	0.49	3.53	1.22	0.22	1.78
13	2.31	0.53	2.79	1.61	0.26	1.43
14	1.27	0.45	1.05	0.72	0.22	0.42
15	0.79	0.46	1.85	0.6	0.2	0.48
16	2.13	1.08	3.71	1.69	0.43	1.68
17	1.86	0.62	3.67	1.14	0.33	1.72
18	1.84	0.62	3.65	1.2	0.4	1.73
19	2.28	0.57	4.44	1.42	0.3	2.07
20	2.18	0.54	4.01	1.18	0.28	1.79

in brain MR images. The landmark specific low-dimensional manifolds were learned using image patches (around the vicinity of the landmark) from 100 brain MR images belonging to the ADNI dataset. Prior knowledge of the spatial distribution of the landmarks was used to reduce the search space. Our results show that the proposed method significantly outperforms a 3D sliding window, B-spline registration and a naive average, in the landmark localization task.

Anatomical landmarks of the brain have a constrained position within the brain and to the relative position of other landmarks. The inferior tip, inner and outer aspect of the splenium of the corpus callosum, are part of the same anatomical structure, so their locations with respect to each other is strongly correlated. A graphical model of the joint spatial distribution probabilities of related landmarks, in the form of a Markov random field, could be used to represent spatial dependencies. Joint spatial probabilities between related landmarks can be modeled as a continuous multivariate Gaussian. The most probable configuration of all the landmarks, according to the graphical model, where the marginal spatial probabilities of the landmarks (possibly a Gaussian fit to the output from all the predictions made from the proposed method or a SW approach) could act as a weight on the joint spatial probabilities, would yield sub-pixel accuracy.

Other possibilities for improvement could be the use of more powerful regression techniques such as support vector regression or a multiple output regression. An interesting path to explore, would be to learn a manifold of the whole brain, this would enable to locate any landmark (or pseudo landmark) within the brain, allowing to determine the position of hundreds or thousands of pseudo landmarks in every new image, allowing for a good fast first stage registration.

References

1. Rohr, K.: On 3D differential operators for detecting point landmarks. Image and Vision Computing 15(3), 219–233 (1997)
2. Gladilin, E., Goetze, S., Mateos-Langerak, J., Van Driel, R., Rohr, K., Eils, R.: Geometrical probability approach for analysis of 3D chromatin structure in interphase cell nuclei. In: Computational Intelligence and Bioinformatics and Computational Biology, pp. 127–134 (2007)
3. Rosten, E., Drummond, T.: Machine learning for high-speed corner detection. In: Leonardis, A., Bischof, H., Pinz, A. (eds.) ECCV 2006. LNCS, vol. 3951, pp. 430–443. Springer, Heidelberg (2006)
4. Li, J., Allinson, N.M.: A comprehensive review of current local features for computer vision. Neruocomputing 71, 1771–1787 (2008)
5. Viola, P., Jones, M.: Robust Real-time Face Detection. International Journal of Computer Vision 57, 137–154 (2004)
6. Lampert, C.H., Blaschko, M.B., Hofmann, T.: Beyond sliding windows: Object localization by efficient subwindow search. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 1–8. IEEE Computer Society, Los Alamitos (2008)
7. Lu, X., Georgescu, B., Littmann, A., Mueller, E., Comaniciu, D.: Discriminative Joint Context for Automatic Landmark Set Detection from a Single Cardiac MR Long Axis Slice. In: Ayache, N., Delingette, H., Sermesant, M. (eds.) FIMH 2009. LNCS, vol. 5528, pp. 457–465. Springer, Heidelberg (2009)
8. Zhou, S.K., Georgescu, B., Zhou, X.S., Comaniciu, D.: Image Based Regression Ussing Boosting Method. In: 10th IEEE International Conference on Computer Vision, vol. 1, pp. 541–548. IEEE Computer Society, Los Alamitos (2005)
9. Zhou, S.K., Zhou, J., Comaniciu, D.: A boosting regression approach to medical anatomy detection. In: IEEE Conference on Computer Vision and Pattern Recognition. IEEE Computer Society, Los Alamitos (2007)
10. Belkin, M., Nigoyi, P.: Laplacian Eigenmaps and spectral techniques for embedding and clustering. In: Advances in Neural Information Processing Systems 14, vol. 14, pp. 585–591. MIT Press, Cambridge (2002)
11. Anderson, W.N., Morley, T.D.: Eigenvalues of the Laplacian of a graph. Linear and Multilinear Algebra 18, 141–145 (1985)
12. Muja, M., Lowe, D.G.: Fast approximate nearest neighbors with automatic algorithm configuration. In: VISAPP International Conference on Computer Vision Theory and Applications, pp. 331–340. INSTICC Press (2009)
13. Fast Library for Approximate Nearest Neighbors, http://www.cs.ubc.ca/~mariusm/index.php/FLANN/FLANN
14. Platt, J.C.: FastMap, MetricMap, and Landmark MDS are all Nyström algorithms. In: Cowell, R.G., Ghahramani, Z. (eds.) 10th International Workshop on Artificial Intelligence and Statistics, pp. 261–268. Society for Artificial Intelligence and Statistics (2005)
15. Bengio, Y., Paiement, J.F., Vincent, P., Delalleau, O., Le Roux, N., Ouimet, M.: Out-of-sample extensions for LLE, Isomap, MDS, eigenmaps, and spectral clustering. In: Thrun, S., Saul, L., Schölkopf, L. (eds.) Advances in Neural Information Processing Systems 16, vol. 16, pp. 177–184. MIT Press, Cambridge (2004)
16. Mueller, S.G., Weiner, M.W., Thal, L.J., Petersen, R.C., Jack, C., Jagust, W., Trojanowski, J.Q., Toga, A.W., Beckett, L.: The Alzheimers Disease Neuroimaging Initiative. Neuroimaging Clinics of North America 15(4), 869–877 (2005)
17. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L.G., Leach, M.O., Hawkes, D.J.: Nonrigid Registration Using Free-Form Deformations: Application to Breast MR Images. IEEE Transactions on Medical Imaging 18(8), 712–721 (1999)

Automatic Alignment of Brain MR Scout Scans Using Data-adaptive Multi-structural Model

Ting Chen[1,⋆], Yiqiang Zhan[2], Shaoting Zhang[3,⋆], and Maneesh Dewan[2]

[1] Department of CISE, University of Florida, Gainesville, FL USA
tichen@cise.ufl.edu
[2] SYNGO US R&D, Siemens Healthcare, Malvern, PA, USA
{yiqiang.zhan,maneesh.dewan}@siemens.com
[3] Department of Computer Science, Rutgers University, Piscataway, NJ, USA
shaoting@cs.rutgers.edu

Abstract. Accurate slice positioning of diagnostic MR brain images is clinically important due to their inherent anisotropic resolution. Recently, a low-res fast 3D "scout" scan has become popular as a pre-requisite localizer for the positioning of these diagnostic high-res images on relevant anatomies. Automation of this "scout" scan alignment needs to be highly robust, accurate and reproducible, which can not be achieved by existing methods such as voxel-based registration. Although recently proposed "Learning Ensembles of Anatomical Patterns (LEAP)" framework [4] paves the way to high robustness through redundant anatomy feature detections, the "somewhat conflicting" *accuracy* and *reproducibility* goals can not be satisfied simultaneously from the single model-based alignment perspective. Hence, we present a data adaptive multi-structural model based registration algorithm to achieve these joint goals. We validate our system on a large number of clinical data sets (731 adult and 100 pediatric brain MRI scans). Our algorithm demonstrates $> 99.5\%$ robustness with high accuracy. The reproducibility is $< 0.32^{\circ}$ for rotation and $< 0.27mm$ for translation on average within multiple follow-up scans for the same patient.

1 Introduction

The inherent MR imaging characteristics of high "in-plane" and low "out-of-plane" resolution warrant high accuracy requirements on the positioning of diagnostic imaging slices. Moreover, there is significant intra- and inter-patient variation in the starting orientations and axes of the scanning volumes. Hence, the use of low-res isotropic 3D "scout" scans as a pre-requisite sequence to improve the positioning accuracy has become a necessity prior to all clinical brain studies. In the current MR brain workflow (Fig.1(a)), a well-trained technician positions the imaging planes for the following high-res scans based on anatomy information in this fast low-res "scout" scan. For example, the axial plane of a standard high-res brain scan should be positioned parallel to the bicommissural line linking the landmarks of anterior and posterior commissure.

⋆ This research was conducted while the authors were at Siemens Healthcare, USA.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 574–581, 2011.

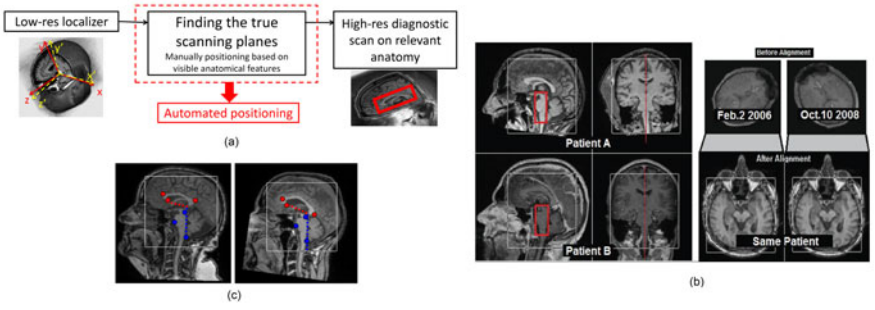

Fig. 1. (a) The 3 steps of the current MRI examination workflow. *The goal of our is to develop an automatic algorithm for the* 2^{nd} *step.* **(b)** *Accuracy* (left): a precise and standardized anatomy for all patients characterized by vertical brainstem and correct middle sagittal plane (MSP). *Reproducibility* (right): a consistent alignment for the follow-up scans of the same patient characterized by "no motion" in rigid structures when switching between aligned re-scan "scout" images. **(c)** Atlas and subject scans with different relations between brain stem and corpus callosum.

The inaccuracy, irreproducibility and time consumption limitations of manual alignment make automatic alignment algorithms greatly desirable. For these automatic alignment algorithms to be clinically acceptable, they need to be not only robust and accurate, but also *reproducible*, i.e., the alignment of the follow-up scans of the same patient should be identical with respect to relevant anatomies. Fig.1(b) illustrates the *accuracy* and *reproducibility* requirements for standardized brain alignment.

In [1], Andre *et al.* proposed an automatic slice positioning technique by registering the scout scan with a pre-aligned one. This method is not robust when the images contain large growing tumor, field of view change, severe noise, image artifacts or missing structures, etc. In addition, it is well known that the voxel-based image registration methods show high reproducibility in registering scans of the same patient, but are not as good for aligning scans from different patients, especially under the constraints of rigid or affine transformation for brain images. We experimentally validated this by registering several pair of re-scans (A_1 and A_2) from the same patient to a model scan (B) from a different patient using mutual information (MI) based image registration. Unfortunately, when both scans (A_1 and A_2) achieve "best" matching with the model scan (B), i.e., MI is maximized, the reproducibility is not good - larger motion was observed when switching between the two aligned scans (A_1' and A_2').

Sharp *et al.* [2] proposed another approach jointly based on feature landmarks and image registration that relies on the availability of previous scans of the *same* patient. This method obtains remarkable accuracy and efficiency, however, their assumption on previous scan availability may not be applicable in real practice. In [3], Zhang *et al.* presented an auto-align system using image registration and active shape model, however, the "reproducibility" issue is not well addressed.

Recently, "Learning Ensembles of Anatomical Patterns (LEAP)" [4] has shown high robustness in anatomy detection through learning redundant local appearance cues and sparse configurational models. Hence, an auto-align system based on LEAP [7] achieves extremely high robustness in the presence of severe artifacts/diseases. However, as [7] rigidly aligns scout scans to a *single* atlas/model, it fails to fully accomplish the "somewhat conflicting" "accuracy" and "reproducibility" goals simultaneously. To illustrate this, a schematic example is shown in Fig. 1(c). Here the brain stem and corpus callosum have different articulation angles (between blue and red dash lines) in the atlas and subject scans. If the alignment relies on all landmarks (3 from brain stem and 3 from corpus callosum), higher *reproducibility* is expected as small variations of landmark detections can be averaged out by robust point set registration techniques.

However, the accuracy requirement would be hard to meet (the aligned brain stem is not vertical) owing to the compromise of corpus callosum (note that only rigid alignment is allowed in MR auto-align). On the contrary, if the alignment relies only on the 3 landmarks of brain stem, one might achieve accurate alignment (vertical brain stem); but *reproducibility* is sacrificed as variations of brain stem landmarks have larger impact on the final alignment. Clearly, this conflict will not exist if the articulations between brain stem and corpus callosum are similar in atlas and subject scan or in other words, if the atlas constructed is adaptive to the detected anatomical feature landmarks in the subject scan. Indeed, this is the motivatation behind our Data-adaptive Multi-structural Model (DMM). It aims to compose a *virtual* atlas using learned exemplars of different local brain structures, based on the auto-detected anatomies in the subject scan. The constructed virtual atlas might not exist in real world, but it is the best approximation of the subject brain structure. Therefore it is a much better solution than picking from multiple models directly, given the limited data in practice. Additionally, multiple models would not be able to capture the local structure variations as well as the DMM. The philosophy behind our method is similar to [8]. However, our method differs significantly from [8] in terms of multi-structural model (landmark based vs. image based) and objectives (alignment reproducibility vs. segmentation accuracy).

2 Methodology

Our auto-align system consists of three major steps. *First*, a set of anatomical landmarks are detected using LEAP [4] algorithm. Briefly put, LEAP is a learning-based algorithm that exploits "redundant" local appearance cues and sparse configuration models to achieve highly robust anatomy detections. *Second*, by querying the Data-adaptive Multi-structural Model (DMM) database, a virtual atlas is composed to optimally match the detected landmark set. *Finally*, a landmark-based rigid alignment is performed to register the landmark set with the virtual atlas, which brings the scout scan to the standard position. Since the major contribution of this work lies in the Data-adaptive Multi-structural Model (DMM), we will focus on it in the remainder of this section.

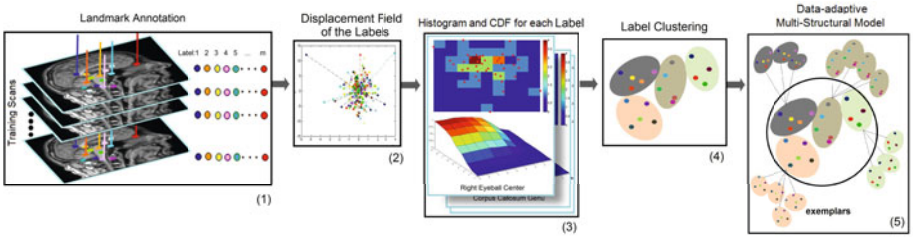

Fig. 2. Flowchart for learning multi-structural clusters. We demonstrate the images in 2D while the data is actually in 3D.

Notations. Define $\mathcal{I} = \{1, 2, \ldots, M\}$ be the labels for M anatomical landmarks. Given N training scans with all these M landmarks annotated, assume x_i^j is the ith landmark in the j^{th} image, $\mathcal{X}^j = \{x_i^j | i = 1, 2, \ldots, M\}$ denotes a set including *all* landmarks from the j^{th} image and $\mathcal{X}_i = \{x_i^j | j = 1, 2, \ldots, N\}$ denotes a set including the i^{th} landmark from *all* N images.

Construction of Multi-structural-Model. The basic idea of multi-structural-model is to group anatomical landmarks into different clusters. During the runtime, the "optimal" exemplars from different cluster can be composed to form a virtual model which has the similar brain structure spatial configuration as the subject's. To leverage the spatial correlations between landmarks, the construction of multi-structural-model starts from calculating the spatial variations.

Calculate spatial variations of each landmark: (Fig.2 (2)) To remove the variations from global transformation, we first register landmarks from different training images, i.e., $\mathcal{X}^j, j = 1, \ldots N$, to a common canonical space. The spatially normalized point set is defined by $\widetilde{\mathcal{X}}^j = R_j(\mathcal{X}^j)$, where R_j is a similarity transformation estimated using group-wise point set registration [5], which is a pre-step to align the training subjects. Given the average position of the i^{th} landmark across all subjects $\bar{\mathcal{X}}_i = E(\widetilde{\mathcal{X}^j})$, the spatial variation of the i^{th} landmark can be captured by a displacement field $U_i = \{u_i^j = \widetilde{x}_i^j - \bar{\mathcal{X}}_i | j = 1, \ldots, N\}$.

Build $3D$ histogram of displacement fields: (Fig.2 (3)) To reveal the statistical insight of the spatial variations, a $3D$ histogram h_i is built on each displacement field U_i. More specifically, the whole space of displacement field is divided into $9 \times 9 \times 9$ bins along x, y and z dimensions and $h_i(x, y, z)$ is obtained by counting the occurrences of u_i^j in bin (x, y, z). As cumulative density function (cdf) is more robust than histogram, we convert $h_i(x, y, z)$ to cdf as $F_i(s, t, l) = \Sigma_{-\infty}^{s,t,l} h_i(x, y, z)$. It is used in the similarity measure in the next step.

Landmark clustering using affinity propagation: (Fig.2 (4)) The clustering of anatomy primitives(landmarks) is finally accomplished by Affinity Propagation [6] method. Compared to other clustering methods, affinity propagation shows advantages in its less sensitivity to bad initialization and no requirement for a preset cluster number which is often unknown in our problem. Here, each landmark is considered as a data point in the affinity network and the similarity between label pairs v and w are defined as:

$$A(v,w) = -\alpha \sum_{s,t,l} ||F_v(s,t,l) - F_w(s,t,l)||_2^2 - (1-\alpha)||\bar{\mathcal{X}}_v - \bar{\mathcal{X}}_w||_2^2 \qquad (1)$$

The second term measures the average distance between two landmarks. It ensures that two nearby landmarks have higher probability to be clustered together (in accordance with our observations of the brain anatomy).

As shown in Fig.3, we now have a multi-structural database that consists of a set of clusters. In other words, brain is divided into multiple pseudo sub-structures (here "pseudo" indicates that a cluster might not strictly correspond to the well defined sub-structures in the medical atlas). Each sub-structure has N instances coming from the N training images. More specifically, each $\mathcal{X}^j$ is now partitioned into a set of $\{\mathcal{S}_k^j\}, k \in \{1, 2, \ldots K\}$, each of which contains a subset of landmarks from $\mathcal{X}^j$. These pseudo-sub-structures provide the flexibility of composing a virtual atlas while still maintaining the connections of closely related anatomy primitives.

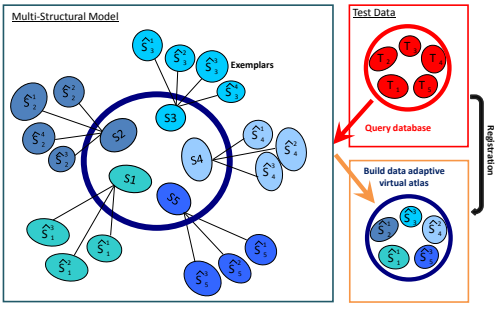

Fig. 3. This figure illustrates the Data-adaptive Multi-structural Model Database

Extraction of representative exemplars: Given the landmarks detected in a testing scout image, a naive way to compose an "optimal" virtual atlas from the multi-structural-model is to go through *all* instances in *all* clusters. Clearly, when the number of training samples become large, the runtime efficiency will be greatly degraded. Therefore, we propose to extract representative exemplars for each cluster $\mathcal{S}_k$. Principal Component Analysis is applied to all instances belonging to $\mathcal{S}_k$, i.e., $\{\mathcal{S}_k^1 ... \mathcal{S}_k^N\}$. The first few principal variation modes $\widehat{\mathcal{S}}_k^j, j \in \{1, 2, \ldots L\}; L \ll N$, are used as representative exemplars of $\mathcal{S}_k$.

Adaptive Composition of Virtual Atlas. During run-time, multi-structural model is employed to derive optimal virtual atlas. Given a testing scout scan, our auto-align system starts by detecting a set of landmarks $T = \{t_i\}, i = 1 \ldots N$ using LEAP. As LEAP algorithm has specific detectors for every anatomical landmark, the correspondences between detected landmarks and our multi-structural model are automatically built. T is then decomposed into pseudo-sub-structures as $\{T_k\}, k = 1 \ldots K$ according to the built-in landmark label clusters in the multi-structural model. For each T_k, it queries the multi-structural model and finds the most similar exemplars as $\widehat{S}_k^{opt} \leftarrow \min_j \|\mathcal{R}(T_k) - \widehat{S}_k^j\|_2$. Here $\mathcal{R}(\cdot)$ defines the transformation that brings T_p to the multi-structural model space.

Finally, a virtual atlas $\mathcal{V}$ is obtained by integrating all these optimal sub-structural exemplars as $\mathcal{V} = \bigcup_{k=1 \ldots K} \widehat{S}_k^{opt}$. As a result, each sub-structure of this atlas has the most similar spatial configurations with the detected anatomy

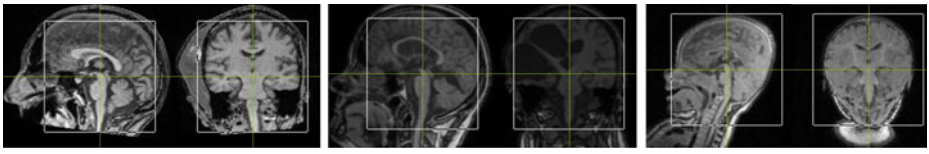

Fig. 4. Auto-align results on challenging cases, i.e., severe artifacts (Left), tumors (Center) and pediatric (Right). Please refer to Fig. 1(b) for alignment requirements.

primitives in the same local regions. As discussed in Section 1, if spatial configurations between different anatomy primitives are similar in atlas and subject, the accuracy and reproducibility requirements are no longer in conflict. Thus the virtual atlas derived from the multi-structural model paves the way to achieve accurate and reproducible alignment simultaneously.

3 Experiments

Experimental Setting. Our validation data set includes a large number of T1-weighted $3D$ scout scans: 731 adult (age: 15-87 years) and 100 pediatric (age: months to 14 years). Some adult volunteers are scanned multiple times to test the reproducibility of the algorithm. 90 adult and 40 pediatric scans were selected as training data. 22 anatomical landmarks are annotated for constructing the DMM model and training the LEAP algorithm. The regularization parameter α in Eqn. (1) was set to 0.7. We compare our method with the landmark-based alignment using Single Fixed Model (SFM), wherein the model is the average constructed by group-wise point set registration [5].

Accuracy. Our method provides *accurate* alignment on 690 cases out of the 701 testing cases. It was verified qualitatively by the experts who checked the correctness of the MSP and the verticality of the brainstem (as indicated in Fig. 1(b)). Some example alignment results on challenging cases in Fig.4 to demonstrate the robustness and accuracy of our algorithm. In Fig. 5, our method shows superiority accuracy with little ($< 1°$) or no tilt, as compared to SFM, wherein brains brainstems are visibly slanted by 3°-5°.

Reproducibility. We validate the reproducibility of our method on 5 volunteers, each with 14-18 re-scans. For better visualization, we overlayed two aligned re-scans in a checkerboard pattern in Fig.6. The edges of the brain structure (e.g.,

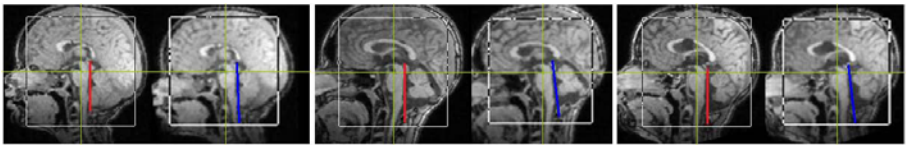

Fig. 5. Comparisons of DMM(red) and SFM(blue) methods (in coronal views)

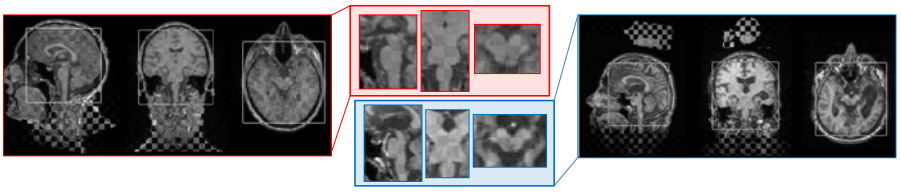

Fig. 6. Checkerboard visualization of aligned re-scans. Brainstem regions are zoomed-in for a better visualization of reproducibility.

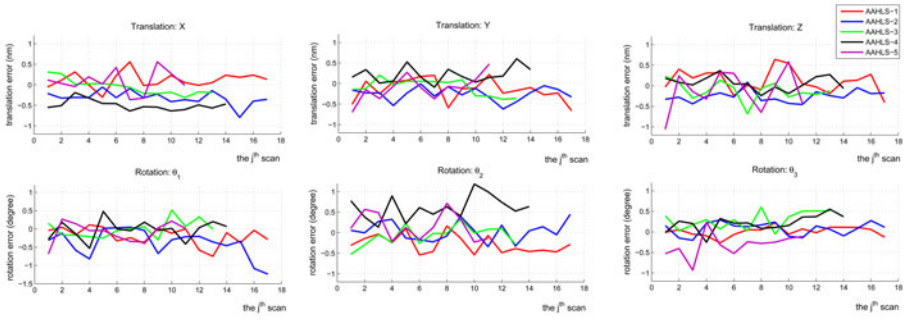

Fig. 7. Each figure shows the rotation or translation errors for each re-scan subject from the 5 data sets. The errors are estimated by registering each re-scan MRI localizer to a selected template from the same data set.

Table 1. Percentage of the re-scan data set with rotation error $\leq 0.1, 0.2$ and 0.5 degrees, and the translation error $\leq 0.1, 0.2$ and 0.5 mm (first three rows). The last three rows show the average rotation and translation errors along with the standard deviation and the improvements of DMM w.r.t SFM method.

	Rotation Error (degree)						Translation Error (mm)					
	DMM			SFM			DMM			SFM		
Statistics	θ_1	θ_2	θ_3	θ_1	θ_2	θ_3	t_x	t_y	t_z	t_x	t_y	t_z
$\leq 0.1^\circ$/mm	34.7%	20.8%	27.8%	22.0%	10.3%	22.1%	25.6%	29.6%	26.8%	15.3%	26.4%	31.9%
$\leq 0.2^\circ$/mm	52.8%	38.9%	52.8%	38.2%	30.9%	36.8%	39.4%	53.5%	50.7%	37.5%	47.2%	59.7%
$\leq 0.5^\circ$/mm	86.1%	80.6%	89.0%	75.0%	67.7%	69.1%	85.9%	90.2%	93.0%	83.3%	93.0%	91.6%
Abs. Mean	0.238	0.318	0.231	0.313	0.580	0.379	0.268	0.226	0.244	0.309	0.253	0.277
Improvement	23.96%	45.17%	39.05%				13.27%	10.67%	11.91%			
STD	0.093	0.174	0.108	0.074	0.260	0.107	0.143	0.037	0.092	0.157	0.098	0.134

brain stems in the zoomed up images) crossing checkerboards are well preserved, which shows the good reproducibility of our method. To quantitatively evaluate the reproducibility of our algorithm, we obtain the "ground truth" transformation between re-scans using MI-based image registration, since it is well known that MI-based image registration is almost perfect in aligning scans from the same person. The quantitative "reproducibility" is shown in Fig.7. Table.1 shows the reproducibility of our method compared to SFM. For each re-scan dataset,

we listed the percentage of scans at different error levels, the average translation and rotation errors, and the percentage of improvement from SFM to DMM.

Runtime Efficiency. Our method is implemented by Python2.5 and C++. Running on an Intel(R) Xeon(R) machine with 2.33GHz CPU and 3GB RAM, it takes about 5 seconds to align one $192 \times 192 \times 144$ brain MRI. This serves to anecdotally illustrate the computational time involved.

4 Conclusions

In this paper, we presented a method to automatically align MR brain "scout" scans with high accuracy, reproducibility and robustness. The joint goals of accuracy and reproducibility were met by using a data-adaptive multi-structural model. Our system, validated on a large number of clinical cases, automates the manually positioning procedure, which can highly improve the quality and speed of the brain MR examination workflow.

Acknowledgement. The authors would like to thank Martin Harder for a lot of valuable discussions and data collection.

References

1. van der Kouwe, A.J.W., Benner, T., Fischl, B., Schmitt, F., Salat, D.H., Harder, M., Sorensen, A.G., Dale, A.M.: On-line Automatic Slice Positioning for Brain MR Imaging. NeuroImage 27(1), 222–230 (2005)
2. Sharp, G.C., Kollipara, S., Madden, T., Jiang, S.B., Rosenthal, S.J.: Anatomic Feature-based Registration for Patient Set-up in Head and Neck Cancer Radiotherapy. Phys. Med. Biol. 50, 4667 (2005)
3. Zhang, L., Xu, Q., Chen, C., Novak, C.L.: Automated Alignment of MRI Brain Scan by Anatomic Landmarks. In: SPIE Med. Imag., vol. 7258 (2009)
4. Zhou, X.S., Peng, Z., Zhan, Y., Dewan, M., Jian, B., Krishnan, A., Tao, Y., Harder, M., Grosskopf, S., Feuerlein, U.: Redundancy, Redundancy, Redundancy: the Three Keys to Highly Robust Anatomical Parsing in Medical Images. In: Int. Conf. MIR, pp. 175–184 (2010)
5. Chen, T., Vemuri, B.C., Rangarajan, A., Eisenschenk, S.J.: Group-wise Point-set Registration using a Novel CDF-based Havrda-Charvat Divergence. Int. J. Comput. Vision 86(1), 111–124 (2010)
6. Frey, B.J., Dueck, D.: Clustering by Passing Messages Between Data Points. Science 315, 972–976 (2007)
7. Dewan, M., Zhan, Y., Peng, Z., Zhou, X.S.: Robust Algorithms for Anatomic Plane Primitive Detection in MR. In: SPIE Med. Imag., vol. 7260 (2009)
8. Shi, F., Yap, P.T., Fan, Y., Gilmore, J.H., Lin, W., Shen, D.: Construction of Multi-Region-Multi-Reference Atlases for Neonatal Brain MRI Segmentation. NeuroImage 51(2), 684–693 (2010)

Reconstruction of 3-D Histology Images by Simultaneous Deformable Registration

Marco Feuerstein, Hauke Heibel, José Gardiazabal,
Nassir Navab, and Martin Groher

microDimensions and Computer Aided Medical Procedures (CAMP),
Technische Universität München, Germany

Abstract. The reconstruction of histology sections into a 3-D volume receives increased attention due to its various applications in modern medical image analysis. To guarantee a geometrically coherent reconstruction, we propose a new way to register histological sections *simultaneously* to previously acquired reference images and to neighboring slices in the stack. To this end, we formulate two potential functions and associate them to the same Markov random field through which we can efficiently find an optimal solution. Due to our simultaneous formulation and the absence of any segmentation step during the reconstruction we can dramatically reduce error propagation effects. This is illustrated by experiments on carefully created synthetic as well as real data sets.

1 Introduction

Today, histology is still the gold standard for assessing anatomical information on a cellular level. Tissue samples are cut into ultra thin slices, stained, and viewed under a microscope. While traditional histology involves only a few slices to be analyzed, there is an increasing need to reassemble consecutive slices into a 3-D volume. Given such a volume, novel high-resolution in-vivo imaging techniques (e.g. micro-CT, high-field MRI, or phase contrast X-ray CT) can be validated, atlases can be created on a micron level, or 3-D micro-structures can be quantified for analysis. However, the geometrically coherent creation of such a 3-D histological volume is difficult to achieve, since the histological sectioning process introduces artifacts and distortions like holes, foldings, and tears.

There are currently two major approaches to create 3-D histological volumes: registration between consecutive sections [12,9,10,4,11] and registration of sections to external reference images coming from e.g. 3-D in-vivo imaging or 2-D block-face images acquired during histological sectioning [2,6,5]. A comprehensive overview of recent techniques is given by Cifor *et al.* [4]. When registering solely consecutive sections, the reconstructed structures are homogeneous but the aperture problem leads to drifts in the stack direction. These drifts cancel out *real* changes between neighboring sections, in particular for curved anatomical structures. This can be avoided by external reference images, but usually those images feature a smaller resolution or contrast than their corresponding

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 582–589, 2011.

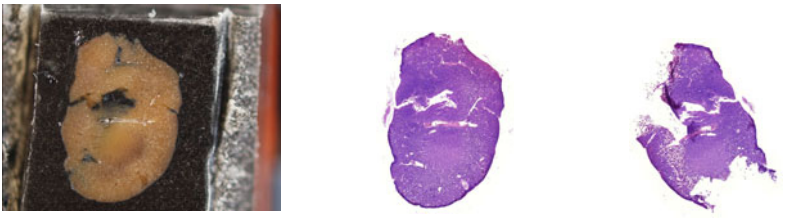

(a) (b) (c)

Fig. 1. Sample block-face image (a), its corresponding histology section (b), and a histology section that was heavily disrupted during cutting (c).

histological sections. Furthermore, the structural homogeneity between consecutive sections can hardly be guaranteed when aligning histological to reference images slice-by-slice.

To benefit from the advantages of both approaches, their sequential utilization has been proposed [13,3]. This can result in more homogeneous histological volumes of higher resolution, but comes at the cost of multiple sequential processing steps and may require a large number of empirically determined parameters. Moreover, the nature of sequential processing can easily annihilate improvements of a previous step and again worsen the overall structural homogeneity or continuity.

Our approach tackles this problem by registering histological sections simultaneously to their corresponding reference images *and* to their neighboring sections. We explicitly avoid any segmentation step during the registration process to keep propagated errors low. In order to efficiently solve the simultaneous alignment problem, we employ discrete optimization techniques for dense deformable registration using Markov random fields (MRFs) [7]. In our MRF, unary potentials account for the registration to reference images, and two distinct pairwise potentials account for the registration to neighboring slices and for regularizing transformations, respectively. This model allows us to jointly register all sections to their respective reference while maintaining structural homogeneity. To the best of our knowledge, this is the first time where histology re-stacking is performed in a single process and solved efficiently by intensity-based registration using discrete optimization.

2 Method

Given a set of histology images $\mathcal{I} = \{\mathcal{I}_1, \ldots, \mathcal{I}_n\}$ and their corresponding block-face images $\mathcal{J} = \{\mathcal{J}_1, \ldots, \mathcal{J}_n\}$ (cf. Fig. 1a, b) we seek a set of sufficiently smooth transformations $\mathbf{T} = \{T_1, \ldots, T_n\}$, which align each $\mathcal{I}_i$ to $\mathcal{J}_i$ and to its adjacent neighbors $\mathcal{I}_{i-1}, \mathcal{I}_{i+1}$. This can be modeled by an energy minimization as

$$\mathbf{T}^* = \arg\min_{\mathbf{T}} \mathbf{E}_R(\mathcal{I}, \mathcal{J}, \mathbf{T}) + \mathbf{E}_C(\mathcal{I}, \mathbf{T}) + \mathbf{E}_S(\mathbf{T}), \tag{1}$$

where $\mathbf{E}_R(\mathcal{I}, \mathcal{J}, \mathbf{T}) = \sum_{i=1}^{n} E_R(\mathcal{I}_i \circ T_i, \mathcal{J}_i)$ computes the energy between histology images and block-face images, $\mathbf{E}_C(\mathcal{I}, \mathbf{T}) = \sum_{i=0}^{n-1} E_C(\mathcal{I}_i \circ T_i, \mathcal{I}_{i+1} \circ T_{i+1})$ computes the energy between consecutive histology slices, and $\mathbf{E}_S(\mathbf{T}) = \sum_{i=1}^{n} E_S(T_i)$ acts as an independent *in-plane* regularizer on each transformation T_i.

In the remainder of this section we will first explain the steps taken to pre-align the given image stacks. Subsequently, we will focus on how to solve the optimization problem in a discrete framework.

2.1 Rigid Pre-alignment

To initialize and speed up our discrete optimization framework, we first rigidly align all pairs of histological sections and their corresponding block-face images. This can be fully automated by extracting and aligning the 2-D contours of our sample for each image pair, followed by a rigid registration. As our tissue sample is embedded in black paraffin as in [3] and hence clearly distinguishable from its background in both the histological and block-face images, we can use Otsu's automatic thresholding method to obtain all 2-D contours. For each pair of contours we compute their semi-major axes and centers based on moments and align them to obtain an initial rotation and translation for rigid registration. To be robust against the variability of intensities and visible structures between block-face and histological images, our rigid registration uses normalized mutual information.

2.2 MRF Formulation of the Deformable Stack Registration

We will now explain our MRF model for optimizing Eq. (1). In order to build an MRF model we first parameterize the transformations $\mathbf{T}$. We use a set of uniform free-form deformation (FFD) grids, i.e. the displacement field representing each T_i is parameterized using 2-D FFDs based on cubic B-splines. Please note that we choose a stack of 2D FFD grids to model the *independent* transformations that happen while cutting each slice individually. The MRF is then constructed by assigning a node to each control point $\mathbf{p}^i$ of an FFD grid $\mathcal{G}^i$. We create two types of links between nodes: (a) between neighboring *in-plane* control points, which are located in the same FFD grid $\mathcal{G}^i$, and (b) between neighbors in consecutive FFD grids $\mathcal{G}^i, \mathcal{G}^{i+1}$, see Fig. 2.

We define a *labeling* $\mathbf{l}$ as the assignment of discrete values to nodes. We associate each label assignment $l_{\mathbf{p}}$ to a corresponding displacement $\mathbf{d}_{l_{\mathbf{p}}}$ of control point $\mathbf{p}$. For a given labeling $\mathbf{l}$ we can then cast $\mathbf{E}_R$ as a sum of unary terms, and $\mathbf{E}_C, \mathbf{E}_S$ as sums of pairwise terms, respectively, leading to an overall MRF energy $\mathbf{E}(\mathbf{l})$:

$$\mathbf{E}(\mathbf{l}) = \sum_{i=1}^{n} \sum_{\mathbf{p} \in \mathcal{G}^i} \Theta_R^i(l_{\mathbf{p}}) + \gamma \sum_{i=1}^{n-1} \sum_{\substack{\mathbf{p} \in \mathcal{G}^i, \\ \mathbf{q} \in \mathcal{G}^{i+1}}} \Theta_C^{i,i+1}(l_{\mathbf{p}}, l_{\mathbf{q}}) + \lambda \sum_{i=1}^{n} \sum_{\mathbf{p} \in \mathcal{G}^i} \sum_{\mathbf{r} \in \mathcal{N}(\mathbf{p})} \Theta_S^i(l_{\mathbf{p}}, l_{\mathbf{r}}), \tag{2}$$

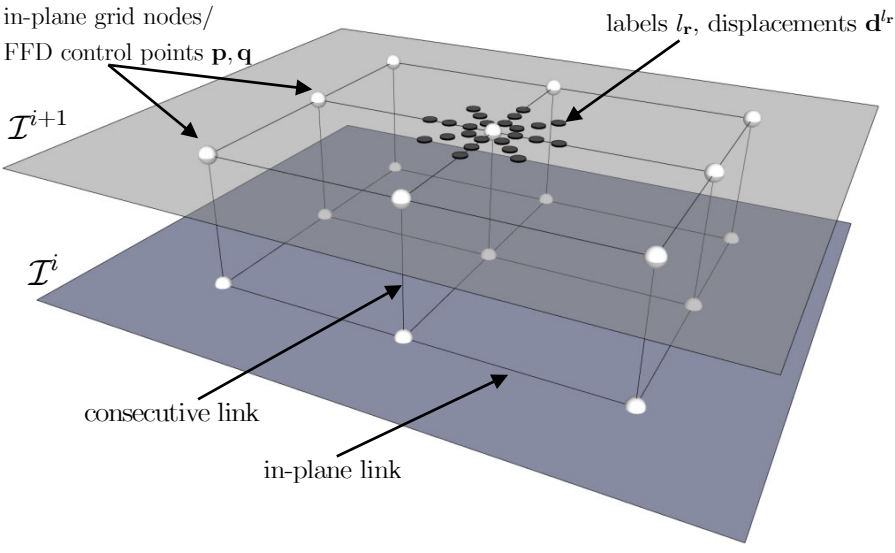

Fig. 2. MRF for two section images

where γ and λ are weighting factors to relate the three terms and $\mathcal{N}(\mathbf{p})$ denotes the set of in-plane neighbors of control point $\mathbf{p}$.

We can now compute costs for each term:

$$\Theta_R^i(l_\mathbf{p}) = \int_{\Omega_i} \eta(\mathbf{x},\mathbf{p}) D_1(\mathcal{I}_i(\mathbf{x}+\mathbf{d}_{l_\mathbf{p}})), \mathcal{J}_i(\mathbf{x}))d\mathbf{x} \tag{3}$$

$$\Theta_C^{i,i+1}(l_\mathbf{p}, l_\mathbf{q}) = \int_{\Omega_i} \eta(\mathbf{x},\mathbf{p})\eta(\mathbf{x},\mathbf{q}) D_2(\mathcal{I}_i(\mathbf{x}+\mathbf{d}_{l_\mathbf{p}}), \mathcal{I}_{i+1}(\mathbf{x}+\mathbf{d}_{l_\mathbf{q}}))d\mathbf{x} \tag{4}$$

$$\Theta_S^i(l_\mathbf{p}, l_\mathbf{r}) = \int_{\Omega_i} R(\mathbf{d}_{l_\mathbf{p}}, \mathbf{d}_{l_\mathbf{r}})d\mathbf{x}, \tag{5}$$

where $D_1(.,.), D_2(.,.)$ compute intensity-based dissimilarity measures and $R(.,.)$ computes a regularizing penalty to achieve *in-plane* smoothness of transformations. $\eta(\mathbf{x},\mathbf{p})$ is a weighting factor, which controls the influence of control point $\mathbf{p}$ to pixel $\mathbf{x}$.

Given the discrete energy formulation $\mathbf{E}$, we can perform the simultaneous deformable registration using two specific dissimilarity measures D_1, D_2, a regularizer R, and a discrete optimization algorithm. In our particular application, we use normalized mutual information for D_1 and normalized cross correlation for D_2 assuming only linear intensity changes between consecutive slices. We choose to penalize the squared difference between neighboring displacement vectors, i.e. $R(\mathbf{d}_{l_\mathbf{p}}, \mathbf{d}_{l_\mathbf{r}}) = ||\mathbf{d}_{l_\mathbf{p}} - \mathbf{d}_{l_\mathbf{r}}||^2$. For solving the discrete labeling problem (involving unary and pairwise terms only) we use the iterative quadratic pseudo-boolean optimization (QPBO) algorithm [8].

3 Results

We evaluated our method on synthetic data as well as histological sections of a rat kidney. To demonstrate the effect of our data term, we compare our method to a sole histology-to-block-face registration utilizing normalized mutual information, a registration only using consecutive histological sections utilizing normalized cross correlation, and a sequential, but not simultaneous combination of the two.

During deformable registration, we set most required parameters to the default values proposed in [7], i.e. the maximum allowed displacement of each level of the multi-scale approach is bound to the grid resolution, the sampling rate from the zero-displacement to the maximum displacement is 5, sparse sampling is used, and 5 optimization cycles are performed on each pyramid level. We however use two image and control point resolution levels. Our label set scaling factor is set to 2/3. We use 12 bins for the histograms needed for the computation of mutual information. Our weighting factors γ and λ are set to 0.2 and 2, respectively. These factors are determined experimentally and altering them did not vastly change the reconstruction results.

3.1 Synthetic Data

Our synthetic ground truth data resembles an ellipsoidal tissue sample embedded with skew tubular structures of varying diameter. To simulate the histological cutting process, we arbitrarily tear some of the sections and apply random FFDs to each section.

For realistic tearing we model each tear by its center, direction (pointing towards the closest point on the contour of the sectioned sample), and its apex angle. A tear symmetrically opens the sample within the apex angle and describes an elliptic sector whose pixels are all nullified. It also linearly affects its neighboring pixels within a specified influence angle, the further away from the direction vector, the less influence a pixel gets. The corresponding deformation field is stored as backward warping and has its maxima along the two vectors defining the elliptic sector. Deformations are linearly decreasing towards the direction vector as well as towards the vectors defined by the influence angle. Figure 3a shows a teared and deformed section.

We model staining and inherent intensity variabilities between block-face and histological images by inverting all original voxel intensities. We also add Gaussian noise to the ground-truth images ($\equiv$*block-face images*) and to the deformed inverted images ($\equiv$*histology sections*) to imitate the image acquisition process. Eventually we store the 2-D ground truth deformation fields for all pairs of deformed and ground truth images and run our method.

In order to quantitatively evaluate our method, we apply the performance measures commonly used for optical flow evaluation by the vision community [1]. In detail, we compute the *absolute* endpoint error (EE) and the *relative* angular error (AE) between the ground truth and resulting deformation fields as well as the interpolation error (IE) and normalized IE (NE) between the ground truth and deformed images. Table 1 shows the measured errors for all compared

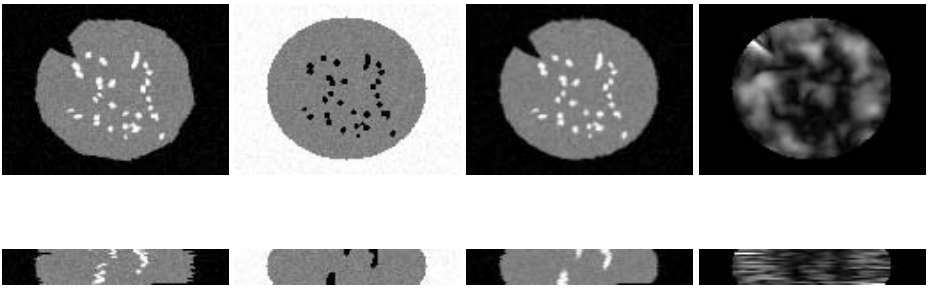

(a) source (b) reference (c) result (d) endpoint error

Fig. 3. Synthetic experimental data showing one section (top) and a cut through the stack (bottom) of the ground truth data (a, b) and our registration results in terms of the deformed image (c) and its corresponding endpoint error (d)

methods, where relative improvements compared to the initial situation (where sections are assumed to be rigidly aligned) are included. Fig. 3d shows the EE of our simultaneous approach.

It can be seen that our proposed method achieves the largest relative improvement of each error measure. The registration only comprising consecutive sections performs worst - this is due to a drift of the stack images, which increases errors in the deformation fields (AE, EE) and hence differences between the ground truth reference and deformed images (IE, NE). Moreover, it should be noted that a sequential application of histology-to-block-face and consecutive registration can decrease the quality of the alignment: the results are worse than a registration based on block-face images only, which is again related to the global drift mentioned above.

3.2 Rat Kidney

The rat kidney sample was cut into 9 μm thin sections (see Fig. 1b for a sample image). Before each cut we acquired a block-face image (cf. Fig. 1a) using an Olympus E-620 SLR camera with a 50 mm 1:2 macro objective. About 5% of our sections were disrupted during cutting, which is usually unavoidable in the histological sectioning process (cf. Fig. 1c for an example). After staining all sections with hematoxylin and eosin, we digitized them using a MIRAX MIDI whole slide scanner by Carl Zeiss. Using a grid spacing of 1 mm, we exemplarily run our registration algorithm on a stack of 580 sections (including the disrupted sections) of 1008 x 756 pixels and a pixel spacing of 33.38 μm, which took about 2.5 hours to complete on a workstation with 24 GB memory and 8 cores.

Figure 4 shows a cut in the middle of the volume orthogonal to slicing direction. Out of all techniques, ours produces the most similar structures compared to the uncut situation. In the close-ups, it can be observed that our method (Fig. 4h) performs satisfactory even in the presence of disrupted slices, while the consecutive (Fig. 4f) and sequential (Fig. 4g) methods fail and histology-to-block-face registration (Fig. 4e) produces large jitter.

Table 1. Experiments on synthetic data. The values in brackets are the relative improvements compared to the initial alignment.

method	error measure			
	AE	EE	IE	NE
initial alignment	50.84	1.57	26.06	13.28
histology-to-block-face registration	36.26 (+28.68%)	1.22 (+22.29%)	17.70 (+32.08%)	9.13 (+31.25%)
consecutive registration	57.41 (-12.92%)	1.91 (-21.66%)	25.76 (+1.15%)	12.35 (+7.00%)
sequential approach	36.56 (+28.09%)	1.23 (+21.66%)	24.97 (+4.18%)	11.72 (+11.75%)
simultaneous approach	**34.30 (+32.53%)**	**1.17 (+25.48%)**	**16.69 (+35.96%)**	**8.66 (+34.79%)**

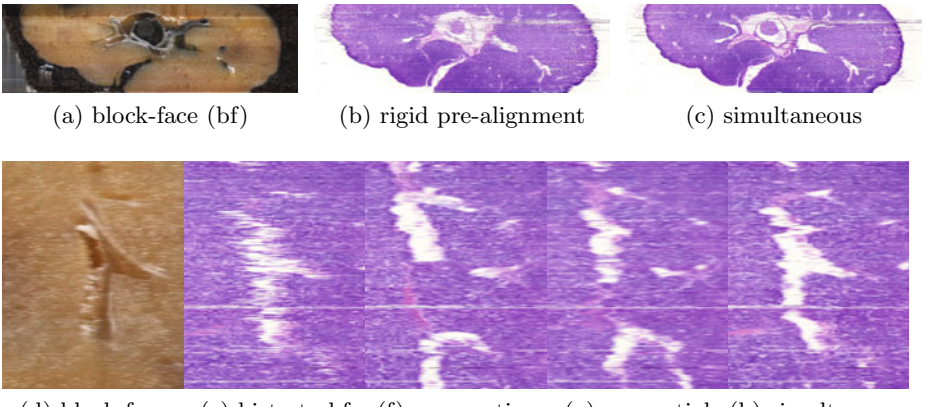

(a) block-face (bf) (b) rigid pre-alignment (c) simultaneous

(d) block-face (e) histo-to-bf (f) consecutive (g) sequential (h) simultaneous

Fig. 4. First row: sagittal cuts through the kidney. The first image is generated from block-face images, middle and right show histology images before and after simultaneous registration. Second row: exemplary close-up of a vessel structure.

4 Conclusion

In this paper we develop a method for the fully automatic reconstruction of 3-D histology stacks. Our approach guarantees geometrical coherence by combining the registration of sections to block-face images with a registration between neighboring slices. This allows us to undo deformations induced by cutting while aligning anatomical structures that are not apparent in the block-face images. Experimental results show that our simultaneous registration algorithm is not only accurate, but also robust against artifacts produced during the cutting process. Moreover, its sensitivity to misalignment due to grossly corrupted slices

is almost negligible. Motivated by our excellent results, we intend to extend our current implementation with adequate streaming techniques to be able to increase the resolution of input and output data in the future.

Acknowledgments. We want to thank F. Pfeiffer and M. Bech from the physics department (E17), Technische Universität München, for providing the kidney data set. Parts of this research were supported by the Alexander von Humboldt Foundation providing a Feodor Lynen Research Fellowship and the Excellence Cluster Initiative, Munich Center for Advanced Photonics.

References

1. Baker, S., Scharstein, D., Lewis, J., Roth, S., Black, M.J., Szeliski, R.: A database and evaluation methodology for optical flow. IJCV 92(1), 1–31 (2011)
2. Bardinet, E., Ourselin, S., Dormont, D., Malandain, G., Tandé, D., Parain, K., Ayache, N., Yelnik, J.: Co-registration of histological, optical and MR data of the human brain. In: Dohi, T., Kikinis, R. (eds.) MICCAI 2002. LNCS, vol. 2488, pp. 548–555. Springer, Heidelberg (2002)
3. Chakravarty, M.M., Bedell, B.J., Zehntner, S.P., Evans, A.C., Collins, D.L.: Three-dimensional reconstruction of serial histological mouse brain sections. In: ISBI, pp. 987–990 (2008)
4. Cifor, A., Pridmore, T., Pitiot, A.: Smooth 3-D reconstruction for 2-D histological images. In: Prince, J.L., Pham, D.L., Myers, K.J. (eds.) IPMI 2009. LNCS, vol. 5636, pp. 350–361. Springer, Heidelberg (2009)
5. Dauguet, J., Delzescaux, T., Cond, F., Mangin, J., Ayache, N., Hantraye, P., Frouin, V.: Three-dimensional reconstruction of stained histological slices and 3D non-linear registration with in-vivo mri for whole baboon brain. Journal of Neuroscience Methods 164, 191–204 (2007)
6. Gefen, S., Tretiak, O., Nissanov, J.: Elastic 3D alignment of rat brain histological images. IEEE TMI 22(11), 1480–1489 (2003)
7. Glocker, B., Komodakis, N., Tziritas, G., Navab, N., Paragios, N.: Dense image registration through MRFs and efficient linear programming. MedIA 12(6), 731–741 (2008)
8. Kolmogorov, V., Rother, C.: Minimizing nonsubmodular functions with graph cuts-a review. IEEE PAMI 29(7), 1274–1279 (2007)
9. Schmitt, O., Modersitzki, J., Heldmann, S., Wirtz, S., Fischer, B.: Image registration of sectioned brains. IJCV 73(1), 5–39 (2006)
10. Tan, Y., Hua, J., Dong, M.: Feature curve-guided volume reconstruction from 2D images. In: ISBI, pp. 716–719 (April 2007)
11. Bagci, U., Bai, L.: Automatic best reference slice (BRS) selection for smooth volume reconstruction of a mouse brain from histological images. IEEE TMI 29(9), 1688–1696 (2010)
12. Wirtz, S., Papenberg, N., Fischer, B., Schmitt, O.: Robust and staining-invariant elastic registration of a series of images from histologic slices. In: Medical Imaging: Image Processing, vol. 5747, pp. 1256–1262 (2005)
13. Yushkevich, P.A., Avants, B.B., Ng, L., Hawrylycz, M., Burstein, P.D., Zhang, H., Gee, J.C.: 3D mouse brain reconstruction from histology using a coarse-to-fine approach. In: Pluim, J.P.W., Likar, B., Gerritsen, F.A. (eds.) WBIR 2006. LNCS, vol. 4057, pp. 230–237. Springer, Heidelberg (2006)

Spatially Adaptive Log-Euclidean Polyaffine Registration Based on Sparse Matches

Maxime Taquet[1,2], Benoît Macq[1], and Simon K. Warfield[2]

[1] ICTEAM Institute, Université Catholique de Louvain, Louvain-La-Neuve, Belgium
[2] Computational Radiology Laboratory, Harvard Medical School, USA
maxime.taquet@uclouvain.be

Abstract. Log-euclidean polyaffine transforms have recently been introduced to characterize the local affine behavior of the deformation in principal anatomical structures. The elegant mathematical framework makes them a powerful tool for image registration. However, their application is limited to large structures since they require the pre-definition of affine regions. This paper extends the polyaffine registration to adaptively fit a log-euclidean polyaffine transform that captures deformations at smaller scales. The approach is based on the sparse selection of matching points in the images and the formulation of the problem as an expectation maximization iterative closest point problem. The efficiency of the algorithm is shown through experiments on inter-subject registration of brain MRI between a healthy subject and patients with multiple sclerosis.

Keywords: Registration, Polyaffine, Log-Euclidean, Structure Tensor.

1 Introduction

Medical image registration is used in a variety of applications, from atlas construction to intraoperative navigation. The goal is to find a transform T that maps an image I onto another image J. A valuable property of the transform is *diffeomorphism* which guarantees invertibility and topology preservation [9].

Different models can be chosen for T [4]. Dense field models, as returned by diffeomorphic demons [9], are very flexible. However, models with fewer parameters are typically more robust. Log-euclidean polyaffine transforms (LEPT) [2,6] are compact and diffeomorphic transforms. They are built by composing affine transforms in the log-domain, and have been successfully used to register predefined anatomical structures [6].

The nature of tissues suggests that higher order deformations could also be modeled by LEPT with more degrees of freedom. However, predefining affine regions at a lower scale would be too cumbersome or impossible.

This paper introduces a registration method that adaptively fits a LEPT based on matching points sparsely selected in the image. The main contribution is the extension of the expectation-maximization iterative closest point (EM-ICP [7]) to use LEPT. This scheme accounts for matching ambiguities due to regularities in the image and naturally integrates regularization as a prior on T.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 590–597, 2011.

The rest of this paper is organized as follows. Section 2 introduces the elements of the algorithm. Section 3 shows results on brain inter-subject registration. Section 4 concludes and presents some directions for future work.

2 Methods

The proposed method, summarized in Table 1, aims at optimizing the parameters of a LEPT to map image J onto I. The location of the anchors of the affine components are first defined. The parameters of the affine components are then estimated in a multi-scale approach. At each scale, corresponding points are selected and their prior matching probabilities are recorded. The optimization then alternates between updating the posterior matching probabilities and optimizing the transform parameters. The image structure tensor is used in the definition of these probabilities to account for matching ambiguities.

Table 1. Summary of the proposed method

1:	Define the anchor locations $\boldsymbol{a}_k$ in I.
2:	**for** $r = 0$ **to** $R - 1$
3:	Compute the structure tensor $\mathcal{S}$ of image I at scale r.
4:	Compute the LEPT weights $\boldsymbol{w}(\boldsymbol{x})$ at scale level r.
5:	Select the K best matches $C = \{(\boldsymbol{m}_i, \boldsymbol{s}_{i,j}, \pi_{ij})\}$ between I and $T^{(r-1)} \circ J$.
6:	**for** $i = 1$ **to** N_{it} **do**
7:	*E-Step* Update the probabilities: $C^{(r)} \leftarrow \text{Update}(C^{(r-1)}; T^{(r-1)}, \mathcal{S})$.
8:	*M-Step* Optimize the affine corrections $\delta L^{(r)} \leftarrow \text{Optimize}(T^{(r)}; C^{(r)}, \mathcal{S})$.
9:	Composition of the affine components: $\exp(L_k^{(r)}) \leftarrow \exp(L_k^{(r-1)}) \exp(\delta L_k^{(r)})$.
10:	Interpolation of the transform: $T^{(r)} \leftarrow \text{Interpolate}(L^{(r)}, \boldsymbol{w}(\boldsymbol{x}))$.
11:	**end for**
12:	**end for**

2.1 Log-Euclidean Polyaffine Transforms

Log-euclidean polyaffine transforms (LEPT) are defined as a weighted composition of N affine transforms in the logarithm domain [2]. In other words, under the action of an LEPT, the point $\boldsymbol{x}$ is transformed to $\boldsymbol{x} + T(\boldsymbol{x})$ with

$$T(\boldsymbol{x}) = \exp\left(\sum_{k=1}^{K} w_k(\boldsymbol{x}) L_k \tilde{\boldsymbol{x}}\right), \tag{1}$$

where $\exp(.)$ is the exponential-map, $L_k \in \mathbb{R}^{3\times 4}$ is the principal logarithm of an affine matrix in homogenous coordinates (with the last line ignored), $w_k(\boldsymbol{x})$ are the weights which depend on the anchor location $\boldsymbol{a}_k$, and $\tilde{\boldsymbol{x}} = [\boldsymbol{x}, 1]^T$. LEPT have remarkable properties. They are invertible and their inverse is a LEPT [2].

In this paper, the location of the anchors are determined a priori based on the local intensity heterogeneity of the image. More precisely, a measure of heterogeneity $h_I(\boldsymbol{x})$ is computed at every $\boldsymbol{x}$ (see Sec. 2.3). The anchors are defined as

the K-means centroids of the point cloud spanned by all $\boldsymbol{x}$ s.t. $h_I(\boldsymbol{x}) > \bar{h}_I(\boldsymbol{x})$, where $\bar{h}_I(\boldsymbol{x})$ is a sliding mean of heterogeneity. Hence, the anchors will tend to concentrate more in areas with high local contrast. Given K anchors $\boldsymbol{a}_k$, we want to optimize the $12K$ parameters corresponding to the L_k's.

2.2 Block Matching

Block matching is used to establish a dense correspondence between points in image I and J. As a similarity measure, the correlation coefficient (CC) is used because of its invariance under any linear mapping of intensities. Let B_m (resp. B_s) be the block centered at $\boldsymbol{m}_i$ (resp. $\boldsymbol{s}_{i,j}$) in image I (resp. J), the CC is:

$$\rho(\boldsymbol{m}_i, \boldsymbol{s}_{i,j}) = \frac{\mathrm{Cov}(B_m, B_s)}{\sqrt{\mathrm{Var}(B_m)\mathrm{Var}(B_s)}}.$$

For each block in I, the CC of the best match $\rho_{\max}(\boldsymbol{m}_i)$ is compared to the mean of all CC for this block $\bar{\rho}(\boldsymbol{m}_i)$. The N points with the highest ratio $\rho_{\max}/\bar{\rho}$ are selected. The prior matching probability is modeled as a normal distribution of the root mean squared error (MSE_n=2−2ρ) between normalized blocks (zero mean, unit variance):

$$\pi_{ij} = \frac{1}{\sqrt{2\pi\sigma_e^2}} \exp\left(-\frac{2 - 2\rho(\boldsymbol{m}_i, \boldsymbol{s}_{i,j})}{2\sigma_e^2}\right).$$

Block matching is not robust to noise and suffers from the aperture problem. To deal with this issue, M matches in J are recorded for each of the N selected points in I, and we will choose $N \gg K$. The remaining of this paper explains how the LEPT can be robustly estimated based on this set C of matching pairs.

2.3 Image Structure Tensor

The location of the matches can be ambiguous due to the regularity of the intensity profile around the point. This problem can be addressed by anisotropically weighting the error with the structure tensor [5].

The structure tensor at $\boldsymbol{x}_k$, $S(\boldsymbol{x}_k) \in \mathbb{R}^3$ is defined as the autocorrelation of the intensity gradient $\boldsymbol{\nabla} I(\boldsymbol{x})$ in a neighborhood Ω_k:

$$S(\boldsymbol{x}_k) = E_{\Omega_k}\left\{\boldsymbol{\nabla} I(\boldsymbol{x}) \boldsymbol{\nabla} I(\boldsymbol{x})^T\right\} = \frac{1}{|\Omega_k|} \sum_{\boldsymbol{x} \in \Omega_k} \boldsymbol{\nabla} I(\boldsymbol{x}) \boldsymbol{\nabla} I(\boldsymbol{x})^T.$$

The normalized structure tensor, $\mathcal{S}(\boldsymbol{x})$=$S(\boldsymbol{x})/||S(\boldsymbol{x})||_2$, is positive semidefinite with a maximum eigenvalue of 1 and an eigenvector aligned with $\nabla I(\boldsymbol{x})$ in Ω_k.

Consequently, the norm of $\boldsymbol{e}' = \mathcal{S}\boldsymbol{e}$ is more affected by the component of $\boldsymbol{e}$ parallel to $\nabla I(\boldsymbol{x})$. In other words, the overall error is less affected by matching ambiguities due to regular structures than by errors made in the direction of $\nabla I(\boldsymbol{x})$.

Besides weighting the errors, the structure tensor is also used to define anchor locations. Indeed, its highest eigenvalue λ_3 is significantly higher than zero only for heterogenous areas, making it natural to define $h_I(\boldsymbol{x}) = \lambda_3(\boldsymbol{x})$.

2.4 Transformation Estimation

Given the set C of matching pairs with prior probabilities π_{ij}, we propose to estimate the transform T by maximizing the joint log-likelihood of C and T [7,1]:

$$T^* = \arg\max_T E\left\{\log P\left(C,T\right)\right\} = \arg\max_T E\left\{\log P\left(C|T\right)\right\} + E\left\{\log P\left(T\right)\right\}. \quad (2)$$

The first term of (2) tends to honor the detected correspondences. The second term is a prior on T. This term favors some transforms over others, based on intrinsic properties of the transforms only. This is a statistical interpretation of the regularization energy used in [4,6,9]. We will use the following prior:

$$\log P(T) = -\lambda^2 \sum_{k,l} a_{k,l} ||L_k - L_l||^2 + \text{cst},$$

where $||.||^2$ is the Frobenius norm and $a_{k,l} = \sum_{\boldsymbol{x}} w_k(\boldsymbol{x}) w_l(\boldsymbol{x}) \left(\frac{1}{\sum_{\boldsymbol{x}} w_k(\boldsymbol{x})} + \frac{1}{\sum_{\boldsymbol{x}} w_l(\boldsymbol{x})}\right)$ is an overlapping coefficient. This expression states that transforms are more likely if nearby affine components (components that share a common affecting area) are close to each other. Interestingly, this term is equivalent to the regularization energy defined in [6].

An efficient method to optimize (2) is EM-ICP [7]. This algorithm consists in alternatively optimizing the criterion for C considering T fixed (*E-step*) and for T considering C fixed (*M-step*). If enough matching pairs are selected, EM-ICP is very robust to noise. It is thus well suited to cope with the block matching issues. The E-step simply results in computing the matching probabilities:

$$\overline{(C)_{ij}} = \frac{\pi_{ij} p(\boldsymbol{s}_{i,j}|\boldsymbol{m}_i, T)}{\sum_k \pi_{ik} p(\boldsymbol{s}_{i,j}|\boldsymbol{m}_k, T)}, \quad (3)$$

where the expression of $p(\boldsymbol{s}_{i,j}|\boldsymbol{m}_i, T)$ is accounts for the matching ambiguities:

$$p(\boldsymbol{s}_{i,j}|\boldsymbol{m}_i, T) = \exp\left(-\frac{||\mathcal{S}(\boldsymbol{m}_i)(\boldsymbol{m}_i - T * \boldsymbol{s}_{i,j}||^2}{2\sigma_n^2}\right).$$

The M-step then consists in optimizing the parameters of T:

$$T^* = \arg\max_T \sum_{i=1}^{N}\sum_{j=1}^{M} \overline{(C)_{ij}} \log p(\boldsymbol{s}_{i,j}|\boldsymbol{m}_i, T) - \lambda^2 \sum_{k,l=1}^{K} a_{k,l} ||L_k - L_l||^2. \quad (4)$$

Unlike affine transforms, this equation does not have an obvious solution for the parameters L_k, due to the exponential map of (1). Therefore, we propose a first order approximation in which $T(\boldsymbol{x})$ is approximated by:

$$T(\boldsymbol{x}) \approx \mathbb{I} + \sum_{k=1}^{N} w_k(\boldsymbol{x}) L_k \tilde{\boldsymbol{x}}.$$

In that case, optimizing (4) amounts to the quadratic programming (QP):

$$l^* = \arg\min_l (Hl - D)^T S(Hl - D) + \lambda^2 l^T A l, \tag{5}$$

where:

- $l \in \mathbb{R}^{12K}$ is the vector of the elements of L_k taken row-wise and concatenated,
- $H \in \mathbb{R}^{3N \times 12K}$ is the interpolation matrix. Each 3×12 block $[H]_{ij}$ corresponds to point $\boldsymbol{m}_i = (x_i, y_i, z_i)$ and the j^{th} anchor $\boldsymbol{a}_j$ and is equal to $I_3 \otimes (w_j(\boldsymbol{m}_i)\tilde{\boldsymbol{m}}_i)^T$, where I_3 is the 3×3 identity matrix and $\otimes$ stands for the Kronecker product,
- $D \in \mathbb{R}^{3N}$ is the vector obtained by taking, for each $\boldsymbol{m}_i$, the barycenter of its matches $\overline{\boldsymbol{s}}_i$ weighted by $\overline{(C)_{ij}}$ (the equivalence between criterion (4) and the use of barycenters is justified in [7].),
- $S \in \mathbb{R}^{3N \times 3N}$ is the block diagonal matrix of the structure tensor. Each 3×3 block on the diagonal corresponds to the structure tensor at point $\boldsymbol{m}_i$,
- $A \in \mathbb{R}^{12K \times 12K}$ is the prior matrix with $(A)_{k \neq l} = -a_{k,l}$ and $(A)_{kk} = \sum_l a_{k,l}$.

In practice, the weights w_i and overlapping coefficients $a_{k,l}$ are thresholded out, so that H, S and A are sparse. Taking the derivative of (5) w.r.t. l and setting it to 0 yields the linear system:

$$(H^T S H + \lambda^2 A) l = H^T S D. \tag{6}$$

Interestingly, the regularization term is optimized simultaneously with the similarity term, unlike [6] where an ad-hoc one step gradient descent is performed on the regularization energy. Alternating the estimation of matching probabilities (3) and the estimation of parameters (6) until convergence results in a globally optimal LEPT that best fits the observations of the block matching.

2.5 Weights: The Kriging Estimator

In [6,2], the weights $w_k(\boldsymbol{x})$ are simply a normalized Gaussian function of the distance between the point and the anchor. Here we use the Kriging estimator (KE) to define these weights, as in [8]. The KE has the advantage of adapting the weights to the spatial distribution of anchors in a statistically sound way.

Let us interpret LEPT as random fields of matrix logarithms $L(\boldsymbol{x})$ for which L_k are observations at locations $\boldsymbol{a}_k$. In this interpretation, the weights are the coefficients of a linear estimator of L at $\boldsymbol{x}$. KE is a best linear unbiased estimator for the field $L(\boldsymbol{x})$, given a (presumably valid) model of its spatial correlation encoded in a variogram: $\gamma(\boldsymbol{x}, \boldsymbol{y}) = E\left\{|L(\boldsymbol{x}) - L(\boldsymbol{y})|^2/2\right\}$. Given $(\boldsymbol{\gamma}_a(\boldsymbol{x}))_i = \gamma(\boldsymbol{a}_i, \boldsymbol{x})$, the weights $\boldsymbol{w}(\boldsymbol{x}) = (w_1(\boldsymbol{x}) \ldots w_K(\boldsymbol{x}))^T$ are obtained by solving the linear system:

$$\begin{bmatrix} \boldsymbol{w}(\boldsymbol{x}) \\ \mu \end{bmatrix} = \Gamma(\boldsymbol{x})^{-1} \boldsymbol{\gamma}_a(\boldsymbol{x}), \text{ with } \Gamma(\boldsymbol{x}) = \begin{bmatrix} (\gamma(\boldsymbol{a}_i, \boldsymbol{a}_j)) & \mathbf{1} \\ \mathbf{1}^T & 0 \end{bmatrix} \in \mathbb{R}^{(K+1)\times(K+1)}$$

where μ is the Lagrange multiplier ensuring unbiasedness of the estimate. In this paper an exponential isotropic variogram is used: $\gamma(\boldsymbol{x}, \boldsymbol{y}) = 1 - e^{-||\boldsymbol{x}-\boldsymbol{y}||/t}$.

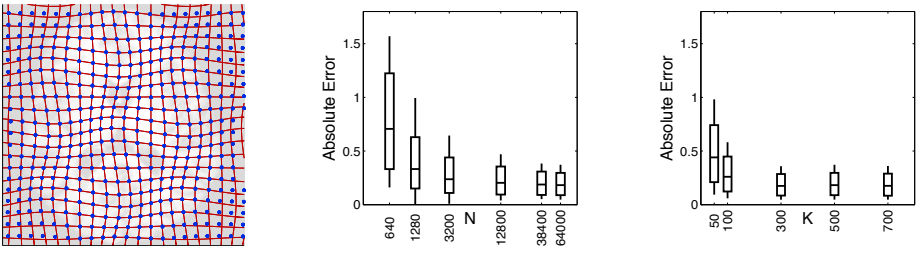

Fig. 1. (Left) LEPT (displayed as checkpoints) are able to accurately recover the synthetic field (deformed grid) within the brain volume (Middle) Influence of N on the recovering accuracy (the $x-$axis is logarithmic) (Right) Influence of K on the recovering accuracy.

3 Experiments and Results

The method was tested on a dataset of ten brain T1-MRI (resolution: 256×256 ×176). A synthetic experiment was first carried out. Inter-subject registration was then performed. Unless otherwise mentioned, parameter values are: K=500, N=$(20+5r)^3$ at scale r, R=5, M=20, N_{it}=5, t=40, λ^2=0.3, σ_e=σ_n=5. For block matching, blocks of 5^3 voxels and searching region of 9^3 voxels are used.

3.1 Synthetic Experiments

The knowledge of a ground true deformation helps understanding how the algorithm behaves. A synthetic field $T(\boldsymbol{x}) = \sin(\frac{\pi x}{50})\cos(\frac{\pi y}{50})(3,3,3)^T$ was applied to the image. Both adaptive LEPT and diffeomorphic demons accurately recovered T within the brain volume, with mean absolute error of 0.21 and 0.12 respectively (Fig. 1) and were invertible (min. jacobian of 0.67 and 0.65 respectively).

The registration was then performed for different number N of sparse matches and K of anchors. The evolution of the accuracy with N (Fig. 1) suggests that the method is not sensitive to the number of selected matches as long as this number is high compared to K. For lower N, the affine transforms can no more be robustly estimated and the performances collapse. The evolution of the accuracy with K (Fig. 1) tells us that no loss of accuracy is observed between $K = 700$ and $K = 300$, suggesting that the number of selected matches is the bottleneck here. For lower K, the performance decreases, but the method does not diverge, since all affine components can still be robustly estimated.

3.2 Inter-subject Registration

The method was then applied to register ten multiple-sclerosis (MS) patients to one healthy subject. Each brain was skull stripped and affinely registered to the subject. Images were manually segmented by an expert in a validated protocol, providing an external validation criterion for the registration. More precisely, the Dice's coefficient between the subject label image and the label image of the patients after alignment were computed for each tissue (Fig. 2).

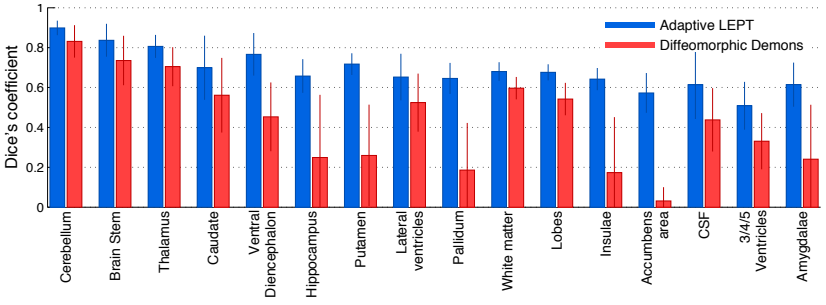

Fig. 2. Dice's coefficient with a 95% CI. Adaptive LEPT aligns structures better than diffeomorphic demons. The difference is especially significant for low contrast structures such as the putamen and the insluae.

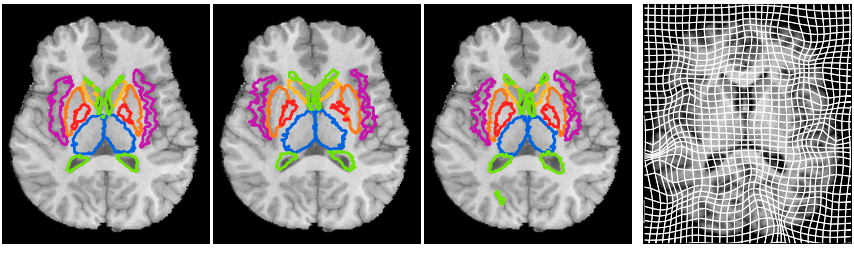

Fig. 3. (Left to right) Subject with labeled tissues, projected labels of the patient's tissues after diffeomorphic demons registration and after adaptive LEPT registration, patient image with the deformed grid overlaid

As a whole, adaptive LEPT aligns tissues better than diffeomorphic demons. For high contrast regions (*e.g.* lateral ventricles), both methods achieve comparable results. However, diffeomorphic demons tend to excessively favor these regions at the cost of a very poor alignment of low contrast structures (*e.g.* putamen). For these structures, the difference of performance between the two algorithms is strongly significant.

Fig. 3 depicts the aligned contours for one slice, along with the deformed grid. Again, the Jacobian of the field never fell under 0 (min. jacobian of 0.12). However, a strong pinching effect appeared in the lobe regions. These regions are typically subject specific and, while good pairs of local matches can be detected, they may be misleading in the transform estimation. Finally, note that the non-rigid deformation of the structures (as seen *e.g.* by the bending of the lateral ventricles) would not be captured if a single affine region was defined for them as in the previous LEPT registration framework.

4 Conclusion and Future Work

This paper introduced a registration algorithm that adaptively fits a LEPT based on a set of sparse matches. Results on inter-subject registration show that LEPT

are able to capture the local affine deformations occurring at small scales. In a future work, we want to investigate an adaptive way of incrementally defining the anchor locations such as in [3]. The choice of location could include a cost related to the confidence of the estimation in order to avoid the pinching effect observed in Fig. 3.

Acknowledgements. MT thanks the B.A.E.F. and the F.R.S.-F.N.R.S for their financial support and the reviewers for their wise advices. This investigation was supported in part by NIH grants R01 RR021885, R01 EB008015, R03 EB008680 and R01 LM010033.

References

1. Akselrod-Ballin, A., Bock, D., Reid, R., Warfield, S.: Improved registration for large electron microscopy images. In: IEEE ISBI, pp. 434–437. IEEE, Los Alamitos (2009)
2. Arsigny, V., Commonwick, O., Ayache, N., Pennec, X.: A Fast and Log-Euclidean Polyaffine Framework for Locally Linear Registration. JMIV 33(2), 222–238 (2009)
3. Buerger, C., Schaeffter, T., King, A.: Hierarchical adaptive local affine registration for fast and robust respiratory motion estimation. MedIA (2011)
4. Cachier, P., Bardinet, E., Dormont, D., Pennec, X., Ayache, N.: Iconic feature based nonrigid registration: the pasha algorithm. CVIU 89(2-3), 272–298 (2003)
5. Clatz, O., Delingette, H., Talos, I., Golby, A., Kikinis, R., Jolesz, F., Ayache, N., Warfield, S.: Robust nonrigid registration to capture brain shift from intraoperative mri. IEEE TMI 24(11), 1417–1427 (2005)
6. Commowick, O., Arsigny, V., Isambert, A., Costa, J., Dhermain, F., Bidault, F., Bondiau, P., Ayache, N., Malandain, G.: An efficient locally affine framework for the smooth registration of anatomical structures. MedIA 12(4), 427–441 (2008)
7. Granger, S., Pennec, X.: Multi-scale EM-ICP: A fast and robust approach for surface registration. In: Heyden, A., Sparr, G., Nielsen, M., Johansen, P. (eds.) ECCV 2002. LNCS, vol. 2353, pp. 418–432. Springer, Heidelberg (2002)
8. Ruiz-Alzola, J., Westin, C., Warfield, S., Alberola, C., Maier, S., Kikinis, R.: Nonrigid registration of 3d tensor medical data. MedIA 6(2), 143–161 (2002)
9. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Diffeomorphic demons: Efficient non-parametric image registration. NeuroImage 45(1), S61–S72 (2009)

Personalized X-Ray Reconstruction of the Proximal Femur via Intensity-Based Non-rigid 2D-3D Registration

Guoyan Zheng

Institute for Surgical Technology and Biomechanics,
University of Bern, CH-3014, Bern, Switzerland
Guoyan.Zheng@ieee.org

Abstract. This paper presents a new approach for reconstructing a patient-specific shape model and internal relative intensity distribution of the proximal femur from a limited number (e.g., 2) of calibrated C-arm images or X-ray radiographs. Our approach uses independent shape and appearance models that are learned from a set of training data to encode the *a priori* information about the proximal femur. An intensity-based non-rigid 2D-3D registration algorithm is then proposed to deformably fit the learned models to the input images. The fitting is conducted iteratively by minimizing the dissimilarity between the input images and the associated digitally reconstructed radiographs of the learned models together with regularization terms encoding the strain energy of the forward deformation and the smoothness of the inverse deformation. Comprehensive experiments conducted on images of cadaveric femurs and on clinical datasets demonstrate the efficacy of the present approach.

1 Introduction

Constructing a personalized three-dimensional (3D) model from a limited number of calibrated C-arm images or X-ray radiographs and a statistical model has drawn more and more attention. The reported techniques can be split into two main categories: those based on statistical models of shape [1,2,3] and those based on statistical models of shape and intensity [4,5,6,7,8,9]. The methods belonging to the former category typically require an implicit or explicit image segmentation which is error-prone and hard to achieve automatically. The errors in segmentation may lead to errors in the final reconstruction. In contrast, shape-intensity statistical model based methods directly compare the input reference images with the floating simulation images called digitally reconstructed radiographs (DRR), which are obtained by ray casting of a volume data instantiated from the learned model. No segmentation is required.

Most of the shape-intensity statistical model based methods are closely related with the Active Appearance Models (AAM) pioneered by Cootes et al. [10]. According to Matthews and Baker [11], there are basically two types of linear shape and appearance models, those model shape and appearance independently, and those which parameterize shape and appearance with a single

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 598–606, 2011.

set of linear parameters. They referred the first set as *independent shape and appearance models* and the second as *combined shape and appearance models*. One good property associated with the independent formulation is that given a hypothesized shape the associated optimal appearance/intensity parameters may be determined analytically.

The contribution of this paper is a new approach for reconstructing a patient-specific shape model and internal relative intensity distribution of the proximal femur from sparse calibrated X-ray images. Unlike existing approaches, where they used either combined shape and appearance models [7] or a shape statistical model together with a density model approximated by Bernstein polynomials [4] [6] [8] [9], our approach uses independent shape and appearance models that are learned from a set of training data. An intensity-based non-rigid 2D-3D registration algorithm is then proposed to deformably fit the learned models to the input images. The fitting is conducted iteratively by minimizing the dissimilarity between the input images and the associated DRRs of the learned models together with regularization terms encoding the strain energy of the forward deformation and the smoothness of the inverse deformation.

2 Construction of the Statistical Models

We constructed two independent shape and appearance models of the proximal femur, one for the left side and the other for the right side. Each independent shape and appearance model was constructed directly from CT data of 20 dry cadaveric proximal femurs of the associated side based on a two-stage procedure as described below.

Our two-stage independent shape and appearance model construction procedure follows the idea introduced by Rueckert et al. [12]. In the first stage, we chose one of the proximal femur volumes as the reference volume $\mathbf{I}_0^{1st}$ and the diffeomorphic Demons algorithm [13] was used to establish the dense correspondences between the reference volume and each one of the 19 floating volumes. The output from the diffeomorphic Demons algorithm include: (1) the dense deformation fields $\{\mathbf{d}_i^{1st}; i = 1, ..., 19\}$; and (2) the 19 non-rigidly deformed floating volumes $\{\mathbf{I}_i^{1st}; i = 1, ..., 19\}$. Each deformation field $\mathbf{d}_i^{1st}$ was expressed as a concatenation of 3D vectors describing the deformation at each voxel of the reference volume and each non-rigidly deformed floating volume $\mathbf{I}_i^{1st}$ as a concatenation of gray values of each voxel in the reference volume. From these data, we computed the average deformation field $\bar{\mathbf{d}}^{1st} = (n^{-1}) \cdot \sum_{i=1}^{n} \mathbf{d}_i^{1st}$ and the average intensity distribution $\bar{\mathbf{I}}^{1st} = ((n+1)^{-1}) \cdot \sum_{i=0}^{n} \mathbf{I}_i^{1st}$.

The purpose of the second stage is to remove the possible bias introduced by the reference volume selection. To achieve this goal, we applied the average deformation field $\bar{\mathbf{d}}^{1st}$ to the reference volume $\mathbf{I}_0^{1st}$ to create a new volume $\mathbf{s}_0$ and assigned the average intensity distribution $\bar{\mathbf{I}}^{1st}$ to this newly created volume. This new volume $\mathbf{s}_0$ was then used as the new reference volume in the second stage and all input 20 proximal femur volumes were regarded as the floating volumes. The diffeomorphic Demons algorithm [13] was used again to establish

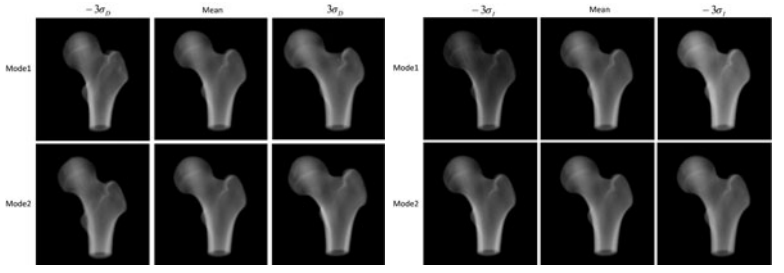

Fig. 1. Projections of the mean and the first two modes of variations of the shape (left) and the intensity (right) statistical models

the dense correspondences between the new reference volume and the other 20 floating volumes.We thus obtained a set of 20 new dense deformation field $\{\mathbf{d}_i; i = 1, ..., 20\}$ and a set of 20 non-rigidly deformed floating volumes $\{\mathbf{I}_i; i = 1, ..., 20\}$. We could then separately construct the shape statistical model and the intensity statistical model. The shape statistical model was constructed using following equations.

$$\begin{aligned} &\mathbf{S}_D = ((m-1)^{-1}) \cdot \textstyle\sum_{i=1}^{m} (\mathbf{d}_i - \bar{\mathbf{d}})(\mathbf{d}_i - \bar{\mathbf{d}})^T \\ &\bar{\mathbf{d}} = (m^{-1}) \cdot \textstyle\sum_{i=1}^{m} \mathbf{d}_i \\ &\mathbf{P}_D = (\mathbf{p}_D^1, \mathbf{p}_D^2, ...); \mathbf{S}_D \cdot \mathbf{p}_D^i = (\sigma_D^i)^2 \cdot \mathbf{p}_D^i \\ &\mathbf{d} = \bar{\mathbf{d}} + \textstyle\sum_{k=1}^{M_D} \alpha_D^k \sigma_D^k \mathbf{p}_D^k \end{aligned} \tag{1}$$

where $m = 20$ is the number of training samples; $\bar{\mathbf{d}}$ and $\mathbf{S}_D$ are the average and the covariance matrix of the displacement fields, respectively; $\{(\sigma_D^i)^2\}$ and $\{\mathbf{p}_D^i\}$ are the eigen values and the eigen vectors of the shape statistical model, respectively; $\{\alpha_D^i\}$ are the model parameters; M_D is the cut-off points.

Fig. 1 shows the projections of the mean and the first two modes of variations of the shape (left) and the intensity (right) statistical models. Each instance of the shape statistical model was generated by evaluating $\mathbf{I}_s = \bar{\mathbf{I}}(\mathbf{s}_0 + \bar{\mathbf{d}} + \alpha_D \sigma_D^i \mathbf{p}_D^i)$, with $\alpha_D \in \{-3, 0, 3\}$. Each instance of the intensity statistical model was obtained by evaluating $\mathbf{I}_I = \bar{\mathbf{I}} + \alpha_I \sigma_I^i \mathbf{p}_I^i$, with $\alpha_I \in \{-3, 0, 3\}$, where $\{(\sigma_I^i)^2\}$ and $\{\mathbf{p}_I^i\}$ are the eigen values and the eigen vectors of the intensity statistical model, respectively; $\bar{\mathbf{I}} = (m^{-1}) \cdot \sum_{i=1}^{m} \mathbf{I}_i$ is the average intensity distribution.

From the instances of the shape statistical model (Fig. 1, left), it can be observed that although the instantiated intensity model is the same (here it is the average intensity distribution $\bar{\mathbf{I}}$), the change of shapes already encode the change of the relative intensity distribution inside the instantiated shape model (e.g., the thickness of the cortical bone along the shaft is different across the instantiated volumes). Thus, we argue that it is sufficient to search only in the space of shapes in order to estimate the personalized shape model and its internal relative intensity distribution. Based on this argument, we developed an intensity-based non-rigid 2D-3D registration algorithm as presented below.

3 Intensity-based Non-rigid 2D-3D Registration

In this work, we assume that we have a set of $Q \geq 2$ X-ray images and that all images are calibrated and co-registered to a common coordinate system called $\mathbf{c}$. Given an initial estimation of the registration parameters, our algorithm iteratively generates a volume and update the parameter estimation by minimizing the dissimilarity between the input images and the associated DRRs that are created from the instantiated volume.

Volume Instantiation and Alignment. The volume instantiation and alignment process is parameterized by two sets of parameters, i.e., the set of shape parameters $\mathbf{b} = (\alpha_D^1, \alpha_D^2, ..., \alpha_D^{M_D})^T$ determining the forward deformation from the reference volume $\mathbf{s}_0$ to an instantiated volume $\mathbf{s}$ and the set of parameters $\mathbf{a} = (\Lambda_x, \Lambda_y, \Lambda_z, \beta, \gamma, \theta, t_x, t_y, t_z)^T$ determining the scaled rigid transformation from the space of the instantiated volume $\mathbf{s}$ to the common coordinate system $\mathbf{c}$, where the first three are scaling parameters; the middle three are rotational parameters and the last three are translational parameters. An instantiated volume that is aligned to the common coordinate system $\mathbf{c}$ is defined by following equation:

$$\bar{\mathbf{I}}(x_c(\mathbf{a}, \mathbf{b})) = \bar{\mathbf{I}}(\mathbf{A}(\mathbf{a}) \circ \mathbf{W}(\mathbf{b}) \circ x_0) \tag{2}$$

where $\mathbf{A}(\mathbf{a})$ is the scaled rigid transformation and $\mathbf{W}(\mathbf{b})$ is the forward deformation.

Eq. (2) describes a forward warping that should be interpreted as follows. Given a voxel x_0 in $\mathbf{s}_0$, the destination of this voxel under the forward transformation is $x_c(\mathbf{a}, \mathbf{b}) = \mathbf{A}(\mathbf{a}) \circ \mathbf{W}(\mathbf{b}) \circ x_0$. The aligned instance at voxel $x_c(\mathbf{a}, \mathbf{b})$ is set to the intensity $\bar{\mathbf{I}}(x_0)$, which then allows creating DRRs by simulating X-ray projection.

Implementing this forward warping to generate the model instance without holes is tricky and is best performed by backward warping. More specifically, as all the input images are calibrated, we can compute a back-projection ray for each pixel in the input images. We then do a sampling to get a set of discrete points along each ray. The intensity values at these points are then obtained by the backward warping with following inverse transformation:

$$x_0 = (\mathbf{W}(\mathbf{b}))^{-1} \circ ((\mathbf{A}(\mathbf{a}))^{-1} \circ x_c(\mathbf{a}, \mathbf{b})) \tag{3}$$

It is straightforward to compute the inverse of the scaled rigid transformation $\mathbf{A}(\mathbf{a})$. However, it is tricky to compute the inverse of the forward deformation $\mathbf{W}(\mathbf{b})$. In this paper, given an estimation of the shape parameters $\mathbf{b} = (\alpha_D^1, \alpha_D^2, ..., \alpha_D^{M_D})^T$, we use the fixed-point approach as proposed by Chen et al. [14] to invert the forward deformation field.

Energy formulation. Given an instantiated volume, the two sets of parameters are updated iteratively by minimizing following energy function:

$$E(\mathbf{a}, \mathbf{b}) = E_{Image}(\mathbf{a}, \mathbf{b}) + \rho_1 \cdot \sum_{k=1}^{M_D} (\alpha_D^k)^2 (\sigma_D^k)^2 + \rho_2 \cdot \int_{\mathbf{s}} ||\nabla((\mathbf{W}(\mathbf{b}))^{-1})||_2 dx \tag{4}$$

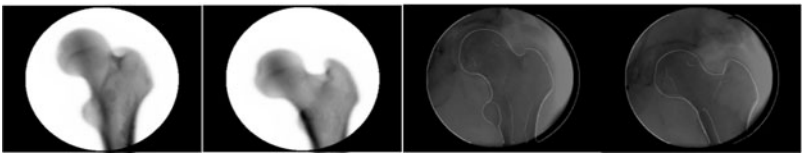

Fig. 2. Validation on a clinical dataset. The left two images show the projections of the reconstructed volume while the right two images show the original C-arm images superimposed by the contours detected from the left two images.

where the first term is the image dissimilarity energy term; the second term is the strain energy of the forward deformation, which, according to Cootes and Taylor [15], can be used to enforce the prior probabilities on the shape's deformations; the last term is a diffusion-like regularization term; ρ_1 and ρ_2 are regularization parameters.

We have chosen to use the robust dissimilarity measure that we introduced in [16] to compare the floating DRRs to the associated reference X-ray images. This dissimilarity measure is defined by modeling the qth observed difference image as a Markov random field with respect to the rth order neighborhood system $N = \{N^r_{i,j}\}$ and is described by following equation:

$$\begin{aligned} E_{Image}(\mathbf{a},\mathbf{b}) = \sum_{q=1}^{Q} [\lambda \sum_{i,j}^{I,J} D^2_{q;(i,j)}(\mathbf{a},\mathbf{b}) + \\ (1-\lambda) \sum_{i,j}^{I,J} \frac{1}{card(N^r_{i,j})} \sum_{(i',j')\in N^r_{i,j}} (D_{q;(i,j)}(\mathbf{a},\mathbf{b}) - D_{q;(i',j')}(\mathbf{a},\mathbf{b}))^2] \end{aligned} \tag{5}$$

where $I \times J$ is the size of each X-ray image; $D_q = \{D_{q;(i,j)}\}$ is the qth observed difference image. We refer interesting readers to [16] for the details about how the difference images are computed and about the details of above equation. The advantage of using such an energy function is that it has a least-squares form and can be effectively minimized by a Levenberg-Marquardt optimizer.

The registration Algorithm. Considering the fact that regularization on a vector field can be implemented as a kernel convolution [18], we developed following intensity-based non-rigid 2D-3D registration algorithm.

Algorithm (Intensity-based non-rigid 2D-3D Registration). The following two stages are executed until the convergence of the algorithm.

- Scaled rigid registration stage: The shape parameters are fixed to the current estimation $\mathbf{b}_t$ and the Levenberg-Marquardt optimizer is used to iteratively minimize the image dissimilarity energy $E_{image}(\mathbf{a},\mathbf{b}_t)$ in order to obtain a new estimation of the scaled rigid transformation parameters $\mathbf{a}_{t+1}$.
- Non-rigid registration stage: The scaled rigid transformation parameters are fixed to $\mathbf{a}_{t+1}$ and the Levenberg-Marquardt optimizer is used again to iteratively estimate the new shape deformation parameters $\mathbf{b}_{t+1}$. At each iteration, following two steps are performed.

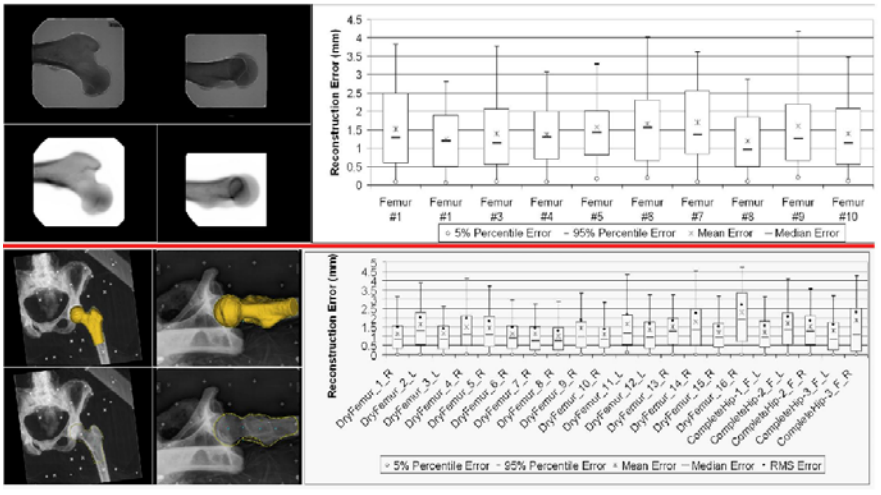

Fig. 3. Results of applying the present approach on sparse calibrated X-ray images. Top row: a reconstruction example as well as the errors of reconstructing all 10 femurs in Group I; bottom row: a reconstruction example as well as the errors of reconstructing all 21 femurs in Group II. In this example, the surface model (yellow) extracted from the reconstructed volume and its silhouettes are superimposed on the X-ray images.

- Step 1: Following the Levenberg-Marquardt optimizer, compute the gradient and the regularized Hessian of the first two terms of Eq. (4) with respect to the shape parameters, and then calculate an additive update $\Delta\mathbf{b}_t$ of the shape parameters to get a new estimation $\mathbf{b'}_t = \mathbf{b}_t + \Delta\mathbf{b}_t$.
- Step 2: Compute the forward deformation based on $\mathbf{b'}_t$ and then invert the forward displacement fields using the fixed-point approach [14]. To implement the third term of Eq. (4), the inverse displacement fields are further regularized by applying a Gaussian kernel convolution. The regularized inverse displacement fields will then be used to instantiate a new volume to generate DRRs for the next iteration.

For all the experiments reported below, the scaled rigid transformation parameters $\mathbf{a}$ were initialized with an anatomical landmarks based registration. All the shape parameters $\mathbf{b}$ were initialized to zeros.

4 Experiments and Results

We conducted comprehensive experiments on 31 cadaveric femurs with different shape (none of them belongs to training population) and 3 clinical datasets to

validate the present approach. The 31 cadaveric femurs were divided into two groups. Group I consists of 10 left dry cadaveric femurs. For each femur in this group, we acquired two calibrated C-arm images around the proximal femur region. Group II contains the rest 21 femurs (7 left and 14 right) where 5 of them are connected with hips. For each femur in this group, we acquired two calibrated X-ray radiographs. Each one of the 3 clinical datasets contains two calibrated C-arm images of a patient. The cut-off point M_D was chosen to be 9 such that more than 90% of the total amount of variations is explained. The registration algorithm was implemented with CPU computation and no attempt was made to optimize its performance. Thus, one registration can take about 15 minutes. It can be accelerated by using modern GPU technology.

Experiment on 3 Clinical Datasets. In this experiment, due to the lack of ground truth, we mainly used these datasets to demonstrate qualitatively the performance of the present approach in clinical settings. Fig. 2 shows a reconstruction example when the present approach was applied to one of the three clinical datasets.

Experiment on 10 Cadaveric Femurs in Group I. In this experiment, the reconstruction accuracies were evaluated by randomly digitizing dozens points from the surface of each femur and then computing the distances from those digitized points to the associated surface model which was segmented from the reconstructed volume. A reconstruction example and the errors of reconstructing volumes of all 10 femurs in this group are shown in the top row of Fig.3. An average mean reconstruction accuracy of 1.5 mm was found.

Experiment on 21 Cadaveric Femurs in Group II. In this experiment, we used the method proposed by Aspert et al.[17] to compute the reconstruction errors, where the ground truth surface models were either obtained with a CT-scan reconstruction method or with a hand-held laser-scan reconstruction method (T-SCAN, Steinbichler, Neubeuern, Germany). The surface models segmented from the reconstructed volumes were then compared to the associated ground truth models to evaluate the reconstruction accuracies. A reconstruction example and the errors of reconstructing volumes of all 21 femurs in this group are shown in the bottom row of Fig. 3. An average mean reconstruction accuracy of 1.4 mm was found.

5 Conclusions

We presented a new approach to reconstruct a personalized shape model and internal relative intensity distribution of the proximal femur from sparse calibrated X-ray images. Our approach used independent shape and appearance models to encode the *a priori* information. An intensity-based non-rigid

2D-3D registration algorithm was then developed to fit the learned models to the input images. Results from our comprehensive experiments demonstrated its efficacy.

References

1. Fleute, M., Lavallée, S.: Nonrigid 3-D/2-D registration of images using statistical models. In: Taylor, C.J., Colchester, A.C.F. (eds.) MICCAI 1999. LNCS, vol. 1679, pp. 138–147. Springer, Heidelberg (1999)
2. Zheng, G., Gollmer, S., Schumann, S., Dong, X., Feilkas, F., Gonzalez Ballester, M.A.: A 2D/3D correspondence building method for reconstruction of a patient-specific 3D bone surface model using point distribution models and calibrated X-ray images. Med. Image Anal. 13, 883–899 (2009)
3. Baka, N., Niessen, W.J., Kaptein, B.L., van Walsum, T., Ferrarini, L., Reiber, J.H.C., Lelieveldt, B.P.F.: Correspondence free 3D statistical shape model fitting to sparse X-ray projections. In: Dawant, B.M., Haynor, D.R. (eds.) SPIE Medical Imaging 2010, vol. 7623, pp. 76230D-1–76230D9 (2010)
4. Sadowsky, O., Chintalapani, G., Taylor, R.H.: Deformable 2D-3D registration of the pelvis with a limited field of view, using shape statistics. In: Ayache, N., Ourselin, S., Maeder, A.J. (eds.) MICCAI 2007, Part II. LNCS, vol. 4792, pp. 519–526. Springer, Heidelberg (2007)
5. Hurvitz, A., Joskowicz, L.: Registration of a CT-like atlas to fluoroscopic X-ray images using intensity correspondences. Int. J. CARS 3, 493–504 (2008)
6. Chintalapani, G., Taylor, R.H.: Integrating statistical models of bone density into shape based 2D-3D registration framework. In: PMMIA 2009, pp. 1–11 (2009)
7. Humbert, L., Whitmarsh, T., De Craene, M., del Rio Barquero, L.M., Fritscher, K.D., Schubert, R., Eckstein, F., Link, T.M., Frangi, A.F.: 3D reconstruction of bone shape and bone mineral density distribution of the femur from DXA images. In: ISBI 2010, pp. 456–459. IEEE, Los Alamitos (2010)
8. Ahmad, O., Ramamurthi, K., Wilson, K.E., Engelke, K., Prince, R.L., Taylor, R.H.: Volumetric DXA (VXA) - A new method to extract 3D information from multiple in vivo DXA images. J. Bone Miner. Res. 25, 2468–2475 (2010)
9. Sadowsky, O., Lee, J., Sutter, E.G., Wall, S.J., Prince, J.L., Taylor, R.H.: Hybrid cone-beam tomographic reconstruction: incorporation of prior anatomical models to compensate for missing data. IEEE T. Med. Imaging 30, 69–83 (2011)
10. Cootes, T.F., Edwards, G.J., Taylor, C.J.: Active appearance models. IEEE T. Pattern Anal. 23, 681–685 (2001)
11. Matthews, I., Baker, S.: Active appearance models revisited. Int. J. Comput. Vision 60, 135–164 (2004)
12. Rueckert, D., Frangi, A.F., Schnabel, J.A.: Automatic construction of 3D statistical deformation models using non-rigid registration. In: Niessen, W., Viergever, M. (eds.) MICCAI 2001. LNCS, vol. 2208, p. 77–84. Springer, Heidelberg (2001)
13. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Non-parametric diffeomorphic image registration with the demons algorithm. In: Ayache, N., Ourselin, S., Maeder, A.J. (eds.) MICCAI 2007, Part II. LNCS, vol. 4792, pp. 319–326. Springer, Heidelberg (2007)

14. Chen, M., Lu, W., Chen, Q., Ruchala, K.J.: A simple fixed-point approach to invert a deformation field. Med. Phys. 35, 81–88 (2008)
15. Cootes, T.F., Taylor, C.J.: Combining point distribution models with shape models based on finite element analysis. Image Vis. Computing 13, 403–409 (1995)
16. Zheng, G.: Effective incorporating spatial information in a mutual information based 3D-2D registration of a CT volume to X-ray images. Comput. Med. Imag. Grap. 34, 553–562 (2010)
17. Aspert, N., Santa-Cruz, D., Ebrahimi, T.: Mesh: measuring error between surfaces using the hausdorff distance. In: ICME 2002, pp. 705–708. IEEE, Los Alamitos (2002)
18. Nielsen, M., Florack, L., Deriche, R.: Regularization, scale-space, and edge detection filters. J. Math. Imaging Vis. 7, 291–307 (1997)

2D Image Registration in CT Images Using Radial Image Descriptors

Franz Graf[1], Hans-Peter Kriegel[1], Matthias Schubert[1], Sebastian Pölsterl[1], and Alexander Cavallaro[2]

[1] Institut für Informatik, Ludwig-Maximilians-Universität München
Oettingenstr. 67, D-80538 München, Germany
{graf,kriegel,schubert}@dbs.ifi.lmu.de, poelsterl@cip.ifi.lmu.de
[2] Radiologisches Institut, Universtätsklinikum Erlangen
Friedrich-Alexander-Universität Erlangen-Nürnberg,
Maximiliansplatz 1 D-91054 Erlangen, Germany
alexander.cavallaro@uk-erlangen.de

Abstract. Registering CT scans in a body atlas is an important technique for aligning and comparing different CT scans. It is also required for navigating automatically to certain regions of a scan or if sub volumes should be identified automatically. Common solutions to this problem employ landmark detectors and interpolation techniques. However, these solutions are often not applicable if the query scan is very small or consists only of a single slice. Therefore, the research community proposed methods being independent from landmark detectors which are using imaging techniques to register the slices in a generalized height scale. In this paper, we propose an improved prediction method for registering single slices. Our solution is based on specialized image descriptors and instance-based learning. The experimental evaluation shows that the new method improves accuracy and stability of comparable registration methods by using only a a single CT slice is required for the registration.

Keywords: Computer Tomography, Similarity Search, Retrieval, Localization.

1 Introduction

CT scans play an important role in the field of medical imaging. Even though CT scans often comprise the complete thorax or even the complete body of a patient, the clinician often only requires a small sub volume of a scan. For example, if a body region should be compared between two scans for a differential diagnosis. In the state of the art work flow both volume scans (each up to more than 1 GB) have to be loaded completely from the clinical PACS (Picture Archiving and Communication System) over the network before they can be aligned manually by the clinician. If instead the clinician could load just the relevant part of a scan from the PACS, this would reduce loading time and also the time needed for navigating through the scans and aligning the positions.

In this paper, we want to address this problem and apply a radial image descriptor [2] for the identification of similar body regions. In the above use

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 607–614, 2011.

case, the clinician then only needs to load the relevant sub volume. Also he would not have to navigate through the complete scan as this task could be sped up by the similarity search using the radial descriptor.

The key contribution of this paper is the application of a radial image descriptor for the identification of similar body regions. The descriptor is independent from any landmark detectors and provides a smaller prediction error as well as an improved stability of the similarity search compared to the reference algorithm [4] that deals with the same problem. Also the descriptor outperforms the algorithm proposed in [5] which requires volumes as query objects. Finally, the descriptor is more robust in body regions where the reference algorithm shows large errors.

The rest of this paper is organized as follows: We describe related work in Sec. 2, then the process of the feature extraction is described in Sec. 3. In Sec. 4 the localization is described. The evaluation is shown in Sec. 5. The paper ends with a conclusion and an outlook on future work (Sec. 6).

2 Related Work

The localization problem could be solved by using the meta data of the DICOM header of a CT image. However, the available information is often error-prone. Gueld et al [6] report that several entries in the DICOM header are often imprecise or even completely wrong. To find a more appropriate solution to the problem, the research community proposed some methods for registering slices to a general atlas with standardized height. The authors of [3] propose to predict the body region from a topogram based on landmarks with invariant positions. A similar approach is proposed in [7] where the mapping is based on a look-up table using 8 landmarks which are detected in various fashions. Seifert et al [9] propose a method to detect up to 19 invariant slices and single point landmarks in full body scans by using PBT and HAAR features. Nevertheless, applying this method requires the availability of a sufficient part of the landmarks and the corresponding landmark detectors. Recently the authors of [4] published a new method using a single slice of a CT scan for registration by using a multi-represented feature descriptor (MR-Descriptor). Feulner et al [5] deal with a similar localization problem. They propose a solution using SURF features with visual words and require sub volumes with a minimum height of 44 mm for the registration. Both [4] and [5] are used for the evaluation of the radial descriptor proposed in this paper.

3 Feature Extraction

The process of generating the compound radial image descriptor consists of the following steps: unifying the image resolutions, extracting the patient's body and combining the two image descriptors to a single radial descriptor.

The resolution of a CT image is determined by the setting of the associated recording device and may vary depending on several external factors. Thus it

is needed to scale the image I to a common resolution (1.5 px/mm) to obtain scale invariance between different scans. The resulting image is defined as I_S. In order to separate the body from the rest of the image, a compound region detection is performed on I_S. A compound region is defined as an area of pixels which is enclosed by a contour of pixels with $p(x,y) > \tau$. $p(x,y)$ defines the Hounsfield Unit (HU) value of a pixel at the position (x,y) and τ defines the according threshold (-500 HU). The resulting compound regions are extracted by starting a contour tracing algorithm from each pixel $\in I_S$ with $p(x,y) > \tau$. The applied algorithm is implemented by using the *analyze particles* function of ImageJ [1] which adapts the well known contour tracing algorithm of Pavlidis [8]. Afterwards the bounding box of the largest compound region defines the region of interest (ROI) represented by the area of the patient's body on the image I_S(cf. Fig. 1(b)). I_S is then cropped to this ROI, building the image I_{ROI}.

Next, a radial sector/shell model comparable to [2] (c.f. Fig. 1(c)) is created from which the two descriptors representing dense structures (bones) and soft tissues (like organs etc.) are extracted. Both descriptors are represented by the circumcircle of I_{ROI} with radius r. They consist of n_y shells and n_x sectors resulting in $n_x \cdot n_y = i$ bins. For each bin i, both the number of pixels of interest (POI) p_i and the number of other pixels (NPOI) n_i is calculated. A POI is defined as a pixel with $p(x,y) \geq \psi_1$ or $p(x,y) \leq \psi_2$ (depending on the type of descriptor, which are described subsequently). The values of bins $\notin I_{ROI}$ are set to a penalty value (-0.25) to achieve a larger difference between descriptors from regions with different aspect ratios.

The first descriptor represents the distribution of bones. Thus, ψ_1 is set 300 HU and the set of all POIs is defined by $p_i = |\{p(x,y) \in I_{ROI} | p(x,y) \geq \psi_1\}|$. Regarding the spatial distribution of the bones, it can be said that the outer shells of the descriptor are more relevant than the shells in the middle of the body where hardly any bones are detected. Thus each bin of the descriptor is weighted w.r.t. the shell index $i \in [1, n_x]$, so that $p_i = p_i \cdot shell(i)^2$. Evaluating the parameters n_x and n_y showed the slightly best results with $n_x = 24$ and $n_y = 11$. Nevertheless, the impact to $\varepsilon_{\text{mean}}$ and σ (if n_x and n_y are changed by ± 4) are less than 2 mm in case of $\varepsilon_{\text{mean}}$ and less than 6 mm for the standard deviation σ, so that the choice of these parameters is not very critical.

Some areas in the human body like the abdomen show a comparatively small amount of dense structures. Therefore, a descriptor representing the location and arrangement of soft tissues is created. The threshold for this descriptor is set to $\psi_2 = -500$ HU. The set of POIs is defined by $p_i = |\{p(x,y) \in I_{ROI} | p(x,y) \leq \psi_2\}|$ in this case. For this descriptor the amount of shells and sectors were set to $n_x = 18$ and $n_y = 8$. Same as above, the impact of changing the parameter values on the accuracy is very small. In contrast to the previous descriptor, the weighting is fixed for this descriptor, so that $p_i = p_i \cdot n_y^2$.

Finally, both descriptors are concatenated to a single feature vector q. An additional step is the application of a principal component analysis to reduce the dimensionality of the feature vectors. In our experiments, the dimensionality could be reduced to 50 dimension without losing accuracy.

Table 1. Parameter setting for both descriptors

Type	ψ	Sectors n_x	Shells n_y	Angle ϕ	Weighting	Bins
Bones	$\geq 300\,\mathrm{HU}$	24	11	15°	quadratic	240
Soft	$\leq -500\,\mathrm{HU}$	18	8	20°	equal	144

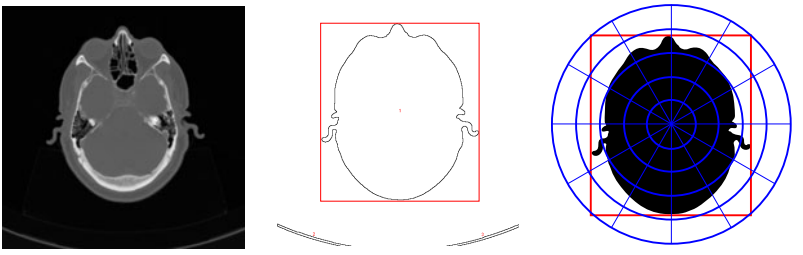

(a) (b) (c)

Fig. 1. Visualization of the feature extraction process for a neck scan image Fig. 1(a): the image is rescaled and the body (in this case the head) is detected Fig. 1(b) and approximated by a bounding box. Afterwards the sector/shell model is created Fig. 1(c) from which the features are extracted.

4 Localization

The task of the prediction method is to localize the query vector q representing the query slice with unknown position q_z to a value $z \in [0, 1]$ in the standardized height model. For this task, we use a database DB containing all feature vectors fv to the corresponding CT slices of n volume scans. DB can thus be regarded as the atlas of the method. Additionally, each $fv \in DB$ is annotated with the position $fv_z \in [0, 1]$ of the according CT slice. For the prediction, we propose a two level $k$nn search, in order to avoid overfitting to a single CT scan: First, the $k_1$nn to q are computed for each of the n scans $\in DB$, so that $k_1 \cdot n$ results are returned, building the set S. This limitation of k_1 results per volume scan is done to avoid that all results origin from a single CT scan. Next, the $k_2$nn to q in S are computed and stored in set T. The result q_z of the prediction is defined by the mean of all position labels in T's feature vectors: $q_z = \frac{1}{k_2} \sum_{fv \in T} fv_z$.

5 Evaluation

All experiments were conducted on a database consisting of 98 CT scans (38 neck and 60 thorax scans) with the collection covering the complete area between the top of the head to the end of the coccyx. The scans were recorded from 71 different patients (43 male, 31 female, 1 unknown) with the age of 4 to 86 years. Each of the patients contributed no more than 1 thorax and 1 neck scan. The data set was recorded with 120 kVp from 5 different types of Siemens CT

scanners[1] and comprises 53 437 DICOM images using more than 27 GB disk space. The resolution along the z-axis varies between 66 and 1 749 slices per scan (= 0.5 - 5 mm/slice). The resolution along the x- and y-axis varies between 1.09 - 2.32 px/mm.

The ground truth was created by a manual annotation of all scans with the following landmarks which were selected by a medical expert. In neck scans: cranial crista galli, cranial sella turcica, cranial dens axis, caudal plate of cervical vertebrae #4, caudal plate of cervical vertebrae #7. In thorax scans: cranial sternum, caudal xiphoid process, caudal plate of thoracic vertebrae #12, sacral promontory, caudal os coccygis. The positions of slices are defined in a standard model $\in [0; 1]$.

All experiments were conducted using a leave-one-out scheme on each of the 97 scans in the database. Each slice of a scan was used as test object while DB did not contain the volume scan which q belongs to. Therefore, overfitting to the same patient could be excluded. The error is determined by the difference between the predicted height q_z and the annotated position of the slice, from which q was extracted. In order to obtain an interpretable result, a height of 1.80 m for each patient was assumed so that all error measurements are relative to this height. We used $k_1 = 1$ and $k_2 = 3$ which proved the best results. Larger values for $k1$ and k_2 only result in a very minor decrease in performance.

The feature extraction takes 55 ms/70 ms for the bone/soft descriptors. Cross validation was accomplished in less than 39 minutes (40 ms/slice). More advanced feature reduction techniques combined with spatial index structures could further decrease the search time, yet we mainly focused on minimizing the prediction error in this paper.

To measure the prediction quality of the radial descriptor, we measured the mean error $\varepsilon_{\text{mean}}$ and standard deviation σ of the prediction. Applying the prediction method introduced in Sec. 4, we were able to reduce the mean error to $\varepsilon_{\text{mean}} = 1.8\,\text{cm}$ and the standard deviation to $\sigma = 2.81\,\text{cm}$. Compared to the MR-Descriptor proposed in [4], this means a reduction of the mean error and standard deviation by a factor of almost 2.5 (cf. $\varepsilon_{\text{mean}} = 4.45\,\text{cm}$, $\sigma = 9.15\,\text{cm}$).

The improvement of σ is visualized by the box plots of Fig. 2: It can be seen that the upper whiskers in the lower diagram of Fig. 2 are clearly lower throughout the whole dataset than in the case of the MR-Descriptor. The improvement of $\varepsilon_{\text{mean}}$ is most significant in the areas between 0-10 cm (representing the head) and $> 70\,\text{cm}$ (region of the hips). Only in the small region of [66 cm, 71 cm], the radial descriptor is outperformed by about 1.8 cm w.r.t. $\varepsilon_{\text{mean}}$ by the MR-Descriptor. Nevertheless the large errors values are still much smaller in this region using the radial descriptor.

In order to evaluate the distribution of the error value, we calculated the cumulative distribution function (CDF) of ε, $F_E(\varepsilon) = F(E \leq \varepsilon)$. Comparing to the MR-Descriptor, we observed an improvement on the complete CDF: The probability for errors less than 1 cm ($F_E(\leq 1\,cm)$) was 0.51 compared to 0.3. $F_E(\varepsilon) \geq 0.9$ was hit at $\varepsilon = 4\,\text{cm}$ for the radial descriptor and $\varepsilon = 8\,\text{cm}$ for the

[1] Definition AS+, Sensation 10/16/64, Volume Zoom.

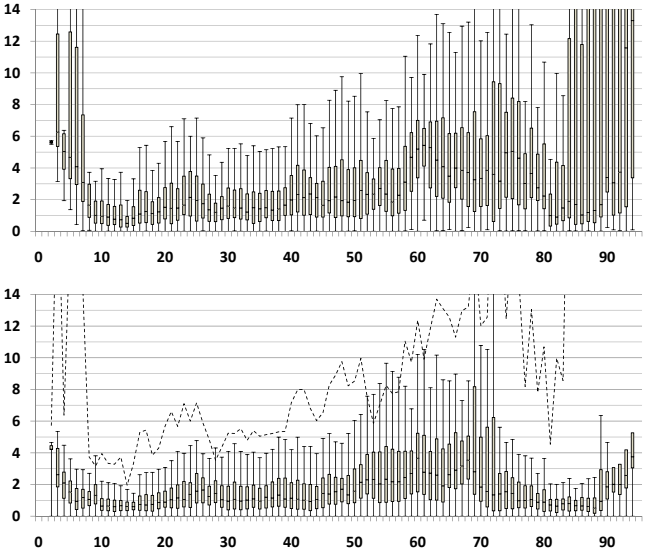

Fig. 2. Comparison of $\varepsilon_{\text{mean}}$ of the MR-descriptor (upper plot) to the radial descriptor (lower plot). The x-axis displays body regions of 1 cm width starting with the head at $x = 0$. The y-axis displays the amount of errors in cm. The dashed line in the lower diagram indicates the upper whiskers of the upper diagram for easier comparison.

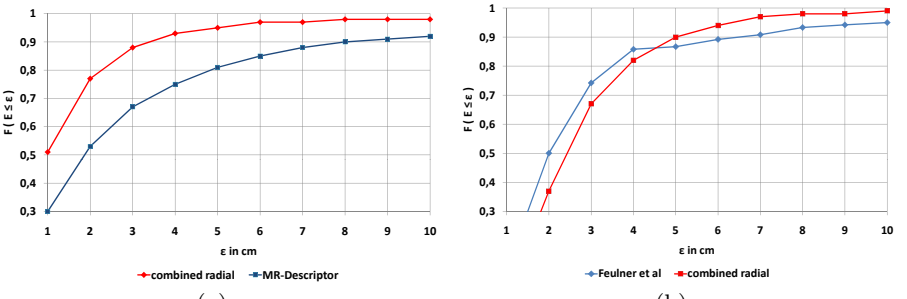

Fig. 3. CDF of errors from the radial descriptor (red) compared to the MR-Descriptor Fig. (a). Fig. (b) compares our descriptor (1 slice/query) to the work of [5] (4.4 cm query volumes) on all volumes containing the required landmarks

MR-Descriptor. This means that we could guarantee 90% of the errors to be less than ±4 cm, while the MR-Descriptor could only garantee ε to be less than ±8 cm with the same confidence. The complete CDF can be seen in Fig. 3(a).

Comparing the scatter plots of Fig. 4 it can clearly be seen that the result of the localization is much more stable in Fig. 4(b) than in Fig. 4(a). Especially large errors in the region [0 cm, 20 cm] and > 75 cm could almost be eliminated. The problematic regions [20 cm, 30 cm] (shoulder) and [60 cm, 75 cm] (abdomen)

can still be identified as a source for larger errors but the overall amount of errors in these regions was also lowered significantly (c.f. CDF in Fig. 3(a)).

We also compared the approach to the work of Feulner et al [5]. As their proposed algorithm is designed for query volumes instead of query slices, we chose sub volumes with a size of 4.4 cm for their algorithm and single slice queries for the radial descriptor. This means of course, that the radial descriptor is using less slices and thus less information for the retrieval. Also, the approach of Feulner et al is based on landmarks, so that we had to reduce the data set to 17 volumes (12 male, 5 female, 6 547 slices) which contained the according landmarks. In Fig. 3 the CDFs for both approaches can be compared. It can be seen that the break even for the radial descriptor is at 4.5 cm. In case of $\varepsilon \leq 4.5$ cm, the work of Feulner et al performs better. Nevertheless $\varepsilon \leq 5$ cm are measured with a confidence of 0.9, whereas Feulner et al yield the same confidence at 6.5 cm. This means that the amount of larger errors is smaller in the case of the radial descriptor even though only a single query slice is used.

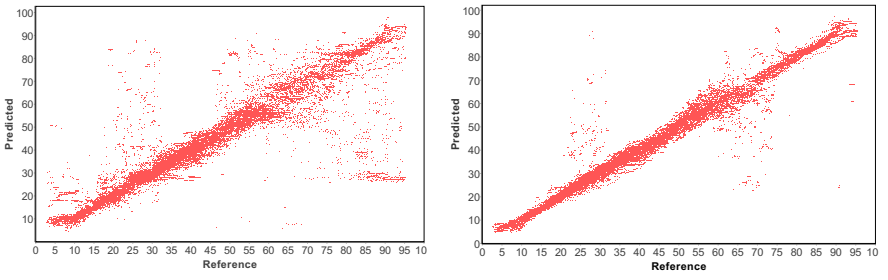

(a) (b)

Fig. 4. Scatter plots displaying the results of the prediction. (a) shows the plot of the MR-Descriptor, (b) shows the plot of the Radial Descriptor. The x-component of a pixel denotes the true position of an image, the y-component describes the prediction.

6 Conclusion

In this paper, we proposed to apply a combined radial descriptor for registering single CT slices in a generalized height scale of the human body. The descriptor is independent of any landmark detectors and combines a descriptor regarding bone structures with a descriptor regarding soft tissues. Both descriptors are based on a sector/shell partitioning schema to provide locality sensitivity. Experiments were conducted on a large real world dataset, which was manually annotated by consulting a clinical expert. The descriptor was compared against two state of the art methods. One, based on landmarks and volume queries, another based on single slice queries. Using the radial descriptor significantly reduces the average prediction errors and the likelihood of large prediction errors by up to 2.5× for the method having the identical setting. For future work, we plan to extend

the prediction method in a way to support query volumes as well as single slice queries. We also plan to apply boosting and instance selection for further improvement of the feature representation.

Acknowledgments. This research has been supported in part by the THESEUS program in the MEDICO and CTC projects. They are funded by the German Federal Ministry of Economics and Technology under the grant number 01MQ07020. The responsibility for this publication lies with the authors.

References

1. Abramoff, M., Magelhaes, P., Ram, S.: Image processing with Image. J. Biophotonics International 11(7), 36–42 (2004)
2. Belongie, S., Malik, M., Puzicha, J.: Shape context: A new descriptor for shape matching and object recognition. In: NIPS, pp. 831–837 (2000)
3. Bürger, C.: Automatic Localisation of Body Regions in CT Topograms. VDM, Saarbrücken (2008)
4. Emrich, T., Graf, F., Kriegel, H.P., Schubert, M., Thoma, M., Cavallaro, A.: CT slice localization via instance-based regression. In: Proceedings of the SPIE Medical Imaging 2010: Image Processing (SPIE), San Diego, CA, USA, vol. 7623 (2010)
5. Feulner, J., Zhou, S.K., Seifert, S., Cavallaro, A., Hornegger, J., Comaniciu, D.: Estimating the body portion of CT volumes by matching histograms of visual words. In: Proceedings of the SPIE Medical Imaging 2009 Conference (SPIE), Lake Buena Vista, FL, USA, vol. 7259, p. 72591V (2009)
6. Güld, M.O., Kohnen, M., Keysers, D., Schubert, H., Wein, B.B., Bredno, J., Lehmann, T.M.: Quality of DICOM header information for image categorization. In: Proceedings of the SPIE Medical Imaging 2002: PACS and Integrated Medical Information Systems: Design and Evaluation (SPIE), San Diego, CA, USA, pp. 280–287 (2002)
7. Haas, B., Coradi, T., Scholz, M., Kunz, P., Huber, M., Oppitz, U., André, L., Lengkeek, V., Huyskens, D., van Esch, A., Reddick, R.: Automatic segmentation of thoracic and pelvic CT images for radiotherapy planning using implicit anatomic knowledge and organ-specific segmentation strategies. Physics in Medicine and Biology 53(6), 1751–1771 (2008)
8. Pavlidis, T., Liow, Y.: Integrating region growing and edge detection. IEEE Transactions on Pattern Analysis and Machine Intelligence, 225–233 (1990)
9. Seifert, S., Barbu, A., Zhou, S., Liu, D., Feulner, J., Huber, M., Suehling, M., Cavallaro, A., Comaniciu, D.: Hierarchical parsing and semantic navigation of full body CT data. In: Proceedings of the SPIE Medical Imaging 2009 Conference (SPIE), Lake Buena Vista, FL, USA, vol. 7259, p. 725902 (2009)

Point-to-Volume Registration of Prostate Implants to Ultrasound*

Ehsan Dehghan[1,4], Junghoon Lee[1], Pascal Fallavollita[2], Nathanael Kuo[1], Anton Deguet[1], E. Clif Burdette[3], Danny Song[1], Jerry L. Prince[1], and Gabor Fichtinger[4]

[1] Johns Hopkins University, USA
[2] Technical University Munich, Germany
[3] Accoustic MedSystems Inc., USA
[4] Queen's University, Canada

Abstract. Ultrasound-Fluoroscopy fusion is a key step toward intra-operative dosimetry for prostate brachytherapy. We propose a method for intensity-based registration of fluoroscopy to ultrasound that obviates the need for seed segmentation required for seed-based registration. We employ image thresholding and morphological and Gaussian filtering to enhance the image intensity distribution of ultrasound volume. Finally, we find the registration parameters by maximizing a point-to-volume similarity metric. We conducted an experiment on a ground truth phantom and achieved registration error of 0.7 ± 0.2 mm. Our clinical results on 5 patient data sets show excellent visual agreement between the registered seeds and the ultrasound volume with a seed-to-seed registration error of 1.8 ± 0.9 mm. With low registration error, high computational speed and no need for manual seed segmentation, our method is promising for clinical application.

1 Introduction

Low dose rate prostate brachytherapy is a treatment for prostate cancer involving permanent implantation of radioactive seeds of ^{125}I or ^{103}Pd inside the prostate and periprostatic tissue. The seeds are implanted using needles that pass through a guiding template, according to a plan, to create an appropriate dose distribution. The procedure is performed under transrectal ultrasound (TRUS) visual guidance. C-arm fluoroscopy images are frequently used for gross visualization of the implant. The treatment quality depends on accurate seed placement which is a challenging task due to problems such as prostate motion and deformation during insertion, needle deflection and edema. Seed misplacement can cause excessive radiation to the healthy tissue that leads to consequent

* E. Dehghan was supported as an Ontario Ministry of Research and Innovation Fellow. G. Fichtinger was supported as Cancer Care Ontario Research Chair. This work was also supported by National Institutes of Health/National Cancer Institute (NIH/NCI) under Grants 2R44CA099374 and 1R01CA151395.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 615–622, 2011.

complications, or can result in insufficient radiation to the cancerous prostate (producing "cold spots") that leads to treatment failure.

Intraoperative dosimetry and planning can improve the treatment quality by intermittently calculating the delivered dose and optimizing the treatment plan in order to compensate for the emerging cold spots [1].

Although TRUS enables visualization of prostate boundary, seed segmentation in TRUS is not robust due to significant number of seed-like artifacts (false positives) created by calcifications and needle tracks, and also missing seeds [2]. However, C-arm images can be used to localize the seeds in 3D space (henceforth, called reconstructed seeds) [3,4]. Spatial registration of reconstructed seeds to the prostate – delineated in the TRUS - can combine the benefits of these two imaging modalities and provide intraoperative dose evaluation.

Lead markers attached to the probe [5] or radio-opaque fiducials attached to the guiding template [6] have been proposed in the past for ultrasound-fluoroscopy registration. However, markers or fiducials need segmentation and their images may overlap with the seeds. Moreover, radio-opaque fiducials cannot compensate for the prostate motion caused by probe retraction. As a solution, Su *et al.* [7] and Tutar *et al.* [8] used point-to-point registration and registered the reconstructed seeds to a set of manually segmented seeds from TRUS images. Manual seed selection in TRUS is tedious as TRUS images are rife with false positives. Fallavollita *et al.* [9] proposed intensity-based registration of CT or fluoroscopy to TRUS. Their method showed successful registration between CT and TRUS in a phantom study and qualitative agreement between the reconstructed seeds and TRUS for a single patient data set.

In this paper, we introduce a point-to-volume intensity-based rigid registration method with application to prostate brachytherapy. We use image thresholding combined with morphological and Gaussian filtering to enhance the quality of TRUS images – without removing the false positives. Except for manual selection of the region of interest, our algorithm is fully automatic and eliminates the need for seed segmentation.

Our registration results on phantom and patient data sets not only show excellent visual agreement between the reconstructed seeds and TRUS images, but also show quantitative registration errors below clinically acceptable levels. In contrast to the work of Fallavollita *et al.* [9], we use different preprocessing steps, similarity metric, and optimizer. In addition, our trials on patient data show smaller registration error and faster computational speed. Considering its low registration error, robustness, and high computational speed, our method is suitable for intraoperative dosimetry.

2 Methods

The following work-flow is envisioned for data acquisition for intraoperative dosimetry. The physician acquires several slices of TRUS images by retracting

the probe from the prostate base to its apex (41-57 slices with spacing of 1 mm in this work). In a preprocessing step, these slices are processed and compounded into a volume. The probe is fully retracted and several C-arm images are taken from different angles. The seeds are reconstructed from 5-6 images in 3D using available seed reconstruction methods [3,4]. The reconstructed seeds comprise a set of 3D points that should be registered to the ultrasound volume.

2.1 Preprocessing

In the preprocessing phase, we follow several steps to enhance the quality of TRUS images (see Fig.1). A region of interest is manually selected from a mid-gland slice of the TRUS volume to limit the search region during optimization and increase the likelihood of convergence. This is the only manual intervention needed in our registration method. Corresponding regions are cropped in all the slices and compounded together to create the volume of interest (VOI).
Although calcifications and air bubbles trapped in the needle tracks have strong reflections, most of the bright areas in TRUS images belong to seeds. Based on this intuition, we apply a threshold (T) to the images to enhance seed visibility (see Fig.1(b)). We define:

$$T = \mu + \alpha\sigma, \tag{1}$$

where, μ and σ are the mean and standard deviation of intensity in the VOI, respectively, and α is a parameter chosen based on the characteristics of the TRUS images ($\alpha = 3$ in this work). Note that we do not try to remove false positives such as calcifications or air bubbles. The thresholded images are then dilated using a disk structural element ($r = 3$ pixels in this work) to increase the size of the bright areas (see Fig.1(c)). Finally, the dilated images are filtered using a Gaussian filter (standard deviation = 10 pixels) in order to spread the bright areas (see Fig.1(d)). The Gaussian filter is applied to provide a smooth change of intensity in the image in order to increase the capture range and enhance the convergence of the optimization algorithm (details in Sec.2.2). The image dilation and Gaussian filtering are applied slice by slice. We sub-sample the filtered TRUS slices with a factor of 2 to gain computational speed.

2.2 Intensity-Based Point-to-Volume Registration

Since transformation of a set of points is computationally faster than the transformation of a volume, we consider the ultrasound VOI as the fixed volume and register the reconstructed seeds to this volume.

The transformation from the C-arm homogeneous coordinate system to the TRUS homogeneous coordinate system is defined as $\mathbf{T}(\underline{\theta}, \underline{\delta})$, where $\underline{\theta} = [\theta_R, \theta_P, \theta_Y]$ represents the roll, pitch and yaw angles, and $\underline{\delta} = [\delta_x, \delta_y, \delta_z]$ represents the translation along x, y and z axes, respectively. We assume that the x axis is parallel to the horizontal axis of the template, the y axis is parallel to the vertical axis of the template from bottom to the top, and the z axis is parallel to the long axis of the probe from the base of the prostate to its apex. Now, consider a mapping $\mathbf{\Psi}$ from every point in the TRUS coordinate system to the indices of

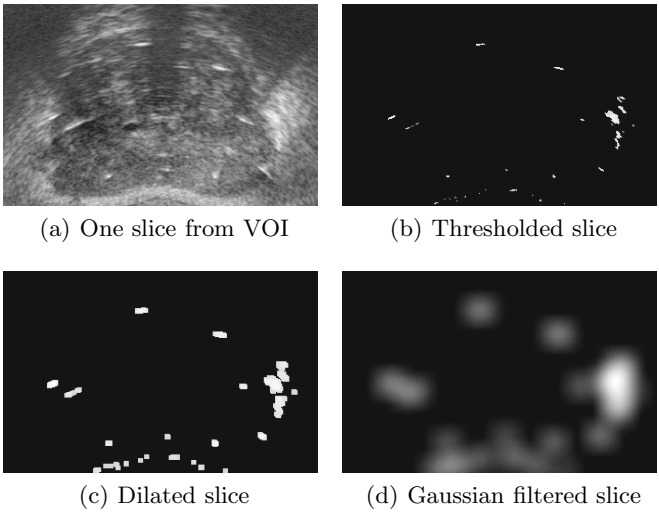

(a) One slice from VOI (b) Thresholded slice

(c) Dilated slice (d) Gaussian filtered slice

Fig. 1. Preprocessing steps

its corresponding voxel in the VOI. The indices of the voxel corresponding to each reconstructed seed (henceforth a seed voxel) can be calculated as:

$$\underline{v}_n = \begin{bmatrix} i_n \\ j_n \\ k_n \end{bmatrix} = \mathbf{\Psi}\left(\mathbf{T}p_n^{\mathrm{C}}\right), \qquad n \in \{1, \cdots, N\}, \tag{2}$$

where p_n^{C} is the coordinates of the n^{th} seed in the C-arm homogeneous coordinate system and N is the number of implanted seeds. We assume rectangular cuboids with dimensions of $(2q_i + 1) \times (2q_j + 1) \times (2q_k + 1)$ voxels centered at each seed voxel. The integers q_i, q_j and q_k are calculated so that each cuboid has dimensions of approximately $2 \times 2 \times 6\mathrm{mm}^3$ (slightly thicker and longer than a seed). The similarity metric is evaluated as:

$$S(\underline{\theta}, \underline{\delta}) = \sum_{n=1}^{N} \sum_{i=-q_i}^{q_i} \sum_{j=-q_j}^{q_j} \sum_{k=-q_k}^{q_k} I\left(\mathbf{\Psi}\left(\mathbf{T}p_n^{\mathrm{F}}\right) + \begin{bmatrix} i \\ j \\ k \end{bmatrix}\right), \tag{3}$$

where, $I(\cdot)$ is the VOI intensity at given indices. In other words, the similarity metric is the summation of the intensities of the voxels inside all cuboids around all the seed voxels. This similarity metric quantifies the overlap between the cuboids and bright regions in the VOI and hence, guides the reconstructed seeds toward the center of the bright regions.

We employ the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [10] in order to maximize the similarity metric. CMA-ES is a stochastic optimization method suitable for nonlinear and non-convex problems and was previously used for registration in [11]. This algorithm samples the search region using a normal distribution, the covariance matrix of which is adapted iteratively.

If the bright regions in the VOI are not expanded and spread using image dilation and Gaussian filtering, the changes in the optimization parameters may result in insignificant or no change in the similarity metric since the bright regions are sparse. Therefore, it is difficult for the optimization algorithm to select the path to improve the similarity metric. However, image dilation and Gaussian filtering provide a smooth and discernible change of similarity metric over the VOI and help the optimization algorithm to hone in on the optimal parameters.

3 Results

3.1 Phantom Study

First, we tested our registration method on a phantom implanted with 48 dummy seeds. Both TRUS and CT images of the phantom were acquired. The registration ground truth between the TRUS and CT volumes was established using a tracked probe and fiducials that were attached to the phantom box [9]. The seeds in the CT volume were segmented using thresholding to yield a set of points similar to the outcome of seed reconstruction using C-arm imaging. These seed positions were assumed to be the ground truth. We applied independent perturbations of -10 to 10 mm, with steps of 1 mm, along each axis and rotations of -10° to 10°, with steps of 1°, around each axis to the ground truth seeds and tried the registration algorithm. The registration algorithm successfully converged close to the ground truth for all of the perturbations. The registration error defined as the distance between the registered seeds and the ground truth was on average 0.7 mm (STD = 0.2 mm). Successful convergence of the algorithm to the global optimum despite the applied perturbations shows its robust performance and wide capture range. Figure 2(a) shows the ground truth and registered seeds overlaid on a slice of ultrasound volume.

3.2 Study on Clinical Data

We also applied our algorithm to clinical data sets. We collected data from 5 patients implanted with 64 to 105 ^{103}Pd seeds. The transverse images – acquired using a BK Pro Focus (BK Medical, Peabody, MA) ultrasound machine - were automatically captured at 1 mm intervals by reading the TRUS stepper position from the encoder while the surgeon continuously retracted the TRUS probe from the prostate base to apex. Several C-arm images were acquired from different angles within a 20° cone around the anterior-posterior axis (AP-axis) using a pre-calibrated GE OEC 9600 mobile C-arm. The C-arm poses were computed using a radio-opaque fiducial [12] that was attached to the guiding template. The C-arm images were preprocessed to correct the image distortion, segment the 2D seed locations and estimate the image poses. The seeds were reconstructed in 3D from 5-6 images by solving a seed matching problem using a dimensionality reduced linear programming algorithm (called REDMAPS) [4].

The registration algorithm was initialized by coinciding the center of mass of the reconstructed seeds with the center of the VOI. The initial rotation angles

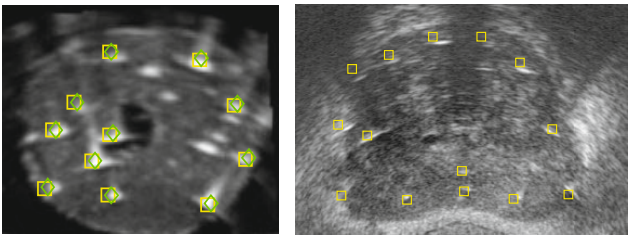

(a) (b)

Fig. 2. Registered seeds overlaid on TRUS images. (a) Phantom result. Green diamonds are ground truth seeds, yellow squares are registered seeds. (b) Patient result. Yellow squares are registered seeds.

were provided by the radio-opaque fiducial which was attached to the guiding template. The registration results showed excellent visual agreement between the reconstructed seeds and the TRUS images as it can be seen in Fig. 2(b). Since no ground truth was available at this stage, we manually identified several seeds from the TRUS images (henceforth called the selected seeds) and measured their distances from the closest registered seed. The average and standard deviation of these seed-to-seed distances are reported in Table 1. The registration error had an overall average of 1.8 mm (STD = 0.9 mm). We also reported the mean and standard deviation of the magnitude of the registration error vectors projected along each axis as shown in Table 1. The registration error along the long axis of the probe (the z axis) is the most significant error.

4 Discussion

Su *et al.* showed that the deviation in D90 (the minimum dose delivered to 90% of the prostate volume) is less than 5% for seed localization uncertainties of less than 2 mm [13]. Our registration errors for 4 of the 5 patients studied are below this limit. Patient 3 had an average registration error slightly greater than 2 mm. This patient had a relatively large prostate that was significantly deformed by the probe pressure. Such patients require a deformable registration between the reconstructed seeds and the TRUS volume. We expect to describe the statistical pattern of deformation from a handful of patients, as the boundary conditions are fairly similar across cases. We also suspect that a simple model, such as a 1D deformation model along the AP-axis suffices to compensate for the primary effect of probe pressure. Alternatively, as recommended in [14], the physician can lower the probe posteriorly in order to decrease the pressure on the prostate. The prostate deformation for the other patients was negligible as the small registration errors along x and y axes confirm.

As it can be seen in Table 1, the error along the z axis is on average 1.2 mm, while the error along x and y axes are on average less than 1 mm. It should be noted that our TRUS volume has a slice spacing of 1 mm along the z axis.

Table 1. Clinical results, showing the mean and standard deviation of registration error, and mean and standard deviation of absolute error along each axis

				Abs. Proj. Error (mm)		
Patient Num.	Total Seeds	Selected Seeds	Reg. Err. (mm) Mean±STD	x Mean±STD	y Mean±STD	z Mean±STD
1	81	12	1.3 ± 0.4	0.5 ± 0.3	0.7 ± 0.4	0.7 ± 0.7
2	76	33	1.7 ± 0.7	0.5 ± 0.5	0.6 ± 0.4	1.3 ± 0.8
3	90	37	2.1 ± 1.2	0.6 ± 0.5	1.3 ± 1.1	1.1 ± 0.9
4	64	20	1.9 ± 1.2	0.5 ± 0.3	0.4 ± 0.3	1.7 ± 1.3
5	105	23	1.5 ± 0.5	0.7 ± 0.4	0.7 ± 0.5	0.8 ± 0.5
Total	416	125	1.8 ± 0.9	0.6 ± 0.4	0.8 ± 0.7	1.2 ± 0.9

In addition, it is difficult to accurately select the center of a 5 mm long seed image which is usually elongated by needle tracks. Therefore our manual seed segmentation can have an error of the same order of magnitude along the z axis, that contributes to the measured registration error.

The algorithm was programmed in MATLAB on a computer with an Intel Core 2 CPU (2 GHz) with 2 GB of RAM. The registration – excluding the manual VOI selection - runs, on average, in approximately 40 s. Our seed-to-seed registration error is less than or equal to the results reported in [7,8] which are based on manual seed segmentation. Our registration error is also smaller than 2.8 mm reported by Fallavollita *et al.* [9]. The registration speed, accuracy and robustness are vital in a clinical setting.

We reported our results based on 5 patient data sets. We expect to get similar results for a larger number of patients. However, we will test our algorithm on a statistically more significant number of patient data sets in the future as data collection is currently underway.

5 Conclusions and Future Work

We presented an intensity-based method for registration of ultrasound to fluoroscopy for intraoperative dosimetry in prostate brachytherapy. Our method obviates the need for tedious manual seed segmentation required for seed-based registration. We applied thresholding, and morphological and Gaussian filtering to the TRUS images to enhance the quality of the images and increase the capture range of the algorithm without removing the false positives or identifying the missing seeds. On a ground truth phantom, the algorithm converged to an average registration error of 0.7 mm despite perturbations of -10 to 10 mm along each axis and -10° to 10° around each axis. This demonstrates the wide capture range and robustness of our algorithm. In a clinical study on 5 patient data sets, we achieved average registration error of 1.8 mm in approximately 40 s. Our registration method with low registration errors, wide capture range and fast computational speed is promising for clinical application.

Extensive tests on more clinical data sets, automatic selection of region of interest and accommodation of prostate deformation as well as sensitivity analysis to the chosen parameters are parts of the future work.

References

1. Polo, A., Salembier, C., Venselaar, J., Hoskin, P.: Review of intraoperative imaging and planning techniques in permanent seed prostate brachytherapy. Radiother. Oncol. 94(1), 12–23 (2010)
2. Han, B.H., Wallner, K., Merrick, G., Butler, W., Sutlief, S., Sylvester, J.: Prostate brachytherapy seed identification on post-implant TRUS images. Med. Phys. 30(5), 898–900 (2003)
3. Dehghan, E., Lee, J., Moradi, M., Wen, X., Fichtinger, G., Salcudean, S.: Prostate brachytherapy seed reconstruction using C-arm rotation measurement and motion compensation. In: Jiang, T., Navab, N., Pluim, J., Viergever, M. (eds.) MICCAI 2010. LNCS, vol. 6361, pp. 283–290. Springer, Heidelberg (2010)
4. Lee, J., Labat, C., Jain, A.K., Song, D.Y., Burdette, E.C., Fichtinger, G., Prince, J.L.: REDMAPS: Reduced-dimensionality matching for prostate brachy-therapy seed reconstruction. IEEE Trans. Med. Imaging 30(1), 38–51 (2011)
5. Todor, D., Zaider, M., Cohen, G., Worman, M., Zelefsky, M.: Intraoperative dynamic dosimetry for prostate implants. Phys. Med. Biol. 48(9), 1153–1171 (2003)
6. Jain, A., Deguet, A., Iordachita, I., Chintalapani, G., Vikal, S., Blevins, J., Le, Y., Armour, E., Burdette, C., Song, D., Fichtinger, G.: Intra-operative 3D guidance and edema detection in prostate brachytherapy using a non-isocentric C-arm. Med. Image Anal. (2010) (in press), doi:10.1016/j.media.2010.07.011
7. Su, Y., Davis, B.J., Furutani, K.M., Herman, M.G., Robb, R.A.: Seed localization and TRUS-fluoroscopy fusion for intraoperative prostate brachytherapy dosimetry. Comput. Aided Surg. 12(1), 25–34 (2007)
8. Tutar, I., Gong, L., Narayanan, S., Pathak, S., Cho, P., Wallner, K., Kim, Y.: Seed-based transrectal ultrasound-fluoroscopy registration method for intraoperative dosimetry analysis of prostate brachytherapy. Med. Phys. 35, 840–848 (2008)
9. Fallavollita, P., Karim-Aghaloo, Z., Burdette, E., Song, D., Abolmaesumi, P., Fichtinger, G.: Registration between ultrasound and fluoroscopy or CT in prostate brachytherapy. Med. Phys. 37, 2749–2760 (2010)
10. Hansen, N.: The CMA evolution strategy: a comparing review. In: Lozano, J.A., Larranaga, P., Inza, I., Bengoetxea, E. (eds.) Towards A New Evolutionary Computation. Advances On Estimation Of Distribution Algorithms. Studies in Fuzziness and Soft Computing, vol. 192, pp. 75–102. Springer, Heidelberg (2006)
11. Gong, R.H., Abolmaesumi, P.: 2D/3D registration with the CMA-ES method, vol. 6918, p. 69181M. SPIE, San Jose (2008)
12. Jain, A.K., Mustufa, T., Zhou, Y., Burdette, C., Chirikjian, G.S., Fichtinger, G.: FTRAC - A robust fluoroscope tracking fiducial. Med. Phys. 32(10), 3185–3198 (2005)
13. Su, Y., Davis, B.J., Furutani, K.M., Herman, M.G., Robb, R.A.: Dosimetry accuracy as a function of seed localization uncertainty in permanent prostate brachytherapy: increased seed number correlates with less variability in prostate dosimetry. Phys. Med. Biol. 52(11), 3105–3119 (2007)
14. Wallner, K., Blasko, J.C., Dattoli, M.: Prostate Brachytherapy Made Complicated, 2nd edn. Smartmedicine Press (2001)

3D Organ Motion Prediction for MR-Guided High Intensity Focused Ultrasound

Patrik Arnold, Frank Preiswerk, Beat Fasel, Rares Salomir, Klaus Scheffler, and Philippe C. Cattin

Medical Image Analysis Center, University of Basel, 4000 Basel, Switzerland
{patrik.arnold,philippe.cattin}@unibas.ch

Abstract. MR-guided High Intensity Focused Ultrasound is an emerging non-invasive technique capable of depositing sharply localised energy deep within the body, without affecting the surrounding tissues. This, however, implies exact knowledge of the target's position when treating mobile organs. In this paper we present an atlas-based prediction technique that trains an atlas from time-resolved 3D volumes using 4DMRI, capturing the full patient specific motion of the organ. Based on a breathing signal, the respiratory state of the organ is then tracked and used to predict the target's future position. To additionally compensate for the non-periodic slower organ drifts, the static motion atlas is combined with a population-based statistical exhalation drift model. The proposed method is validated on organ motion data of 12 healthy volunteers. Experiments estimating the future position of the entire liver result in an average prediction error of 1.1 mm over time intervals of up to 13 minutes.

1 Introduction

Respiratory organ motion is a complicating factor in the treatment of pathological tissue with MR-guided High Intensity Focused Ultrasound (MRgHIFU). Focused ultrasound has the unique capability to deposit sharply localised energy deep into the tissues, producing thermal ablation. Accurate spatial and rapid temporal beam spot focusing in the range of millimetres and within milliseconds, respectively, is reachable and hence increasing the demand of more exact knowledge about the organ's position. Accurate tracking of pathological tissue in mobile organs would not only increase patient safety, but also reduce the treatment time, as the gating window can be increased without sacrificing precision. Although the patient is located within the MR system during sonication, the scan-time is mainly required for the temperature feedback control of the HIFU system to determine the thermal dose given to a tumour. Non MR-based external surrogate markers would thus be ideal for the prediction of the respiratory-induced organ motion. Several techniques have been proposed in the literature to handle respiratory organ motion. Existing approaches that ensure comprehensive target coverage with minimal damage to the surrounding tissue include the optimisation of safety margins, voluntary or forced breath-hold, respiratory

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 623–630, 2011.

gating or full tracking of the target. These methods were discussed in depth and compared in recent publications [1,2]. The holy grail of three-dimensional motion compensation in free-breathing awake patients is still out-of-reach, though. On the one hand, this goal could be reached by using ultrasound for real-time tracking [3]. On the other hand, Ries *et al.* [4] proposed only recently a real-time tracking method that observes the target on a 2D image plane combined with a perpendicular acquired pencil beam navigator, finally obtaining 3D information of the targets trajectories. The future target position is then estimated by a 3D Kalman filter. The method was tested in phantom experiments on human kidneys and in vivo with kidneys of ventilated pigs, both following a regular and stable breathing pattern. The tracking quality is evaluated by comparing temperature distribution obtained after 60 s of HIFU application with and without motion compensation, resulting in higher final temperatures in the target area with enabled motion compensation.

In this paper, we present a novel atlas-based respiratory motion prediction method for free breathing patients. In contrast to the state-of-the-art, the slower modes of non-periodic organ deformation, that occur in addition to the fitful respiratory motion and that are not detectable by external sensors, are compensated by means of a population-based statistical drift model. Although the proposed generic framework is applicable to any abdominal organ, *e.g.* the kidney, the prediction technique is evaluated on real 4DMRI motion data of the liver.

2 Materials and Methods

2.1 Data Acquisition

To learn the patient setup specific breathing characteristics and organ motion, 4DMRI [5] sequences of 12 healthy volunteers (6 female, 6 male, average age 31, range 17-75) were acquired. During roughly one hour acquisition sessions, 14-26 minutes of time-resolved motion data was captured. MR volumes consisting of 25-30 slices covering the right liver lobe with a voxel size of $1.4 \times 1.4 \times 4\ \mathrm{mm}^3$ and with a temporal resolution of 300-400 ms were obtained. The retrospectively reconstructed stacks cover the entire range of observed breathing depths. The vector fields describing the motion between the different respiratory states of the liver were estimated by means of 3D non-rigid registration [5,6] between the reconstructed volumes. In order to estimate the liver's future position we need to keep track of the current respiratory state on the basis of a breathing signal (surrogate marker). Regardless whether we measure the breathing signal by a breathing belt, by an optical chest wall tracker or by a pencil beam navigator placed on the diaphragm, the different respiratory states of the liver can reliable be tracked over a short period of time, as has been shown in [7]. In this work, we extracted a pencil beam navigator by tracking a manually defined region (Fig. 1(a)) on the navigator slice. The inferior-superior motion of the diaphragm was persistently tracked by template matching the dedicated region with all subsequent navigator frames throughout the acquisition sequence, providing one respiratory position per reconstructed volume (Fig. 1(b)).

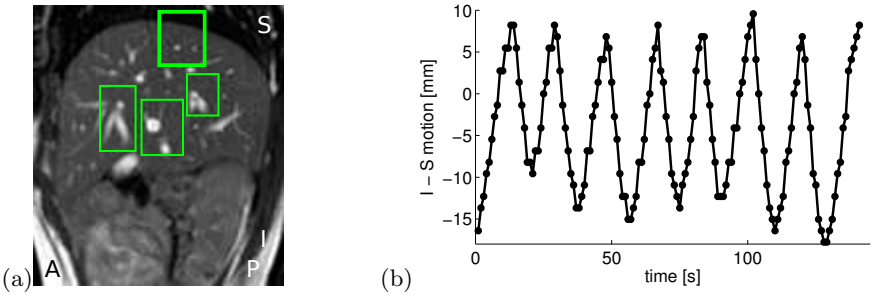

Fig. 1. (a) Sagittal view of a navigator slice with 4 marked regions used for slice stacking and drift compensation. The region indicated by the thicker frame is dedicated to track the diaphragm's position, providing the respiratory signal (b).

2.2 Atlas Creation

Using the 4D organ motion data of all 12 volunteers, we simulated a realistic MRgHIFU scenario. In particular, the first 7-13 minutes of 4DMRI scan time were used to build the motion atlas. This initial training time was long enough to cover all typical respiratory cycles. The remaining 4D motion data was used as ground truth to validate our prediction scheme. In order to keep correspondence between the acquisition of the atlas and the final treatment, the volunteers were asked not to move over the course of the entire sessions.

For each patient, a specific atlas and ground truth dataset is created, wherein both the breathing signal and corresponding 3D vector fields are stored pair-wise for each time step. An example of such an atlas is illustrated on the left side of Fig. 2 as well as the ground truth data for the validation on the right, both containing the respiratory signal and the associated organ displacement depicted by the black arrows between the reconstructed volumes. The breathing signal and the 3D vector fields describing the motion covering 150-400 breathing cycles or 1200-2000 time steps, respectively, serve as the atlas' database.

2.3 Motion Prediction

To readjust the treatment focus, any breathing-controlled tracking method must be able to estimate the target's position at some future time. This estimation must be based on measurements of the past breathing signal. Since our approach deals with a rather low sample rate, we use the atlas as combined breath and 3D motion look-up-table instead of an on-line learning based algorithm. However, only realistic, already seen motion patterns are being generated. Let S be the respiratory signal given as the sequence $S = s_j|_{j=1,\ldots,m}$, with the indices denoting the running time steps t. At a given point in time j, the prediction provides an estimate $s_p = s'_{j+\Delta}$ of S and of the corresponding 3D motion vector field $\mathbf{u}_p = \mathbf{u}'_{j+\Delta}$, describing the future displacement of the organ for a later time point (Fig. 2). In the following experiments, we predicted $\Delta = 1$ time step into the future. One time step corresponds to roughly 300 ms, given by the

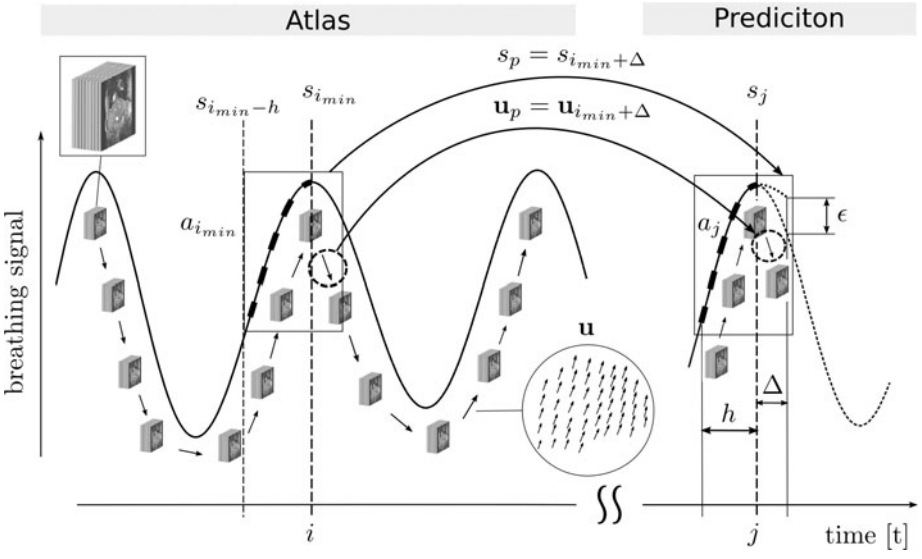

Fig. 2. Schematic illustration of the combined breath and motion atlas. Based on the signals history length h, the prediction yields the 3D displacement field $\mathbf{u}_p$ that estimates the organ's future position, Δ time steps ahead.

4DMRI sequence. Note, that although the experiments are performed and validated on the temporal resolution of 3-4 Hz, Δ can be chosen arbitrary. In that case, the breathing signals and the vector fields are interpolated, allowing any predictive time gap and smooth beam re-focusing. For the prediction, we propose using the last h values of the breathing signals history denoted by the vector $\mathbf{a}_j = (s_{j-h}, \ldots, s_j)$. The reference signal S_{ref} basically serving as the atlas is represented by similar vectors $\mathbf{a}_i$ for prior time points $\mathbf{a}_i = (s_{i-h}, \ldots, s_i)$. The prediction of S at the time point j is chosen by finding the best match of the current breathing signal vector $\mathbf{a}_j$ within the reference signals from the atlas:

$$i_{min} = \arg\min_i |\mathbf{a}_i - \mathbf{a}_j| \,. \tag{1}$$

The future run of $\mathbf{a}_{i_{min}}$ with minimum aberration from the history $\mathbf{a}_j$ is chosen to estimate the organ's prospective respiratory state and corresponding displacement field, Δ time steps ahead:

$$s_p = s_{i_{min}+\Delta} \quad \text{and} \quad \mathbf{u}_p = \mathbf{u}_{i_{min}+\Delta} \,. \tag{2}$$

Finally, the task of predicting the organ's motion is handled by estimating the breathing signal's future evolution, yielding the well adapted displacement fields for the prediction. Since the algorithm is continuously adjusting to new input data, it can quickly adapt to the irregularity of the periods and amplitudes of the respiratory signal of a free breathing person.

2.4 Drift Compensation

Besides the displacements of the liver caused by respiratory motion, additional deformation independent of the fitful breathing motion can occur within a few minutes. Since the proposed prediction method base on collecting the patient-specific liver motion during the initial training phase of a couple of minutes, these organ drifts can not be captured during the short acquisition of the atlas. These drifts can quickly invalidate the applicability of the static atlas with the consequence of increasing systematic prediction errors. However, during sonication within the MR system, the functionality of the scanner is used for temperature feedback of the HIFU device and therefore, a scan time intensive 3D drift tracking is hard to achieve. On the other side, measuring a one dimensional breathing signal only, tracking the inferior-superior motion of the diaphragm respectively, is not sensitive to drifts in the inferior part of the liver (Fig. 3(a)). We propose to acquire one update-navigator slice after every 60 s to capture the exhalation position of the liver based on the breathing signal. Comparing the displacements of the tracked regions with the regions on the actually acquired slice provides the needed information used for the correction of the previously acquired static atlas. In order to compensate these drifts, we introduce a population-based statistical drift model describing the inter-subject variations of exhalation positions in a shared shape-free coordinate system [8]. Shape-free means, that only the relative differences to the first exhalation position of each subject, the drifts, respectively, are used for modelling. Thereby, we assume that the drift is independent of the respiratory motion and is similar for all subjects. From each subject, 200 exhalation positions ($m = 11 \times 200$) with $N = 290$ corresponding points per liver ($n = 3N$), placed on a 3D regular grid with a 15 mm resolution, are mean-free concatenated in a data matrix $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_m) \in \mathbb{R}^{n \times m}$ with $\mathbf{x}_k = \mathbf{v}_k - \bar{\mathbf{v}}$ and sample mean $\bar{\mathbf{v}} = \frac{1}{m} \sum_{k=1}^{m} \mathbf{v}_k$. Applying Principal

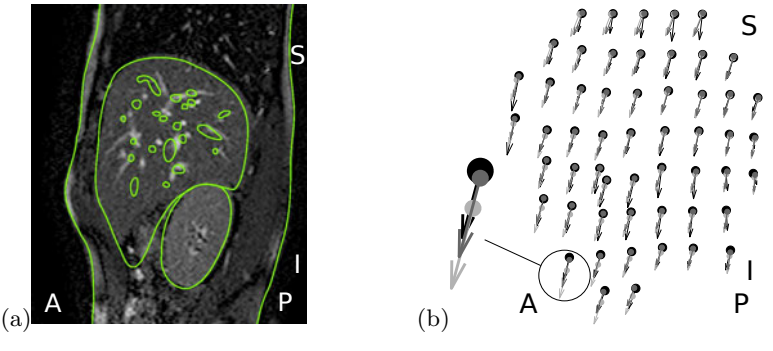

Fig. 3. (a)View of a sagittal placed navigator slice before (*edges*) and 20 minutes later after an exemplary drift displacement. The position of the diaphragm remains almost constant while the inferior part of the liver is drifting. (b) Ground truth motion field (*black*), prediction with static atlas (*light grey*) and with updated atlas (*dark grey*).

Component Analysis on the data, the vectors $\mathbf{x}$ are defined by the coefficients c_k and the eigenvectors $\mathbf{s}_k$ of $\mathbf{S} = (\mathbf{s}_1, \mathbf{s}_2, \dots)$ of the covariance matrix of the data:

$$\mathbf{x} = \sum_{k=1}^{m-1} c_k \sigma_k \mathbf{s}_k = \mathbf{S} \cdot \mathrm{diag}(\sigma_k)\, \mathbf{c}\,. \tag{3}$$

Hereby, σ_k are the standard deviations within the data along each eigenvector $\mathbf{s}_k$. As elaborated in [9], the full vector $\mathbf{x}$ can be found by an incomplete measurement $\mathbf{r} \in \mathbb{R}^l, l < n$ that minimises

$$E(\mathbf{x}) = \| \mathbf{L}\mathbf{x} - \mathbf{r} \|^2\,, \tag{4}$$

where $\mathbf{L}$ represents a subspace mapping $\mathbf{L} : \mathbb{R}^n \mapsto \mathbb{R}^l$ such that $\mathbf{r} = \mathbf{L}\mathbf{x}$. The reduced version of $\mathbf{S}$ can be written as $\mathbf{Q} = \mathbf{L}\mathbf{S} \cdot \mathrm{diag}(\sigma_k) \in \mathbb{R}^{l \times m-1}$, yielding eigenvectors of the form $\mathbf{q}_k = \sigma_k \mathbf{L}\mathbf{s}_k \in \mathbb{R}^l$. The most probable organ deformation $\mathbf{v}$ given the incomplete measurements $\mathbf{r}$ is then

$$\mathbf{v} = \mathbf{S} \cdot \mathrm{diag}(\sigma_k)\, \mathbf{c} + \bar{\mathbf{v}}, \text{ where } \mathbf{c} = \mathbf{Q}^{+}\mathbf{r}\,. \tag{5}$$

Hereby $\mathbf{Q}^{+}$ is the pseudoinverse of $\mathbf{Q}$. The vector $\mathbf{r}$ describes the relative differences from a few grid points at the beginning of the data acquisition to the actual exhalation position. These displacements are captured again by template matching the defined regions (Fig. 1(a)) with the update-navigator slice, measuring the distinct distances between the matching regions. Tracking 4 individual regions enables the detection of non-rigid deformations. As the centers of the templates may rarely coincide with the grid points of the model, the shifts of the templates have been adopted to the 3 closest points of the grid (12 out of 290 points), used as inputs for the drift model. The prediction from the static atlas $\mathbf{u}_p$ (Eq. 2) is updated by the non-rigid correction field (Fig. 3(b)) provided by Eq. (5):

$$\hat{\mathbf{u}}_p = \mathbf{u}_p + \mathbf{v}. \tag{6}$$

3 Results

3.1 Motion Prediction without Drift Compensation

In a first approach we evaluated the capability of organ motion compensation by means of a static atlas without drift compensation. Based on the data mentioned in Sec. 2.2, motion prediction experiments were performed on 4DMRI datasets of all 12 volunteers. The parameter h introduced in Sec. 2.3 was optimised and found to work best for $h = 3$ time steps (≈ 1 s). The prediction experiments were evaluated for all subjects, covering 75-200 full respiratory breathing cycles. The error in prediction for $\Delta = 1$ time step (≈ 300 ms) was calculated point-wise over all grid points and time steps. The results are plotted in Fig. 4(a), characterised by the median, 5th and 95th percentiles. The dashed line is set to 2 mm, marking an acceptable precision limit for HIFU treatments [10]. The impact of the liver's

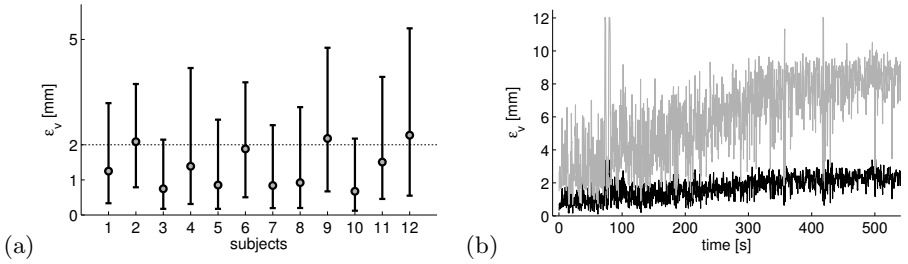

Fig. 4. (a) Resulting deviations between predicted and ground truth liver motions for 12 different subject over time intervals up to 13 minutes. Error bars around the median show the 5th and 95th percentile deviation. (b) Mean (*black*) and maximum error (*grey*) of motion prediction based on the static atlas for the drifting liver of subject 4.

drift is clearly visible in Fig. 4(b), when monitoring the prediction performance over several minutes. The average error over all subjects is 1.6 mm.

3.2 Motion Prediction with Drift Compensation

With equal settings as in Sec. 3.1, the same experiments but with drift compensation as elaborated in Sec. 2.4, were realised (Fig. 5(a)). The statistical drift models were built from 11 of 12 livers in leave-one-out experiments. Although the residual MR time during HIFU treatment is rather sparse, we allowed the acquisition of one 2D navigator slice every 60 s, capturing the actual exhalation position. This time interval is based on the maximal observed drifting speed of 0.5 mm/min. Following Eq. (5) and (6), the most probable drift deformation of the left-out liver is provided by the model and used as a drift-update of the previously acquired static atlas. Taking the drift into account, the prediction performance remains constant over time as shown in Fig. 5(b). The error averaged over all subjects improved by 30 % to 1.1 mm, with a notable impact for

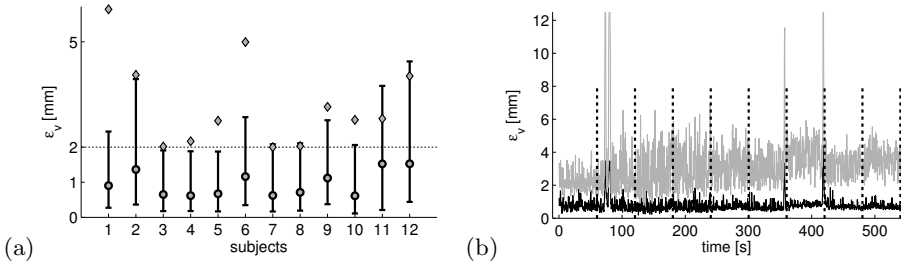

Fig. 5. (a) Residual error of motion prediction with drift compensation and median error without any motion compensation (◊). (b) Mean (*black*) and maximum error (*grey*) of the prediction with drift compensation every 60 s (*black-dotted lines*).

the subjects 4,9 and 12. Without any motion compensation, the mean prediction error would be 4.7 mm.

4 Discussion and Outlook

Despite frequently occurring organ drifts, our proposed method proved to be reliable enough for the application in MRgHIFU systems. Using the R software package (Version 2.11.1), we used the Kolmogorof-Smirnov test to test the mean errors of both experiments for normality. Assuming a significance level of 0.05, the t-tests showed that the statistical drift model significantly improved the prediction accuracy ($p<0.05$). By replacing the pencil beam navigator with a faster low lag signal, such as the breathing belt or a spirometer, the prediction quality should further improve, as the time span the system has to predict into the future decreases.

Acknowledgments. This work has been supported by the research network of the Swiss National Science Foundation-Project Nr. CR32I3_125499.

References

1. Keall, P., Mageras, G., Balter, J., Emery, R., Forster, K., Jliang, S., Kapatoes, J., Low, D., Murphy, M., Ramsey, B.M.C., Herk, M.V., Vedam, S., Wong, J., Yorke, E.: The management of respiratory motion in radiation oncology report of AAPM task group. Med. Phys. 33, 3874–3900 (2006)
2. Webb, S.: Motion effects in (intensity modulated) radiation therapy. Phys. Med. Biol. 51, R403–R425 (2006)
3. Pernot, M., Tanter, M., Fink, M.: 3-D real-time motion correction in high-intensity focused ultrasound therapy. Ultrasound in Medicine and Biology 30(9), 1239–1249 (2004)
4. Ries, M., de Senneville, B.D., Roujol, S., Berber, Y., Quesson, B., Moonen, C.: Real-time 3D target tracking in mri guided focused ultrasound ablations in moving tissues. Magnetic Resonance in Medicine 64, 1704–1712 (2010)
5. von Siebenthal, M., Szekely, G., Gamper, U., Boesiger, P., Lomax, A., Cattin, P.: 4D MR imaging of respiratory organ motion and its variability. Phys. in Med. Biol. 52, 1547–1564 (2007)
6. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L.G., Leach, M.O., Hawkes, D.J.: Nonrigid registration using free-form deformations: Application to breast MR images. Transactions on Medical Imaging 18, 712–721 (1999)
7. von Siebenthal, M., Szekely, G., Lomax, A.J., Cattin, P.C.: Systematic errors in respiratory gating due to intrafraction deformations of the liver. Med. Phys. 34, 3620–3629 (2007)
8. von Siebenthal, M., Székely, G., Lomax, A., Cattin, P.C.: Inter-subject modelling of liver deformation during radiation therapy. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) MICCAI 2007, Part I. LNCS, vol. 4791, pp. 659–666. Springer, Heidelberg (2007)
9. Blanz, V., Vetter, T.: Reconstructing the complete 3D shape of faces from partial information. Informationstechnik und Technische Informatik 44, 295–302 (2002)
10. Vedam, S., Kini, V.R., Keall, P.J., Ramakrishnan, V., Mostafavi, H., Mohan, R.: Quantifying the predictability of diaphragm motion during respiration with a noninvasive external marker. Med. Phys. 30, 505–513 (2003)

Geometry-Aware Multiscale Image Registration via OBBTree-Based Polyaffine Log-Demons

Christof Seiler[1,2], Xavier Pennec[2], and Mauricio Reyes[1]

[1] Institute for Surgical Technology and Biomechanics, University of Bern, Switzerland
[2] Asclepios Research Group, INRIA Sophia Antipolis, France

Abstract. Non-linear image registration is an important tool in many areas of image analysis. For instance, in morphometric studies of a population of brains, free-form deformations between images are analyzed to describe the structural anatomical variability. Such a simple deformation model is justified by the absence of an easy expressible prior about the shape changes. Applying the same algorithms used in brain imaging to orthopedic images might not be optimal due to the difference in the underlying prior on the inter-subject deformations. In particular, using an un-informed deformation prior often leads to local minima far from the expected solution. To improve robustness and promote anatomically meaningful deformations, we propose a locally affine and geometry-aware registration algorithm that automatically adapts to the data. We build upon the log-domain demons algorithm and introduce a new type of OBBTree-based regularization in the registration with a natural multiscale structure. The regularization model is composed of a hierarchy of locally affine transformations via their logarithms. Experiments on mandibles show improved accuracy and robustness when used to initialize the demons, and even similar performance by direct comparison to the demons, with a significantly lower degree of freedom. This closes the gap between polyaffine and non-rigid registration and opens new ways to statistically analyze the registration results.

1 Introduction

In orthopaedic research, and particularly in reconstructive trauma of mandibles, there are a number of surgical interventions such as tumor resection, fracture reconstructions, osteomyelitis or other bone loss repair, which requires implantation of a reconstructive plate. Conventional plate designs, based on shape analysis of a typically small cohort of cadaver specimens or computerized tomography (CT) derived data, fail at capturing the anatomical complexity of the mandible and its shape variation as encountered in a population. In many cases, the results are unsatisfactory and complications such as plate exposure, plate fracture, and screw failure may occur. For instance, it has been reported that suboptimal plate design leads to plate fracturing in 2.9 to 11% of implantations [4]. Additionally intra-operative bending and re-bending of a plate leads to residual stresses,

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 631–638, 2011.

which affect the mean stress in fatigue loading. The solution is to manufacture the plate as close to the human anatomy as possible. For this the morphometry of the mandible should be analyzed in 3D. In the last years, methodologies based on computational anatomy techniques have been explored for shape analysis of the mandible [8], and for population-based orthopaedic femural plate design [7]. A common key step in these works is the need to capture the shape variability as encountered in a population, which is typically performed through non-rigid image registration.

Many existing non-linear registration algorithms for medical images require target deformations to be smooth and invertible. In absence of more precise priors, this could be a reasonable assumption, but it often leads to many local minima far from the optimal deformation. The complexity of the shape (e.g. condyle and coronoid process in the mandible) and the encapsulated population shape variability makes the registration a challenging task, which calls for improvements in non-rigid registration for orthopeadic research. Whenever additional knowledge about the anatomy is known, we should exploit it to guide the search of the deformation towards anatomically more meaningful ones.

We hypothesized that a hierarchical anatomically-based characterization of the structure under study can improve the quality of non-rigid registration results. In this paper we present a novel method, which incorporates such a description of the geometry of the anatomy into the registration process. The method builds on the work of [9], where the authors presented a polyaffine-regularized demons algorithm, however we believe that defining the regions for the locally affine deformations should not be left to the user. This becomes even more evident in the case of a multiscale representation of the geometry, where the definition and division process is not clear, not to mention the time aspect. Thus we propose an automatic and anatomically motivated hierarchical scheme.

Our contributions are three-fold: First, we demonstrate that registration accuracy and robustness can be improved by incorporating a locally affine and anatomy-aware hierarchical scheme. Second, we demonstrate that the complexity of the target deformation can be reduced, which is of great interest in statistics for computational anatomy. The complexity reduction stems from the fact that in contrast to standard multilevel schemes we introduce a higher degree of freedom per region by allowing the region to undergo an affine transformation as opposed to a mere translation. Furthermore, the affine transformation provides important information about the anatomy since it is implicitly defined by the geometry of the anatomical structure. Third, we visualize that our approach opens new ways for multiscale statistics in medical imaging.

2 An OBBTree-Based Polyaffine Log-Demons

Our main contribution in this paper is to show that the demons algorithm can be enhanced in orthopedics applications with an appropriate model. The main idea of the proposed algorithm is to introduce a multiscale regularization implicitly defined via the geometry of the anatomical structure under study. In the following

sections we review three methods that will be important to formulate our registration: Log-demons, polyaffine registration and polyaffine log-demons. Then we present the new multiscale scheme and the integration into the log-demons.

Log-Demons Registration. To setup correspondences between anatomical images, a set of images are registered to a reference. We use the recent diffeomorphic log-domain demons registration approach described in [10]. One interesting point of this registration framework is the efficient optimization in the domain of stationary velocity fields. These velocity fields can be looked at as generators for diffeomorphic deformations through the group exponential map that can be very efficiently computed using the scaling and squaring method [2]. This property explains the denomination of log-domain (or simply log-demons) registration.

Polyaffine Registration. Polyaffine transformations were introduced in [3] to fuse locally rigid and affine transformations into a diffeomorphism. The basic idea is to consider each local affine transformation as the flow of a speed vector field obtained through the log of the affine transformation. Then, instead of averaging the affine matrices, one averages these vector fields with spatial weights describing the influence of each region. The flow of the resulting vector field automatically gives a diffeomorphic transformation. In [5] the authors introduced an efficient approximation of that framework. However, the regions are manually defined on the reference prior to the registration procedure.

Polyaffine Log-Demons Registration. The polyaffine framework was later enhanced in [1] to work with stationary vector fields. This new formulation suggested that polyaffine transformations could be compatible in some sense with the log-demons. In [9] the authors present a marriage of the two approaches and showed a specific application to femur bones with three fixed regions, head, shaft and condyles. However manually choosing the number of regions is a difficult process which is not easy for new applications. In case of hierarchical schemes, this is an intractable task, due to the need to find regions for each level, and the fact that the number of regions usually grows at least quadratically with levels.

Now let us introduce a general formulation of [9] for n regions. Let M_i be the 3×4 non null components of the log of the affine transformation and $\tilde{v}(x)$ be the polyaffine velocity field model:

$$\tilde{v}(x) = \frac{\sum_{i=1}^{n} w_i(x).M_i.x}{\sum_{i=1}^{n} w_i(x)} \quad \text{with} \quad \log\left(\begin{bmatrix} A_i & t_i \\ 0 & 1 \end{bmatrix}\right) = \begin{bmatrix} M_i \\ 0 \end{bmatrix}, \tag{1}$$

where x, w_i, n, are spatial position, weight for region i and number of regions, respectively. Now given v_c, the correspondence velocity field computed by the demons (not regularized), we can solve for M_i using linear least squares, i.e. minimizing $C_{\text{poly}}(M_1, \ldots, M_n) = \int \lambda(x) \left\| v_c(x) - \tilde{v}(x) \right\|^2 dx$, where λ is a binary mask indicating background voxels. This problem has an explicit solution given by $\sum_j M_i.\Sigma_{ij} = B_i$, with $\Sigma_{ij} = \int \lambda(x).w_i(x).w_j(x).x.x^T dx$ and $B_i = \int \lambda(x).w_i(x).v(x).x^T dx$. To estimate M_i we need to solve the system $M.\Sigma = B$, where Σ is symmetrical and thus diagonalizable and the minimal norm solution is given using the pseudo inverse $M = B.\Sigma^{+}$.

OBBTree-Based Hierarchical Scheme for the Polyaffine Log-Demons. The concept of oriented bounding boxes (OBB) has been used extensively in computer graphics to speed up ray tracing and interference detection computations. In [6] the authors presented a hierarchical version and an algorithm to compute it efficiently. An OBBTree is a hierarchy of OBB's in 3D space. Let us consider first a surface (in our case an iso-surface of the CT image). The algorithm computes OBB's via principle component analysis of the vertex coordinates, which give the orientation (principle component directions) and the extent (outmost point on the principle component). A refinement to avoid bias towards densely populated patches is to sample the convex hull of the vertex coordinates and approximate the analytic surface by a linear sum of all triangle areas. There are two ways of calculating the hierarchy, bottom-up and top-down. Top-down approaches start with all vertices and subdivide the points into two groups at every subsequent hierarchical level, whereas bottom-up approaches start by assigning one box per vertex and combine vertices until one box contains all vertices. We used the top-down approach, which divides the vertices into two groups by projecting the vertex coordinates onto the principle components, and using the mean point as the group boundary. The algorithm stops once there are no more divisions possible along any component. For images, the point set could result from a random sampling weighted by the importance of points (e.g. norm of the gradient).

The weights $w_i(x)$ are defined on the OBBTree using multidimensional Gaussian functions as follows,

$$w_i(x) = \exp\left(-0.5.(x-\bar{x}_i)^T.[R_i.\,\mathrm{diag}(\sigma_{i1}^2, \sigma_{i2}^2, \sigma_{i3}^2).R_i^T]^{-1}.(x-\bar{x}_i)\right), \quad (2)$$

where $\bar{x}_i$ and R_i are center point and orientation of the ith OBB, and σ_{i1}, σ_{i2}, σ_{i3} are the extent of the weights along each axis of the OBB. The parameter σ_{ij}^2 can be set by the user to enforce different smoothing behavior between regions.

To obtain an intuitive understanding of the smoothing parameter σ_{ij}^2 we introduce the concept of relative regional mass (RRM). The RRM defines the relative mass for the most prominent weight in a given region: $\mathrm{RRM}_i = \frac{\int_{\Omega_i} w_i(x)dx}{\int_{\Omega} w_i(x)dx}$, where Ω_i and Ω are the volume given by the ellipsoid that fits inside the ith OBB and the volume of the entire image, respectively. In the special case of RRM equal to 1, there is no overlap between regions. In Fig. 2 (right), a range of possible values are shown with corresponding mean squared error and harmonic energy. In practice, we adjust σ_{ij}^2 to obtain the same RRM for all regions.

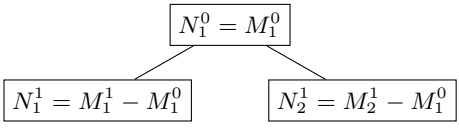

Fig. 1. Tree of log affine transformations for the first two levels

Algorithm 1. OBBTree Log-Demons (OBB-LD)

Sequentially register OBBTree levels $l = 0, \ldots, k$

- Initialize demons with previous level $v^l = v^{l-1}$ (for starting level $v^0 = 0$)
- Iterate until convergence
  - Compute correspondence stationary velocity field v_c^l (not regularized)
  - Solve linear least square problem $M^l = B^l.\Sigma^{l+}$
  - Combine $M_1^l, \ldots, M_n^l$ using the polyaffine model (1) to $\tilde{v}^l$
  - Let $v^l = \tilde{v}^l$

The integration of our hierarchical model into the demons (after rigid alignment) is shown in Algorithm 1. The OBBTree-based polyaffine log-demons regularizes the velocity fields during each update step. The final result is a velocity field that includes all registration steps allowing for statistics on diffeomorphisms as described in [2]. In addition we obtain M_i^l, which is the ith log affine transformation at level l, providing us with a low-parametric representation of the anatomical structure and allowing further hierarchical statistical analysis and modeling. The hierarchical structure is shown in Fig. 1. Since we are working in the log-Euclidean framework we can subtract the previous log transform from the current level to obtain N_i^l, which describes the remaining transformation at that level. The importance of such a hierarchy lies in its power of decomposing features into different scales, this will be elaborated in the next section by performing a hierarchical PCA for mandibles.

3 Experiments on Mandible CT's

To evaluate the performance of our new method we conducted registrations on 47 CT images of mandibles. After rigid alignment of the OBB at level 0, we

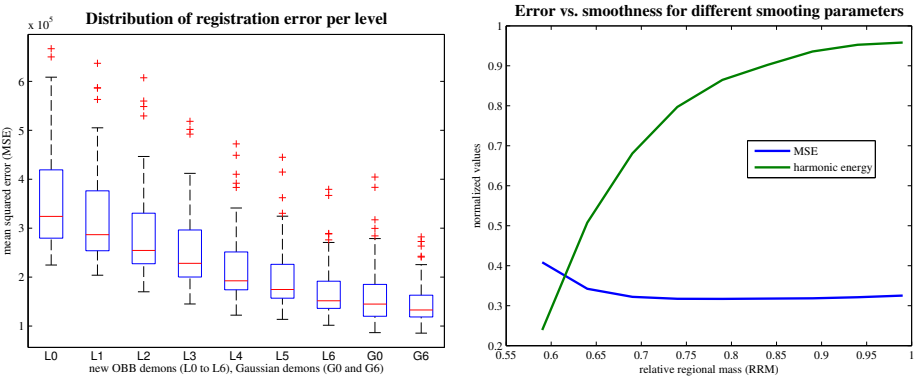

Fig. 2. (left) Boxplot of mean squared error (MSE) of intensities for each level. G0 and G6 are standard log-demons initialized with L0 and L6, respectively. L0 is one region only, i.e. standard affine initialization. The red crosses indicate outliers. (right) The relative regional mass (RRM) represents the amount of smoothing between regions, where 1 means no overlap between regions.

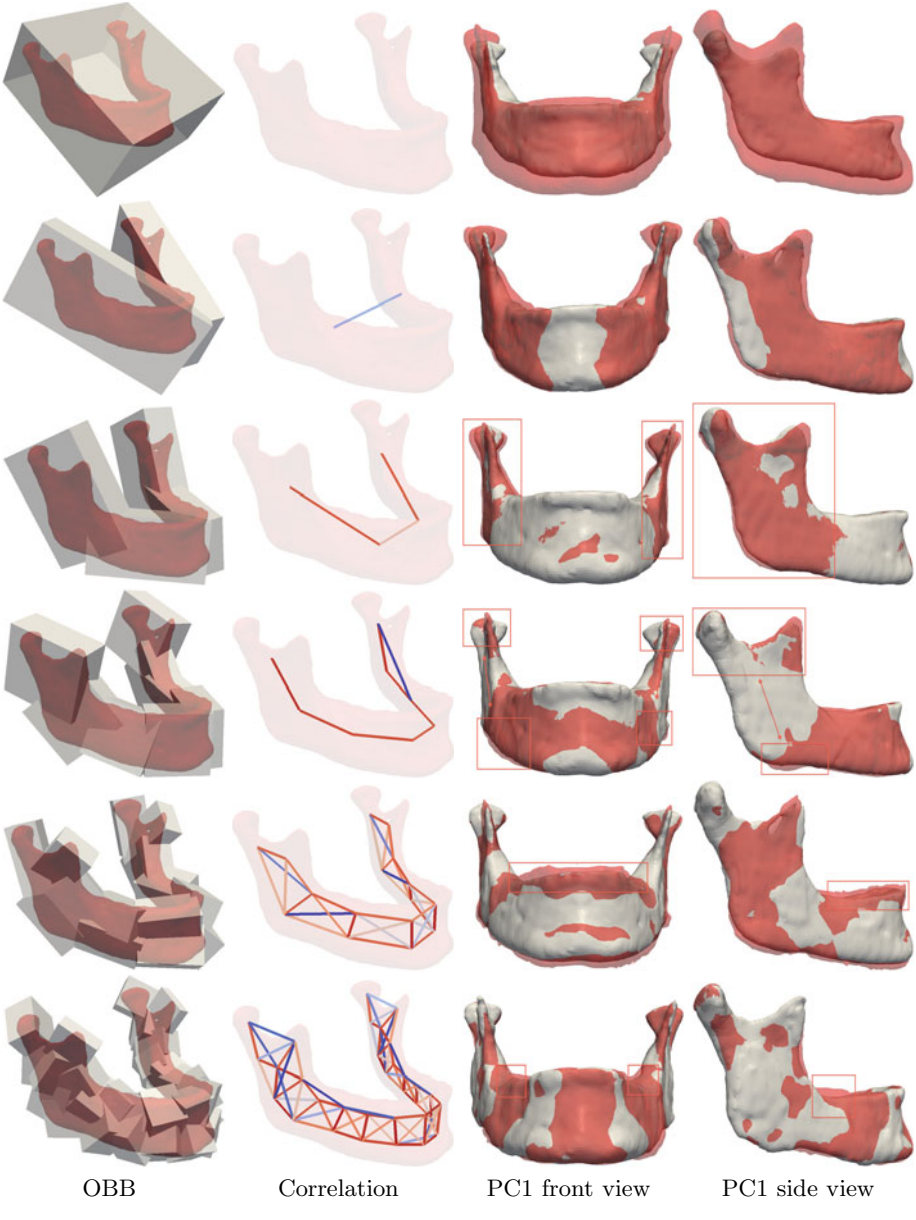

Fig. 3. (OBB) 6 levels of the OBBTree visualized on the reference mandible, starting with level 0 up to level 5. (Correlation) Major axis of correlations between regions, color coded from cold (blue=0.4) to warm (red=1.0). (PC1 front and side view) Red and white surface are -2 and $+2$ standard deviations from the mean.

performed an OBBTree-based registration with 0 to 6 levels (L0 to L6), each level being initialized by the previous one. Additionally, we initialized two standard log-demons with Gaussian regularization (G0 and G6) at the first and last level (L0 and L6), respectively. Since L0 is one region, it is equivalent to a standard affine initialization.

Fig. 2 (left) shows three important results: (1) Decrease of mean squared error (MSE) with increasing amount of regions (levels). (2) G0 shows higher median, variability and more outliers than G6, which indicates the improvements in terms of robustness, when used as an initialization for standard demons. (3) L6 shows similar performance as G0 and G6, this indicates that by using $2^6 = 64$ regions, we can model a standard demons (relative degrees of freedom: three times the number of voxels). For example, given the dimensions of our mandible CT's our method uses only $\approx 0.01\%$ degrees of freedom to reach the same MSE compared to a standard demons.

Fig. 2 (right) depicts MSE and harmonic energy as a function of different RRM's. The MSE converges at RRM=0.74, whereas the harmonic energy increases with increasing RRM values. For the experiments we chose a RRM of 0.74 to obtain the best MSE with highest smoothness.

Fig. 3 shows 5-level OBBTree, weight correlation structure and the first principle component for each level. The weight correlation structure is computed as $\Gamma_{ij} = \Sigma_{ii}^{-1/2}.\Sigma_{ij}.\Sigma_{jj}^{-1/2}$, and Γ_{ij} is decomposed using singular value decomposition to extract the major axis of correlation. The correlation of the major axis is colored coded from cool (blue=0.4) to warm (red=1). The graph structure clearly reveals the intrinsic underlying dimensionality of the object at each scale, going from a curve to a ribbon and finally in some areas locally to a 3D volume. It would be interesting to study if this could be a robust alternative to the medial axis or surface representation.

The per level PCA can be interpreted as follows: (L0) global scaling; (L1) thickness; (L2) reorientation in the region of the masseter; (L3) relative displacement of condyles and coronoid processes; (L4) change in teeth region; (L5) change in back teeth region. This gives a visual validation of the usefulness of per level hierarchical statistical analysis, clearly distributing feature to different scales.

4 Conclusions

In this work we presented a geometry-aware multiscale approach for registration. We showed that our method significantly robustifies the standard demons and increases its repeatability when the full range of scales is used. When stopping at a reasonable scale (here only 64 components), it even performs similarly while presenting significant reduction of degrees of freedom of the registration. We further visualized a per hierarchical level PCA, which suggested a clear division of shape features into the different levels, allowing for a better interpretation, as opposed to a PCA on the entire field with potential mixtures of features.

This new method opens a large number of potential research opportunities: How to analysis tree-like data objects? How to enforce sparsity on these objects?

What is the correlation between levels? How to measure anatomical meaningfulness of deformations quantitatively (in this work a qualitative approach visualizing the modes of variation is used)? How to measure the explained variance in trees? All these questions will be addressed in future work.

Acknowledgements. This work has been funded by the Swiss National Science Foundation.

References

1. Arsigny, V., Commowick, O., Ayache, N., Pennec, X.: A Fast and Log-Euclidean Polyaffine Framework for Locally Linear Registration. J. Math. Imaging Vis. 33(2), 222–238 (2009)
2. Arsigny, V., Commowick, O., Pennec, X., Ayache, N.: A Log-Euclidean Framework for Statistics on Diffeomorphisms. In: Larsen, R., Nielsen, M., Sporring, J. (eds.) MICCAI 2006. LNCS, vol. 4190, pp. 924–931. Springer, Heidelberg (2006)
3. Arsigny, V., Pennec, X., Ayache, N.: Polyrigid and Polyaffine Transformations: A Novel Geometrical Tool to Deal with Non-rigid Deformations Application to the Registration of Histological Slices. Med. Image Anal. 9(6), 507–523 (2005)
4. Coletti, Ord, R., Liu, X.: Mandibular Reconstruction and Second Generation Locking Reconstruction Plates: Outcome of 110 Patients. Int. J. Oral Max. Surg. 38(9), 960–963 (2009)
5. Commowick, O., Arsigny, V., Isambert, A., Costa, J., Dhermain, F., Bidault, F., Bondiau, P.-Y., Ayache, N., Malandain, G.: An Efficient Locally Affine Framework for the Smooth Registration of Anatomical Structures. Med. Image Anal. 12(4), 427–441 (2008)
6. Gottschalk, S., Lin, Manocha, D.: OBBTree: A Hierarchical Structure for Rapid Interference Detection. In: SIGGRAPH, pp. 171–180 (1996)
7. Kozic, N., Weber, S., Büchler, P., Lutz, C., Reimers, N., Ballester, M.A., Reyes, M.: Optimisation of Orthopaedic Implant Design Using Statistical Shape Space Analysis Based on Level Sets. Med. Image Anal. 14(3), 265–275 (2010)
8. Metzger, M.C., Vogel, M., Hohlweg-Majert, B., Mast, H., Fan, X., Rüdell, A., Schlager, S.: Anatomical Shape Analysis of the Mandible in Caucasian and Chinese for the Production of Preformed Mandible Reconstruction Plates. J. Cranio. Maxill. Surg. (2010)
9. Seiler, C., Pennec, X., Ritacco, L., Reyes, M.: Femur Specific Polyaffine Model to Regularize the Log-Domain Demons Registration. In: SPIE Medical Imaging, vol. 7962 (2011)
10. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Symmetric Log-Domain Diffeomorphic Registration: A Demons-Based Approach. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part I. LNCS, vol. 5241, pp. 754–761. Springer, Heidelberg (2008)

Geometric Metamorphosis

Marc Niethammer[1,2], Gabriel L. Hart[3], Danielle F. Pace[3], Paul M. Vespa[5], Andrei Irimia[4], John D. Van Horn[4], and Stephen R. Aylward[3]

[1] University of North Carolina (UNC), Chapel Hill NC 27599-3175, USA
[2] Biomedical Research Imaging Center, UNC Chapel Hill NC 27599-7515, USA
[3] Kitware, Inc., Carrboro NC 27510, USA
[4] Laboratory of Neuro Imaging, University of California, Los Angeles CA 90095, USA
[5] Brain Injury Research Center, University of California, Los Angeles CA 90095, USA

Abstract. Standard image registration methods do not account for changes in image appearance. Hence, metamorphosis approaches have been developed which jointly estimate a space deformation and a change in image appearance to construct a spatio-temporal trajectory smoothly transforming a source to a target image. For standard metamorphosis, geometric changes are not explicitly modeled. We propose a *geometric metamorphosis* formulation, which explains changes in image appearance by a global deformation, a deformation of a geometric model, and an image composition model. This work is motivated by the clinical challenge of predicting the long-term effects of traumatic brain injuries based on time-series images. This work is also applicable to the quantification of tumor progression (e.g., estimating its infiltrating and displacing components) and predicting chronic blood perfusion changes after stroke. We demonstrate the utility of the method using simulated data as well as scans from a clinical traumatic brain injury patient.

1 Introduction and Background

Image registration is based on structural similarity between a source and a target image. Similarity is often measured either by comparing image intensities directly or using indirect intensity measures like mutual information or cross correlation. However, for images with pathologies, assumptions of structural and intensity similarity may not hold.

In traumatic brain injury (TBI) cases, one clinical challenge is distinguishing permanent from transient changes in the brain in order to prescribe effective treatment and rehabilitation plans. Scans are acquired upon initial presentation in the clinic as well as after four to eight months (see Fig. 1). The geometry of the pathology, the deformation of the brain, and infiltration of the pathology into the brain change drastically between these scans. Determining the regions in which the infiltration has receded can be particularly useful in predicting long-term outcome. Similarly, in tumor cases, post-treatment assessment requires determination of changes in tumor geometry, tumor infiltration, scarring, and overall brain morphology. In stroke cases, there is a clinical need to predict chronic changes in blood perfusion from acute scans. In general, these cases are

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 639–646, 2011.

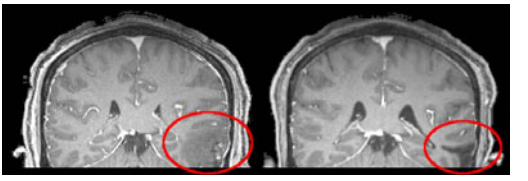

Fig. 1. MP-RAGE, post-contrast MRI scans from TBI case. *Left*: Initial scan. *Right*: Eight months after initial scan. Rigidly registered. 1x1x1mm voxels.

characterized by global tissue deformations, local changes in the geometry of a pathology, and local changes in the composition of the tissue and the pathology. We refer to these changes as "geometric metamorphosis."

While geometric metamorphosis changes may be tolerated by registration methods with low-dimensional image transformation models, direct application of a classical deformable registration method will likely produce unrealistic estimates of deformation. To address geometric metamorphosis changes, deformable registration approaches with weak and strong models of appearance change have been proposed. For example, methods having strong models of brain tumor mass effects and infiltration have been developed [4,7] and have been used to simulate tumors in atlas images to allow for spatial normalization of subjects with brain tumors [10]. While highly sophisticated, these methods are application-specific and rely on a good match of the tumor model to the observed tumor. On the other hand, image metamorphosis methods [9] use weak models to smoothly transform a source to a target image exactly. However, the transformations estimated by image metamorphosis do not explicitly model the deformation or composition of the pathologies and instead compromise between a globally smooth spatial transformation and the interpolation of image intensities along individual point trajectories. Hence, image metamorphosis models will have difficulty quantifying effects such as tumor infiltration or tissue recovery in stroke. Consider also that the approach proposed in this paper is, in spirit, related to the methods proposed in [5,8,2] in which areas that cannot be matched (because no correspondence exists) are masked out. However, in those methods registration results inside these masked-out areas are only driven by the spatial regularity terms of the deformable registration algorithms. Our method explicitly includes a deformable geometric model of the extent of the appearance change in order to capture pathology deformations in conjunction with underlying image deformations.

Sec. 2 discusses the geometric metamorphosis model. Its numerical solution method is discussed in Sec. 3. Results are presented in Sec. 5. The paper concludes with a summary and outlook on future work in Sec. 6.

2 Geometric Metamorphosis Model

Taking tumor growth as an example, changes in image appearance can be caused by a mixture of tissue deformation, tumor growth displacing healthy tissue, and tumor infiltration into healthy tissue. Tissue deformation and displacement due to tumor growth could be captured (between time-points) using a standard

registration method[1]. However, infiltration does not imply spatial changes. A registration method should be able to distinguish image appearance changes arising from the *composition* of *background deformations* of the image and *foreground deformations* of an embedded geometric object, e.g., a tumor.

We model these transformations through a fluid-registration formulation. In large displacement diffeomorphic metric mapping (LDDMM) [6] one minimizes

$$E = \int_0^1 \|v\|_L^2 \, dt + \frac{1}{\sigma^2}\|I(1) - I_1\|^2, \quad \text{s.t.} \quad I_t + \nabla I^T v = 0; \; I(0) = I_0,$$

where v is a sought-for time-dependent velocity field which induces a spatial transformation warping the source image I_0 to the target image I_1. Typical LDDMM formulations register from source I_0 to target I_1 on the time interval $[0, 1]$, thus $I(1)$ represents the warped source image. L is a differential operator (here, $L = \gamma - \alpha\nabla^2$, $\alpha, \gamma > 0$) controlling spatial regularity of v; $\sigma > 0$ controls the influence of the image match term. This is an inexact matching that only allows for the deformation of the source image I_0, but not for a change of its appearance. In order to explicitly model appearance changes, we augment this standard image registration model with an extra control to model foreground deformation of a geometric model (Sec. 2.1) which induces image changes through an image composition model (Sec. 2.2).

2.1 Deformation Model

Fig. 2 illustrates the principle of the geometric metamorphosis model. Since we model geometric metamorphosis as the composition of background and foreground deformations, we introduce the (smooth) indicator functions T_1 and T_2 as models of the geometric object, $T_1(x)$ and $T_2(x) \in [0, 1]$. We then register the background global deformation on time $[0, 1]$ and the foreground geometry change on time $(1, 2]$, solving for the time-dependent velocity fields v and v^τ, respectively. We define the geometric metamorphosis problem as the minimization of

$$\begin{aligned} E = (1-w)\int_0^1 \|v\|_L^2 \, dt + w\int_1^2 \|v^\tau\|_L^2 \, dt \\ + \frac{1}{\sigma_1^2} Sim(I^c(I_1, I^\tau(1), T_2), I^c(I(1), I^\tau(1), T_2)) + \frac{1}{\sigma_2^2} Sim(I^\tau(2), T_2), \end{aligned}$$

$$\text{s.t.} \quad I_t + \nabla I^T v = 0; \; I(0) = I_0, \quad \begin{cases} I_t^\tau + \nabla (I^\tau)^T v = 0, & I^\tau(0) = T_1, \; t \in [0, 1], \\ I_t^\tau + \nabla (I^\tau)^T v^\tau = 0, & t \in (1, 2], \end{cases} \tag{1}$$

where I^τ is the image of the geometric object and $w \in (0, 1)$ controls the trade-off between background and foreground deformations. Note that the geometric

[1] We assume for simplicity that the geometric object causing image change is present in both the source and target image, although it may undergo significant distortions.

model T_1 and its image I^τ are subject to both deformations, whereas the source image is only subjected to the background deformation. $I^c(\cdot,\cdot,\cdot)$ denotes the image composition model (see Sec. 2.2). Sim denotes a similarity measure of choice. For simplicity, we use the L^2 distance measure, $Sim(I,J) = \|I-J\|^2$. Two similarity terms are used to assure matching of (i) the regions which correspond in both images and (ii) the geometric models.

2.2 Image Composition Model

To accommodate local expansions and contractions of the geometric model affecting image appearance, the image composition model needs to preserve regions where the source and target image can be reliably matched. It needs to disregard areas where no matching image information can be found due to the shape change of the geometric model. The composition model

$$I^c(I, I^\tau(1), T_2)(x) := I(x)(1 - I^\tau(1,x))(1 - T_2(x)), \tag{2}$$

achieves this by zeroing out regions defined by the smoothed indicator functions $I^\tau(1)$ and T_2. Since this happens for both arguments of the similarity function in Eq. 1 the image match is effectively disregarded in these regions. This definition is reminiscent of cost function masking as for example used when registering images with and without lesions [2]. Here, we use regions in the source *and* target image to alter the energy function and estimate the regions which should be excluded in a joint optimization process. This allows for a combined estimation of foreground and background deformation.

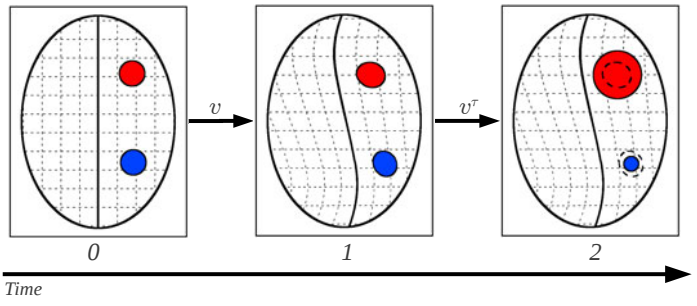

Fig. 2. Geometric Metamorphosis. An image is explained by a global deformation (via v) and a geometric model deformation (via v^τ). Corresponding structures in the source and target guide the estimation of v and v^τ addresses additional appearance differences at the pathology. To avoid faulty evaluation of image similarities, a suitable image composition method is required (Sec. 2.2). Regions which carry no matchable information are set to 0 in the image composition model. For a shrinking geometric model (blue) this region is specified by $I^\tau(1)$ (which already includes the background deformation) and for a growing geometric model (red) by T_2. Defining the composition model as Eq. 2 allows localized growing and shrinking simultaneously.

3 Numerical Solution

We follow the solution method of [3] to solve the registration problem. To compute the optimality conditions, we add the dynamic constraints through the Lagrange multipliers λ and λ^τ. Note that λ^τ is allowed to be discontinuous at $t = 1$ due to the energy term depending on $I^\tau(1)$. After some computations we obtain the optimality conditions for $t \in [0, 1)$

$$
\begin{aligned}
0 &= 2(1-w)L^\dagger Lv + \lambda \nabla I + \lambda^\tau \nabla I^\tau \\
I_t + \nabla I^T v &= 0,\ I(0) = I_0, \\
-\lambda_t - div(\lambda v) &= 0,\ \lambda(1) = \frac{2}{\sigma_1^2}(I_1 - I(1))(1 - T_2)^2(1 - I^\tau(1))^2, \\
I_t^\tau + \nabla (I^\tau)^T v &= 0,\ I^\tau(0) = T_1, \\
-\lambda_t^\tau - div(\lambda^\tau v) &= 0,\ \lambda^\tau(1-) = \lambda^\tau(1+) + \frac{2}{\sigma_1^2}(I(1) - I_1)^2(1 - T_2)^2(1 - I^\tau(1)),
\end{aligned}
$$

and for $t \in (1, 2]$

$$
\begin{aligned}
2wL^\dagger Lv^\tau + \lambda^\tau \nabla I^\tau &= 0, \\
I_t^\tau + \nabla (I^\tau)^T v^\tau &= 0, \\
-\lambda_t^\tau - div(\lambda^\tau v^\tau) &= 0,\ \lambda^\tau(2) = -\frac{2}{\sigma_2^2}(I^\tau(2) - T_2).
\end{aligned}
$$

The final conditions for λ and λ^τ in $[0, 1)$ reflect the "don't care" areas of the registration: areas where $T_2 = 1$ or $I^\tau(1) = 1$ are zeroed out. This is sensible, because the Lagrangian multipliers represent the image-matching error. We obtain a solution fulfilling the optimality conditions through the following adjoint solution method:

---

0) Initialize v, v^τ to zero.
1) Solve $I_t + \nabla I^T v = 0,\ I(0) = I_0$ and $I_t^\tau + \nabla (I^\tau)^T v = 0,\ I^\tau(0) = T_1$ forward in time in $[0, 1]$.
2) Continue solving for I^τ for $t \in (1, 2]$ but with velocity field v^τ.
3) Compute the adjoint solution λ backward for $t \in [0, 1]$.
4) Compute the adjoint solution λ^τ backward for $t \in (1, 2]$.
5) Apply the jump condition to λ^τ at $t = 1$.
6) Compute the adjoint solution λ^τ backward for $t \in [0, 1)$.
5) Compute for every point and time-point the gradients

$$
\begin{aligned}
\nabla_v(t)E &= 2(1-w)L^\dagger Lv + \lambda \nabla I + \lambda^\tau \nabla I^\tau,\ t \in [0, 1], \\
\nabla_{v^\tau}(t)E &= 2wL^\dagger Lv^\tau + \lambda^\tau \nabla I^\tau,\ t \in (1, 2].
\end{aligned}
$$

6) Do a gradient descent step to update the velocities (using line search).
7) Repeat steps 1 - 6 until convergence.

---

4 Estimating Geometric Deformation

Once the foreground and background velocity fields v and v^τ have been estimated, they can be used to represent the geometric deformation modulo the background deformation. This allows for visualization and quantification, for example of tumor growth. Computing (backward in time) the mapping

$$-\Phi^r_t - D\Phi^r v = 0,\ \Phi^r(1) = id,\ t \in [0,1]$$

where id is the identity map and D the Jacobian, shape change is computed as

$$S(0) = T_2 \circ \Phi^r(0) - T_1, \quad S(1) = T_2 - I^\tau(1)$$

in the coordinate system of the source and the target image respectively. Here positive values indicate expansion and negative values contraction with respect to the source image.

5 Experimental Results

We test the geometric metamorphosis model on two sets of synthetic images, and a TBI image pair. The first synthetic set (Fig. 3) illustrates four different scenarios: (i) all change caused by infiltration, (ii) all change caused by global deformation, (iii) global deformation and local infiltration, and (iv) global deformation and local recession.

The second synthetic set consists of ten different global and object warps applied to the same source image and geometric object. After registering, we compute the mean and standard deviation of the percent overlap for the geometric object as,

$$Overlap(T_2, I^\tau(2)) = sum((I^\tau_{\geq 0.5}(2)) \cap (T_{2,\geq 0.5}))/sum((T_{2,\geq 0.5})),$$

where $I_{\geq x}$ is a binary mask of all pixels in I greater than or equal to x.

We also compute the background registration accuracy using six manually selected landmarks in the tissue region as ground truth. Landmark locations are calculated on the pixel grid and are accurate to +/- 0.5 pixels. We compare the results of our method against the B-Spline and LDDMM registration methods. Given the expected similar results in the regions of the image away from the geometric object we look at the extreme percentiles of the landmark distance mismatch values. Both LDDMM and geometric metamorphosis compute 95% of the landmarks to within 0.5 pixels of their correct location, but our method is able to achieve a significantly higher overlap accuracy (Fig. 3).

The TBI test case contains considerable deformation as well as object recession around the pathology site. To illustrate those changes, we provide manual segmentations of the pathology sites in both images (Fig. 4)[2]. Of clinical importance, Fig. 5 shows the progression of the pathology label map over the entire time solution interval [0,2].

[2] Note that it is expected that a high segmentations accuracy is not required [1].

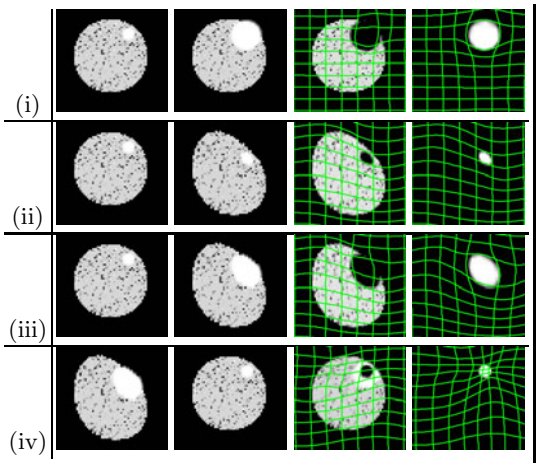

Overlap Accuracy

	Mean	Std. Dev.
B-Spline	0.394	0.033
LDDMM	0.540	0.013
Geo. Met.	0.975	0.039

Landmark Pixel Distance Mismatch Percentiles

	90^{th}	95^{th}	100^{th}
B-Spline	2.74	3.33	3.66
LDDMM	0.5	0.5	0.5
Geo. Met.	0.5	0.5	1.5

Algorithm Comparison: Ten synthetic cases registered using B-Spline, LDDMM, and geometric metamorphosis.

Fig. 3. Synthetic Results. Each row, left to right: I_0, I_1, $I(1)$ and global deformation, $I^\tau(2)$ and composite deformation. (i) Local infiltration. (ii) Image deformation. (iii) Image deformation and local infiltration. (iv) Image deformation and local recession.

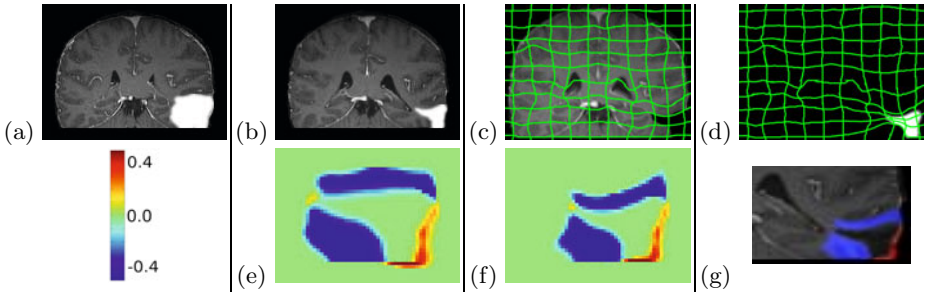

Fig. 4. TBI Results. *Top*: (a) Initial scan, T_1 overlaid. (b) Second scan, T_2 overlaid. (c) Image deformation and $I(1)$. (d) Retraction area deformation and $I^\tau(2)$. *Bottom*: (e) $S(0)$: Shape change in the source image coordinate frame (f) $S(1)$: Shape change in the target image coordinate frame, (g) Incursion (red) and retraction (blue) in I_2.

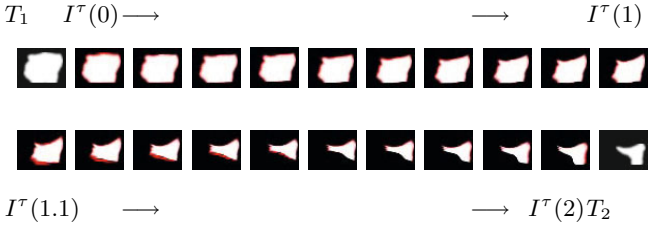

Fig. 5. TBI Pathology Label Map. Progression of TBI pathology label map on [0,2]. Red portion of each frame shows the change from the previous time point. *Top*: [0,1] Changes in global deformation. *Bottom*: (1,2] Changes in infiltration and recession.

6 Conclusions and Future Work

We proposed an image registration method allowing for background deformation of a source image, foreground deformation of a geometric object, and their composition to match a target image. The method can thereby account for processes causing images to change due to pathology infiltration/recession and image deformation. Since it makes minimal assumptions about the underlying change, it is generally applicable. We demonstrated its behavior for the registration of simulated data and traumatic brain injury cases. If desired, the registration framework can be augmented with an application-specific model, for example, of tumor growth, as in [4]. In future work we will investigate adaptations wherein a geometric model is only available in one of the two images and a transformed model either needs to be estimated from the image or does not exist, e.g., when registering a tumor patient to a healthy atlas. This work was sponsored in part by the following grants: NIH 1R01CA138419-01, NIH 1U54EB005149-01, NIH 1R01MH091645-01A1, NIH 2P41EB002025-26A1 and NSF EECS-0925875.

References

1. Anderson, S.M., Rapcsak, S.Z., Beeson, P.M.: Cost function masking during normalization of brains with focal lesions: still a necessity? Neuroimage 53(1), 78–84 (2010)
2. Brett, M., Leff, A., Rorden, C., Ashburner, J.: Spatial normalization of brain images with focal lesions using cost function masking. Neuroimage 14(2), 486–500 (2001)
3. Hart, G.L., Zach, C., Niethammer, M.: An optimal control approach for deformable registration. In: MMBIA Workshop, CVPR (2009)
4. Hogea, C., Davatzikos, C., Biros, G.: An image-driven parameter estimation problem for a reaction–diffusion glioma growth model with mass effects. Journal of Mathematical Biology 56(6), 793–825 (2008)
5. Lamecker, H., Pennec, X.: Atlas to image-with-tumor registration based on demons and deformation inpainting. In: Proc. MICCAI Workshop on Computational Imaging Biomarkers for Tumors - From Qualitative to Quantitative, CIBT 2010 (2010)
6. Miller, M.I., Younes, L.: Group actions, homeomorphisms, and matching: A general framework. International Journal of Computer Vision 41(1/2), 61–84 (2001)
7. Prastawa, M., Bullitt, E., Gerig, G.: Simulation of brain tumors in MR images for evaluation of segmentation efficacy. Medical Image Analysis 13(2), 297–311 (2009)
8. Stefanescu, R., Commowick, O., Malandain, G., Bondiau, P., Ayache, N., Pennec, X.: Non-rigid atlas to subject registration with pathologies for conformal brain radiotherapy. In: Barillot, C., Haynor, D.R., Hellier, P. (eds.) MICCAI 2004. LNCS, vol. 3216, pp. 704–711. Springer, Heidelberg (2004)
9. Trouve, A., Younes, L.: Metamorphoses through Lie group action. In: Foundations of Computational Mathematics, pp. 173–198 (2005)
10. Zacharaki, E., Hogea, C., Shen, D., Biros, G., Davatzikos, C.: Non-diffeomorphic registration of brain tumor images by simulating tissue loss and tumor growth. Neuroimage 46(3), 762–774 (2009)

Longitudinal Brain MRI Analysis with Uncertain Registration

Ivor J.A. Simpson[1,2,⋆], Mark W. Woolrich[2,3], Adrian R. Groves[2], and Julia A. Schnabel[1]

[1] Institute of Biomedical Engineering, University of Oxford, Oxford
{ivor.simpson,julia.schnabel}@eng.ox.ac.uk
[2] Oxford Centre for Functional MRI of the Brain, University of Oxford, Oxford
adriang@fmrib.ox.ac.uk
[3] Oxford Centre for Human Brain Activity, University of Oxford, Oxford
mark.woolrich@ohba.ox.ac.uk

Abstract. In this paper we propose a novel approach for incorporating measures of spatial uncertainty, which are derived from non-rigid registration, into spatially normalised statistics. Current approaches to spatially normalised statistical analysis use point-estimates of the registration parameters. This is limiting as the registration will rarely be completely accurate, and therefore data smoothing is often used to compensate for the uncertainty of the mapping. We derive localised measurements of spatial uncertainty from a probabilistic registration framework, which provides a principled approach to image smoothing. We evaluate our method using longitudinal deformation features from a set of MR brain images acquired from the Alzheimer's Disease Neuroimaging Initiative. These images are spatially normalised using our probabilistic registration algorithm. The spatially normalised longitudinal features are adaptively smoothed according to the registration uncertainty. The proposed adaptive smoothing shows improved classification results, (84% correct Alzheimer's Disease vs. controls), over either not smoothing (79.6%), or using a Gaussian filter with $\sigma = 2$mm (78.8%).

1 Introduction

Alzheimer's disease (AD) is a progressive neurological disease, and the most common to be associated with the symptoms of dementia. The Alzheimer's Disease Neuroimaging Initiative (ADNI) [7] is a large multi-site study whose primary goal is to test whether serial magnetic resonance imaging (MRI) and other biological and imaging markers can be combined to measure the progression of mild cognitive impairment (MCI) and early AD. Determination of sensitive and specific markers of very early AD progression is intended to aid researchers and clinicians to develop new treatments and monitor their effectiveness, as well as lessen the time and cost of clinical trials. ADNI's initial goal was to recruit 800

⋆ This work was funded by the EPSRC through the LSI DTC.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 647–654, 2011.

adults, ages 55 to 90, to participate in the research, approximately 200 cognitively normal older subjects to be followed for 3 years, 400 people with MCI to be followed for 3 years and 200 people with early AD to be followed for 2 years.

Longitudinal studies of this scale provide a platform to analyse the progression of the anatomical effects of neurodegenerative diseases. The use of imaging tools on longitudinal data allows quantification of the rate of brain atrophy, e.g. SIENA [11]. Feature maps can be derived from the deformation fields required to warp between follow-up images and baseline scans, providing high-resolution information illustrating the anatomical changes taking place over time [9]. The use of deformation derived features is known as tensor-based morphometry (TBM). A popular feature one can choose to use is the determinant of the Jacobian of the deformation fields, which is interpreted as the local expansion/contraction at a particular location and in the case of longitudinal brain imaging this would describe atrophy. Once these longitudinal features have been created, they need to be spatially normalised to enable comparison. This spatial normalisation is required to be sufficiently accurate to ensure robust inference. Accurate spatial normalisation can be provided using non-rigid registration methods [6]. A drawback in current approaches is that a point-estimate of the registration parameters is used, e.g. in [2], which assumes that the mapping is exactly correct. This however, in the case of inter-subject brain registration, is highly unlikely. In a recent comparative study of volumetric brain registration algorithms, it was shown that no algorithm was capable of entirely correctly registering large labelled regions [5]. This situation can be improved upon by using a probabilistic registration method, which includes a level of uncertainty on the inferred transformation. The concept of spatial uncertainty in low-level vision has been previously explored [12]. More recently, probabilistic registration methods have been shown to produce a map of registration uncertainty in intra-subject brain registration following tumour resection [8]. However, although their registration model is capable of estimating the weighting between the similarity and regularisation terms, in this work they prefer to use an ad-hoc weighting of these factors. This is probably because their approach employs Markov chain Monte Carlo (MCMC) to infer the registration model parameters, which is particularly computationally expensive. The significance of this weighting is that it will influences both the registration result, and the distribution of the spatial uncertainty.

In this work we propose the use of a generic and adaptive approach for probabilistic non-rigid registration to provide a more principled estimate of the spatial uncertainty of a transformation. Based on these estimates, local Gaussian smoothing kernels can be automatically estimated and used to smooth image features over the set of probable locations, rather than just the most likely. This approach is demonstrated using longitudinal data from the ADNI dataset, where Jacobian maps are used as image features. These are spatially normalised using a principled probabilistic approach. This transformed data is adaptively smoothed based on the registration derived uncertainty. Our adaptive approach to image smoothing is compared, and performs favourably to an ad-hoc Gaussian smoothing, or no image smoothing for multi-variate disease state classification.

2 Methods

Image registration can be described probabilistically using a generative model, where it is assumed that the target image data $\mathbf{Y}$ can be generated from a source image $\mathbf{X}$, which is deformed by a transformation, where $\mathbf{T}(\mathbf{X},\mathbf{w})$ is the transformed source image, and $\mathbf{w}$ parametrises the transformation. The specific form of $\mathbf{T}$ used in this implementation is a Free-Form Deformation (FFD) model, where $\mathbf{w}$ is the set of control point displacements in each direction.

As the model will have residual error throughout the registration process, this error estimate needs to be included. The noise is assumed to have zero mean and be independent and identically distributed (i.i.d.) across image voxels. In this case, the noise is assumed to be normally distributed $\mathbf{e} \approx N(0, \phi^{-1}\mathbf{I})$, where $\mathbf{I}$ represents the matrix identity. The noise distribution is assumed to have global precision (inverse variance) across the image, ϕ. The generic generative model for registration is therefore given as $\mathbf{Y} = \mathbf{T}(\mathbf{X}, \mathbf{w}) + \mathbf{e}$. Using a normally distributed noise model, as we use here, is equivalent to using the sum of squared differences (SSD) as the image similarity term.

2.1 Priors

In a probabilistic model for registration, regularisation can be incorporated as a prior on the transformation parameters, which is modelled using a Multivariate Normal Distribution (MVN). The prior on $\mathbf{w}$ is described in eq. 1 where Λ encodes the regularisation as a spatial kernel matrix providing bending energy regularisation. λ is the spatial precision parameter, controlling the level of spatial regularisation. λ is modelled as an unknown parameter, and determined adaptively from the data resulting in an automated approach to regularisation. Where λ is constant, this method of regularisation is seen in other probabilistic approaches to registration [2]. In a related approach in groupwise registration, the covariance matrix of the deformations used to warp a template to a set of observed images is calculated from the data [1]. This provides an alternative principled approach to inferring the covariance between deformation parameters, however the model inference is computationally infeasible on full 3D volumes.

Non-informative priors on the spatial precision (eq. 2) and noise precision (eq. 3) are specified using Gamma (Ga) distributions, where the subscript $_0$ denotes initial parameter estimates. We use wide priors over λ and ϕ with initial hyperparameter for scale s $s_0 = 10^{10}, a_0 = 10^{10}$ and shape $c_0 = 10^{-10}, b_0 = 10^{-10}$.

$$P(\mathbf{w}|\lambda) = MVN(\mathbf{w}; 0, (\lambda\Lambda)^{-1}) \tag{1}$$

$$P(\lambda) = Ga(\lambda; s_0, c_0) \tag{2}$$

$$P(\phi) = Ga(\phi; a_0, b_0) \tag{3}$$

2.2 Model Inference

The model parameters are inferred upon using Variational Bayes [3] which uses an objective function of the Variational Free Energy (VFE). The VFE

measures model fit and complexity. Model fit can be approximately considered to be minimising the sum of squared difference, and model complexity as the level of bending energy of the transformation. Using the VB framework, analytic updates are derived for the approximate posterior distributions of the transformation, regularisation and noise parameters, which maximise the VFE.

VB uses the mean-field approximation, hence the posterior parameter distribution $p(\mathbf{w}, \lambda, \phi|\mathbf{Y})$ is approximated as $q(\mathbf{w})q(\lambda)q(\phi)$. The functional forms of the approximate posterior distributions are constrained to be conjugate to the priors, and are given as: $q(\mathbf{w}) = MVN(\mathbf{w}; \mu, \Upsilon)$, $q(\lambda) = Ga(\lambda; s, c)$ and $q(\phi) = Ga(\phi, a, b)$. The hyper-parameter updates (eqs. 4-9) are derived by integrating out the factorised parameters from the log posterior model distribution.

$$\Upsilon = (\alpha\bar{\phi}\mathbf{J}^{\mathsf{T}}\mathbf{J} + \bar{\lambda}\Lambda)^{-1} \quad (4) \qquad \mu_{new} = \Upsilon\left[\alpha\bar{\phi}\mathbf{J}^{\mathsf{T}}(\mathbf{J}\mu_{old} + \mathbf{k})\right] \quad (7)$$

$$c = c_0 + \frac{N_c}{2} \quad (5) \qquad \frac{1}{s} = \frac{1}{s_0} + \frac{1}{2}\left(Tr(\Lambda\Upsilon^{-1}) + \mu^{\mathsf{T}}\Lambda\mu\right) \quad (8)$$

$$b = b_0 + \frac{N_v\alpha}{2} \quad (6) \qquad \frac{1}{a} = \frac{1}{a_0} + \frac{1}{2}\alpha(\mathbf{k}^{\mathsf{T}}\mathbf{k} + Tr(\Upsilon^{-1}\mathbf{J}^{\mathsf{T}}\mathbf{J})) \quad (9)$$

Here, $\mathbf{J}$ is the matrix of first order partial derivatives of the transformation parameters with respect to the transformed image $\mathbf{T}(\mathbf{X}, \mu_{old})$, centred about the previous estimate of the mean μ_{old}. $\mathbf{k}$ is the vector representing the residual image $\mathbf{Y}-\mathbf{T}(\mathbf{X}, \mathbf{w})$. μ_{new} describes the current estimated transformation parameters, and is dependent on the old estimated values. The approximate posterior covariance matrix of the set of transformation parameters is given by Υ. $\bar{\lambda}$ is the expectation of the posterior spatial precision distribution and $\bar{\phi}$ is the expectation of the estimated noise precision. N_c is the number of active control points in the model and N_v is the number of active voxels, α is the virtual decimation factor which accounts for the correlation in the image noise [4].

2.3 Spatial Uncertainty

The probabilistic registration method intrinsically provides a measurement of uncertainty on the posterior distribution of the model parameters. This is particularly interesting for the transformation parameters, given by Υ, which has units of mm^2. The uncertainty of $\mathbf{w}$ represents how certain we are that a given point in the source image should be transformed to a particular point in the target image. By accurately estimating Υ, and interpolating the variance and cross-directional covariance across the image, a multivariate normal distribution illustrating the spatial uncertainty can be calculated at each voxel. The spatial uncertainty of a particular control point is governed by 5 factors:

- The local image information which is affected by this control point's movement ($\mathbf{J}^T\mathbf{J}$).
- The noise precision: how much the image data is trusted, related to SSD.
- The spatial precision: how similar the transformation is to the spatial prior.

– The form of the spatial prior, e.g. bending energy, membrane energy.
– The uncertainty of neighbouring control points.

This distribution provides both the magnitude and direction of the uncertainty at a given point. As the uncertainty is dependent on the image information, it is lower across an image boundary than along it. This results in an anisotropic measure of spatial uncertainty, which varies across the image. The scale of uncertainty will also vary across individual registrations. A straightforward approach to compensate for the spatial uncertainty in the mapping of any given voxel is to smooth the data according to the local uncertainty distribution, which is an anisotropic Gaussian kernel. Where the uncertainty of the mapping is high, smoothing the data compensates for this, but still retains potentially discriminating information. We propose this novel method to help reduce the inter-subject variability due to mis-registration of spatially normalised image data.

3 Materials

Data was provided from the Alzheimer's Disease Neuroimaging Initiative (ADNI) [7], a large multi-side longitudinal study containing MR imaging data. For this study 125 random control subjects, and AD patients, of both sexes and a range of ages AD (mean age 76.6 std.dev 7.67, 67M, 58F) and Normal controls (mean age 78.6 std.dev 5.61, 66 M, 59F) were drawn from the database. Pre-processed images which have been corrected for geometric distortions, bias fields and geometric scaling are present on the ADNI website, and were used in this work. Subjects were chosen with at least 2 scans with a minimum interval of 1 year.

4 Experiments

Tools from the publicly available open-source software library, FSL[1] were used to pre-process the images. To correct for difference in size and location, each of the scans was registered to the MNI 152 template using 9 degrees of freedom. Each scan was also re-sampled to have 1mm isotropic voxels and was processed to remove non-brain tissue. An initial atlas was created by averaging 40 healthy individual control scans after initial affine alignment to the MNI152 template. To create a sharper atlas, each of these scans was then non-rigidly registered to the affine atlas using the probabilistic registration tool with a 5mm knot spacing.

The most recent follow-up scan was registered to the baseline scan of each individual using the probabilistic registration tool with a 5mm knot spacing, which yielded individual maps of local atrophy. As the interval between scans varied across subjects, the Jacobian values were linearly scaled to a single year.

The probabilistic non-rigid registration algorithm was used to provide accurate spatial normalisation, which allows measurement of spatial uncertainty. Each baseline scan was registered to the atlas, an example is given in figure 2. Each registration result yields a map of the variance and cross-directional covariance of the transformation parameters. An illustration of the uncertainty in the mapping between an AD subject to the atlas is given in figure 1.

[1] `http://www.fmrib.ox.ac.uk/fsl/`

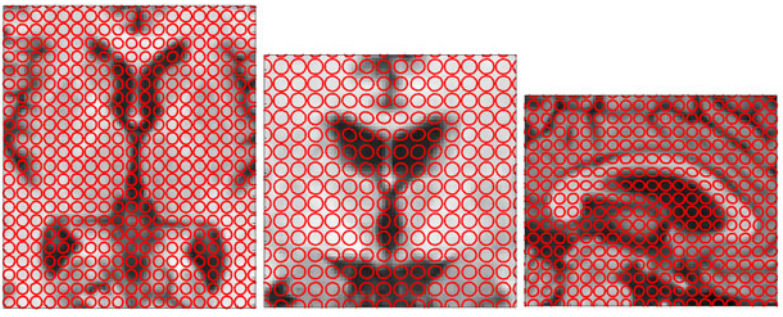

Fig. 1. An example image showing the distribution of spatial uncertainty in the mapping between an AD patient and the atlas. Ellipsoids are plotted every 5mm showing the full-width at half-maximum of the uncertainty distribution at each point. The underlying image is the patient image transformed to the atlas.

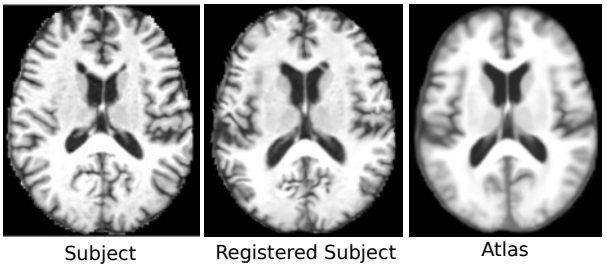

Fig. 2. An example registration from the AD patient shown in figure 1 to the atlas

Each individual subject's longitudinal Jacobian map was transformed to the atlas space. It is then adaptively smoothed using a unique 3-D Gaussian kernel for each spatial location derived from the uncertainty of the registration from the baseline to atlas space. As control experiments, we compare this strategy against not smoothing the images, and Gaussian smoothing with $\sigma = 2mm$.

5 Results

Data decomposition and feature selection was required to provide computationally tractable classification on the 250 spatially normalised and smoothed Jacobian maps, each of which had 2068979 voxels within the anatomical mask. The concatenated data was processed using a principal component analysis (PCA) algorithm to reduce the data dimensionality. The feature maps were masked prior to the PCA using a voxelwise t-test between the two populations with a threshold of $p < 0.05$ uncorrected. The percentage of voxels within the anatomical atlas region included in the mask was 26.82% for the unsmoothed data, 31.91% for the Gaussian smoothed data and 29.74% for the proposed method. Thresholded z-stat maps of the 3 smoothing methods are presented in figure 3. The lowest

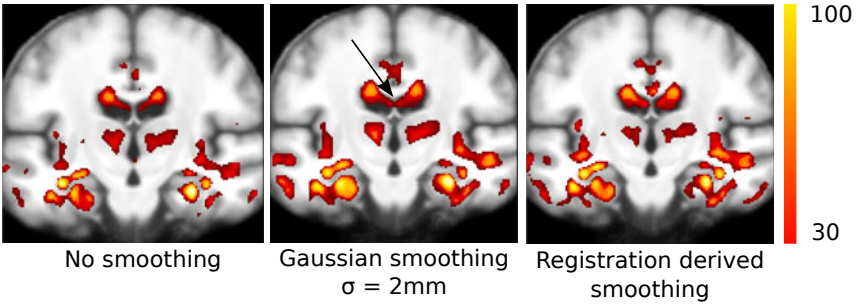

Fig. 3. Z-stat maps showing the results from an uncorrected, two-tail t-test on the spatially normalised Jacobian maps between AD patients and control subjects. The registration derived smoothing does not blur the boundaries between the ventricles, unlike Gaussian smoothing, and it increases the signal to noise over no smoothing.

Table 1. Classification results using an SVM with a radial basis function kernel for the different smoothing methods

Smoothing method	Correct rate	Sensitivity	Specificity
No smoothing	0.796	0.704	0.888
2mm Gaussian Smoothing	0.788	0.904	0.67
Adaptive Smoothing	0.84	0.968	0.71

uncorrected p-value observed for the individual methods is 6.8736×10^{-20} with no smoothing, 1.1289×10^{-20} for Gaussian smoothing and 1.5601×10^{-21} when using the proposed adaptive smoothing. In each case the most significant voxel was located in the left hippocampus which is consistent with previous work [6].

The components which explained 99.0% of the sample variance were used to project the data. The projected data were classified using a support vector machine with a radial basis function kernel using $\sigma_{RBF} = 2$. A leave-two-out methodology was used, where an instance of each class was used for testing at each iteration, while the rest of the data were used for training. The average classification results are given in table 1 and show an improvement in the correct classification rate when using the proposed approach.

6 Conclusions and Future Work

Adaptive smoothing of spatially normalised image feature data, based on registration derived uncertainty, has been demonstrated to provide an increase in the ability to classify between patients with Alzheimer's Disease and normal controls using longitudinal Jacobian maps as features. Although the uncertainty measurements we derive from the image registration process are only a surrogate indication of the true anatomical uncertainty, we have shown that they still provide useful information for reducing the inter-subject variability due

to mis-registration. As the methods presented here are generic, they could be employed in alternative applications, including probabilistic segmentation propagation [10]. Future work includes a more rigorous experimental design and including MCI subjects. Additionally, we will consider the use of the independent uncertainty distribution for each FFD control point, as opposed to the unfactorised estimates used in this work. This may provide more accurate estimates of the spatial uncertainty for use as a smoothing kernel.

References

1. Allassonniére, S., Amit, Y., Trouvè, A.: Toward a coherent statistical framework for dense deformable template estimation. Journal of the Royal Statistical Society, Series B 69(2) (2007)
2. Ashburner, J.: A fast diffeomorphic image registration algorithm. Neuroimage 38(1), 95–113 (2007)
3. Attias, H.: A variational Bayesian framework for graphical models. In: Leen, T., Dietterich, T., Tresp, V. (eds.) NIPS 2000, vol. 12, pp. 209–215. MIT Press, Cambridge (2000)
4. Groves, A.R., Beckmann, C.F., Smith, S.M., Woolrich, M.W.: Linked independent component analysis for multimodal data fusion. NeuroImage 54(3), 2198–2217 (2011)
5. Klein, A., Andersson, J., Ardekani, B., Ashburner, J., Avants, B., Chiang, M., Christensen, G., Collins, D., Gee, J., Hellier, P., Song, J., Jenkinson, M., Lepage, C., Rueckert, D., Thompson, P., Vercauteren, T., Woods, R., Mann, J., Parsey, R.: Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration. Neuroimage 46(3), 786–802 (2009)
6. Leow, A., Yanovsky, I., Parikshak, N., Hua, X., Lee, S., Toga, A., Jack Jr., C., Bernstein, M., Britson, P., Gunter, J., Ward, C., Borowski, B., Shaw, L., Trojanowski, J., Fleisher, A., Harvey, D., Kornak, J., Schuff, N., Alexander, G., Weiner, M., Thompson, P.: Alzheimer's disease neuroimaging initiative: a one-year follow up study using tensor-based morphometry correlating degenerative rates, biomarkers and cognition. Neuroimage 45(3), 645–655 (2009)
7. Mueller, S., Weiner, M., Thal, L., Petersen, R., Jack, C., Jagust, W., Trojanowski, J., Toga, A., Beckett, L.: Alzheimer's Disease Neuroimaging Initiative. Advances in Alzheimer's and Parkinson's Disease, 183–189 (2008)
8. Risholm, P., Pieper, S., Samset, E., Wells, W.: Summarizing and visualizing uncertainty in non-rigid registration. In: Jiang, T., Navab, N., Pluim, J., Viergever, M. (eds.) MICCAI 2010. LNCS, vol. 6362, pp. 554–561. Springer, Heidelberg (2010)
9. Scahill, R., Schott, J., Stevens, J., Rossor, M., Fox, N.: Mapping the evolution of regional atrophy in Alzheimer's disease: unbiased analysis of fluid-registered serial MRI. Proceedings of the National Academy of Sciences of the United States of America 99(7), 4703 (2002)
10. Simpson, I., Woolrich, M., Schnabel, J.: Probabilistic segmentation propagation. In: Medical Image Understanding and Analysis 2011 (2011)
11. Smith, S., Zhang, Y., Jenkinson, M., Chen, J., Matthews, P., Federico, A., De Stefano, N.: Accurate, robust, and automated longitudinal and cross-sectional brain change analysis. Neuroimage 17(1), 479–489 (2002)
12. Szeliski, R.: Bayesian modeling of uncertainty in low-level vision. International Journal of Computer Vision 5(3), 271–301 (1990)

Geodesic Regression for Image Time-Series

Marc Niethammer[1,2], Yang Huang[1], and François-Xavier Vialard[3]

[1] UNC Chapel Hill
[2] BRIC
[3] Imperial College, London

Abstract. Registration of image-time series has so far been accomplished (i) by concatenating registrations between image pairs, (ii) by solving a joint estimation problem resulting in piecewise geodesic paths between image pairs, (iii) by kernel based local averaging or (iv) by augmenting the joint estimation with additional temporal irregularity penalties. Here, we propose a *generative model* extending least squares linear regression to the space of images by using a second-order dynamic formulation for image registration. Unlike previous approaches, the formulation allows for a compact representation of an approximation to the full spatio-temporal trajectory through its initial values. The method also opens up possibilities to design image-based approximation algorithms. The resulting optimization problem is solved using an adjoint method.

1 Introduction

The analysis of image-time series is important to study brain development, aging processes, or tumor growth to name but a few application areas. Establishing image correspondences and localizing change is essential if global measures are insufficient for analysis. While image registration between image pairs has been extensively researched for decades, considering populations of images is more recent. Here, joint-alignment procedures [6] have become standard tools for cross-sectional population-based image analysis. Lately, methods for longitudinal data analysis have been proposed based on the extension of large-deformation-diffeomorphic-metric-mapping (LDDMM) registration to image time-series [4,8] or series of point clouds [9]. Here, only [9] provides a true generative model, since the approach is based on an initial value formulation for the registration of shapes. Initial value formulations have been theoretically discussed for images [7], but only recently solved as initial value problems [11,1].

Statistics in the LDDMM setting are most naturally performed on momenta with respect to a mean image [10]. If a piecewise geodesic estimation for a time-series is used, this requires transporting the set of momenta (for each measurement point of a time-series) to a reference coordinate frame (as proposed in [9] for points sets)[1]. Instead of reformulating [9] for images using the initial-value formulation for image registration, we investigate the behavior of an approximative time-series model, which generalizes least-square linear regression to the

[1] All LDDMM models for image time-series so far estimate piecewise geodesic paths.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 655–662, 2011.

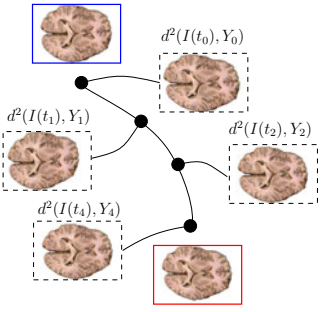

Fig. 1. Principle: The geodesic is determined by the blue and the red images and is closest to the dashed images

image-valued case. The method (i) is generative, describing a full time-trajectory by an initial image and momentum, (ii) will allow for compact statistical analyses of time series based on sets of initial momenta (one for each time series), (iii) opens up the possibility to design approximative algorithms for image time-series (e.g., an approximative spline, useful for random-design data), and (iv) handles non-uniform sampling in time. We use a second order image-based formulation and minimize the sum of squared distances of a set of measured images to a geodesic (Fig. 1). Once a distance measure is defined, we only need to estimate the change of the sum of squared distances with respect to the initial conditions of the dynamical system. This is accomplished by an adjoint solution method (Sec. 4) motivated by a scalar-valued formulation (Sec. 2) and generalized to the image-valued case (Sec. 3). Sec. 5 presents results, Sec. 6 conclusions.

2 Dynamic Formulation of Least-Squares Line Fitting

Consider least-squares linear regression from an optimal control viewpoint: Let $\{y_i\}$ be a set of M measurements at time points $\{t_i\}$ *not necessarily distinct* and $\dot{x}_1 = x_2$; $\dot{x}_2 = 0$ be the dynamical system where the states denote y-intercept and slope respectively. The goal is to find initial conditions $x_1(t_0)$, $x_2(t_0)$ s.t.

$$E = \int_{t_0}^{t_{M-1}} \lambda_1(\dot{x}_1 - x_2) + \lambda_2(\dot{x}_2)\ dt + \sum_{i=0}^{M-1} (x_1(t_i) - y_i)^2 \tag{1}$$

is minimized, where the Lagrangian multipliers λ_1 and λ_2 may be discontinuous. The variation yields the state equation and a boundary value problem for λ_1, λ_2

$$\begin{cases} -\dot{\lambda}_1 = 0, & \lambda_1(t_0^-) = 0,\ \lambda_1(t_{M-1}) = -2(x_1(t_{M-1}) - y_{M-1}), \\ -\dot{\lambda}_2 = \lambda_1, & \lambda_2(t_0) = \lambda_2(t_{M-1}) = 0, \end{cases}$$

with jump conditions $\lambda_1(t_i^-) = \lambda_1(t_i^+) - 2(x_1(t_i) - y_i)$. The gradients of the energy with respect to the initial conditions are $-\lambda_1(t_0^-) = \nabla_{x_1(t_0)} E$, $-\lambda_2 = \nabla_{x_2(t_0)} E$. We can explicitly solve the equations and obtain the conditions

$$\sum_{i=0}^{M-1} (x_1(t_i) - y_i) = 0, \quad \sum_{i=0}^{M-1} (t_i - t_0)(x_1(t_i) - y_i) = 0.$$

The first condition is a force balance (of model residuals) and the second a moment balance. The dynamic formulation extends to the space of images (Sec. 3).

We obtain the mechanical interpretation that λ_1 is the (backward in time) running sum of forces and λ_2 is the (backward in time) running sum of moments. Both need to vanish at optimality. The second-order constraint ($\dot{x}_2 = 0$) is necessary to obtain a straight line. Relaxing this constraint, the method falls back to a piecewise geodesic model as currently used in the LDDMM framework.

3 Geodesic Regression on the Space of Images

The LDDMM framework provides a convenient Riemannian setting where geodesics corresponds to straight lines in the scalar case. Geodesics on the space of the deformed images are obtained by minimizing the functional

$$\mathcal{E}(v) = \int_0^1 \|v\|_V^2 \, dt + d^2(I(1), Y)\,, \tag{2}$$

where v is a time-dependent velocity field in V, a Reproducing Kernel Hilbert Space (RKHS) of smooth velocity fields; d^2 denotes a general squared-distance(-like) term. Extending (2) to multiple timepoints leads to piecewise geodesic interpolation. For a single timepoint however, it gives the geodesicity that needs to be enforced for our least squares generalization. The Euler-Lagrange equation for $\mathcal{E}$ is a special case of the EPDiff equation and parametrizes a geodesic in image space given an initial image $I(t_0)$ and an initial momentum $p(t_0)$ [11]:

$$I_t + \nabla I^T v = 0,\ p_t + div(pv) = 0,\ v + K \star (\nabla I p) = 0. \quad \text{(EPDiff)} \tag{3}$$

where K is the (translation invariant) smoothing kernel of the RKHS and $\star$ the convolution operator. Weighted ($w_i \geq 0$) geodesic regression is to minimize

$$E = \langle p(t_0)\nabla I(t_0), K \star (p(t_0)\nabla I(t_0)\rangle + \sum_{i=0}^{M-1} w_i d^2(I(t_i), Y_i) \tag{4}$$

wrt. the initial conditions ($I(t_0)$, $p(t_0)$) subject to the EPDiff equation (3) that replaces the dynamic line model $\dot{x}_1 = x_2$, $\dot{x}_2 = 0$ in Sec. 2. More importantly, the first term (not present in the scalar case) ensures the well-posedness of the model by preventing high frequencies in the time-dependent velocity field.

3.1 Optimality Conditions

Evolution equations for the adjoints are valid piece-wise with jump conditions at measurement instants. Jumps depend on how much a measured image "pulls" at the geodesic. Weights w_i for the measurements allow for the equivalent of locally linear regression [5] on the space of images[2]. We obtain the state equations and the optimality conditions for the adjoints λ^I, λ^p

[2] This can be seen as an alternative to kernel based methods (as proposed in [3]) and is expected to have improved performance at the boundaries of the time interval.

$$\begin{cases} -\lambda_t^I - div(v\lambda^I) - div(pK * \lambda^v) & = 0, \\ \lambda^I(t_{M-1}) & = -w_{M-1}\nabla_{I(t_{M-1})} d^2(I(t_{M-1}, Y_{M-1})), \\ -\lambda_t^p - v^T \nabla \lambda^p + \nabla I^T K * \lambda^v & = 0,\ \lambda^p(t_{M-1}) = 0, \\ \lambda^I \nabla I - p\nabla \lambda^p + \lambda^v & = 0,\ t \in [t_i^+, t_{i+1}^-],\ i = 0(1)M-2. \end{cases}$$

subject to the compatibility conditions

$$\begin{cases} \lambda^I(t_i^-) = \lambda^I(t_i^+) - w_i \nabla_{I(t_i)} d^2(I(t_i), Y_i), & i > 0, \\ \lambda^p(t_i^-) = \lambda^p(t_i^+), & i > 0, \\ 0 = \lambda^I(t_0^+) - w_i \nabla_{I(t_0)} d^2(I(t_0), Y_0) - 2div(p(t_0)K * (p(t_0)\nabla I(t_0 0)), & i = 0, \\ 0 = -\lambda^p(t_0^+) + 2\nabla I(t_0)^T K * (p(t_0)\nabla I(t_0)), & i = 0. \end{cases}$$

For notational convenience define $\lambda^I(t_0) := \lambda^I(t_0^+) - w_0 \nabla_{I(t_0)} d^2(I(t_0), Y_0)$, $\lambda^p(t_0) := \lambda^p(t_0^+)$. We obtain the gradients

$$\begin{aligned} \nabla_{I(t_0)} E &= -\lambda^I(t_0) - 2div(p(t_0)K * (p(t_0)\nabla I(t_0))), \\ \nabla_{p(t_0)} E &= -\lambda^p(t_0) + 2\nabla I(t_0)^T K * (p(t_0)\nabla I(t_0)). \end{aligned}$$

Note that both initial momentum *and* the initial image are unknowns. To fully define the problem we need to specify the distance measure d^2 and its gradient.

3.2 Choices for d^2

Selecting d^2 is a design choice. The gradients (wrt. $I(t_i)$) can be viewed as forces pulling on the geodesic. We discuss the gradients for distances based on the L^2 metric and LDDMM registration; a metamorphosis approach could also be used.

L^2, Interpolation-Based Image-Match Term. The squared L^2 distance between images and its (infinite-dimensional) derivative is

$$d^2(J, Y) := \|J - Y\|^2; \quad \nabla_J d^2(J, Y) = 2(J - Y).$$

Note that other similarity measure such as cross-correlation or mutual-information could also be used. This definition simplifies computations. It is only meaningful if the geodesic is close to the measured images. If large distances are admissible the squared distances can be defined by registration themselves.

Approximation-Based Inexact Image-Match Term. We use the same second order model as for the regression line for image-to-image registration to define:

$$d^2(J, Y) = \underset{p(0)}{\operatorname{argmin}} \langle p(0)\nabla I(0) K \star p(0)\nabla I(0) \rangle_{L^2} + \frac{1}{\sigma^2} \|I(1) - Y\|^2, \qquad (5)$$

subject to the EPDiff equation with initial image given by $I(0) = J$. This is a special case of the geodesic regression problem with two images with an L^2

distance measure. For an optimal set $\{I^*, p^*, v, \lambda^{I*}, \lambda^{p*}, \lambda^{v*}\}$) the gradient of the squared distance measure with respect to J is hence given by

$$\nabla_J d^2(J, Y) = -\lambda^I(0) - 2div(p(0)K * (p(0)\nabla I(0))).$$

Note the slight abuse of notation, since here λ^I is not the Lagrangian multiplier for geodesic regression, but for the geodesic of the registration problem.

3.3 Mechanical Interpretation

A similar mechanical interpretation as for the scalar-valued case holds. Since the state and the adjoint equations are more involved, we can no longer explicitly solve the optimality conditions. However, the Lagrangian multipliers λ^I can be considered as the generalized running sum of forces and λ^p as the generalized running sum for the moment. The gradients of the squared distance measures with respect to the respective source images can be considered generalized forces. The additional terms appearing in the energy gradients with respect to the initial conditions result from penalizing the length of the geodesic.

4 Numerical Solution

To obtain a solution fulfilling the optimality conditions, we compute the gradients with respect to $I(t_0)$ and $p(t_0)$ through the adjoints. We use a multi-scale approach to speed up convergence. The algorithm proceeds as follows

0) Specify an initial ($I(t_0)$, $p(t_0)$).
1) Solve the state equation forward, while saving computed values $I(t)$ and $p(t)$.
2) Solve the adjoint equations backward while applying jump conditions at every time-point with an available measured image.
3) Compute $\nabla_{I(t_0)}E$ and $\nabla_{p(t_0)}E$ from $\lambda^I(t_0)$, $\lambda^p(t_0)$, $I(t_0)$ and $p(t_0)$.
4) Use the gradients to update $I(t_0)$ and $p(t_0)$ through a line search.
5) Repeat from step 1 until converged.

We solve the advection equation for I and the scalar conservation law for p using a map-based approach as proposed in [2] to minimize numerical dissipation. We use a similar approach to solve for the adjoints λ^I and λ^p, but treat all terms which cannot be explained by advection or scalar conservation as source terms which are added to the solutions at each time step as in [11,8]. If it is desired to obtain a least squares fit such that $I(t_0) = I_0$ (where I_0 is a given fixed initial image) the gradient with respect to $I(t_0)$ can simply be disregarded. In the scalar case this amounts to fixing the y intercept at t_0 and searching for the best slope.

5 Experimental Results

We tested the geodesic regression model using synthetic and real images. Note that this paper focuses on the formulation and solution of the geodesic regression model. Validation in the context of population studies will be future work.

5.1 Synthetic Images

We applied the method to a translating circle with $w_i = 0.1$ and a Gaussian kernel K with $\sigma = 4$ pixels. Fig. 2 (right) shows the original images and the geodesic regression results when updating initial momentum and the initial image. Since the displacements between consecutive images is small we used the L^2 distance. The geodesic regression captures the translation well. As expected for a fluid registration non-uniform compressions and dilations occur. Fig. 2 (left) shows an example geodesic regression trajectory for a fixed initial template through three distinct geometric objects. Since this is an approximative algorithm, the shapes are not perfectly recovered, but instead an intermediate solution is obtained, which approximates the square as a shape in-between the circle and the diamond-shape. The size of all synthetic images is 64×64 pixels.

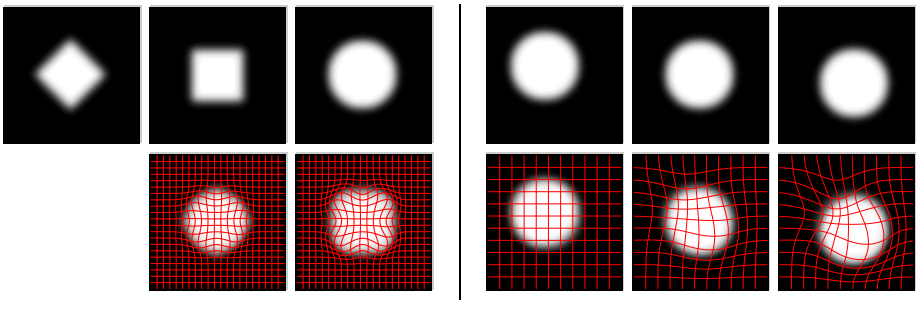

Fig. 2. Left: Geodesic regression result with initial image fixed. The trajectory is a compromise between the shapes. A perfect match of the square and the circle would require a local contraction (with respect to the diamond shape) followed by an expansion, which cannot be expressed. Right: Translation experiment and results for geodesic regression with initial image and momentum free. The translation is well captured.

5.2 Real Images: Brain Slices from the OASIS Database

To illustrate the behavior for real images, we took 5 brain slices (176×208 pixels) from 5 subjects at different ages from the OASIS database and applied geodesic regression with $w_i = 10$ and a Gaussian kernel with $\sigma = 5$ pixels using the L^2 distance. Cases were selected to exhibit an expansion of the ventricles. Since the ventricle topology is different for the subjects, perfect matching cannot be achieved. Even though large scale deformations were present, geodesic regression approximates the temporal evolution of the ventricles. This experiment illustrates that reasonable approximations can be obtained. Computing the geodesic from young to old is numerically challenging, because of the large expansions occurring, which need to be represented by relatively few grid cells in the image. Fig. 3 (left) therefore also shows the estimation results when the geodesic is represented in the space of the oldest subject. Computed deformations are better behaved, because they are easier to represent numerically. In comparison to

time-series approaches estimating piece-wise geodesics [8,4], geodesic regression results in a smooth temporal evolution. This is shown in Fig. 3 (right) which compares a coronal cross section (through the ventricles) for the slices over time. Kinks are visible for the piece-wise geodesic approach (at the time-points of the images), but not for the geodesic regression approach.

Much smaller changes are expected for longitudinal data. Fig. 4 shows four slices (128×161 pixels) for a longitudinal dataset from the OASIS database. Changes are subtle and most easily seen around the ventricles. The top row shows overlays with respect to the youngest image and illustrates the increase in ventricle size. The ventricle expansion is well captured by geodesic regression.

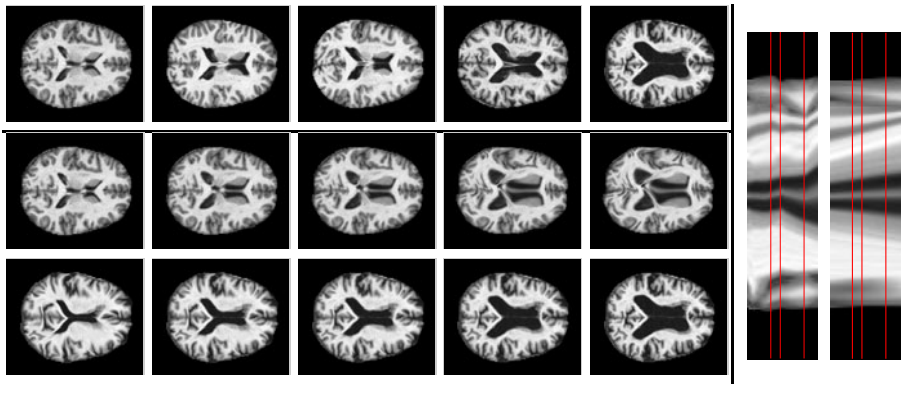

Fig. 3. Left: Brain slices of subjects with increasing age (left to right, 38, 52, 58, 73, 81 [years]) and geodesic regression results. Results with initial conditions in the space of the 38 year old (middle row) and the 81 year old (bottom row) subject. Right: Coronal cross sections through ventricles for piecewise-geodesic approach (left) and for geodesic regression (right) with fixed initial image.

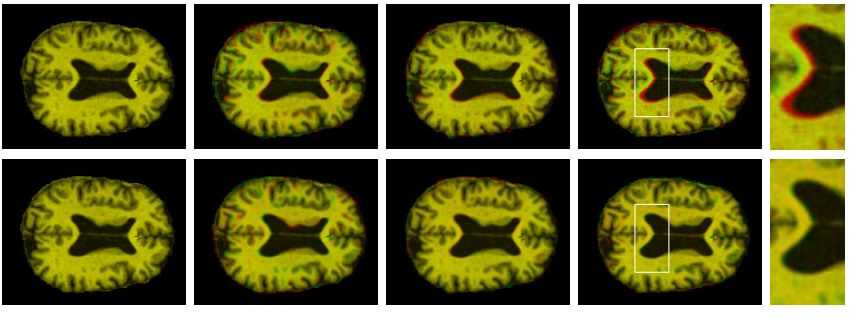

Fig. 4. Brain slices for longitudinal subject data (left to right, age: 67, 68, 71, 73 [years]). Top: original images overlaid with image at age 67. Bottom: geodesic regression results overlaid with original images. Yellow indicates agreement. Zoom in to view.

6 Discussion and Conclusions

We proposed a generative model for image time-series, where trajectories are fully parametrized by their initial conditions. To measure distances from the regression geodesic we proposed an L^2 and a registration-based distance and integrated them into an adjoint solution method. Geodesic regression is an approximative estimation method, which opens possibilities for other approximation methods on the space of images (e.g., approximating splines). We addressed how to estimate an individual trajectory. To perform statistical analysis for populations requires comparing the initial momenta in a common coordinate system, which can be achieved by transporting the initial momenta to a common atlas-space [12]. Geodesic regression simplifies statistical analysis, because of its compact representation of image time-series. It generalizes to piecewise geodesic approximations by concatenating geodesic regressions at *specified* time-points. This allows standardized representations to compare time-series data with non-uniform temporal sampling. This work was sponsored by NIH 1R01MH091645-01A1, NIH 2P41EB002025-26A1 and NSF EECS-0925875.

References

1. Ashburner, J., Friston, K.J.: Diffeomorphic registration using geodesic shooting and Gauss-Newton optimisation. Neuroimage 55(3), 954–967 (2011)
2. Beg, M., Miller, M., Trouvé, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. IJCV 61(2), 139–157 (2005)
3. Davis, B., Fletcher, P., Bullitt, E., Joshi, S.: Population shape regression from random design data. In: ICCV, pp. 1–7 (2007)
4. Durrleman, S., Pennec, X., Trouvé, A., Gerig, G., Ayache, N.: Spatiotemporal atlas estimation for developmental delay detection in longitudinal datasets. In: Yang, G.Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (eds.) MICCAI 2009. LNCS, vol. 5761, pp. 297–304. Springer, Heidelberg (2009)
5. Fan, J.: Local linear regression smoothers and their minimax efficiencies. The Annals of Statistics 21(1), 196–216 (1993)
6. Joshi, S., Davis, B., Jomier, M., Gerig, G.: Unbiased diffeomorphic atlas construction for computational anatomy. Neuroimage 23, S151–S160 (2004)
7. Miller, M.I., Trouve, A., Younes, L.: Geodesic shooting for computational anatomy. Journal of Mathematical Imaging and Vision 24, 209–228 (2006)
8. Niethammer, M., Hart, G., Zach, C.: An optimal control approach for the registration of image time series. In: Conference on Decision and Control (CDC), pp. 2427–2434. IEEE, Los Alamitos (2009)
9. Qiu, A., Albert, M., Younes, L., Miller, M.I.: Time sequence diffeomorphic metric mapping and parallel transport track time-dependent shape changes. NeuroImage 45(1), S51–S60 (2009)
10. Vaillant, M., Miller, M.I., Younes, L., Trouvé, A.: Statistics on diffeomorphisms via tangent space representations. Neuroimage 23, S161–S169 (2004)
11. Vialard, F.X., Risser, L., Rueckert, D., Cotter, C.J.: Diffeomorphic 3D image registration via geodesic shooting using an efficient adjoint calculation (2011) (preprint)
12. Younes, L., Qiu, A., Winslow, R.L., Miller, M.I.: Transport of relational structures in groups of diffeomorphisms. Journal of Mathematical Imaging and Vision 32, 41–56 (2008)

Mapping the Effects of Aβ_{1-42} Levels on the Longitudinal Changes in Healthy Aging: Hierarchical Modeling Based on Stationary Velocity Fields

Marco Lorenzi[1,2], Nicholas Ayache[1], Giovanni B Frisoni[2], Xavier Pennec[1], and The Alzheimer's Disease Neuroimaging Initiative*

[1] Project Team Asclepios, INRIA Sophia Antipolis, France
[2] LENITEM, IRCCS San Giovanni di Dio, Fatebenefratelli, Brescia, Italy

Abstract. Mapping the effects of different clinical conditions on the evolution of the brain structural changes is of central interest in the field of neuroimaging. A reliable description of the cross-sectional longitudinal changes requires the consistent integration of intra and inter-subject variability in order to detect the subtle modifications in populations. In computational anatomy, the changes in the brain are often measured by deformation fields obtained through non rigid registration, and the stationary velocity field (SVF) parametrization provides a computationally efficient registration scheme. The aim of this study is to extend this framework into an efficient and robust multilevel one for accurately modeling the longitudinal changes in populations. This setting is used to investigate the subtle effects of the positivity of the CSF Aβ_{1-42} levels on brain atrophy in healthy aging. Thanks to the higher sensitivity of our framework, we obtain statistically significant results that highlight the relationship between brain damage and positivity to the marker of Alzheimer's disease and suggest the presence of a presymptomatic pattern of the disease progression.

1 Introduction

The ability to map the different areas involved in the neurodegenerative processes is of primary importance for the formulation of new clinical hypotheses on the pathological mechanisms. Moreover, the availability of a longitudinal model of the disease progression would provide a reliable standard for diagnostic purposes. The problem is particularly relevant in the field of Alzheimer's disease (AD) which is characterized by the progressive abnormal configuration of the

* Data used in preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (www.loni.ucla.edu/ADNI). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: `www.loni.ucla.edu/ADNI/Collaboration/ADNI_Authorship_list.pdf`

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 663–670, 2011.

biochemical, functional and structural markers in the brain which may occur up to decades before the clinical assessment [5]. Among the earliest potential markers, the pathological configuration of the CSF $A\beta_{1-42}$ was shown to be associated with a general increased predisposition to clinical conversion to AD. It is therefore of great interest to model the subtle differential evolution from normal aging of the brain changes in subjects who are not affected by the disease but present lower $A\beta_{1-42}$ levels. For this purpose, robust, sensitive, accurate and reproducible modeling techniques are required.

The non-rigid registration is a candidate instrument to quantify the structural differences between brain images and the new generation registration algorithms provide diffeomorphic registration ([14], [9]). Among them, the Log-Demons algorithm provides an accurate and computationally efficient approach, by using stationary velocity fields (SVF) as parametrization of the deformations.

The analysis of longitudinal data requires to go one step further and to integrate the temporal dimension into the registration procedure. The main complexity of the problem lies in the different levels of variation introduced by the different nature of the small intra (*longitudinal*) and large inter-subject (*cross-sectional*) changes: the measurements from time series of a specific subject must be normalized into a comprehensive spatio-temporal atlas. Although different approaches have been proposed in the past for the group-wise analysis of longitudinal dataset ([2], [15]), a consensus on the optimal strategy to handle the different levels of information is still missing, for instance for the choice of the different metrics for intra and inter subject normalization.

We believe that the reliable quantification of the group-wise longitudinal changes should independently address the different sources of variability with proper methods, and consistently integrate the different levels into a general framework. In previous works the SVF setting was shown to provide:

1. An *efficient* pairwise-registration scheme with Log-Demons [14];
2. A *straightforward* way to model the subject-specific deformation trend from time series with a spatio/temporal regularization procedure [7];
3. A *stable* way to transport the subject-specific trends in the atlas geometry using the parallel transport given by the Schild's Ladder procedure [8].

The goal of this paper is 1) to combine these previous contributions in a robust, efficient and precise tool for modeling group-wise deformation, and 2) to use the framework to analyze and model the subtle effects of the CSF $A\beta_{1-42}$ levels on longitudinal brain atrophy in healthy elders.

2 Modeling Changes in Time Series of Images with the SVF Framework

We assume that the subject specific evolutions are random realizations of an underlying ideal population trend. The hierarchical generative model is therefore composed of the following levels:

1. We model the *population trend* as the deformation $\mu^G(t)$ of a template T_0 over time. The (spatially normalized) deformation trend of subject K in the template space is assumed to be a random realization of a Gaussian process $\mu^K(t) = \mu^G(t) + \epsilon_K$. It is the goal of step 4 to estimate the population trend $\mu^G(t)$ from the spatially normalized subject's longitudinal trends.
2. To account for the spatial variability of the anatomy across the population, the subjects specific coordinate system is defined by a spatial changes of coordinates $\phi^{K(-1)}$ from the template to the subject. The *subject specific* deformation trend $\mathbf{v}^K(t)$ is then modeled as the parallel transport of the spatially normalized subject's longitudinal trend $\mu^K(t)$ along the template-to-subject spatial change of coordinates $\phi^{K(-1)}$. Step 3 is taking care of solving the reverse problem in a discrete time setting.
3. Subject specific longitudinal trends are then sampled in time (modeling the discrete acquisition times) and a deformation noise accounting for the influence of random confounding factors (hydratation, vasodilation, etc) is added independently at each time point to obtain the *subject-specific* deformation $\mathbf{v}_i^K = \mathbf{v}^K(t_i) + \epsilon_i$ at time point t_i. Step 2 aims at solving the inverse problem.
4. Last but not least, the subject time series of images is generated by deforming the subject baseline image I_0^K with an acquisition noise on intensities: $I_i^K = \exp(\mathbf{v}_i^K) * I_0^K + \epsilon_i^I$. Step 1 is solving the inverse problem using non-linear registration.

Let us now address the inverse problem: estimating the population trend from the time series of patient images. We detail below step by step the solution we propose to solve each level of the generative model (in the reverse order).

Step 1: Robust pairwise registration with the Log-Demons algorithm. For each subject K, the longitudinal changes along the time series of images I_i^K, $i = 0, \ldots, n$ acquired at time $t_0 = 0, \ldots, t_n$, are evaluated by non rigid registration with respect to the reference time point, here the baseline I_0^K.

The Log-Demons algorithm aims at matching the images I_0 and I_i by looking for the deformation φ which maximises their similarity. The deformation φ belongs to the one-parameter subgroup generated by an optimal vector field $\mathbf{v}$, and the parametrisation is defined by the group exponential map $\varphi = \exp(\mathbf{v})$ [1].

In the standard log-Demons algorithm the "unregularized" correspondence field $\mathbf{v_x}$ is given by the minimization of the sum of squared differences (SSD) between the intensities of the two images, which is not robust to the intensity biases. In order not to mistake spurious intensity variations for morphological differences, we first propose to resort to the local correlation coefficient, introduced in [3]:

$$E(I_0, I_i, \mathbf{v_x}, \mathbf{v}) = \min_{(a,b)} \int G_{\sigma_S} * \|(a(x) \cdot I_0(x) + b(x)) - I_i(x) \circ \exp(\mathbf{v_x})(x)\|^2 + \\ + \frac{1}{\sigma_x^2} \| \log(\exp(-\mathbf{v})(x) \circ \exp(\mathbf{v_x})(x))\|_{L_2}^2 \quad (1)$$

The spatially varying coefficients $a(x), b(x)$ account for the additive and multiplicative biases for the intensities. Moreover the bias estimation is local, thanks

to the Gaussian weights on the error norm. In practice, the standard correspondence energy of the Log-Demons is replaced by $E(I_0, I_i, \mathbf{v_x}, \mathbf{v})$, while preserving the remaining structure of the algorithm. As proposed in [3], the minimization of (1) is operated through a two step procedure: a first step evaluates the optimal scaling factors a and b voxel-wise, that are then reintroduced for the optimization of $\mathbf{v}_x$ through a Gauss-Newton scheme. Experiments on both synthetic and real data showed that the local similarity criteria allows to robustly compute deformations in presence of bias and generally provides smoother estimation of the anatomical differences (data not shown due to space constraints). The important robustness improvements came at the price of a reasonable increase of the computational time (around 25 minutes on a Pentium Intel Core Duo 2.4Ghz for registering images with resolution 182x182x218, voxel size 1x1x1) .

Step 2: Modeling the subject specific longitudinal trends. In order to obtain smoother estimations of the subject specific trajectory and to reduce the intra-subject variability given by possible confounding factors, the Step 2 consists in introducing a temporal correlation into the estimated serial deformations through a 4D registration scheme [7]. The procedure is particularly indicated here, since we are going to investigate the subtle morphological changes occurring in the brain of cognitively healthy subjects on a small number of time points (around 4 for the ADNI dataset), and we do not expect to model sharp variations or sudden modification of the longitudinal series.

The subject specific trend $\bar{\mathbf{v}}^K(t) = L(\mathbf{v}_i^K, t_i, t)$ is estimated with a linear model in time (which is a non-linear deformation model) from the time series of static velocity fields $\mathbf{v}_i^K$ evaluated in the Step 1 for the pairs I_0^K, I_i^K. The $4D$ registration integrates the $\bar{\mathbf{v}}^K(t)$ in a new registration step in order to provide a temporal prior for finally estimate the spatio-temporal regularized sequence of the static velocity fields $\mathbf{v}_i^{'K}$.

The solution at each time point t_i is represented by the weighted average between the temporal prior $\bar{\mathbf{v}}^K(t_i)$ and the spatial correspondence $\mathbf{v_x}$ provided by the similarity measure. Previous experiments showed that the 2:1 trade-off between spatial and temporal weights defines sufficiently smooth trajectories while not biasing the changes towards a completely linear model.

Step 3: Transporting the subjects trajectories in the atlas geometry. In order to compare the longitudinal trajectories between the different subjects and to perform statistical analysis, we need to transport the series of velocity fields $\mathbf{v}_i^{'K}$ in a common reference. For this purpose, we base the transport on the Schild's Ladder method [8]. The method relies on the technique introduced in the field of theoretical physics for computing the parallel transport of tangent vectors on a general manifold without requiring the knowledge of the global geometrical properties of the space. It is based on the construction of a "geodesic parallelogram" for transporting vectors along *any* curve (and not just the geodesics of a specific choice of metric)[1]. More precisely, the parallel

[1] In the case of SVF, the geodesic parallelogram is based on the one-parameter subgroups which are the geodesics of the Cartan connections [12].

transport of the trajectory $\mathbf{v}_i'^K$ from Step 2 along $\phi^K = \exp(t\mathrm{u}^K)$ connecting I_0^K and T_0 is the field $\mathbf{v}_i^{*K} = \Pi^{\phi^K}(\mathbf{v}_i'^K) \simeq \mathbf{v}_i'^K + [\mathrm{u}^K, \mathbf{v}_i'^K] + \frac{1}{2}[\mathrm{u}^K, [\mathrm{u}^K, \mathbf{v}_i'^K]]$.

Step 4: Longitudinal group-wise modeling. The transported time series of SVF $v_i^{*K} = \Pi^{\phi^K}(\mathbf{v}_i'^K)$ belonging to different subjects can now be easily compared in the reference space T_0. In order to develop a group-wise model for the trajectories, we propose here a random effect analysis based on the longitudinal transported trends. Let $\mu^K(t) = L(\mathbf{v}_i^{*K}, t_i, t)$ be the spatially normalized subject trend modeled in the reference space with a linear model in time[2]. The different subject trends $\mu^K(t)$ characterize the trajectories across the populations and by comparing them it is possible to provide a description of the group-wise evolutions. In the following, the different evolutions across the groups (say + and -) will be statistically assessed on the group-wise mean deformation trends $\mu^+(t)$ and $\mu^-(t)$. However, the visual differences between the trends will be illustrated by applying the longitudinal evolutions to the template image: $T^+(t) = \exp(\mu^+(t)) * T_0$ and $T^-(t) = \exp(\mu^-(t)) * T_0$.

3 Effects of Aβ_{1-42} Positivity on Healthy Aging

The T1 weighted longitudinal scans (baseline, 6, 12, 24 and 36 months) were selected for 98 healthy subjects from the ADNI dataset [10]. Two subgroups were then defined based on the positivity to the Aβ_{1-42} marker defined by values below the threshold of 192 pg/ml and resulted in 41 subjects Aβ_{1-42} positives and 57 negatives (Aβ^+_{1-42} and Aβ^-_{1-42}). The two groups were similar at baseline for gender (% of women: 45 % for Aβ^+_{1-42} , 51 % for Aβ^-_{1-42}), age (75±5, 75±5) and education (15.8±3.17, 15.5±2.7). For each subject, the time series of images were aligned through an unbiased procedure consisting on the iterative rigid registration to the median image computed voxel-wise. The final median image was linearly registered to the MNI132 template and the affine transformation was then applied to the series.

The 4D registration algorithm was applied to the longitudinal series of each subject, with $\sigma_S = 10mm$ for the local similarity criteria, $\sigma_{fluid} = 0.5mm$ and $\sigma_{elastic} = 1mm$ for the regularization. The Schild's Ladder was used to transport the longitudinal trajectories from the subject to an unbiased population-based Template T, computed as in [6] (Inter-subject registrations were also computed with the log-Demons algorithm).

The mean trends μ^- of the Aβ^-_{1-42} and μ^+ of the Aβ^+_{1-42} groups were computed from the estimated subject-specific trends. Their difference was assessed on a voxel-by-voxel basis by a multivariate analysis based on the Hotelling's two-sample T^2 statistic (Figure 2C). The statistical significance was assessed after correction for multiple comparisons by means of permutation test (1000 permutations). Moreover, the trends allowed to compute the mean evolutions for the

[2] We notice that the model fitted in the Log-domain does not imply a linear trend for the parametrized deformations.

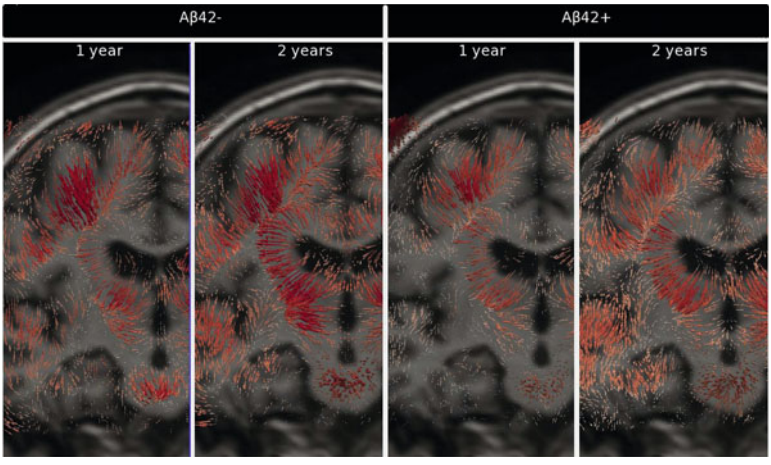

Fig. 1. Average SVF from baseline for the $A\beta^{-}_{1-42}$ (left) and $A\beta^{+}_{1-42}$ (right) groups. For both groups the average forces increase longitudinally, but we can notice an acceleration for the changes across the hippocampus and the temporal regions for the $A\beta^{+}_{1-42}$ group.

Template space and to qualitatively assess the differential progression between the two groups (Figure 2A/B). Finally, a region of interest (ROI) based analysis was performed on the average log-Jacobian values of the estimated trajectories in selected areas of the Template space, segmented through an automated procedure(Ventricles, Hippocampus, Amygdalae, Caudate and Thalamus) [11].

4 Results

Figure 1 shows the average SVF estimated for the two groups from baseline. Althought the two groups show a similar pattern for the ventricular expansion, the $A\beta^{+}_{1-42}$ shows an increased flow of vectors across the temporal regions and hippocampus. Figure 2A highlights the modeled longitudinal changes from baseline for the $A\beta^{-}_{1-42}$ group. The aging effect can be appreciated in the ventricular expansion and in the spread cortical changes. The *additional* changes due to the positivity to the marker $A\beta_{1-42}$ are displayed in Figure 2B. The positivity to $A\beta_{1-42}$ is characterized by increased longitudinal changes located in the temporal areas and by ventricles expansion. We notice that the average progression built from the estimated SVF allowed to extrapolate the expected evolution 2 years after the end of the study. The multivariate statistical assessment of the differences between the evolution of the two groups is shown in Figure 2C. It involves hippocampi, ventricles and the temporal regions. Interestingly, the voxel-by-voxel statistical analysis on the associated log-Jacobian scalar maps showed similar patterns but failed to reach the statistical significance after the correction for multiple comparisons. This suggests a higher sensitivity of

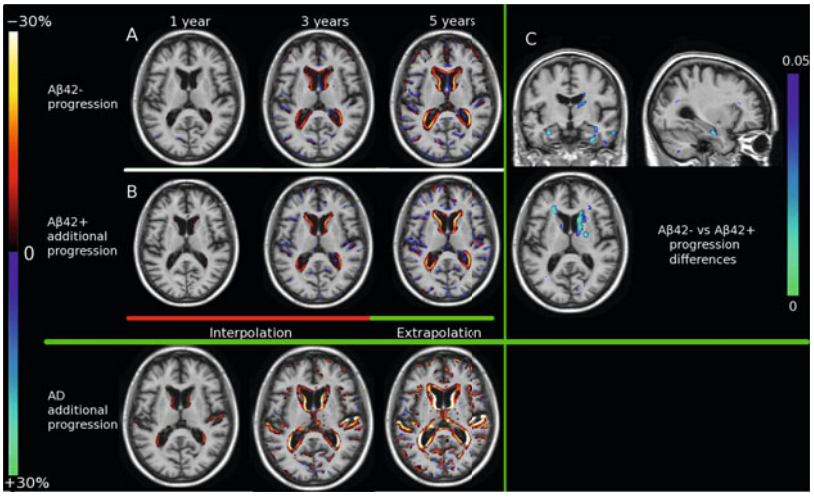

Fig. 2. Modeled longitudinal annual % intensity changes for A) the Aβ^{-}_{1-42} group with respect to the baseline, and for B) the Aβ^{+}_{1-42} group with respect to the Aβ^{-}_{1-42} longitudinal progression. In C) are shown the areas of statistically significant difference between the trends of the Aβ^{-}_{1-42} and the Aβ^{+}_{1-42} groups (p<0.05 corrected). Last row: modeled additional loss with respect to the Aβ^{-}_{1-42} progression for an AD group from the ADNI dataset. We can notice the analogies with the Aβ^{+}_{1-42} trend.

the analysis when performed on the multivariate SVF $\mathbf{v}$ rather than on scalar higher order quantities such $\mathbf{det}(\nabla \mathbf{v})$. Supplementary material can be found in http://www.inria.fr/sophia/members/Marco.Lorenzi/SVF_Framework. The regional differences were confirmed by the ROI based analysis, where significant differences for the volume change/year were found in the ventricles (3.84% for Aβ^{-}_{1-42}, 6.72% for Aβ^{+}_{1-42}, p=0.009) and in the hippocampus (0.14%, 0.24%, p= 0.014) while no significant differences were detected in the other regions.

5 Conclusions

The present work introduces a consistent and effective framework for the analysis of longitudinal data of 3D MRI images. It allowed to model the subtle changes which differentiate the longitudinal evolution of healthy people with abnormal Aβ_{1-42} level from those in the normal range, given by increased ventricular expansion and spread matter loss in the temporal regions ([4], [13]). The resulting trajectories incorporate a wide range of informations (velocities, deformations, volume changes, ...) which could provide new insights for the understanding of the biological phenomenas, like modeling the pathological evolutions (such as in Figure 2). For instance, the extrapolation result is an appealing feature in epidemiology as it enables previsions that could motivate clinical hypothesis. Moreover, the soundness of the extrapolated data indicate the stability and the robustness of the proposed method.

References

1. Arsigny, V., Commowick, O., Pennec, X., Ayache, N.: A log-euclidean framework for statistics on diffeomorphisms. In: Larsen, R., Nielsen, M., Sporring, J. (eds.) MICCAI 2006. LNCS, vol. 4190, pp. 924–931. Springer, Heidelberg (2006)
2. Avants, B., Anderson, C., Grossman, M., Gee, J.: Spatiotemporal normalization for longitudinal analysis of gray matter atrophy in frontotemporal dementia. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) MICCAI 2007, Part II. LNCS, vol. 4792, pp. 303–310. Springer, Heidelberg (2007)
3. Cachier, P.: Recalage non rigide d'images madicales volumiques, contributions aux approches iconiques et geometriques. PhD thesis (2002)
4. Fjell, A., Walhovd, K., Notestine, C., McEvoy, L., Hagler, D., Holland, D., Blennow, K., Brewer, J., Dale, A.: The Alzheimer's Disease Neuroimaging Initiative: Brain atrophy in healthy aging is related to CSF levels of Ab1-42. Cereb. Cortex, 20-9 (2010)
5. Frisoni, G., Fox, N., Jack Jr., C.R., Scheltens, P., Thompson, P.: The clinical use of structural MRI in Alzheimer's disease. Nat. Rev. Neurol. 6, 67–77 (2010)
6. Guimond, A., Meunier, J., Thirion, J.: Average Brain Models: A Convergence Study. Computer Vision and Image Understanding 77-2 (2000)
7. Lorenzi, M., Ayache, N., Frisoni, G., Pennec, X.: 4D registration of serials brain's MR images: A robust measure of changes applyed to Alzheimer's disease. In: Spatio Temporal Image Analysis Workshop (STIA), MICCAI (2010)
8. Lorenzi, M., Ayache, N., Pennec, X.: Schild's ladder for the parallel transport of deformations in time series of images. In: Székely, G., Hahn, H. (eds.) Information Processing in Medical Imaging - IPMI. Springer, Heidelberg (2011)
9. Miller, M., Trouvé, A., Younes, L.: On the metrics and Euler-Lagrange equations of computational anatomy. Annu. Rev. Biomed. Eng. 4(1), 375–405 (2002)
10. Mueller, S., Weiner, M., Thal, L., Peternsen, R., Jack, C., Jagust, W., Trojanowsky, J., Toga, A., Beckett, L.: The Alzheimer's disease neuroimaging initiative. Neuroimaging Clin. N. Am. 15(4), 869–877 (2005)
11. Patenaude, B., Smith, S., Kennedy, D., Jenkinson, M.: A Bayesian Model of Shape and Appearance for Subcortical Brain. NeuroImage (2011) (in press)
12. Postnikov, M.: Geometry VI. Springer, Heidelberg (2001)
13. Schott, J.M., Bartlett, J., Fox, N., Barnes, J., The Alzheimer's Disease Neuroimaging Initiative Investigators: Increased brain atrophy rates in cognitively normal older adults with low cerebrospinal fluid Ab1-42. Annals of Neurology 68-6 (2010)
14. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Symmetric log-domain diffeomorphic registration: A demons-based approach. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part I. LNCS, vol. 5241, pp. 754–761. Springer, Heidelberg (2008)
15. Younes, L., Qiu, A., Winslow, R., Miller, M.: Transport of relational structures in groups of diffeomorphisms. J. Math. Imaging Vis. 32(1), 41–56 (2008)

Consistent Reconstruction of Cortical Surfaces from Longitudinal Brain MR Images

Gang Li, Jingxin Nie, and Dinggang Shen

Department of Radiology and BRIC, University of North Carolina at Chapel Hill, USA
dinggang_shen@med.unc.edu

Abstract. Accurate and consistent reconstruction of cortical surfaces from longitudinal human brain MR images is of great importance in studying subtle morphological changes of the cerebral cortex. This paper presents a new deformable surface method for consistent and accurate reconstruction of inner, central and outer cortical surfaces from longitudinal MR images. Specifically, the cortical surfaces of the group-mean image of all aligned longitudinal images of the same subject are first reconstructed by a deformable surface method driven by a force derived from the Laplace's equation. And then the longitudinal cortical surfaces are consistently reconstructed by jointly deforming the cortical surfaces from the group-mean image to all longitudinal images. The proposed method has been successfully applied to both simulated and real longitudinal images, demonstrating its validity.

Keywords: Cortical surface reconstruction, longitudinal cortical surface.

1 Introduction

The human cerebral cortex is a thin, highly folded sheet of gray matter with the thickness varying between 1 and 5 mm [1-3]. Reconstruction of cortical surfaces from brain MR images plays a vital role in studying structure and function of human brains, and many methods have been proposed in the literature [1-7]. However, most existing cortical surface reconstruction methods were designed for working on a single image. For studying longitudinal change of cortical structures, which is important to normal development, aging, and disease progression of human brains, it requires more accurate and consistent cortical surfaces reconstruction and representation, since longitudinal cortical changes are usually very subtle, especially in normal aging and Alzheimer's disease. Therefore, applying these existing cross-sectional methods independently to reconstruction of cortical surfaces at each time point in a longitudinal imaging study may generate longitudinally-inconsistent cortical surfaces, due to the inconsistency of tissue segmentation, topology correction, and surface tessellation. Accordingly, efforts have been made toward the reconstruction of cortical surfaces from longitudinal images [8, 9], e.g., the longitudinal processing pipeline in FreeSurfer [8], in which cortical surfaces of the mean image of rigidly-aligned longitudinal images or the median image are used as initialization for each longitudinal image and then independently deformed at each time point.

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 671–679, 2011.

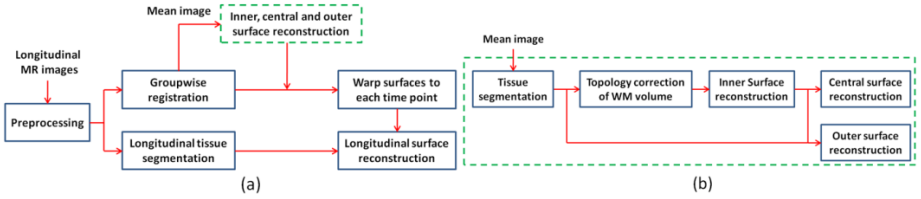

Fig. 1. (a) The flow chart of the longitudinal cortical surface reconstruction method. (b) The flow chart of the cortical surface reconstruction for the group-mean image.

This paper presents a new method for consistent reconstruction of inner, central, and outer cortical surfaces from longitudinal brain MR images. The inner cortical surface is the interface between white matter (WM) and gray matter (GM), and the outer cortical surface is the interface between GM and cerebrospinal fluid (CSF). The central cortical surface is defined as the layer lying in the geometric center of the cortex, approximately corresponding to the cytoarchitechtonic layer four [3, 4, 7]. For consistent reconstruction of cortical surfaces from longitudinal images, in our method, the cortical surfaces of a group-mean image of all non-rigidly aligned longitudinal images are first reconstructed using a deformable surface method, and then the cortical surfaces of the group-mean image are used as the initialization to reconstruct all longitudinal cortical surfaces *simultaneously*. To drive the deformable surfaces towards the target surfaces (inner, central and outer surfaces), a force derived from the Laplace's equation [10] is adopted. Our method has been successfully applied to both simulated and real longitudinal images, demonstrating its validity.

Our proposed method has several advantages over exiting methods. First, temporal constraints are incorporated in longitudinal cortical surface reconstruction by jointly deforming cortical surfaces of all longitudinal images simultaneously in contrast to the existing methods which independently deform the cortical surfaces from the mean image to each time point [8, 9]. Second, a force derived from Laplace's equation [10] is used to drive the deformable surface to help preserve the topology of the deformable surface. Third, the group-mean image is obtained by a nonlinear groupwise registration method in contrast to the rigid alignment used in the existing methods [8, 9]. At last, a longitudinally-consistent tissue segmentation method is adopted to facilitate the longitudinally consistent cortical surface reconstruction.

2 Methods

Given longitudinal brain MR images of a subject, our method for consistent cortical surface reconstruction consists of the following 4 major steps, as shown in Fig. 1. First, longitudinal images are preprocessed. Second, longitudinal images are groupwisely registered to obtain a group-mean image, and also their tissues are consistently segmented into WM, GM, and CSF. Third, the cortical surfaces of the group-mean image are reconstructed using a deformable surface method. Finally, the cortical surfaces of the group-mean image are warped to each longitudinal image and jointly deformed to reconstruct longitudinal cortical surfaces.

2.1 Preprocessing

The preprocessing procedure includes the following steps: (1) intensity inhomogeneity correction, (2) rigid registration of follow-up images onto the baseline image using FLIRT, (3) removing of non-brain tissues of the baseline images, and (4) masking of the brains of follow-up images using the brain mask of the baseline image.

2.2 Groupwise Registration and Consistent Longitudinal Tissue Segmentation

Considering the nonlinear longitudinal changes of brains, instead of averaging rigidly aligned images as did in existing methods [8, 9], a groupwise registration method [11] is adopted to obtain the group-mean image of longitudinal images, as well as the deformation fields from each longitudinal image to the group-mean image. To achieve longitudinally-consistent tissue segmentation, CLASSIC [12] is adopted to perform tissue segmentation on longitudinal images. To recover deep and narrow sulci, we use the ACE method [3] to modify the segmented GM volume to generate a no-more-than-one-voxel thick separation between opposite sulcal GM banks.

2.3 Cortical Surface Reconstruction for Group-Mean Image

To reconstruct cortical surfaces of the group-mean image, the inner cortical surface is reconstructed first in our method, and then the inner surface is deformed to reconstruct both central and outer surfaces. To obtain the topologically-correct inner surface, the topology of the WM volume is first corrected by a graph based method in [5] to ensure a spherical topology, and then the Marching cubes method is used to convert the boundary of the corrected WM volume to an explicit surface representation. Using a deformable surface method, this reconstructed *rough* inner surface is deformed under imposed forces to obtain the refined inner surface, as well as to obtain the reconstruction of the central and outer surfaces one-by-one. Since the deformable surface for cortical surface reconstruction from a single image can be considered as a special case of longitudinal cortical surface reconstruction, the general deformable surface will be detailed in Section 2.4. Fig. 2 shows an example of the reconstructed inner, central, and outer surfaces from an image by our method.

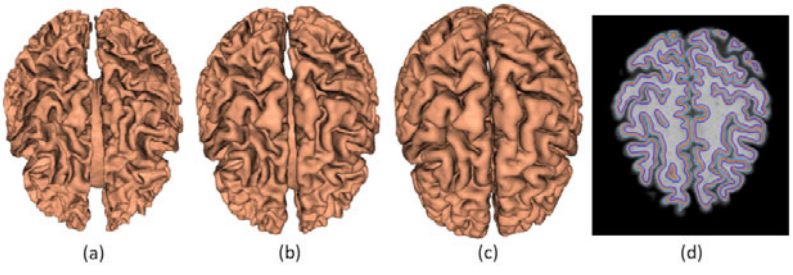

Fig. 2. Reconstructed surfaces of an image. (a) Inner surface. (b) Central surface. (c) Outer surface. (d) Surfaces overlaid on a slice. Dark blue: inner; orange: central; light blue: outer.

2.4 Cortical Surfaces Reconstruction for Longitudinal Images

The cortical surfaces of the group-mean image are warped to each longitudinal image and further jointly deformed for longitudinal surface reconstruction using a deformable surface method. For longitudinal surface reconstruction with n time points, the deformable surfaces at 4D (3D spatial + 1D temporal) domain are parameterized as $\{\mathbf{x}_t(\mathbf{u}) = [x_t(\mathbf{u}), y_t(\mathbf{u}), z_t(\mathbf{u})]^T, \quad \mathbf{u} = (u_1, u_2) \in [0,1] \times [0,1], \quad t \in \{1, \dots, n\}\}$, which can be obtained by minimizing the following energy function:

$$E = \sum_{t=1}^{n} \left(\int \frac{1}{2} \left(\alpha \sum_{i=1}^{2} \left| \mathbf{x}_{t,i} \right|^2 + \beta \sum_{i,j=1}^{2} \left| \mathbf{x}_{t,ij} \right|^2 + \gamma |\mathbf{x}'_t|^2 \right) + E^{ext}(\mathbf{x}_t) d\mathbf{u} \right) \tag{1}$$

where parameters α and β control the tension and rigidity of surfaces, respectively. And the parameters γ controls the temporal smoothness of surfaces. $\mathbf{x}_{t,i}$ and $\mathbf{x}_{t,ij}$ denote the first and second partial derivative of $\mathbf{x}_t$ w.r.t. $\mathbf{u}_i$, respectively. And $\mathbf{x}'_t$ denotes the finite difference of $\mathbf{x}_t$ w.r.t. t. $E^{ext}(\mathbf{x}_t)$ is the external energy derived from the image at time t. The solution to the above energy minimization problem can be obtained by solving the following dynamic equation [4, 7]:

$$\mathbf{x}_{t,\tau}(\tau) = \mathbf{F}^{int}(\mathbf{x}_t(\tau)) + \mathbf{F}^{ext}(\mathbf{x}_t(\tau)) \tag{2}$$

where the internal force $\mathbf{F}^{int}(\mathbf{x}_t(\tau)) = \alpha \Delta_{\mathbf{u}} \mathbf{x}_t(\tau) - \beta \Delta_{\mathbf{u}} (\Delta_{\mathbf{u}} \mathbf{x}_t(\tau)) + \gamma \mathbf{x}''_t$ and the external force $\mathbf{F}^{ext}(\mathbf{x}_t(\tau))$ will be defined later. $\Delta_{\mathbf{u}} = \partial^2/(\partial u_1)^2 + \partial^2/(\partial u_2)^2$ is the Laplacian operator to $\mathbf{u}$, $\mathbf{x}''_t$ denotes the second order finite difference of $\mathbf{x}_t$ w.r.t. t. The first two terms in $\mathbf{F}^{int}(\mathbf{x}_t(\tau))$ is the spatial regularizing force, and the third term is the temporal regularizing force, enforcing the temporal smoothness along the time t. Note that the deformable surface is treated as a function of surface evolution time τ. If only one time point exists, the temporal regularizing force will be 0 and the deformable surface can be used for surface reconstruction for the group-mean image as mentioned above.

The external force driving initial surfaces towards target surfaces is designed as:

$$\mathbf{F}^{ext}(\mathbf{x}_t(\tau)) = \mathrm{G}(\mathbf{x}_t(\tau)) \cdot \mathbf{F}^{GM}(\mathbf{x}_t(\tau)) + \left(1 - \mathrm{G}(\mathbf{x}_t(\tau))\right) \cdot \mathbf{F}^{NonGM}(\mathbf{x}_t(\tau)) \tag{3}$$

where $G(\mathbf{x}_t(\tau))$ is the GM indicator function and $\mathbf{F}^{GM}(\mathbf{x}_t(\tau))$ is a force activated inside of GM and derived from the Laplace's equation [10]. $\mathbf{F}^{NonGM}(\mathbf{x}_t(\tau))$ is a force activated outside of GM and defined as:

$$\mathbf{F}^{NonGM}(\mathbf{x}_t(\tau)) = D(\mathbf{x}_t(\tau)) \cdot (2W(\mathbf{x}_t(\tau)) - 1) \cdot \mathbf{n}(\mathbf{x}_t(\tau)) \tag{4}$$

where $W(\mathbf{x}_t(\tau))$ is the WM indicator function, and $\mathbf{n}(\mathbf{x}_t(\tau))$ is the outward-oriented unit normal vector. $D(\mathbf{x}_t(\tau))$ is the force strength at vertex $\mathbf{x}_t(\tau)$. For inner and outer surface reconstruction, $D(\mathbf{x}_t(\tau))$ is the distance along the direction of $(2W(\mathbf{x}_t(\tau)) - 1) \cdot \mathbf{n}(\mathbf{x}_t(\tau))$ to the WM/GM and GM/CSF interfaces, respectively. For central surface reconstruction, $D(\mathbf{x}_t(\tau))$ is set as the average distance along the direction of $(2W(\mathbf{x}_t(\tau)) - 1) \cdot \mathbf{n}(\mathbf{x}_t(\tau))$ to WM/GM and GM/CSF interfaces.

$\mathbf{F}^{GM}(\mathbf{x}_t(\tau))$ is derived from Laplace's equation of the GM [10], which is a second-order partial differential equation for a scalar field φ enclosed between two interfaces:

$$\Delta\varphi = \frac{\partial^2 \varphi}{\partial x^2} + \frac{\partial^2 \varphi}{\partial y^2} + \frac{\partial^2 \varphi}{\partial z^2} = 0 \tag{5}$$

By setting the WM as the minimal value and the CSF as the maximum value, the Laplace's equation is solved inside of the GM to obtain the harmonic function. The normalized gradient vector field of the harmonic function $\mathbf{T}(\mathbf{x}_t(\tau))$ and the streamlines to both WM/GM and GM/CSF interfaces are computed for each point in GM. The Laplace's equation establishes a one-to-one correspondence between WM/GM and GM/CSF interfaces and the streamlines of the harmonic function never intersect each other. This elegant property helps preserve the topology of the deformable surface. Denote the streamline lengths from a point $\mathbf{x}_t(\tau)$ in GM to WM/GM and GM/CSF interfaces as $L_0(\mathbf{x}_t(\tau))$ and $L_1(\mathbf{x}_t(\tau))$, respectively. $\mathbf{F}^{GM}(\mathbf{x}_t(\tau))$ for inner, central and outer surfaces reconstruction are respectively defined as:

$$\mathbf{F}^{GM}_{inner}(\mathbf{x}_t(\tau)) = -\mathbf{T}(\mathbf{x}_t(\tau)) \cdot L_0(\mathbf{x}_t(\tau)) \tag{6}$$

$$\mathbf{F}^{GM}_{central}(\mathbf{x}_t(\tau)) = \mathbf{T}(\mathbf{x}_t(\tau)) \cdot (L_1(\mathbf{x}_t(\tau)) - L_0(\mathbf{x}_t(\tau))) \tag{7}$$

$$\mathbf{F}^{GM}_{outer}(\mathbf{x}_t(\tau)) = \mathbf{T}(\mathbf{x}_t(\tau)) \cdot L_1(\mathbf{x}_t(\tau)) \tag{8}$$

The central idea is that the direction of $\mathbf{F}^{GM}(\mathbf{x}_t(\tau))$ should point toward the target surface and the magnitude of $\mathbf{F}^{GM}(\mathbf{x}_t(\tau))$ should be directly proportional to the distance to the target surface. Fig. 3 illustrates the streamlines to GM/CSF interface and the force directions in GM.

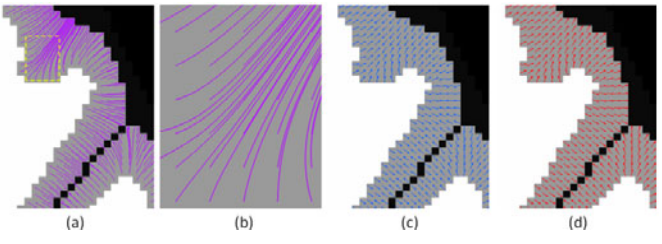

Fig. 3. Illustrations of the streamlines to GM/CSF interface and the force directions in GM. (a) Streamlines. (b) Zooming view of the yellow rectangular region in (a). (c) and (d) Force directions for central and outer surface reconstruction, respectively.

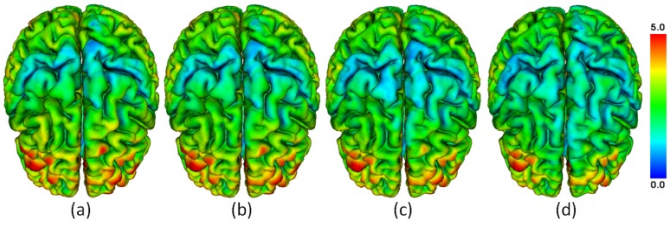

Fig. 4. An example of reconstructed longitudinal outer surfaces, color-coded by thickness

To prevent from self-intersection, a triangle-triangle intersection detection method [13] is used in triangle faces contained in a local sub-volume. Once self-intersection is detected, the deformation is cropped to a valid location. Fig. 4 shows an example of reconstructed longitudinal outer surfaces of a subject with color-coded cortical thickness, calculated as the closest distance between the inner and outer surfaces. We can observe the overall decline trend of the cortical thickness in aging.

3 Results

Data used in experiments are from ADNI [15]. The parameter β controls the rigidity of the deformable surface and is set as 0, as suggested in [4, 7]. Experimentally, α and γ are set as 0.25 and 0.1, respectively. In experiments, we also find that results are not sensitive to subtle changes of parameters.

Real Data. To evaluate the accuracy of the longitudinal inner and outer cortical surface reconstruction results, we compare the GM volume in the tissue-segmented image by CLASSIC [12] (denote as A) with the GM volume enclosed by the corresponding reconstructed inner and outer surfaces (denote as B). Three statistical values are calculated, including: (1) true positive: $(A \cap B)/A$, (2) false negative: $(A \cap \bar{B})/A$, (3) false positive: $(\bar{A} \cap B)/A$, similar to the measurements adopted in [14]. To evaluate the accuracy of the longitudinal central cortical surface, we calculate the percentage of vertices of the reconstructed central surface falling outside the GM (non-GM vertices), as adopted in [7], since the central surface is converged to the inside of GM. The proposed method is applied to 10 normal healthy subjects, each with 4 longitudinal scans. The average true positive, false negative, false positive and percentage of non-GM vertices of 4 time points for each subject are shown in Fig. 5 (a). The above measurements are further compared to the values reported in the literature. Compared to those reported in [14] for validation of inner and outer surfaces, the average true positive of our method is around 0.77, which is higher than the value in [14] which is less than 0.70. And the average false negative is 0.23, compared to the value of around 0.35 in [14]. And the false positive of our method is 0.15, which is similar to the value in [14]. Compared to those reported in [7] for validation of central surfaces, the average percentage of non-GM vertices of the reconstructed central surfaces by our method is around 0.03, which is less than the value of 0.06 in [7]. Although different dataset are adopted in [7] and [14], we believe that this comparison can reflect the accuracy of our method to some degree.

Simulated Data. To further validate the inner and outer surfaces, we simulate images from the reconstructed inner and outer surfaces using the method in [14]. Briefly, the voxels inside the inner surface are labeled as WM, and the voxels between inner and outer surface are labeled as GM, and the voxels between skull and outer surface are labeled as CSF. The longitudinal image sequences are simulated from the corresponding reconstructed inner and outer surfaces. Then the inner and outer surfaces reconstructed from the simulated longitudinal images are compared to those original cortical surfaces, which are treated as the "ground truth". The average distance errors between the two sets of surface are calculated and further averaged for 4 time points of each subject. To validate the central surfaces, we also simulate

images from the reconstructed central surfaces using the method in [7]. Specifically, a thickness value following a Gaussian distribution with the mean 3.0mm and variance 1.0mm is generated for each point of the central surface, and all voxels inside the half thickness range are labeled as GM. Also, all voxels enclosed by GM are defined as WM, and all voxels between the skull and GM are labeled as CSF. The reconstructed central surfaces from the simulated longitudinal images are compared to those original central surfaces. The distance errors of inner, central, and outer surfaces of 10 simulated subjects are shown in Fig 5(b). The average distance errors of inner, central and outer surfaces are all around 0.6mm, indicating the accuracy of our method.

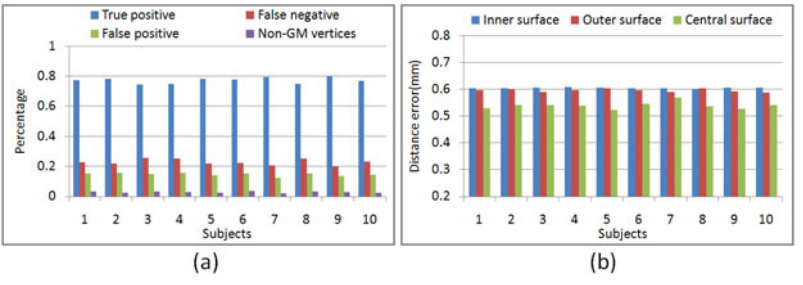

Fig. 5. (a) The average true positive, false negative, false positive and percentage of non-GM vertices of 4 time points for each subject. (b) The average distance errors of inner, central and outer surfaces of 4 time points for each subject compared to the "ground truth".

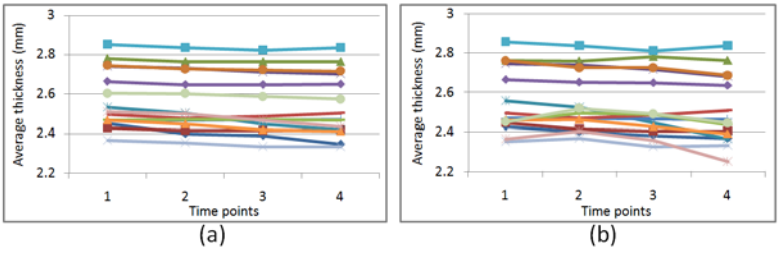

Fig. 6. The longitudinal change of average cortical thickness on 15 normal subjects. (a) Our method with temporal constraint. (b) Our method without temporal constraint.

Longitudinal Thickness Changes. To test the capability of consistently capturing longitudinal cortical changes, we apply the proposed method on 15 normal healthy subjects and calculate the cortical thickness. The trajectory of the average cortical thickness for each subject is shown in Fig. 6 (a). For comparison, the results from our method without temporal constraint (by setting $\gamma = 0$) is shown in Fig. 6 (b). For quantitative comparison, a linear regression is performed on the longitudinal curves of average thickness of each subject, and the residuals are calculated. The average residuals of our method with and without temporal constraint are $0.01 \pm 0.005mm$ and $0.024 \pm 0.02mm$, respectively. As we can see, the results with temporal constraint are more longitudinally consistent and smoother than those without temporal constraint.

Comparison with FreeSurfer. To further demonstrate the consistency of our method, we compare the trajectories of vertices on reconstructed longitudinal outer surface with those obtained from the longitudinal processing pipeline in FreeSurfer [8]. For quantitative evaluation, we calculate the ratio between the length of trajectory and the straight line distance of the first and the last time-point vertices in the longitudinal surfaces. Smaller ratio indicates smoother trajectory. Fig. 7 (a) and (b) show a comparison of the trajectory cropped from longitudinal surfaces obtained by our method and FreeSurfer, respectively. The distributions of ratios of whole surfaces are shown in Fig. 7 (c) and (d). Clearly, the trajectories from our method are much smoother than those obtained from FreeSurfer, since no temporal constraint is incorporated in longitudinal surface reconstruction in FreeSurfer.

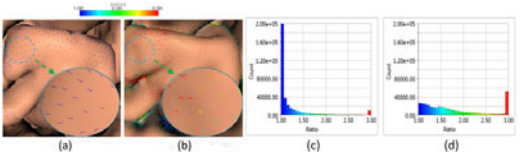

Fig. 7. Comparison of consistency of longitudinal surfaces between our method and FreeSurfer. (a) and (b) The trajectories color-coded by ratios obtained from our method and FreeSurfer, respectively. (c) and (d) The distribution of ratios by our method and FreeSurfer, respectively.

4 Conclusion

This paper presented a new method for consistent reconstruction of cortical surfaces from longitudinal brain MR images and experimental results demonstrate its validity. In future work, we will validate our proposed method using more longitudinal data.

Acknowledgement. This work was supported in part by NIH grants EB006733, EB008374, EB009634 and MH088520.

References

1. Dale, A.M., Fischl, B., Sereno, M.I.: Cortical surface-based analysis. I. segmentation and surface reconstruction. NeuroImage 9, 179–194 (1999)
2. MacDonald, D., Kabani, N., Avis, D., Evans, A.C.: Automated 3-D extraction of inner and outer surfaces of cerebral cortex from MRI. NeuroImage 12, 340–356 (2000)
3. Han, X., Pham, D.L., Tosun, D., Rettmann, M.E., Xu, C., Prince, J.L.: CRUISE: Cortical Reconstruction Using Implicit Surface Evolution. NeuroImage 23, 997–1012 (2004)
4. Xu, C., Pham, D.L., Rettmann, M.E., Yu, D.N., Prince, J.L.: Reconstruction of the human cerebral cortex from magnetic resonance images. IEEE Trans. Med. Imag. 18, 467–480 (1999)
5. Shattuck, D.W., Leahy, R.M.: BrainSuite: an automated cortical surface identification tool. Med. Image Anal. 6, 129–142 (2002)
6. Kim, J.S., Singh, V., Lee, J.K., Lerch, J., Ad-Dab'bagh, Y., MacDonald, D., Lee, J.M., Kim, S.I., Evans, A.C.: Automated 3-D extraction and evaluation of the inner and outer cortical surfaces using a Laplacian map and partial volume effect classification. NeuroImage 27, 210–221 (2005)

7. Liu, T., Nie, J., Tarokh, A., Guo, L., Wong, S.T.: Reconstruction of central cortical surface from brain MRI images: method and application. NeuroImage 40, 991–1002 (2008)
8. FreeSurfer, http://surfer.nmr.mgh.harvard.edu/
9. Nakamura, K., Fox, R., Fisher, E.: CLADA: cortical longitudinal atrophy detection algorithm. NeuroImage 54, 278–289 (2011)
10. Jones, S.E., Buchbinder, B.R., Aharon, I.: Three-dimensional mapping of cortical thickness using Laplace's equation. Hum. Brain Mapp. 11, 12–32 (2000)
11. Wu, G., Jia, H., Wang, Q., Shen, D.: SharpMean: Groupwise registration guided by sharp mean image and tree-based registration. NeuroImage 56, 1968–1981 (2011)
12. Xue, Z., Shen, D., Davatzikos, C.: CLASSIC: consistent longitudinal alignment and segmentation for serial image computing. NeuroImage 30, 388–399 (2006)
13. Moller, T.: A fast triangle – triangle intersection test. J. Graphics Tools 2, 25–30 (1997)
14. Lee, J.K., Lee, J.M., Kim, J.S., Kim, I.Y., Evans, A.C., Kim, S.I.: A novel quantitative cross-validation of different cortical surface reconstruction algorithm using MRI phantom. NeuroImage 31, 572–584 (2006)
15. ADNI, http://adni.loni.ucla.edu/

Inferring 3D Kinematics of Carpal Bones from Single View Fluoroscopic Sequences

Xin Chen[1], Jim Graham[1], Charles Hutchinson[2], and Lindsay Muir[3]

[1] ISBE, School of Cancer and Enabling Sciences, University of Manchester, Oxford Road, Manchester, M13 9PT, UK
{xin.chen,jim.graham}@manchester.ac.uk
[2] Clinical Sciences Research Institute, Clinical Sciences Building, University Hospital -Walsgrave Campus, Clifford Bridge Road, Coventry, CV2 2DX, UK
[3] Consultant Orthopaedic Surgeon, Salford Royal Hospital NHS Foundation Trust, Stott Lane, Salford, M6 8HD, UK

Abstract. We present a novel framework for inferring 3D carpal bone kinematics and bone shapes from a single view fluoroscopic sequence. A hybrid statistical model representing both the kinematics and shape variation of the carpal bones is built, based on a number of 3D CT data sets obtained from different subjects at different poses. Given a fluoroscopic sequence, the wrist pose, carpal bone kinematics and bone shapes are estimated iteratively by matching the statistical model with the 2D images. A specially designed cost function enables smoothed parameter estimation across frames. We have evaluated the proposed method on both simulated data and real fluoroscopic sequences. It was found that the relative positions between carpal bones can be accurately estimated, which is potentially useful for detection of conditions such as scapholunate dissociation.

Keywords: Carpal bones kinematics, 2D 3D registration, Statistical model.

1 Introduction

Chronic pain in the wrist arises due to a number of conditions, such as instability patterns, nonunion or malunion of fractures, primary osteoarthritis and inflammatory arthritis. The result for patients is a severe reduction in quality of life due to impairment of everyday functions, lost work time, increased morbidity and loss of the capacity to live independently. The current method of distinguishing between these conditions is by examining 2D video fluoroscopy sequences showing movement of the hand from full ulnar to full radial deviation and from full flexion to extension in two orthogonal views. From these images clinicians can infer the three-dimensional translations and rotations of the carpal bones that take place during wrist movement, and arrive at a differential diagnosis on the basis of variations from normal bone kinematics. The interpretation is difficult and the accuracy of diagnosis depends wholly on the experience of the practitioner. Currently, accurate diagnosis requires referral to a specialist hand consultant and treatment is often delayed to the detriment of the patient.

The aim of the project is computer interpretation of the fluoroscopy sequences to attain a higher degree of objectivity and quantification in the diagnostic process. During

G. Fichtinger, A. Martel, and T. Peters (Eds.): MICCAI 2011, Part II, LNCS 6892, pp. 680–687.

wrist movement, the eight carpal bones follow a complex, multi-dimensional trajectory, making interpretation of radiographs difficult. For this study we have trained a hybrid statistical model (SM) from a set of CT images from different subjects at different poses. Subsequently, the full 3D carpal bone motions can be recovered by matching the SM with the fluoroscopy sequences through 3D-2D image registration techniques. A number of studies have sought to represent the carpal kinematics using CT or MR data, mainly concentrating on representing 'average' kinematics over a small number of individuals (e.g. [1], [2]). More recently, Van deGiessen et al. [3] presented a 3D rigid registration method based on segmented meshes, which aims to build SM of carpal bones. A study of carpal bone kinematics based on a 4D imaging system was reported in [4]. 3D-2D registration has been the subject of many studies (e.g. [5]), mainly in the field of registration of pre-operative MR or CT images to intra-operative 2D images. Our work differs from the above in that we seek to achieve registration of a 2D image sequence to a 3D model (not derived from the same individual) to derive the kinematics of an individual wrist. Zheng [6] took a similar approach to estimate the orientation of pelvis from a single X-ray image.

The main contributions of this paper, distinguishing it from these studies, are: (1) A hybrid SM is developed representing both the complex kinematics and shape variation of the eight carpal bones plus radius and ulna. (2) The full 3D motion and bone shapes are recovered by matching the SM with a single view fluoroscopy sequence: a difficult ill-posed problem. (3) Our initial results show that the relative positions between the carpal bones can be estimated accurately through the proposed framework. We are not aware of any study which attempts to make a 2D to 3D inference in a system of this level of complexity.

The system consists of a training phase and a 3D-2D image registration phase. We currently have CT data from 10 subjects, each at five poses (neutral pose and two extreme poses in flexion-extension and radial-ulnar deviation). In the training phase, only the data from the neutral pose and two extreme poses in the radial-ulnar movement were used, as the radial-ulnar movement is the current concern of this paper. The segmentation of each bone and rigid registration parameters that align bones at different poses within and across the subjects were obtained using an iterative segmentation and registration algorithm [7]. A hybrid statistical model, representing both the kinematics and shape variation, was built efficiently from the results of the segmentation-registration framework. The kinematic model was built based on the transformation parameters, while the segmentation result was used to build the statistical shape model for each individual bone. In the 3D-2D image registration phase, the 3D rigid transformation, the kinematic motion and bone shapes were estimated in sequence from each frame of the fluoroscopy sequences. Detailed descriptions are given in the following sections.

2 Problem Parameterisation

We use a perspective projection model to represent the relationship between the 2D fluoroscopy image and the 3D configuration of bones. Almost all parameters necessary for this model (pixel size and optical centre) are known. The distance from the X-ray source and the detector needs to be measured for each patient. If this parameter is not

accurate, it will lead to a scale difference of the estimated 3D model. The resulting translation effects on the relative motion between carpal bones at pixels away from the centre of the field is very small.

Three sets of parameters need to be estimated during image registration in order to interpret the true 3D motion of the carpal bones: (1) Rigid transformation parameters of the wrist and a global scale factor, denoted by $\theta=\{t_x, t_y, t_z, \alpha, \beta, \gamma, s_{global}\}$. t_x, t_y and t_z denote the translations, and α, β and γ denote the rotation angles. s_{global} controls the distance between the centroid of each bone to the origin in the radius, and the global size of the bones. (2) Kinematic model parameters M representing the carpal bone poses during movement. (3) Shape model parameters Q_i and scale factor s_i for each bone (i).

3 Training of Kinematic Model and Shape Model

We use the six rigid transformation parameters for each bone to train the kinematic model. The common coordinate system for all pose and scale parameters has an origin at the centroid of the head of radius for one subject. The pose of one subject is described by $(tx_1, ty_1, tz_1, \alpha_1, \beta_1, \gamma_1, ..., tx_{10}, ty_{10}, tz_{10}, \alpha_{10}, \beta_{10}, \gamma_{10})^t$. (8 carpal bones, 1 radius and 1 ulna). The orientation parameters all occupy values distant from the angular discontinuity. Then the kinematic model can be parameterised as,

$$M = \mu^m + \phi^m b^m \tag{1}$$

where the mean pose μ^m (m is a notation indicating the model parameters) and the principal subspace matrix ϕ^m are computed from 3 (poses)$\times$ 10 (subjects) training samples using PCA. The vector b^m represents the kinematic parameters that describe the pose of M along each principal direction. In our experiments, the first 8 significant modes are used, which keeps 98% of variation.

The statistical shape model of each bone is a point distribution model, built using the segmented volume of the same training subjects. The 3D structure of each bone is described by a set of approximately 1000 points on the segmented surface. Correspondence between these points across subjects was established by the minimum description length algorithm [8]. The deformable shape model is then described as,

$$Q_i = \mu_i^q + \phi_i^q b_i^q \tag{2}$$

where μ_i^q and ϕ_i^q (q is a notation indicating the shape parameters) are the mean shape and the principal subspace matrix for the i^{th} bone. b_i^q is the shape model parameter to be estimated. In order to keep the complexity within limits, only the first 3 significant modes are used which keeps 84% of variation.

Based on the point distribution model of each bone and the kinematic model, a hybrid statistical mesh model can be built by using the Crust mesh construction algorithm [9]. Figure 1 shows the poses of the first mode of the kinematic model (represented by the mean shapes of each bone) and the shapes of the first mode of the scaphoid.

4 3D-2D Image Registration

The statistical mesh model from the training data is then used to match with the fluoroscopic sequence to infer the 3D motion and bone shapes. Figure 2(a) summarises

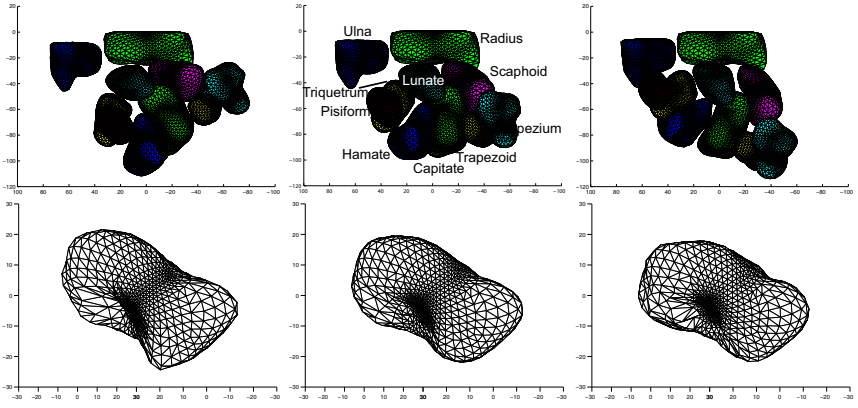

Fig. 1. Top row: The poses of the first mode of kinematic model. Bottom row: the first mode of the shape model of the scaphoid. In each case the mean +/-1.5s.d. are shown.

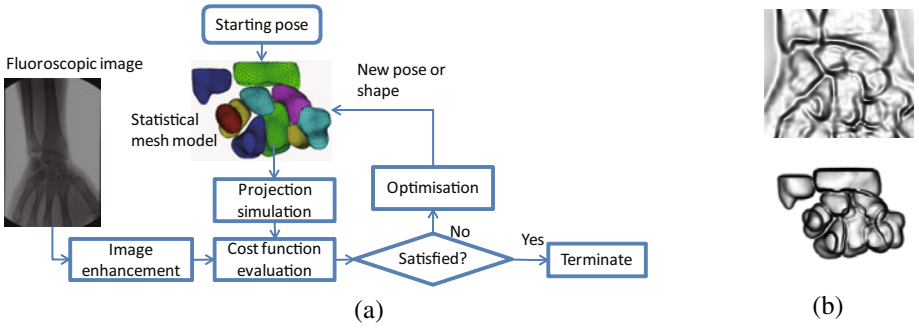

Fig. 2. Overview of the 3D-2D image registration process. (b) The gradient magnitude map of the fluoroscopic image after enhancement (cropped to show the region of interest) (top) and the simulated image from mesh model (bottom).

the registration process, in which the preprocessed fluoroscopic image is iteratively matched with a simulated projection generated from an updated pose of the mesh model. Detailed descriptions are given in the following subsections.

4.1 Fluoroscopic Image Enhancement and Projection Simulation

As the edges are strong features that can be used for image matching, the fluoroscopic image was firstly pre-processed to enhance the edges and reduce noise in homogenous regions. Local intensity normalisation was achieved by subtraction of the local mean intensity and division by the local standard deviation. The anisotropic diffusion [10] filter is then used to smooth the image while preserving the edges. Figure 2(b) shows an example of the gradient magnitude map of the fluoroscopic image after enhancement.

To optimise the pose parameters we iteratively generate projections from the mesh model with updated parameters, using the perspective projection described in section

2. The mesh model is considered to be a binary volume, and the projected intensity is negatively proportional to the sum of binary values along the ray from the source to each pixel in the image plane. Figure 2(b) shows an example.

4.2 Cost Function

To evaluate the similarity between the fluoroscopic image and the simulated image, we investigated several forms of the cost function, achieving best results from the one shown in Eqn. (3), based on the gradient along horizontal and vertical directions as well as the gradient magnitude of the two images. Additionally, the adjacent frames of the current fluoroscopic image were also taken into account in the cost function to make the estimated poses smooth across frames.

Taking C(A,B) as the Normalised Correlation Coefficient between two images A and B, we can write the cost function as:

$$E = C(Om_{k-1}, Om_k) + \sum_{p=k-1,k,k+1} w_p(C(Im_p, Dm_k) + C(Ix_p, Dx_k) + C(Iy_p, Dy_k)) \quad (3)$$

where k is the current frame number of the fluoroscopic sequence. Im_p, Ix_p and Iy_p are the gradient magnitude image, vertical gradient and horizontal gradient of the fluoroscopic image at the p^{th} frame respectively. Dm_k, Dx_k and Dy_k are the corresponding values of the simulated image. The second term calculates a cross-correlation between sets of three adjacent frames with weights w_{k-1}, w_k and w_{k+1}= 0.2, 0.6, 0.2 respectively, making the estimated pose smooth across frames. For the first term of the cost function, the vertices in the statistical mesh model are projected to the image plane, we assume the intensities at those projected points are similar across adjacent frames. Om_{k-1} and Om_k represent the gradient magnitude of the previous frame and the current frame at the projected correspondence positions. The first term makes the shape of the cost function sharper, leading to a faster and more accurate optimisation result. The $(k-1)^{th}$ frame and $(k+1)^{th}$ frame are not evaluated for the first and last frame respectively.

4.3 Optimisation

The optimisation method used is the best neighbour search combined with parabola fitting. The multi-dimensional search space (θ, M and Q) is explored by iteratively individual 1D line search. The cost function is evaluated at the current position, positive and negative neighbour positions (defined by a search range), then an optimum is found by fitting a parabola to the 3 evaluated positions. The optimum is iteratively refined by reducing the search range until convergence.

In our case, the true sizes of the bones are unknown; recovering the 3D pose from a single image is therefore a difficult, ill posed, problem. Any movement along the out-of-plane translation, could be compensated by scaling of the bone. In order to minimise this effect, the optimisation is carefully sequenced. We firstly assume that the wrist is not moving along the out-of-plane direction during radial-ulnar movement (t_y=0), as it is placed on a flat surface. The position of the model is firstly initialised by clicking the centre of the radius in the first frame of the fluoroscopic sequence. In the first step of the optimisation, only the first frame of the fluoroscopic sequence is used, and only the in-plane rigid transformation parameters (t_x, t_z, β) are estimated along with the global scale

factor (s_{global}) and the relative scale parameters of each bone (s_i). The first significant parameter of the kinematic model (b^m) is also estimated to provide an estimate of the overall pose. Other, less significant modes may include components of deviation along the out-of-plane direction that would affect the estimation of the global scale parameter. Inclusion of this first step resulted in significantly lower estimated error along the out-of-plane direction than optimisation without this step. Starting from this initial estimate of pose, the first frame is evaluated again, taking all the parameters into account (except t_y) in the following sequence: t_x, t_z, β, α, γ, b^m, s_{global}, s_i and b_i^q. After convergence, the estimated pose of the current frame is used as the starting pose for the next frame. The shape model parameters b_i^q are only estimated once in the first frame. From our initial experiments, the shape parameters are not improved significantly when we include more frames and the fitting is made significantly more complex and time consuming. At each stage, when t_x, t_z, β, α and γ are estimated, only the region immediately surrounding the radius and ulna are used for cost function evaluation, while the larger region that includes the carpal bones is used for estimating the other parameters. There are about 60-80 frames per sequence. The whole process was performed in a 3-level multi-scale framework at each frame to enhance the robustness of the registration.

5 Evaluation

The ground truth of the recovered 3D pose corresponding to real fluoroscopic sequences is almost impossible to obtain. It would require the synchronisation of 3D imaging with the fluoroscopy. Hence, we evaluated our framework based on a number of simulated fluoroscopic sequences generated from the 3D CT data. All CT volumes have been resampled to an isocubic volume with voxel dimension of 0.5 mm. We linearly interpolated a number of poses between the neutral pose and two extreme poses of radial-ulnar deviation in a full movement cycle containing 50 poses for each of 10 subjects. The ray-casting method was then used to generate a simulated fluoroscopic sequence from those interpolated 3D data. We evaluated the proposed framework in the leave-one-out manner. The 3D pose of the simulated test subject was then calculated as described in section 4, and registration error measured by the 3D Euclidian distance of each corresponding point of the mesh between the target pose and the estimated pose is presented in Table 1. The error of the registration is mainly caused by the ill posed problem (confusion between the scale and translation along the out-of-plane direction), whereas the errors along the in-plane directions are very small with average error of about 2 pixels and maximum error within 4 pixels.

It is important to mention that the relative positions of the carpal bones with respect to each other can be estimated much more accurately than the absolute positions of the individual bones. The registration error of the 3D distance between the centroid of Triquetrum and the centroid of Lunate (dTL), and the distance between the centroid of Lunate and the centroid of Scaphoid (dLS) were also measured. The errors are 1.18 ± 0.74 and 1.82 ± 0.99 pixels for dTL and dLS respectively, compared to a bone size of about 30 pixels. One of the conditions that may be assessed using this method is dissociations, where the distance between the bones is larger than normal. Scapholunate dissociation is one of the most common of these. We normalise the dLS by dividing it by the estimated global scale factor s_{global} and an average of the scale factor s_i for lunate and

Table 1. The average error, measured in 3D, between the target and estimated correspondence points of each carpal bone of 10 subjects: Triquetrum(Tri), Lunate(Lun), Scaphoid(Sca), Pisiform(Pis), Hamate(Ham), Capitate (Cap), Trapezoid (Trd) Trapezium (Trm). The measurement errors of dTL and dLS.

	eTri	eLun	eSca	ePis	eHam	eCap	eTrd	eTrm	Total	eTL	eLS
Err3D	5.4±2.6	5.1±2.5	6.5±3.6	6.8±3.7	6.5±3.8	6.6±4.0	6.5±4.6	7.6±4.3	6.3 ±3.7	1.18±0.74	1.82±0.99
ErrX	1.6±1.3	2.0±1.6	2.1±1.8	2.4±1.9	1.8±1.4	2.1±1.5	1.8±1.4	2.2±1.8	2.1 ± 1.7	/	/
ErrY	3.7±2.8	3.0±2.6	4.9±3.7	4.8±3.9	5.4±4.2	5.3±4.4	5.5±5.0	6.1±4.6	4.6±4.0	/	/
ErrZ	2.5±2.0	2.5±1.9	2.2±1.8	2.6±2.1	1.6±1.3	1.7±1.3	1.5±1.2	2.2±2.0	2.3±1.9	/	/

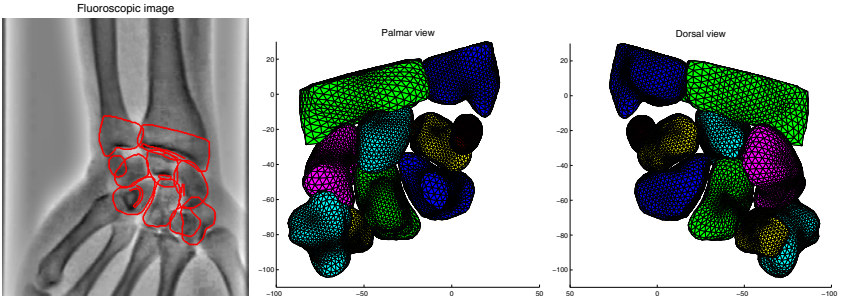

Fig. 3. Registration result of one frame from a real fluoroscopic sequence. The registration result for the whole sequence can be found in [11].

scaphoid. From the tested 10 subjects, we successfully identified the subjects suffering from scapholunate dissociation (dLS=38.78±1.53 pixels) from the normal subjects (dLS=34.49±0.83 pixels). Making this type of measurement without a 3D statistical model would be impossible.

We also tested our framework on real fluoroscopic sequences. Although the matching error cannot be quantified, the registration results show good visual correspondence and have been confirmed by a clinician. A sample frame of the matching result and the corresponding 3D pose are shown in Fig. 3 in which the projected contours from the 3D mesh model are superimposed on the preprocessed fluoroscopy image. The estimated 3D mesh model in the palmar and dorsal views are shown in middle and right respectively. The registration result for the whole sequence can be found in [11].

6 Concluding Remarks

We have presented a complete framework that is able to infer the 3D motion of carpal bones from a single view fluoroscopic sequence. It uses a hybrid statistical model to estimate both the kinematics and bone shapes from the fluoroscopic sequences allowing the motion of carpal bones during radial-ulnar deviation to be estimated. Particularly, the relative positions between carpal bones can be estimated accurately. This is potentially useful for detection of dissociation conditions, such as scapholunate dissociation,

where the underlying pathology is a rupture of one or more ligaments, and the diagnosis rests on a judgement regarding the joint separation.

In further work we will extend the current statistical model with more training data (in progress) and test the framework for the flexion-extension movement.

References

1. Snel, J.G., Venema, H.W., Moojen, T.M., Ritt, M., Grimbergen, C.A., den Heeten, G.J.: Quantitative in vivo analysis of the kinematics of carpal bones from three-dimensional CT images using a deformable surface model and a three-dimensional matching technique. Medical Physics 27, 2037–2047 (2000)
2. Sonenblum, S.E., Crisco, J.J., Kang, L., Akelman, E.: In vivo motion of the scaphotrapeziotrapezoidal (STT) joint. Journal of Biomechanics 37, 645–652 (2004)
3. van de Giessen, M., Streekstra, G.J., Strackee, S.D., Maas, M., Grimbergen, K.A., van Vliet, L.J., Vos, F.M.: Constrained Registration of the Wrist Joint. IEEE Transactions on Medical Imaging 28(12), 1861–1869 (2009)
4. Foumani, M., Strackee, S.D., Jonges, R., Blankevoort, L., Zwinderman, A.H., Carelsen, B., Streekstra, G.J.: In-vivo three-dimensional carpal bone kinematics during flexion-extension and radio-ulnar deviation of the wrist: Dynamic motion versus step-wise static wrist positions. Journal of Biomechanics 42, 2664–2671 (2009)
5. Penney, G.P., Batchelor, P.G., Hill, D.L.G., Hawkes, D.J., Weese, J.: Validation of a two- to three-dimensional registration algorithm for aligning preoperative CT images and intraoperative fluoroscopy images. Medical Physics 28, 1024–1032 (2001)
6. Zheng, G.: Statistically deformable 2D/3D registration for accurate determination of postoperative cup orientation from single standard X-ray radiograph. In: Yang, G., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (eds.) MICCAI 2009. LNCS, vol. 5761, pp. 820–827. Springer, Heidelberg (2009)
7. Chen, X., Graham, J., Hutchinson, C.E.: Integrated framework for simultaneous segmentation and registration of carpal bones. In: Accepted by the 18th ICIP, Belgium (2011)
8. Davies, R.H., Twining, C., Cootes, T.F., Taylor, C.J.: Building 3-D Statistical Shape Models by Direct Optimisation. IEEE Transactions on Medical Imaging 29(4), 961–980 (2010)
9. Amenta, N.: The Crust Algorithm for 3D Surface Reconstruction. In: Proceeding of the Fifteenth Annual Symposium on Computational Geometry (1999)
10. Black, M.J., Sapiro, G.: Edges as Outliers: Anisotropic Smoothing Using Local Image Statistics. In: Nielsen, M., Johansen, P., Fogh Olsen, O., Weickert, J. (eds.) Scale-Space 1999. LNCS, vol. 1682, pp. 259–270. Springer, Heidelberg (1999)
11. http://personalpages.manchester.ac.uk/staff/xin.chen/CarpalReg.htm

Author Index