Medical Image Analysis 16 (2012) 819–830

Contents lists available at SciVerse ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Nonlinear dimensionality reduction combining MR imaging with non-imaging information

Robin Wolz [a,*], Paul Aljabar [a], Joseph V. Hajnal [b], Jyrki Lötjönen [c], Daniel Rueckert [a], The Alzheimer's Disease Neuroimaging Initiative [1]

[a] Medical Image Analysis Group, Department of Computing, Imperial College London, 180 Queen's Gate, London SW7 2AZ, UK
[b] Division of Neuroscience and Mental Health, MRC Clinical Sciences Center, Imperial College at Hammersmith Hospital Campus, Du Cane Road, London W12 0HS, UK
[c] Knowledge Intensive Services, VTT Technical Research Centre of Finland, Tekniikankatu 1, FIN-33101 Tampere, Finland

ARTICLE INFO

Article history:
Received 11 January 2011
Received in revised form 7 December 2011
Accepted 7 December 2011
Available online 22 December 2011

Keywords:
Manifold learning
Laplacian Eigenmaps
Classification
Metadata
Alzheimer's disease

ABSTRACT

We propose a framework for the extraction of biomarkers from low-dimensional manifolds representing inter-subject brain variation. Manifold coordinates of each image capture information about structural shape and appearance and, when a phenotype exists, about the subject's clinical state. Our framework incorporates subject meta-information into the manifold learning step. Apart from gender and age, information such as genotype or a derived biomarker is often available in clinical studies and can inform the classification of a query subject. Such information, whether discrete or continuous, is used as an additional input to manifold learning, extending the Laplacian Eigenmap objective function and enriching a similarity measure derived from pairwise image similarities. The biomarkers identified with the proposed method are data-driven in contrast to a priori defined biomarkers derived from, e.g., manual or automated segmentations. They form a unified representation of both the imaging and non-imaging measurements, providing a natural use for data analysis and visualization. We test the method to classify subjects with Alzheimer's Disease (AD), mild cognitive impairment (MCI) and healthy controls enrolled in the ADNI study. Non-imaging metadata used are ApoE genotype, a risk factor associated with AD, and the CSF-concentration of $A\beta_{1-42}$, an established biomarker for AD. In addition, we use hippocampal volume as a derived imaging-biomarker to enrich the learned manifold. Our classification results compare favorably to what has been reported in a recent meta-analysis using established neuroimaging methods on the same database.

1. Introduction

Identifying imaging biomarkers for diseases such as schizophrenia or Alzheimer's Disease (AD) is the focus of many neuroimaging studies. This research is driven by the potential role of imaging biomarkers as an important adjunct to traditional biomarkers, such as psychological tests, in order to achieve more accurate and earlier diagnoses. Imaging biomarkers are used to classify subjects into different clinical categories (e.g. for differential diagnosis) or to give indicators of disease severity or progression. Many of the well-established biomarkers for dementia from magnetic resonance (MR) images are based on traditional morphometric measures, such as the volume or shape of brain structures (Fox et al., 2000; Lerch et al., 2008; Schuff et al., 2009; Chupin et al., 2009; Wolz et al., 2010b) and their change over time (Freeborough and Fox, 1997; Smith et al., 2002; Boyes et al., 2006; Leow et al., 2007; Wolz et al., 2010c). More recently, models based on machine learning techniques have been proposed which seek discriminating features over the whole brain or within a defined region of interest (Fan et al., 2007, 2008; Vemuri et al., 2008; Gerardin et al., 2009; Wolz et al., 2010a). Finding a low-dimensional representation of complex and high-dimensional data is a central problem in machine learning and pattern recognition. Many methods have been proposed to learn the underlying low-dimensional space of intrinsically low-dimensional data lying in a high-dimensional space. Linear models like principal component analysis (PCA) (Pearson, 1901; Hotelling, 1933; Jolliffe, 1986) and multi-dimensional scaling (MDS) (Torgerson, 1958; Cox and Cox, 1994) have a long history for dimensionality reduction. More recently nonlinear methods like principal curves, (Hastie and Stuetzle, 1989), ISOMAP (Tenenbaum et al., 2000), locally linear embedding (LLE) (Roweis and Saul, 2000) and Laplacian Eigenmaps (LE) (Belkin and Niyogi,

* Corresponding author. Tel.: +44 20 7594 8374; fax: +44 20 7581 8024.
E-mail address: r.wolz@imperial.ac.uk (R. Wolz).

[1] Data used in the preparation of this article were obtained from the 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. ADNI investigators include (complete listing available at http://adni.loni.ucla.edu/wp-content/uploads/how_to_apply/ADNI_Authorship_List.pdf).

1361-8415/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.media.2011.12.003

2003) have been proposed to better model highly non-nonlinear data. Van der Maaten et al. (2007) give a comprehensive overview on dimensionality reduction techniques. Using the concepts of dimensionality reduction, Aljabar et al. (2008) applied spectral analysis (von Luxburg, 2007) to pairwise label overlaps obtained from a structural segmentation to discriminate AD patients from healthy controls. Focusing on intensity-based similarities between MR brain images, Klein et al. (2010) used vectors defined by the similarities of a given query subject with a cohort of training images as features from which to learn a classifier. Computer vision applications, particularly for face recognition, also use pairwise similarities to learn a low-dimensional subspace and to classify unseen images mapped to this space (Chen et al., 2005; He et al., 2005; Zhao et al., 2007). These methods are typically linear, making it easy to transform data between image space into the learned subspace, but this linearity can limit the ability to generalize to complex datasets. Indeed, recent work suggests that the complex natural variation of brain images is best described by nonlinear models (Gerber et al., 2010; Hamm et al., 2010). We aim to learn the manifold[2] structure of brain images in healthy aging and neurodegeneration by considering both clinically labeled and unlabeled image data. Nonlinear dimensionality reduction of a set of brain images with Laplacian Eigenmaps (LE) is based on pairwise image similarities that can be evaluated either over the whole image or in a region of interest (ROI). A weighted similarity graph is built that represents neighborhood information in the image data set. With the Laplacian of the graph, a low-dimensional embedding that respects the input relations is determined. The LE objective function, which is based on edge weights in the similarity graph, places more similar images in the input space closer in the embedded space. Building on this principle, we propose a method to extend the LE objective function in order to learn a manifold not only defined by pairwise image similarities but also by some metadata available for the subjects under consideration. Such metadata in a clinical setting can be discrete (e.g. gender) or continuous (e.g. age). We propose to extend the similarity graph defined in LE by a set of additional nodes representing a number of discrete states or intervals of a continuous variable. The weights from every subject to these nodes are defined based on the subjects' metadata. This groups subjects with similar metadata closer in the manifold. The proposed approach is related to work by Costa and Hero (2005) where binary label information in partially labeled data sets is used to enforce constraints in a low-dimensional manifold representation. Optimizing the extended LE objective function, results in an embedding that incorporates metadata and pairwise image similarities at the same time. The coordinates of a particular subject in the low-dimensional space can then be regarded as encoding information about the shape and appearance of the brain as well as the state of the meta-variable and thus about clinically relevant differences across the population described by both measures. Manifold coordinates can be interpreted as biomarkers defined by pairwise similarities as well as the embedded metadata, and images with clinical labels can be used to infer information about unlabeled images in their neighborhood within the learned geometrical space. In this work we use support vector machines (SVM) to perform classification of unlabeled subjects in the low-dimensional manifold and evaluate the power of the manifold representation to predict clinical variables by fitting a multiple linear regression model of clinical data vs. manifold coordinates. The contribution of this work can be summarized as follows: we propose a method for the extraction of a unified biomarker combining imaging information with non-imaging metadata. Such a unified representation makes its use for data analysis and for visualization in a potential clinical application more powerful. The method can handle discrete and continuous metadata and offers a natural way to deal with incomplete information. Compared to classical biomarkers, the proposed method is data-driven and only requires minimal a priori information. We evaluate the proposed method on brain MR images from healthy controls, subjects with mild cognitive impairment (MCI) and AD taken from the Alzheimer's Desease Neuroimaging Initiative (ADNI)[3] (Mueller et al., 2005). We use the 420 subjects for which a measurement of the cerebrospinal fluid (CSF)-concentration of the $A\beta_{1-42}$ peptide and the subject's ApoE genotype are currently available. A decrease in the concentration of the $A\beta_{1-42}$ peptide that follows its accumulation into plaques is considered a diagnostic biomarker for AD (Trojanowski, 2004) and the ApoE genotype has been shown to be a risk factor for the disease (Lehtovirta et al., 1996). Besides non-imaging metadata, we also test the power of automatically derived hippocampal volumes, a well-known biomarker for AD (Jack et al., 1999), as meta-information. We furthermore test the ability of the proposed method to combine different metadata in a single manifold learning step.

2. Method

2.1. Image distance function

The manifold learning framework described in the following builds on the definition of pairwise image distances between two images I_i and I_j defined on a discrete domain $\Omega \in \mathbb{R}^3$. Souvenir and Pless (2007) propose an image distance function for manifold learning based on a deformation-based metric and intensity differences. The deformation-based part represents the shape difference between anatomical structures in two images and can be estimated from the deformation that is needed to non-rigidly transform one image to another. The intensity measure is based on the appearance differences remaining after transforming one image to the other. Assuming a transformation $\phi_{ij} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ that transforms image I_i to image I_j, the combined distance function is formulated as

$$\mathcal{D}_{ij} = \alpha \mathcal{D}_{ij}^{\text{def}}(I_i, I_j, \phi_{ij}) + (1-\alpha)\mathcal{D}_{ij}^{\text{int}}(I_i, I_j, \phi_{ij}) \quad (1)$$

where α defines the influence of the shape- and appearance based parts. A deformation-based metric to measure pairwise image distances can be obtained in the large deformation dipheomorphic metric mapping (LDDMM) framework (Beg et al., 2005). In recent work on manifold learning from brain MRI, this metric is estimated in a small deformation setting using a single displacement field $u : \mathbb{R}^3 \times \mathbb{R} \rightarrow \mathbb{R}^3$ to represent ϕ (Gerber et al., 2010; Hamm et al., 2010; Aljabar et al., 2011). The deformation-based term in Eq. (1) can then be defined as

$$\mathcal{D}_{ij}^{\text{def}}(I_i, I_j, \phi_{ij}) = \sum_{x \in \Omega} \|u_{ij}(x)\|_2^2 dx \quad (2)$$

where u_{ij} is the mean magnitude of the displacement vector field between images I_i and I_j. The intensity-based term in Eq. (1) is defined as the sum of squared differences between image I_i and the deformed image I_j:

$$\mathcal{D}_{ij}^{\text{int}}(I_i, I_j, \phi_{ij}) = \sum_{x \in \Omega} (I_i(x) - I_j(\phi_{ij}(x)))^2. \quad (3)$$

2.2. Manifold learning using pairwise image similarities

A set of N images $\{I_i : i = 1, \ldots, N\}$ is described by vectors of intensities $\mathbf{x}_i$ as $\mathbf{X} = \{\mathbf{x}_1, \ldots, \mathbf{x}_N\} \in \mathbb{R}^D$, where D is the number of voxels per image or region of interest (typically $D >$ 1,000,000 for

[2] In this paper, we use the terms "manifold learning/embedding" and "dimensionality reduction" interchangeably.

[3] www.loni.ucla.edu/ADNI.

brain MR images). Assuming $\mathbf{x}_1, \ldots, \mathbf{x}_N$ lie on or near an d-dimensional manifold $\mathcal{M}$ embedded in $\mathbb{R}^D$, a low dimensional representation $\mathbf{Y} = \{\mathbf{y}_1, \ldots, \mathbf{y}_N\}$ with $\mathbf{y}_i \in \mathbb{R}^d$ of the input images in $\mathcal{M}$ may be learned. We use Laplacian Eigenmaps (LE) (Belkin and Niyogi, 2003) to achieve the nonlinear dimensionality reduction

$$f : \mathbf{X} \to \mathbf{Y}, \quad \mathbf{y}_i = f(\mathbf{x}_i).$$

An undirected weighted graph $G = \langle V, E \rangle$ with N nodes V representing the images and edges E connecting neighboring nodes is defined on $\mathbf{X}$. Edge weights of E are defined based on pairwise image similarities. A k-nearest neighbor graph is defined, where the weight w_{ij} between nodes i and j is set to their image similarity if they are in a local neighborhood and set to zero otherwise:

$$w_{ij} = \begin{cases} e^{-\frac{\mathcal{D}(I_i I_j, \phi_{ij})^2}{t}} & \text{if } i \in \mathcal{N}_i \text{ or } j \in \mathcal{N}_j \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

where distances are transformed to similarities using the heat kernel as proposed in Belkin and Niyogi (2003) and $\mathcal{N}_x$ describes the k nearest neighbors of subject x. A low-dimensional representation $\mathbf{y}_i = f(\mathbf{x}_i)$ that respects the defined edge weights w_{ij} can be obtained by minimizing the energy function

$$\sum_{ij} \|\mathbf{y}_i - \mathbf{y}_j\|^2 w_{ij}. \tag{5}$$

This energy function ensures that more similar images in the input space are closer together in the embedded space. With the diagonal degree matrix $\mathbf{D} = \sum_j w_{ij}$, this can be reformulated as

$$\begin{aligned} \sum_{ij} \|\mathbf{y}_i - \mathbf{y}_j\|^2 w_{ij} &= \sum_{ij} \left(\|\mathbf{y}_i\|^2 + \|\mathbf{y}_j\|^2 - 2\mathbf{y}_i^T \mathbf{y}_j \right) w_{ij} \\ &= \sum_i \|\mathbf{y}_i\|^2 D_{ii} + \sum_j \|\mathbf{y}_j\|^2 D_{jj} - 2 \sum_{ij} \mathbf{y}_i^T \mathbf{y}_j w_{ij} \\ &= 2\mathbf{Y}^T \mathbf{L} \mathbf{Y} \end{aligned} \tag{6}$$

with the graph Laplacian $\mathbf{L} = \mathbf{D} - \mathbf{W}$. Since $\mathbf{L}$ is positive semidefinite, the minimization problem can be formulated as

$$\begin{aligned} &\operatorname*{argmin}_{\mathbf{Y}} \mathbf{Y}^T \mathbf{L} \mathbf{Y} \\ &\mathbf{Y}^T \mathbf{D} \mathbf{Y} = 1 \end{aligned} \tag{7}$$

where the constraint $\mathbf{Y}^T \mathbf{D} \mathbf{Y} = 1$ removes an arbitrary scaling factor in the embedding and prevents the trivial solution where all $\mathbf{y}$'s are set to zero (Belkin and Niyogi, 2003). Finding the $\mathbf{y}_i$'s that optimize this objective function can be formulated in closed form as finding the eigenvectors associated with the d smallest non-zero eigenvalues of the generalized eigenvalue problem

$$\mathbf{L} v = \lambda \mathbf{D} v. \tag{8}$$

2.3. Manifold learning incorporating non-imaging information

In many settings, an additional variable z_i providing further information on subject i may be available in addition to MR imaging data. Such meta-information can inform judgments such as clinical diagnosis. We propose a method to incorporate such information into the manifold learning process (Section 2.2). The hypothesis is that by using this additional information, a more accurate representation of the population can be obtained leading to a more reliable biomarker in the low-dimensional space. Metadata available in a clinical setting can be defined by discrete categories (two or more), or by a continuous variable. Examples of discrete variables are gender, blood- or genotype. Continuous variables can be, e.g., a subject's age, weight or measurements derived from a phenotype associated with the disease of interest. In (Costa and Hero, 2005), graph G describing the LE objective function in Eq. (5) is extended by two nodes, each representing one of two classes available for training data. Connecting each training subject with its respective class node imposes the class differences in the training data on the manifold structure. Assuming generalizability between labeled and unlabeled nodes, a more accurate classification performance on the test data is expected. Extending this idea, we propose to use metadata available for *all* or a *subset* of subjects in contrast to the class labels itself to enrich the low-dimensional representation. The concept of the proposed method is illustrated in Fig. 1: in classic Laplacian eigenmaps, graph G connects neighboring subjects to enforce similar images to be close in the embedding. In the proposed framework, additional edges are defined that connect every subject to an additional node representing it's metadata class. This enforces subjects in the same class to be closer in the manifold. Metadata can be incorporated from one or more measures, each defining a number of discrete classes or a continuous interval leading to a fuzzy-class membership. Graph G is extended by M nodes $\hat{V}$ representing the metadata variable z and called *support nodes* in the following. By connecting each image $\mathbf{x}_i$ to the support nodes with weights defined according to the value of z_i, the distribution of the meta-variable will influence neighborhoods in the low-dimensional representation. In the discrete setting with $z_i \in Z_{\text{discr}} = \{z^1, \ldots, z^M\}$, the weight $\hat{w}_{im}$ between subject i and the mth support node is defined by

$$\hat{w}_{im} = \begin{cases} 1 & \text{if } z_i = z^m \\ 0 & \text{otherwise} \end{cases} \tag{9}$$

For a continuous metadata variable $z_i \in Z_{\text{cont}} = [z^a, z^b]$, a set of support nodes is defined, each representing a sub-interval of the variable space z. The input space is subdivided into M subintervals $\bar{z}^m \in Z^m_{\text{cont}} = [z^{a,m}, z^{b,m}], m = 1, \ldots, M$. Subintervals are defined by percentiles of the variable interval, giving equal probabilities to each $\bar{z}^m$. Their bounds can then be defined as

$$\begin{aligned} z^{a,m} &= P_z\left((m-1)\frac{100}{M} \right) \\ z^{b,m} &= P_z\left(m\frac{100}{M} \right) \end{aligned} \tag{10}$$

where $P_z(x)$ gives the xth percentile of interval z. With the mean value $\mu^m = \frac{1}{|\bar{z}^m|} \sum_{z \in \bar{z}^m} z$ of interval $\bar{z}^m$, the continuous weight $\hat{w}_{im}$ between subject i and the mth support noe is defined based on the distance between z_i and μ^m, grouping subject i closer to subjects with a similar value of z:

$$\hat{w}_{im} = \begin{cases} \frac{1}{c}(1 + (z_i - \mu^m)^2)^{-1} & , \text{ if } z_i \text{ is available} \\ 0 & , \text{ otherwise} \end{cases} \tag{11}$$

where c is a normalizing constant to ensure $\sum_m \hat{w}_{im} = 1$. The weighting schemes in the discrete and continuous settings for the case where an additional variable z is available for all images are illustrated in Fig. 2. Incorporating the M support nodes $\hat{V}$ and the weights $\hat{w}_{im}$, leads to an extended Laplacian Eigenmaps (E-LE) objective function

$$\gamma \sum_{ij} \|\mathbf{y}_i - \mathbf{y}_j\|^2 w_{ij} + \sum_{im} \|\mathbf{y}_i - \hat{\mathbf{y}}_m\|^2 \hat{w}_{im} \tag{12}$$

with $\hat{\mathbf{y}}_m$ representing the cluster center of state z^m of a discrete variable or of the interval $\bar{z}^m$ of a continuous meta-variable. As in classic Laplacian Eigenmaps, the first term maps subjects with a high similarity w_{ij} close in the embedding space. In addition, subjects are mapped close to the cluster centers of the groups of subjects with similar metadata variable as represented by the weights $\hat{w}_{im}$. The extended low-dimensional embedding space is described by

$$\mathbf{Y}' = \{\hat{\mathbf{y}}_1, \ldots, \hat{\mathbf{y}}_M, \mathbf{y}_1, \ldots, \mathbf{y}_N\}, \quad \hat{\mathbf{y}}_m, \mathbf{y}_i \in \mathbb{R}^d. \tag{13}$$

In this embedding, the proximity of subject i to the mth group (discrete or continuous) and its centroid $\hat{\mathbf{y}}_m$ is determined by the weights $\hat{w}_{im}$ defined by the metadata as well as image-based

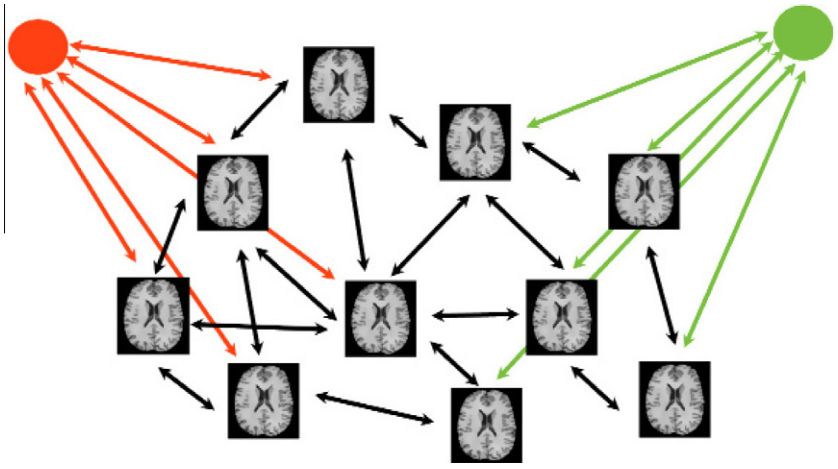

Fig. 1. Concept of the proposed extension to Laplacian Eigenmaps (LE): in LE, a set of images is connected to it's closest neighbors, enforcing similar images to be close in the embedding (black edges). The proposed extension introduces additional edges in this similarity graph, connecting every subject to additional nodes representing metadata classes (red and green) according to the subject's class membership. This enforces subjects in the same metadata class to be closer in the embedding space. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

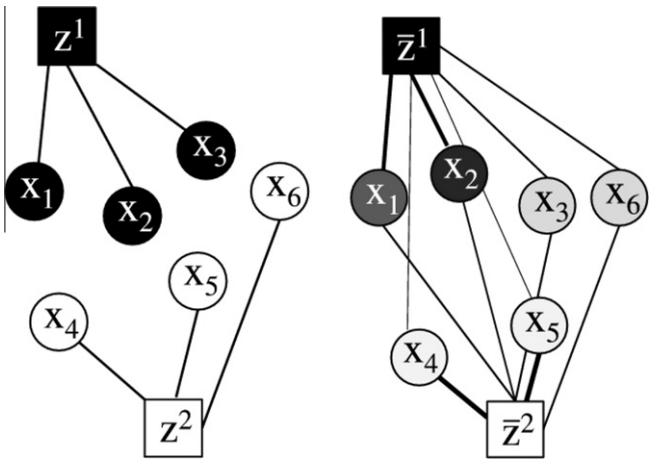

Fig. 2. Weights defined between image nodes $\mathbf{x}_i$ and support nodes representing metadata Z. In the discrete setting (left), equally weighted edges are defined according to Eq. (9). In the continuous setting (right), weights to both additional nodes are defined according to Eqs. (10) and (11). A higher weight is illustrated by a thicker edge.

weights w_{ij}. A low weight of parameter γ arranges the subjects mainly according to the metadata, whereas a high value of γ is closer to the standard embedding with Laplacian Eigenmaps which considers only pairwise image similarities. The influence, γ has on the embedding is illustrated in Fig. 3. In a synthetic example pairwise similarities for 16 nodes are defined from a set of distances between points in 2D to arrange them in a grid-shaped embedding when using standard Laplacian Eigenmaps. Every node is associated with a randomly assigned meta-variable varying between zero and one which is encoded by the color in Fig. 3. In panel (a) with $\gamma = 1$, the embedding is dominated by the value of the meta-variable. Panel (b) shows an embedding influenced by both measures and panel (c) shows an embedding close to the one obtained with LE for $\gamma = 50$. With the $N \times M$ matrix $\widehat{\mathbf{W}}$ defining the weights between subject i and the M support nodes, an extended weight matrix

$$\mathbf{W}' = \begin{pmatrix} \mathbf{I} & \frac{1}{2}\widehat{\mathbf{W}}^T \\ \frac{1}{2}\widehat{\mathbf{W}} & \gamma\mathbf{W} \end{pmatrix} \tag{14}$$

is derived, where $\mathbf{I}$ is an $M \times M$ identity matrix. Following Eqs. (6) and (8) to solve the generalized eigenvector problem associated with the extended weight matrix, yields the embedding coordinates which optimize the objective function in Eq. (12).

2.4. Extraction of biomarkers

Assuming the pairwise image similarities and the metadata variable represent clinically relevant differences between clinical groups of interest, a subject's manifold coordinates $\mathbf{y}_i$ can be used as a biomarker to support inferences about their clinical state.

2.4.1. Classification

When aiming at classifying unlabeled subjects for which no clinical label is available, information from labeled subjects can be used to make a decision. Please note that "unlabeled" in this context refers to the clinical label (e.g. AD, healthy control) that is to be predicted. Every subject (labeled or unlabeled) may or may not have metadata associated with it that can be used to enrich the manifold learning as described in Eqs. (9) and (11). When dealing with a two-class problem, the coordinates of N' labeled training images $\{\mathbf{y}_j,\ d_j\},\ j = 1,\ldots,N' < N,\ \mathbf{y_j} \in \mathcal{R}^d$ with clinical labels $d_j \in \{-1, 1\}$ is used to train a classifier on the derived manifold coordinates $\mathbf{y}_j = y_{j1},\ldots,y_{jd}$. Support Vector Machines (SVMs) minimize a Lagrangian energy function which leads to the hyperplane

$$\mathbf{a} \cdot \mathbf{y} - \mathbf{b} = \mathbf{0} \tag{15}$$

in the manifold space that best separates the two subject groups (Cortes and Vapnik, 1995). The location of embedding coordinates of the $N - N'$ unlabeled images in relation to this plane may then be used to classify them.

2.4.2. Regression

A continuous assignment can be achieved by, e.g., building a linear regression model between a clinical variable $\hat{d}_j$ vs. manifold coordinates $y_{j1},\ldots,y_{jd}$:

$$\hat{d}_j = a_0 + \sum_{i=1}^{d} a_i y_{ji} \tag{16}$$

Learning such a model from a subset of subjects for which clinical labels exist, allows its application to unlabeled subjects and predictions to be made about clinical information associated with those subjects.

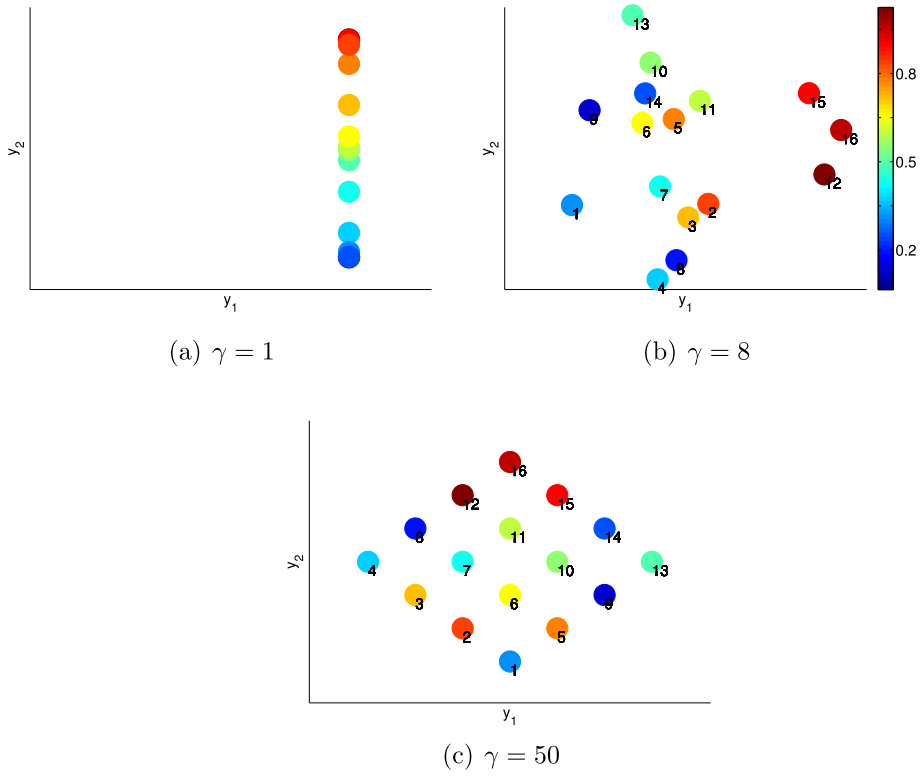

Fig. 3. First two embedding coordinates with varying influence of γ. A high weight leads to an embedding similar to the one obtained with classic Laplacian Eigenmaps (c). A very low weight embeds the images mainly based on metadata (a). A more balanced weighting results in an embedding influenced by metadata and image similarities (b).

3. Data and results

The proposed method was evaluated by performing classification in a large study on Alzheimer's Disease (AD) and mild cognitive impairment (MCI). The results are presented as follows: Section 3.1 describes the image data used. The derived image similarity used for manifold learning is described in Section 3.2. An overview on the different experiments to incoroporate metadata with the proposed method is given in Section 3.3. Sections 3.4 and 3.5 present classification and regression results. Section 3.6 compares the proposed method to other approaches for combining imaging and non-imaging information in classification tasks.

3.1. Subjects

Images used to evaluate the proposed method were obtained from the ADNI database (Mueller et al., 2005). 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 MCI, and 200 subjects with early AD. A more detailed description of the ADNI study is given in Appendix B. ADNI provides the ApoE genotype (determined by the ApoE alleles carried) for all subjects. Humans carry two out of three possible ApoE alleles ($\varepsilon 2$, $\varepsilon 3$, $\varepsilon 4$). Carriers of $\varepsilon 4$ have been shown to have a higher risk of developing AD, while $\varepsilon 2$ carriers have a lower risk (Lehtovirta et al., 1996). In addition an $A\beta_{1-42}$ protein analysis of cerebrospinal fluid (CSF) is available for a subset of ADNI subjects. A decrease in the concentration of this protein has been shown to be associated with a development of AD (Trojanowski, 2004). In this work, we used the 1.5 T T1-weighted baseline images of the 420 subjects for which a CSF analysis was available. Out of 201 MCI subjects, 89 were progressive, i.e. were diagnosed as converting to AD as of October 2010. We therefore independently analyzed stable (S-MCI) and progressive (P-MCI) subjects.[4] Table 1 presents an overview of the subjects studied and their metadata as well as their MMSE scores used for clinical diagnosis. In addition average hippocampal volumes (right + left) for the different subject groups are displayed in the very right column. This information is used as a derived imaging biomarker to enrich manifold learning.

3.2. Pairwise image similarities

The image distance function described in Eq. (1) is used to define pairwise image similarities as defined in Eq. (4). Pairiwse non-rigid registrations were performed using a deformation model based on free-form deformations (FFDs) (Rueckert et al., 1999) in a multi-resolution fashion with B-spline control point spacings of 20 mm, 10 mm, 5 mm and 2.5 mm. The first three resolution levels where applied to all registrations between the 420 study subjects. Based on the resulting deformation fields, the 50 closest neighbors to every subject were identified using the distance measure in Eq. (2). The registrations to these identified subjects were refined using the final control point spacing of 2.5 mm. Based on the resulting deformation fields and the transformed images, pairwise image similarities to define the objective function in Eq. (5) are obtained using Eqs. (1) and (4). To allow an accurate extraction of the appearance-based distance measure in Eq. (3), intensities of transformed images are normalized to the target image using a linear regression model. The parameter t in Eq. (4) used to transform image distances to similarities is set to the median of the measure data so that it corresponds with the densest part of the data. The medial temporal lobe and especially the hippocampus and

[4] Note that since the ADNI study is still ongoing it is likely that some subjects will convert from the S-MCI group to the P-MCI group in the future.

Table 1
Subject data of the study subjects are shown for the different groups. Non-imaging metadata in the form of ApoE genotype and $A\beta_{1-42}$ concentration as well as the derived imaging metadata, hippocampus volume, are presented. Carriers of the ApoE $\varepsilon 2/\varepsilon 4$ alleles are shown. The remaining subjects only carry the $\varepsilon 3$ allele. There is no significant difference in age between the clinical groups with an average age of 74.95 ± 7.03 years.

	Subject data		Non-imaging metadata		Derived metadata
	N (F)	MMSE	$\varepsilon 2/\varepsilon 4$	$A\beta_{1-42}$ (pg/ml)	Hippo. vol. (cm^3)
CN	116 (56)	29.12 ± 1.02	16/28	202.3 ± 57.5	4.53 ± 0.55
S-MCI	112 (36)	27.16 ± 1.75	9/49	178.9 ± 61.6	4.26 ± 0.59
P-MCI	89 (33)	26.64 ± 1.8	1/52	146.3 ± 46.30	3.93 ± 0.65
AD	103 (43)	23.55 ± 1.87	4/63	147.5 ± 45.8	3.92 ± 0.73

amygdala have been shown to be predominantly affected by onset and progression of MCI and AD (Dubois et al., 2007). We therefore restricted the evaluation of pairwise similarities to a region defined around both structures (see Fig. 4). This ROI is based on the definitions of hippocampus and amygdala in the anatomical atlas described in Hammers et al. (2003). Manually defined label maps on 30 subjects were affinely aligned with the MNI152-brain T1 atlas (Mazziotta et al., 1995). A fused and dilated segmentation was then transformed affinely to every target image to serve as an ROI.

3.3. Experiments

Both discrete and continuous metadata were incorporated into the manifold learning process to compare the resulting biomarker to the one achieved with Laplacian Eigenmaps based on image similarities only. ApoE genotype, which has been shown to be a risk factor for AD (Lehtovirta et al., 1996), and the CSF-concentration of $A\beta_{1-42}$, an established biomarker for AD (Trojanowski, 2004), were used as non-imaging information. In addition, we used automatically determined hippocampal volumes (Wolz et al., 2010b) as a derived imaging biomarker to enrich the manifold learning process. Furthermore, we evaluated the impact of adding support nodes to more than just one meta-variable. In particular, we combined CSF with hippocampal volume and CSF with hippocampal volume and ApoE genotype. The following list gives an overview of the different experiments performed:

- I : Laplacian Eigenmaps (LE)
- II : Extended LE (E-LE) with ApoE genotype
- III : E-LE with $A\beta_{1-42}$
- IV : E-LE with hippocampal volume
- V : E-LE with $A\beta_{1-42}$ and hippocampal volume
- VI : E-LE with $A\beta_{1-42}$, hippocampal volume and ApoE genotype

The model parameters α (influence of shape- and appearance based parts on image distance), k (no. of neighbors in the similarity graph), d (dimension of manifold) and γ (influence of metadata) are set globally for all experiments. Their value is determined based on 418 ADNI baseline images that are not used in the evaluation because no CSF measurement is available. A detailed description of these settings is given in Appendix A. In order to correct for age as a confounding variable, a linear regression model was first estimated to predict the manifold coordinates from the age of the subjects. The residual after this prediction was then used for extracting biomarkers. Fig. 5 shows exemplars of the projected embedding onto the first two coordinate directions when using standard Laplacian Eigenmaps (top panel) and the proposed method with hippocampal volume as metadata (bottom panel). A separating hyperplane between AD and control subjects as defined by SVM is displayed in both cases. Better discrimination between the two groups can be observed when using the proposed method especially for subjects close to the separating plane.

3.4. Classification

The manifold representations, obtained from standard LE embedding and from the five experiments using the extended version (E-LE) proposed in this work, are used to perform classification between the different clinical subject groups. For each relevant pairing (AD vs. CN, S-MCI vs. P-MCI, CN vs. P-MCI), a leave-25%-out cross-validation was performed using the image set described in Section 3.1. Average classification rates after 1,000 runs are determined for every dimension $d \in \{5, 6, 7\}$ (see Appendix A). To arrive at a more robust and generalizable result, we report the average classification rates over these three dimensions. Table 2 presents the correct classification rates for all six experiments. For each experiment, the multiple runs provide a distribution estimate for the corresponding classifica-

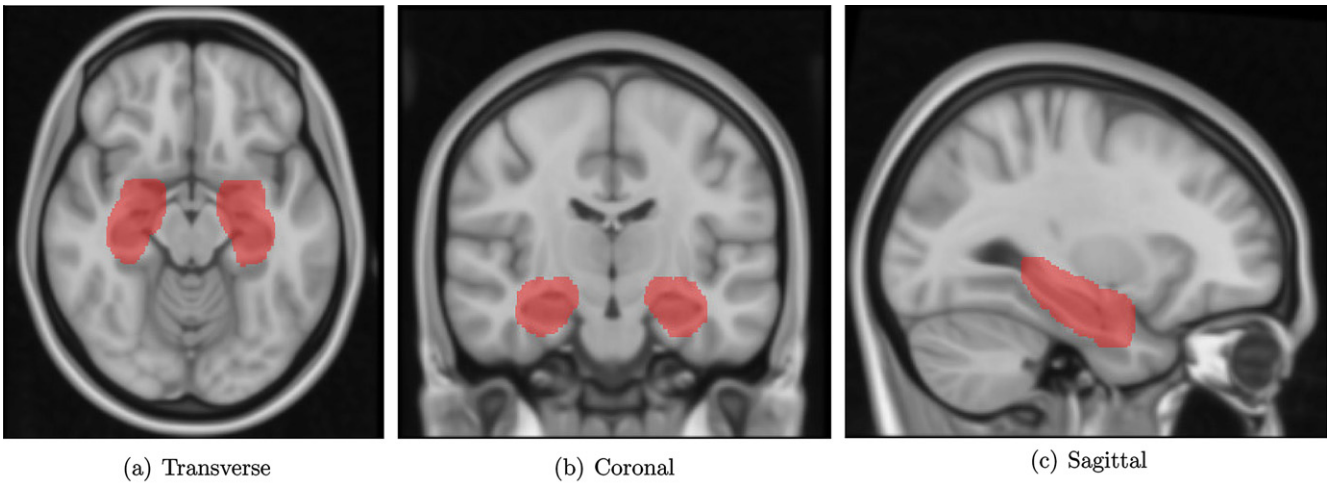

Fig. 4. Orthogonal views of MNI152 space showing the ROI around hippocampus and amygdala used to evaluate pairwise image similarities.

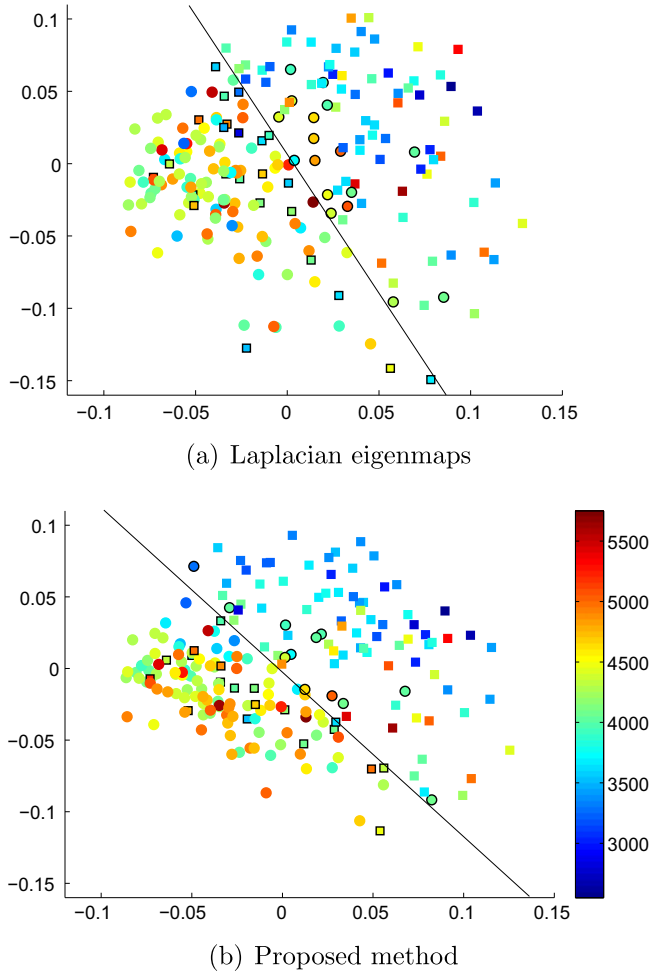

(a) Laplacian eigenmaps

(b) Proposed method

Fig. 5. Standard embedding using Laplacian Eigenmaps based on pairwise image similarities only (panel a). Extended embedding using the proposed method with hippocampal volume as metadata (panel b, $\gamma = 10$). 103 AD patients are represented by squares, 116 healthy controls by circles. Hippocampal volume (cm^3) is encoded in the marker color. A SVM separating hyperplane in 2 dimensions is displayed. Misclassified subjects with both approaches are highlighted by a black outline (41 with LE, 31 with E-LE). An improved separating ability can be observed in the extended embedding especially for subjects close to the separating plane in the original embedding. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

tion rate.[5] For each pair of clinical groups, we used these distributions to carry out unpaired *t*-tests between the results of methods I (LE) and IV (E-LE with ApoE, hippo. vol. and $A\beta_{1-42}$) with the respectively remaining methods in order to estimate the significance of any performance improvements when incorporating metadata. Resulting *p*-values are presented in Table 2. For comparison, correct classification rates when only using the different sources of meta-information are presented in the bottom part of Table 2.

3.5. Regression

The Mini-Mental State Examination is a psychological test to screen for cognitive impairment. To test the ability of manifold coordinates to predict clinical variables, we fit a multiple linear regression model of MMSE vs. manifold coordinates. We use a model that predicts MMSE from the first d manifold coordinates:

$$\mathrm{MMSE} = a_0 + \sum_{i=1}^{d} a_i y_i. \tag{17}$$

We evaluated the model for $d = 1$ to allow comparison with work presented in Gerber et al. (2010) on regressing MMSE vs. manifold coordinates in ADNI. Residuals and correlation of the regressed values for the manifolds defined in experiments I–VI, are displayed in Table 3. It can be observed that adding metadata leads to lower residuals and an improved correlation.

3.6. Alternative approaches to incorporate metadata

There are several alternative approaches to perform classification based on multiple measurements or to learn a manifold based on more than one similarity measure. In this Section we consider two obvious choices of alternatives:

3.6.1. Concatenation of feature vectors in a SVM-based classification

When performing SVM-based classification, an extended feature vector $\mathbf{f}_i = \{\mathbf{y}_i, z_i\}$ concatenating the manifold coordinates of subject i, $\mathbf{y}_i$ with its meta-variable z_i can be defined. Classification can then be performed in the resulting $d+1$ dimensional space. Table 4 shows classification results using this approach for ApoE genotype, $A\beta_{1-42}$ concentration and hippocampal volume. In addition, *p*-values for the differences between these results and the relevant results in Table 2 are presented. While in most comparisons, the proposed method achieves the same accuracy as the concatenation approach, an improved classification accuracy can be observed for the identification of P-MCI subjects. Adding ApoE genotype to the manifold coordinates obtained with LE leads to a better accuracy than the proposed scheme with a unified biomarker.

3.6.2. Learning a low-dimensional manifold from a combined similarity measure

The similarity between two instances $z_i,\ z_j \in Z_c$ of a continuous variable can be defined as $\hat{s}_{ij} = \frac{\mathrm{abs}(z_i - z_j)}{\max(Z_c)}$. With the intensity-based similarity s_{ij}, the edge weight used for manifold learning with LE (Eqs. (5)–(8)) can then be defined using a combined similarity measure:

$$w_{ij} = \begin{cases} s_{ij} + \beta \hat{s}_{ij}, & \text{if } i \in \mathcal{N}_i \text{ or } j \in \mathcal{N}_j \\ 0, & \text{else.} \end{cases} \tag{18}$$

where $\mathcal{N}_x$ describes the n nearest neighbors to subject x and β defines the relative influence of the two similarity measures. We evaluated the classification performance when using this approach to define an LE embedding for $\hat{s}_{ij}$ defined by hippocampal volume and $A\beta_{1-42}$ concentration alternatively. Varying $\beta \in [0, 7]$ and applying SVM-based classification on the resulting manifold coordinates with the procedure described in Section 3.4, results in the classification rates displayed in Fig. 6.

3.7. Summary

The preceding sections describe different experiments performed to evaluate the proposed method. Several metadata variables were used in combination with pairwise image similarities to learn low-dimensional manifold representations of a set of 420 images from the ADNI database. The metadata used include ApoE genotype, the concentration of the $A\beta_{1-42}$ peptide, hippocampal volume, and the combination of all measures. Section 3.4 presents results from using manifold-representations to classify individual subjects according to clinical subject groups. A significantly improved classification accuracy can be observed when incorporating metadata. In Section 3.5 MMSE data is regressed

[5] All estimated distributions passed a normality test using a Kolmogorov–Smirnov test at $\alpha = 0.05$.

Table 2
Correct classification accuracy (ACC), sensitivity (SEN) and specificity (SPE) (%) for classic Laplacian Eigenmaps (LE, I) and the extended version E-LE incorporating different types of meta-information (II–VI). *P*-values for the difference between methods I–VI and method I (p_1) and method VI (p_2) are presented. † stands for $p < 0.001$ The results for method V are significantly different from all other results with $p < 0.001$ – apart from method VI. The bottom rows of the table present classification rates when using different types of metadata only.

	AD vs. CN			P-MCI vs. S-MCI			P-MCI vs. CN			p_1	p_2
	ACC	SEN	SPE	ACC	SEN	SPE	ACC	SEN	SPE		
I: LE	86	82	89	63	55	69	82	73	89	n.a.	†/†/†
II: E-LE: ApoE	83	80	86	69	64	73	81	76	84	†/†/†	†/†/†
III: E-LE: Aβ_{1-42}	87	84	89	68	65	70	84	81	87	†/†/†	†/†/†
IV: E-LE: hippo. vol.	86	82	89	66	60	71	83	77	87	0.94/†/†	†/†/†
V: E-LE: Aβ_{1-42}, hippo. vol.	88	85	90	67	64	70	87	87	88	†/†/†	0.006/†/0.63
VI: E-LE: Aβ_{1-42}, hippo. vol., ApoE	88	85	91	69	68	70	87	87	88	†/†/†	n.a.
ApoE only	67	59	75	57	52	64	68	63	75	†/†/†	†/†/†
Aβ_{1-42} only	75	64	85	64	51	81	73	63	84	†/†/†	†/†/†
Hippo. vol. only	74	74	74	61	58	62	72	73	70	†/†/†	†/†/†

Table 3
Residuals and R^2 statistics obtained from regressing the results of the MMSE vs. the first manifold coordinate using a multiple linear model. An improvement of statistics can be observed when incorporating metadata into the manifold learning process.

	I	II	III	IV	V	VI
Residual	1.86	1.85	1.85	1.85	1.85	1.82
R^2	0.20	0.22	0.22	0.22	0.21	0.24

from the manifold coordinates, observing improved regression statistics when incorporating metadata. Section 3.6 compares the proposed data combination strategy to two approaches that combine manifold learning with metadata: (1) an approach that defines a unified similarity measure from two separate input measures (imaging and metadata) and (2) a method that concatenates a feature vector obtained from manifold learning with some metadata before performing SVM-based classification. The proposed method compares favorably to both other combination strategies.

4. Discussion

We have presented a method to extract biomarkers from MR brain images, combining imaging information with non-imaging metadata. These biomarkers are defined in a low-dimensional manifold space that is learned from image-based similarities and non-imaging metadata. Laplacian Eigenmaps were used to derive

Table 4
Classification accuracy (ACC), sensitivity (SEN) and specificity (SPE) when incorporating metadata for classification into the SVM featurevector. *P*-values for differences with the according methods (III, IV) in Table 2 are presented. †: $p < 0.001$.

	AD vs. CN				P-MCI vs. S-MCI				P-MCI vs. CN			
	ACC	SEN	SPE	*p*	ACC	SEN	SPE	*p*	ACC	SEN	SPE	*p*
ApoE	84	80	88	0.02	64	57	71	†	82	76	87	†
Aβ_{1-42}	87	85	89	0.09	65	60	70	†	84	77	89	0.03
Hippo. vol.	86	81	89	0.76	63.7	52.3	71.9	†	83	76	88	0.43

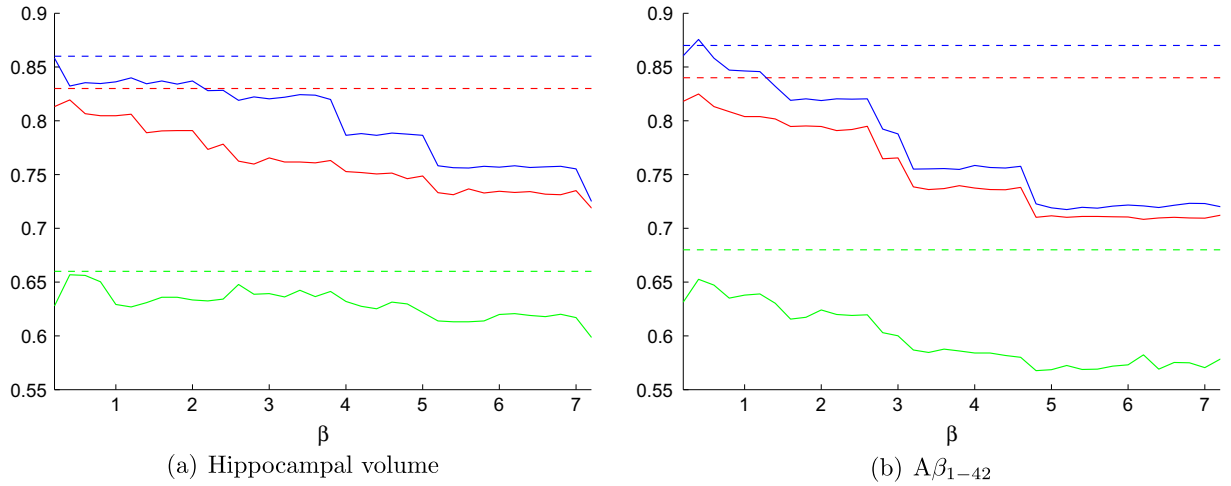

(a) Hippocampal volume (b) Aβ_{1-42}

Fig. 6. Classification accuracy obtained from defining a combined similarity measure incorporating both imaging and non-imaging information before performing manifold learning. Results with hippocampal volume and Aβ_{1-42} are presented over an increasing influence of the metadata. Performance is evaluated for the classification tasks AD vs. CN (blue), P-MCI vs. CN (red) and P-MCI vs. S-MCI (green). The dotted lines indicate the classification accuracy obtained with the proposed method and the respective metadata. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

a nonlinear and low-dimensional representation of a set of images. The graph defined by pairwise similarities and used to define the Laplacian Eigenmaps objective function is extended by *support nodes* representing metadata. Weights defined from all image nodes to all support nodes incorporate the metadata into an extended objective function. Optimizing this target function leads to an embedding that is expressed by both pairwise image similarities and the similarity represented by the metadata. While we used Laplacian Eigenmaps to learn a low-dimensional representation, other manifold learning techniques such as LLE or ISOMAP could be used assuming appropriately defined pairwise image distances are defined. In fact, Yan et al. (2007) showed that most spectral dimensionality reduction techniques can be formulated in one general graph embedding framework. We have validated the proposed method on a large and diverse clinical dataset (ADNI) applying linear SVMs to the learned manifold coordinates to discriminate between different subject groups. Other classification methods such as linear discriminant analysis (LDA) (Krzanowski, 1988) or kernel SVMs could be applied. While a nonlinear method might further improve classification accuracy, it requires an additional parameter or parameters to be optimized. Our results show that the proposed method is able to produce a classification accuracy between clinical groups with an accuracy that compares favorably to established and well-cited methods in neuroimaging. Cuingnet et al. (2011) recently presented the comparison of ten different methods for classification on a subset of ADNI similar to that used in our study. These methods comprise five high dimensional voxel-based approaches, three methods based on cortical thickness and two methods based on the hippocampus. Using only imaging similarities, the proposed manifold-based method outperforms the majority of the ten methods in individual classification experiments. This is witnessed by a much more balanced performance of the proposed method resulting in a significantly improved sensitivity together with an only slightly lower specificity. This effect is particularly striking for the classification between stable and progressive MCI. Combining imaging and non-imaging information in one manifold, yields substantial and significant improvements in classification accuracy compared to the results based on image similarities or metadata alone. The individual use of ApoE genotype, the concentration of $A\beta_{1-42}$ and hippocampal volume, improves classification rates. Using all metadata in one step, further improves results to 88% for AD vs. CN, 69% for P-MCI vs. S-MCI and 87% for P-MCI vs. CN. These results highlight the potential role of such metadata as suitably complementary information to MR image data in future studies and show the ability of the proposed method to combine both measures into one biomarker. Several studies have proposed the combined use of MR imaging with metadata for AD classification. A combined classifier based on hippocampal volume extracted from MRI and the CSF proteins *t*-tau and $A\beta_{1-42}$ has been described by Eckerstrom et al. (2010). The presented results show that on a small dataset of 68 subjects, a significant improvement in identifying progressive MCI subjects can be achieved by combining both measures. Kohannim et al. (2010) show for the classification between AD and CN on ADNI an improvement from 82% to 88% when adding CSF-derived measures to several MRI-based biomarkers. The results based on a similar subset of ADNI and presented in this paper show that combining several measures into one combined biomarker with the proposed method produces comparable classifiacation rates than combining several measures at the classification stage. This can also be observed from the results presented in Section 3.6.1. In addition to classification performance, we also evaluated the ability of the learned manifold to predict clinical variables. Learning a multiple linear regression model of MMSE vs. manifold coordinates, leads to significantly improved results compared to what has been published using similar data. Gerber et al. (2010) report $R^2 = 0.05$ and a residual of 2.37 when regressing MMSE vs. the first manifold coordinate. Incorporating metadata led to further improved regression statistics in this sample dataset. We discussed two alternative approaches to incorporate non-imaging information into a manifold classification setting. The proposed method shows better classification accuracy compared with a method in which image similarities and non-imaging similarities are combined before performing manifold learning. However, tuning the weighting factor between the concatenated similarities on the test images, did lead to results comparable with the proposed method for the sub-comparisons of AD vs. CN and P-MCI vs. CN. Compared to an approach where manifold features are combined with a meta-variable before performing SVM-based classification, the proposed method gave slightly superior performance. The strength of the proposed method, however, lies in the unified representation of information taken from different measurements. This enables not only classification but can also help in visualizing the determined biomarker in a clinical environment. Plots of the form shown in Fig. 5 can potentially enhance interpretation of computer-aided diagnosis (CAD) systems, such as the one developed in PredictAD (www.predictad.eu). A clinician can locate the patient studied relative to all other database cases providing information about the severity of the disease not only the on/off-classification result. Furthermore, the capability to define a single continuous biomarker facilitates the definition of regression models such as the one presented in Section 3.5. We evaluated the influence of the number of embedding dimensions d on training data. Robust results where achieved for $d \in \{5, 6, 7\}$. Assuming normalized weights defined on the metadata and a normalized pairwise similarity measure, the weighting factor γ that dictates the influence of metadata on the manifold coordinates, can be set globally. While individually tuning γ for every type of metadata is expected to lead to better results, we determined a weighting based on training data and used the same parameter for all experiments performed to work with a more realistic setting. Many state-of-the-art methods for the extraction of biomarkers for AD from MR images are computationally expensive (run-time of hours to days) or require complex a priori information (e.g. manual segmentation in atlas-based methods) (Cuingnet et al., 2011). Here, we have proposed a fast and robust alternative to classify subjects that is generic and data-driven. Extending on our previous work (Wolz et al., 2010a), we plan to further explore the joint estimation of intra- and inter-subject variation in a manifold learning setting. While the acquisition of longitudinal brain scans is more expensive and may delay diagnosis, the longitudinal development of brain structures has been shown to give more powerful biomarkers, e.g., (Freeborough and Fox, 1997; Smith et al., 2002; Boyes et al., 2006). Using both longitudinal development and metadata to learn a low-dimensional representation can be expected to further improve the classification results presented in this crossectional study.

Acknowledgments

This project is partially funded under the 7th Framework Programme by the European Commission (http://cordis.europa.eu/ist/). Data collection and sharing for this project was funded by the Alzheimer's Disease Neuroimaging Initiative (ADNI) (National Institutes of Health Grant U01 AG024904). ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and Bioengineering, and through generous contributions from the following: Abbott, AstraZeneca AB, Bayer Schering Pharma AG, Bristol-Myers Squibb, Eisai Global Clinical Development, Elan Corporation, Genentech, GE Healthcare, GlaxoSmithKline, Innogenetics, Johnson and Johnson, Eli Lilly and Co., Medpace, Inc., Merck and Co., Inc., Novartis AG, Pfizer Inc., F. Hoffman-La Roche, Schering-Plough, Synarc, Inc., as well as non-profit partners

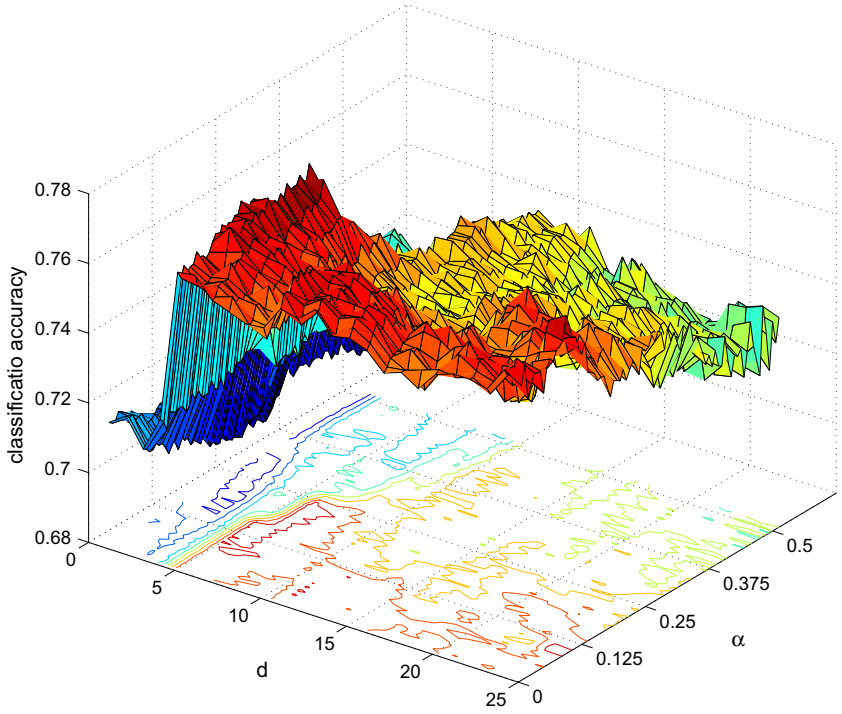

Fig. 7. Displayed is the classification accuracy when increasing the influence of the deformation-based distance α and when changing the embedding dimension d.

the Alzheimer's Association and Alzheimer's Drug Discovery Foundation, with participation from the US Food and Drug Administration. Private sector contributions to ADNI are facilitated by the Foundation for the National Institutes of Health (www.fnih.org). The grantee organization is the Northern California Institute for Research and Education, and the study is coordinated by the Alzheimer's Disease Cooperative Study at the University of California, San Diego. ADNI data are disseminated by the Laboratory for Neuro Imaging at the University of California, Los Angeles. This research was also supported by NIH Grants P30 AG010129, K01 AG030514, and the Dana Foundation.

Appendix A. Parameter settings

This section gives a description of the parameter sets used to evaluate the proposed method.

A.1. Neighborhood size k

In our experiments, the neighborhood size k used to define graph G did not substantially influence classification accuracy when varying between 10 and 50. We set it to $k = 20$ in all experiments following results presented in (Gerber et al., 2010).

A.2. Weighting of distance measures (α) and dimension d

The optimal number of embedding coordinates d is expected to depend on the defined influence of the shape- and appearance measures on the pairwise similarity used. Furthermore different numbers of dimensions have been shown to produce good results for different applications in manifold learning of brain images (Gerber et al., 2010; Wolz et al., 2010b; Jia et al., 2010; Hamm et al., 2010; Wolz et al., 2010a). To get an overview of how the proposed classification framework reacts to varying the influence of both distance measures and dimension of the manifold, we performed classification between clinical groups on the 418 ADNI baseline images not used for testing the classifier (images for which no CSF information is available). We evaluated the classification accuracy for the pairings AD vs. CN, P-MCI vs. S-MCI and P-MCI vs. CN when varying the dimension $d \in [1, \ldots, 25]$ and $\alpha \in [0, \ldots, 0.5]$ as displayed in Fig. 7. When increasing the embedding dimension, a steep increase in classification accuracy can be observed from $d = 5$. Also, the accuracy increases with an increased influence of the deformation-based distance measure before falling down again. Relatively stable classification results are achieved when varying $d \in [5, \ldots, 10]$ and $\alpha \in [0, \ldots 0.25]$ with a maximum around $d = 5$, $\alpha = 0.12$). To arrive at a robust parameter set, we set $\alpha = 0.125$ and evaluated the embedding for $d \in \{5, 6, 7\}$. The reported classification rates are averaged over these dimensions.

A.3. Influence of metadata (γ)

The weighting factor γ defined in Eqs. (12) and (14) determines how much the final embedding is influenced by image similarities and metadata. To evaluate the influence it has on classification accuracy, we evaluated the performance on the images for which no CSF measurement is available ($N = 418$) when using hippocampal volume, which is available for all subjects, as metadata. Fig. 8 shows classification results with $\alpha = 0.125$ and averaged over dimensions $d \in \{5, 6, 7\}$ plotted over varying γ. As γ is increased, initially improved results finally asymptote to the results obtained with standard LE (illustrated with the red line in Fig. 8). Following these results, we globally set $\gamma = 10$ for all types of metadata. Tuning γ individually for different types of metadata is expected to further improve results but requires a more complex training and makes the application to new datasets more difficult.

A.4. Number of support nodes M

For the discrete metadata in experiment II, ApoE genotype, $M = 3$ support nodes are defined, each trivially associated with a

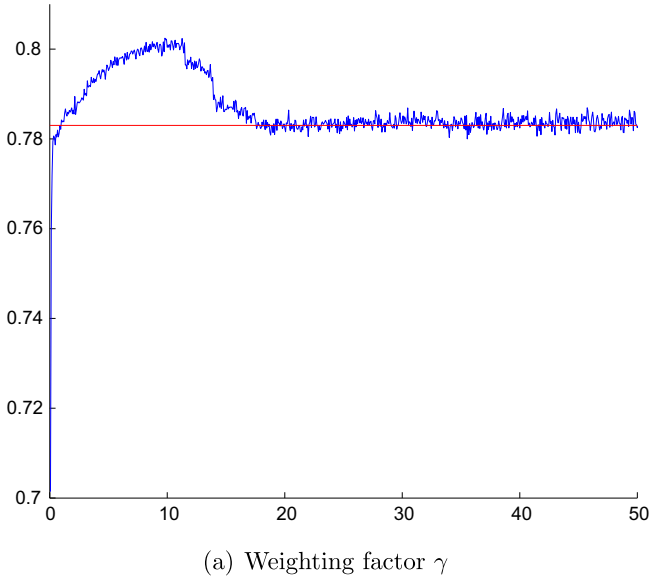

(a) Weighting factor γ

Fig. 8. Classification rate when varying γ between 1 and 50 evaluated for $d \in \{5,6,7\}$ with $\alpha = 0.125$.

possible genotype (z_1: subjects that carry at least one $\varepsilon 2$ allele. z_2: subjects that carry at least one $\varepsilon 4$ allele. z_3: subjects that only carry the $\varepsilon 3$ allele). Following Eq. (9), $\hat{w}_{im}$ is set to one if subject i has a genotype associated with node m, otherwise it is set to zero. For the continuous variables in experiments III and IV, $A\beta_{1-42}$ concentration and hippocampal volume, a continuous weighting $\hat{w}$ is defined as described by Eqs. (10) and (11). To accommodate the four clinical groups (CN, S-MCI, P-MCI, AD), we used $M = 4$ support nodes with subintervals $\bar{z}^m, m = 1,\ldots,4$ to describe the metadata as defined in Eq. (10). For experiments V and VI, that use more than one meta-variable, edges to the support nodes associated with all variables are defined. This results in $M = 8$ and $M = 11$ support nodes for experiments V and VI respectively with weights defined as in experiments II–IV.

Appendix B. The Alzheimer's Disease Neuroimaging Initiative

The Alzheimer's Disease Neuroimaging Initiative (ADNI) was launched in 2003 by the National Institute on Aging (NIA), the National Institute of Biomedical Imaging and Bioengineering (NIBIB), the Food and Drug Administration (FDA), private pharmaceutical companies and non-profit organizations, as a $60 million, 5-year public–private partnership. The primary goal of ADNI has been to test whether serial MRI, positron emission tomography (PET), other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of MCI and 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. The Principle Investigator of this initiative is Michael W. Weiner, M.D., VA Medical Center and University of California – San Francisco. ADNI is the result of efforts of many co-investigators from a broad range of academic institutions and private corporations, and subjects have been recruited from over 50 sites across the US and Canada. The initial goal of ADNI was to recruit 800 adults, ages 55–90, to participate in the research – approximately 200 cognitively normal older individuals 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. For up-to-date information see www.adni-info.org.

References

Aljabar, P., Rueckert, D., Crum, W., 2008. Automated morphological analysis of magnetic resonance brain imaging using spectral analysis. NeuroImage 43 (2), 225–235.

Aljabar, P., Wolz, R., Srinivasan, L., Counsell, S.J., Rutherford, M.A., Edwards, A.D., Hajnal, J.V., Rueckert, D., 2011. A combined manifold learning analysis of shape and appearance to characterize neonatal brain development. IEEE Transactions on Medical Imaging 30 (12), 2072–2086.

Beg, M.F., Miller, M.I., Trouve, A., Younes, L., 2005. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. International Journal of Computer Vision 61 (2), 139–157.

Belkin, M., Niyogi, P., 2003. Laplacian eigenmaps for dimensionality reduction and data representation. Neural Computation 15 (6), 1373–1396.

Boyes, R.G., Rueckert, D., Aljabar, P., Whitwell, J., Schott, J.M., Hill, D.L., Fox, N.C., 2006. Cerebral atrophy measurements using Jacobian integration: comparison with the boundary shift integral. NeuroImage 32 (1), 159–169.

Chen, H.T., Chang, H.W., Liu, T.L., 2005. Local discriminant embedding and its variants. In: Computer Vision and Pattern Recognition, pp. II:846–II:853.

Chupin, M., Hammers, A., Liu, R., Colliot, O., Burdett, J., Bardinet, E., Duncan, J., Garnero, L., Lemieux, L., 2009. Automatic segmentation of the hippocampus and the amygdala driven by hybrid constraints: method and validation. NeuroImage 46 (3), 749–761.

Cortes, C., Vapnik, V., 1995. Support-vector networks. Machine Learning 20 (3), 273–297.

Costa, J., Hero III, A., 2005. Classification constrained dimensionality reduction. International Conference on Acoustics, Speech and Signal Processing, 1077–1080.

Cox, T.F., Cox, M.A.A., 1994. Multidimensional Scaling. Chapman & Hall, London.

Cuingnet, R., Gerardin, E., Tessieras, J., Auzias, G., Lehricy, S., Habert, M.-O., Chupin, M., Benali, H., Colliot, O., 2011. Automatic classification of patients with Alzheimer's disease from structural MRI: a comparison of ten methods using the ADNI database. NeuroImage 56 (2), 766–781.

Dubois, B., Feldman, H.H., Jacova, C., DeKosky, S.T., Barberger-Gateau, P., Cummings, J., Delacourte, A., Galasko, D., Gauthier, S., Jicha, G., Meguro, K., O'Brien, J., Pasquier, F., Robert, P., Rossor, M., Salloway, S., Stern, Y., Visser, P.J., Scheltens, P., 2007. Research criteria for the diagnosis of Alzheimer's disease: revising the NINCDS-ADRDA criteria. The Lancet Neurology 6 (8), 734–746.

Eckerstrom, C., Andreasson, U., Olsson, E., Rolstad, S., Blennow, K., Zetterberg, H., Malmgren, H., Edman, A., Wallin, A., 2010. Combination of hippocampal volume and cerebrospinal fluid biomarkers improves predictive value in mild cognitive impairment. Dementia and Geriatric Cognitive Disorders 29 (4), 294–300.

Lerch, J.P., Pruessner, J., Zijdenbos, A.P., Louis Collins, D., Teipel, S.J., Hampel, H., Evans, A.C., 2008. Automated cortical thickness measurements from MRI can accurately separate Alzheimer's patients from normal elderly controls. Neurobiology of Aging 29 (1), 23–30.

Fan, Y., Batmanghelich, N., Clark, C.M., Davatzikos, C., 2008. Spatial patterns of brain atrophy in MCI patients, identified via high-dimensional pattern classification, predict subsequent cognitive decline. NeuroImage 39 (4), 1731–1743.

Fan, Y., Shen, D., Gur, R.C., Gur, R.E., Davatzikos, C., 2007. COMPARE: classification of morphological patterns using adaptive regional elements. IEEE Transactions on Medical Imaging 26 (1), 93–105.

Fox, N.C., Cousens, S., Scahill, R., Harvey, R.J., Rossor, M.N., 2000. Using serial registered brain magnetic resonance imaging to measure disease progression in Alzheimer disease: power calculations and estimates of sample size to detect treatment effects. Archives of Neurology 57 (3), 339–344.

Freeborough, P.A., Fox, N.C., 1997. The boundary shift integral: an accurate and robust measure of cerebral volume changes from registered repeat MRI. IEEE Transactions on Medical Imaging 16 (5), 623–629.

Gerardin, E., Chetelat, G., Chupin, M., Cuingnet, R., Desgranges, B., Kim, H.-S., Niethammer, M., Dubois, B., Lehericy, S., Garnero, L., Eustache, F., Colliot, O., 2009. Multidimensional classification of hippocampal shape features discriminates Alzheimer's disease and mild cognitive impairment from normal aging. NeuroImage 47 (4), 1476–1486.

Gerber, S., Tasdizen, T., Fletcher, P.T., Joshi, S., Whitaker, R., 2010. Manifold modeling for brain population analysis. Medical Image Analysis 14 (5), 643–653.

Hamm, J., Ye, D.H., Verma, R., Davatzikos, C., 2010. GRAM: a framework for geodesic registration on anatomical manifolds. Medical Image Analysis 14 (5), 633–642.

Hammers, A., Allom, R., Koepp, M.J., Free, S.L., Myers, R., Lemieux, L., Mitchell, T.N., Brooks, D.J., Duncan, J.S., 2003. Three-dimensional maximum probability atlas of the human brain, with particular reference to the temporal lobe. Human Brain Mapping 19 (4), 224–247.

Hastie, T., Stuetzle, W., 1989. Principal curves. Journal of the American Statistical Association 84, 502–516.

He, X., Yan, S., Hu, Y., Niyogi, P., Zhang, H.-J., 2005. Face recognition using laplacianfaces. IEEE Transactions on Pattern Analysis and Machine Intelligence 27, 328–340.

Hotelling, H., 1933. Analysis of a complex of statistical variables into principal components. Journal of Educational Psychology 24, 417–441.

Jack Jr., C.R., Petersen, R.C., Xu, Y.C., O'Brien, P.C., Smith, G.E., Ivnik, R.J., Boeve, B.F., Waring, S.C., Tangalos, E.G., Kokmen, E., 1999. Prediction of AD with MRI-based hippocampal volume in mild cognitive impairment. Neurology 52 (7), 1397–1407.

Jia, H., Wu, G., Wang, Q., Shen, D., 2010. ABSORB: Atlas building by self-organized registration and bundling. NeuroImage 51 (3), 1057–1070.

Jolliffe, I.T., 1986. Principal Component Analysis. Springer-Verlag.

Klein, S., Loog, M., van der Lijn, F., den Heijer, T., Hammers, A., de Bruijne, M., van der Lugt, A., Duin, R., Breteler, M.M.B., Niessen, W., 2010. Early diagnosis of dementia based on intersubject whole-brain dissimilarities. In: IEEE International Symposium on Biomedical Imaging, pp. 249–252.

Kohannim, O., Hua, X., Hibar, D.P., Lee, S., Chou, Y.-Y., Toga Jr., A.W., Jack, C.R., Weiner, M.W., Thompson, P.M., 2010. Boosting power for clinical trials using classifiers based on multiple biomarkers. Neurobiology of Aging 31 (8), 1429–1442.

Krzanowski, W.J., 1988. Principles of Multivariate Analysis: A User's Perspective. Oxford University Press.

Lehtovirta, M., Soininen, H., Laakso, M.P., Partanen, K., Helisalmi, S., Mannermaa, A., Ryynnen, M., Kuikka, J., Hartikainen, P., Riekkinen Sr., P.J., 1996. SPECT and MRI analysis in Alzheimer's disease: relation to apolipoprotein E epsilon 4 allele. Journal of Neurology, Neurosurgery, and Psychiatry 60, 644–649.

Leow, A.D., Yanovsky, I., Chiang, M.C., Lee, A.D., Klunder, A.D., Lu, A., Becker, J.T., Davis, S.W., Toga, A.W., Thompson, P.M., 2007. Statistical properties of jacobian maps and the realization of unbiased large-deformation nonlinear image registration. IEEE Transactions on Medical Imaging 26 (6), 822–832.

Mazziotta, J.C., Toga, A.W., Evans, A.C., Fox, P.T., Lancaster, J.L., 1995. A probabilistic atlas of the human brain: theory and rationale for its development. The international consortium for brain mapping (ICBM). NeuroImage 2 (2a), 89–101.

Mueller, S.G., Weiner, M.W., Thal, L.J., Petersen, R.C., Jack, C., Jagust, W., Trojanowski, J.Q., Toga, A.W., Beckett, L., 2005. The Alzheimer's disease neuroimaging initiative. Neuroimaging Clinics of North America 15 (4), 869–877.

Pearson, K., 1901. On lines and planes of closest fit to systems of points in space. Philosophical Magazine 2 (6), 559572.

Roweis, S.T., Saul, L.K., 2000. Nonlinear dimensionality reduction by locally linear embedding. Science 290 (5500), 2323–2326.

Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L.G., Leach, M.O., Hawkes, D.J., 1999. Nonrigid registration using free-form deformations: application to breast MR images. IEEE Transactions on Medical Imaging 18 (8), 712–721.

Schuff, N., Woerner, N., Boreta, L., Kornfield, T., Shaw, L.M., Trojanowski, J.Q., Thompson, P.M., Jack, C.R., Weiner, M.W., 2009. The Alzheimer's disease neuroimaging initiative. MRI of hippocampal volume loss in early Alzheimer's disease in relation to ApoE genotype and biomarkers. Brain 132 (4), 1067–1077.

Smith, S.M., Zhang, Y., Jenkinson, M., Chen, J., Matthews, P.M., Federico, A., Stefano, N.D., 2002. Accurate, robust, and automated longitudinal and cross-sectional brain change analysis. NeuroImage 17 (1), 479–489.

Souvenir, R., Pless, R., 2007. Image distance functions for manifold learning. Image and Vision Computing 25 (3), 365–373, articulated and non-rigid motio.

Tenenbaum, J.B., de Silva, V., Langford, J.C., 2000. A global geometric framework for nonlinear dimensionality reduction. Science 290 (5500), 2319–2323.

Torgerson, W., 1958. Theory and Methods of Scaling. Wiley, New York.

Trojanowski, J.Q., 2004. Searching for the biomarkers of Alzheimers. Practical Neurology 3, 30–34.

Van der Maaten, L., Postma, E., van der Herik, J., 2007. Dimensionality reduction: a comparative review. Published Online 10, 1–35.

Vemuri, P., Gunter, J.L., Senjem, M.L., Whitwell, J.L., Kantarci, K., Knopman, D.S., Boeve, B.F., Petersen, R.C., Jack Jr., C.R., 2008. Alzheimer's disease diagnosis in individual subjects using structural MR images: validation studies. NeuroImage 39 (3), 1186–1197.

von Luxburg, U., 2007. A tutorial on spectral clustering. Statistics and Computing 17 (4), 395–416.

Wolz, R., Aljabar, P., Hajnal, J., Rueckert, D., 2010a. Manifold learning for biomarker discovery in MR imaging. In: Dee, H. (Ed.), Machine Learning in Medical Imaging. Lecture Notes in Computer Science, vol. 6357. Springer, Berlin/Heidelberg, pp. 116–123.

Wolz, R., Aljabar, P., Hajnal, J.V., Hammers, A., Rueckert, D., 2010b. LEAP: learning embeddings for atlas propagation. NeuroImage 49 (2), 1316–1325.

Wolz, R., Heckemann, R.A., Aljabar, P., Hajnal, J.V., Hammers, A., Lotjonen, J., Rueckert, D., 2010c. Measurement of hippocampal atrophy using 4D graph-cut segmentation: application to ADNI. NeuroImage 52, 1009–1018.

Yan, Xu, Zhang, Zhang, Yang, Lin, 2007. Graph embedding and extensions: a general framework for dimensionality reduction. IEEE Transactions on Pattern Analysis and Machine Intelligence, 29.

Zhao, D.L., Lin, Z.C., Xiao, R., Tang, X., 2007. Linear Laplacian discrimination for feature extraction. In: Computer Vision and Pattern Recognition, pp. 1–7.