



Rigid motion invariant statistical shape modeling based on discrete fundamental forms

Data from the osteoarthritis initiative and the Alzheimer's disease neuroimaging initiative

Felix Ambellan^{a,1,*}, Stefan Zachow^{a,b,2}, Christoph von Tycowicz^{a,c,3}

^a Zuse Institute Berlin, Berlin, Germany

^b 1000shapes GmbH, Berlin, Germany

^c Freie Universität Berlin, Berlin, Germany

ARTICLE INFO

Article history:

Received 9 July 2020

Revised 8 July 2021

Accepted 13 July 2021

Available online 15 July 2021

Keywords:

Statistical shape analysis

Principal geodesic analysis

Lie groups

Classification

Manifold valued statistics

ABSTRACT

We present a novel approach for nonlinear statistical shape modeling that is invariant under Euclidean motion and thus alignment-free. By analyzing metric distortion and curvature of shapes as elements of Lie groups in a consistent Riemannian setting, we construct a framework that reliably handles large deformations. Due to the explicit character of Lie group operations, our non-Euclidean method is very efficient allowing for fast and numerically robust processing. This facilitates Riemannian analysis of large shape populations accessible through longitudinal and multi-site imaging studies providing increased statistical power. Additionally, as planar configurations form a submanifold in shape space, our representation allows for effective estimation of quasi-isometric surfaces flattenings. We evaluate the performance of our model w.r.t. shape-based classification of hippocampus and femur malformations due to Alzheimer's disease and osteoarthritis, respectively. In particular, we outperform state-of-the-art classifiers based on geometric deep learning as well as statistical shape modeling especially in presence of sparse training data. To provide insight into the model's ability of capturing biological shape variability, we carry out an analysis of specificity and generalization ability.

© 2021 Elsevier B.V. All rights reserved.

1. Introduction

Statistical shape models (SSMs) have become an essential tool for medical image analysis with a wide range of applications such as segmentation of anatomical structures, computer-aided diagnosis, and therapy planning. SSMs describe the geometric variability in a population in terms of a mean shape and a hierarchy of major modes explaining the main trends of shape variation. Based on a notion of shape space, SSMs can be learned from a database of consistently parametrized instances from the object class under study. The resulting models provide a shape prior that can be used to constrain synthesis and analysis problems. Moreover, their parameter space provides a compact representation that is amenable to learning algorithms (e.g. classification or clustering), evaluation, and exploration.

Standard SSMs treat the space of shapes as a Euclidean vector space allowing for linear statistics to be applied (see e.g. Heimann and Meinzer (2009) and the references therein). Linear methods, however, are often inadequate for capturing the high variability in biological shapes (Davis et al., 2010). Nonlinear approaches have been developed based on geometric as well as physical concepts such as Hausdorff distance (Charpiat et al., 2006), elasticity (Rumpf and Wirth, 2011; von Tycowicz et al., 2015; Zhang et al., 2015), and viscous flows (Fuchs et al., 2009; Brandt et al., 2016; Heeren et al., 2018). In general, these methods lack numerical robustness as well as fast response rates limiting their practical applicability especially in interactive applications. To address these challenges, one line of work models shapes by a collection of primitives (Fletcher et al., 2003; Freifeld and Black, 2012; Ambellan et al., 2019b) that naturally belong to Lie groups and effectively describe local changes in shape. Performing intrinsic calculus on the uncoupled primitives allows for fast computations while, at the same time, accounting for the nonlinearity in shape variation. However, solving the inverse problem, i.e. mapping from primitives back to surface meshes, is generally non-trivial. Recently, von Tycowicz et al. (2018) presented a physically

* Corresponding author.

E-mail addresses: ambellan@zib.de (F. Ambellan), zachow@zib.de (S. Zachow), vontycowicz@zib.de (C.v. Tycowicz).

¹ orcid 0000-0001-9415-0859.

² orcid 0000-0001-7964-3049.

³ orcid 0000-0002-1447-4069.




Fig. 1. Relations between reference shape $\tilde{S}$ (left) and shape $S = \phi(\tilde{S})$, a deformation thereof (right), s.t. $D_i := \nabla\phi|_{\tilde{T}_i}$. Note that each frame $\tilde{F}_i = R_i \tilde{F}_i$ is defined solely on the respective triangle $\tilde{T}_i$ and all neighboring frames are connected across the shared edge of their underlying triangles via $\tilde{F}_i C_{ij} = \tilde{F}_j$.

motivated approach based on differential coordinates for which the inverse problem is well-known and can be solved at linear cost. Despite their inherent nonlinear structure, the employed representations are not invariant under Euclidean motion and, thus, analysis thereon suffers from bias due to arbitrary choices. While the effect of rigid motions can be removed between pairs of shapes using alignment strategies, non-transitivity thereof prevents true group-wise alignment.

A related concept is to exploit the homogeneous structure of the ambient space and to encode displacements of points in terms of (e.g. rigid or affine) transformations (Gilles et al., 2011; Arsigny et al., 2003, 2009; McLeod et al., 2015). Exploiting the redundancy of such representations present e.g. in articulated motion, these approaches provide low-dimensional encodings of deformations. Considering the limit case of triangle-wise supported polyaffine/rigid deformations is similar to simplicial maps underlying the construction in (Freifeld and Black, 2012; von Tycowicz et al., 2018; Ambellan et al., 2019b) as well as our setup. However, the latter employ differential characterizations of such maps that remove translational components and put local geometric variability into focus.

This work presents a novel shape representation based on discrete fundamental forms that is invariant under Euclidean motion. We endow this representation with a Lie group structure that admits bi-invariant metrics and therefore allows for consistent analysis using manifold-valued statistics based on the Riemannian framework. Furthermore, we derive a simple, efficient, robust, yet accurate (i.e. without resorting to model approximations) solver for the inverse problem that allows for interactive applications. Beyond statistical shape modeling the proposed framework is amenable for surface processing such as quasi-isometric flattening. A publicly available implementation of the proposed model is given in the Morphomatics⁴ library.

Although in computer graphics and vision communities, rotation invariant differential coordinates have also been successfully employed for geometry processing applications, e.g. Kircher and Garland (2008); Hasler et al. (2009); Gao et al. (2016), these approaches fall short of a fully intrinsic treatment (e.g. due to lack of bi-invariant group structure and linearization) and have not been adapted to the field of SSMS. A recent string of contributions investigates functional characterizations of intrinsic and extrinsic geometry (Rustamov et al., 2013; Corman et al., 2017; Wang et al., 2018) to obtain shape descriptors. While the underlying functional map framework alleviates the requirement on point-to-point correspon-

dences, the reduced function spaces are based on low-frequency variations and, thus, prone to insensitivity for localized shape variability (such as osteophyte formation during the course of osteoarthritis).

2. Rotation invariant surface representation

In this section, we derive a discrete surface representation based on concepts from differential geometry of smooth surfaces. This representation's key feature, its invariance under Euclidean motion and hence well-suitedness for shape analysis purposes, arises directly from discretization of surface theoretical results. Finally, the proposed representation setting exhibits a Lie group structure that we endow with a bi-invariant metric in order to ensure structural unity between Riemannian and Lie group framework (see e.g. Pennec and Lorenzi (2020)).

2.1. Fundamental forms and surface theory

To every smooth surface there uniquely exist two smoothly pointwise varying and symmetric bilinear forms on the tangent plane, the so called *fundamental forms*. The *first* fundamental form I (a.k.a. metric tensor) is positive-definite and allows for angle, length and area measurement. The *second* fundamental form II describes the curvature of the surface. A prominent result in classical mathematics, the *Fundamental Theorem of Surface Theory* according to Bonnet (≈ 1860 , e.g. do Carmo (1976) Sec. 4.3), states that if given two symmetric bilinear forms (one of them positive-definite), s.t. for both certain integrability conditions hold (viz. the Gauß–Codazzi equations), then they (locally) determine uniquely, up to global rotation and translation, a surface embedded in three dimensional space with these two as its fundamental forms. Therefore, a discrete description of the fundamental forms is an excellent candidate for a rotation-invariant surface representation. In the following, we will denote our proposed shape model as the *fundamental coordinate model* (FCM).

2.2. Discretization

We consider shapes that belong to a particular population of anatomical structures, s.t. each digital shape S can be described as a left-acting deformation ϕ of a common reference shape $\tilde{S}$ given as triangulated surface. Let deformation ϕ be affine on each triangle $\tilde{T}_i$ of $\tilde{S}$, then the deformation gradient $\nabla\phi$ is the 3×3 matrix of partial derivatives of ϕ and constant on each triangle $D_i := \nabla\phi|_{\tilde{T}_i}$ (see e.g. Botsch et al. (2006) for detailed expressions). Note that transition from deformation to deformation gra-

⁴ morphomatics.github.io.




Fig. 2. Non-flat surface $\bar{S}$ is employed as reference within the deformation setup. Flat surface S is determined via deformation of $\bar{S}$ by ϕ , s.t. metric distortion, i.e. U_i , is close to identity and R_i is determined by means of C_{ij} that are normal vector fixing modifications of $\bar{C}_{ij}$.

dient provides invariance under translations. Assuming ϕ to be an orientation-preserving embedding of $\bar{S}$, we can decompose D_i uniquely into its rotational R_i and stretching U_i components by means of the polar decomposition $D_i = R_i U_i$. Note that U_i furnishes a complete description of the metric distortion of $\bar{T}_i$ and is defined in reference coordinates, hence invariant under rotation of S . Indeed, we can obtain a representation of the first fundamental form by restricting the stretches to the tangent plane. To this end, we define an arbitrary but fixed element-wise field $\{\bar{F}_i\}$ of orthonormal frames on $\bar{S}$, s.t. the last column of each frame is the normal of the respective element. Then, we represent the metric in terms of reduced stretch $\bar{U}_i := [\bar{F}_i^T U_i \bar{F}_i]_{3,3} = I_{\bar{T}_i}^{1/2}$, where $[\cdot]_{3,3}$ denotes the submatrix with the third row and column removed.

As for the second fundamental form, we note that at a point $p \in S$ it is determined by the differential of the normal field N , viz. $\mathbb{I}_p(v, w) = I_p(-dN_p(v), w)$ for tangent vectors v, w . For a triangulated surface, the differential dN is supported along the edges. In order to derive a representation thereof, we induce the frame field $\{F_i\}$ on S consistent to $\{\bar{F}_i\}$ using the rotational part of the deformation gradient, i.e. $F_i = R_i \bar{F}_i$. This allows us to define *transition rotations* $F_i C_{ij} = F_j$ for each inner edge (incident to triangles T_i, T_j) that fully describe the change in normal directions. Note that, while both the frames $\{F_i\}$ and the rotations $\{R_i\}$ are equivariant, the transition rotations $\{C_{ij}\}$ are invariant under global rotations of S and $\bar{S}$. Further details hereon are depicted in Fig. 1.

2.3. Group structure

In order to perform intrinsic statistical analysis, we derive a distance that is compatible with the underlying representation space. In particular, we endow the space with a Lie group structure together with a bi-invariant Riemannian metric for which group and Riemannian notions of exponential and logarithm coincide. This allows us to exploit closed-form expressions to perform geodesic calculus yielding simple, efficient, and numerically robust algorithms. We recommend chapter two of [Alexandrino and Bettiol \(2015\)](#) to readers interested in deeper insight into bi-invariant metrics on Lie groups. Especially regarding their existence and the geometric consequences thereof.

Our shape representation consists of transition rotations $C_{ij} \in \text{SO}(3)$ (one per inner edge) and tangential stretches $\bar{U}_i \in \text{Sym}^+(2)$ (one per triangle), where $\text{SO}(3)$ is the Lie group

of rotations in $\mathbb{R}^3$ and $\text{Sym}^+(2)$ the space of symmetric and positive-definite 2×2 matrices. Following the approach in [Arsigny et al. \(2006\)](#), we equip $U, V \in \text{Sym}^+(2)$ with a multiplication $U \circ V := \exp(\log(U) + \log(V))$, s.t. $\text{Sym}^+(2)$ turns into a commutative Lie group. It now allows for a bi-invariant metric induced by the Frobenius inner product yielding distance $d_{\text{Sym}^+(2)}(U, V) = \|\log(V) - \log(U)\|_F$. Note that this structure and metric do not exhibit the *swelling effect* of determinants in interpolation ([Goh et al., 2011](#)). $\text{SO}(3)$ as a compact Lie group also admits a bi-invariant metric induced by the Frobenius inner product with distance $d_{\text{SO}(3)}(Q, R) = \|\log(Q^T R)\|_F$, s.t. we define our representation space as the product group $G := \text{SO}(3)^n \times \text{Sym}^+(2)^m$ and m, n the number of triangles and inner edges. Finally, we define the distance of two shapes S, T based on the respective group representation $s = s(S), t = t(T) \in G$ as

$$d_\omega^2(s, t) = \frac{\omega^3}{\bar{A}_\mathcal{E}} \sum_{(i,j) \in \mathcal{E}} \bar{A}_{ij} d_{\text{SO}(3)}^2(C_{ij}^s, C_{ij}^t) + \frac{\omega}{\bar{A}} \sum_{i=1}^m \bar{A}_i d_{\text{Sym}^+(2)}^2(\bar{U}_i^s, \bar{U}_i^t), \quad (1)$$

where $\omega \in \mathbb{R}^+$ is a weighting factor, $\mathcal{E}$ is the set of inner edges, $\bar{A}_i$ is the area of triangle $\bar{T}_i$, $\bar{A}_{ij} = 1/3(\bar{A}_i + \bar{A}_j)$, $\bar{A}_\mathcal{E} = \sum_{(i,j) \in \mathcal{E}} \bar{A}_{ij}$, and $\bar{A} = \sum_{i=1}^m \bar{A}_i$. The area terms hereby provide invariance under refinement of the mesh as well as simultaneous scaling of $\bar{S}, S, T$, whereas ω allows for commensuration of the curvature and metric contributions in analogy to the Koiter thin shell model (e.g. [Ciarlet \(2005\)](#) Sec. 4.1).

3. Shape analysis and processing

3.1. Statistical shape modeling

The derived representation carries a rich non-Euclidean structure calling for manifold-valued generalizations for first and second moment statistical analysis. By virtue of the bi-invariant metric, the proposed representation allows for consistent analysis within the Riemannian framework for which statistics are well-developed, while at the same time providing closed-form, group theoretic expressions for geodesic calculus ([von Tycowicz et al., 2018](#)). In particular, we employ the *Riemannian center of mass* that provides a rigorous notion of a mean μ of elements $\{s_i = s_i(S_i)\}$ and




Fig. 3. Left: Example for a flattened femoral articular cartilage region. Right: Flattened femoral cartilage with gray value coded cartilage thickness. The top row shows healthy subjects whereas the subjects in the bottom row exhibit denuded areas within the cartilage region.

can be efficiently computed using the Gauss-Newton descent algorithm (Pennec, 2006; Arsigny et al., 2006):

$$\mu^{k+1} = \exp \left(\sum_i \log(s_i \cdot (\mu^k)^{-1}) \right) \cdot \mu^k.$$

As our representation space comprises a symmetric positive-definite and a rotational part the algorithm's respective behavior can be assessed separately. Since $\text{Sym}^+(2)$ is abelian and flat (indeed a vector space) the algorithm converges after exactly one step (Pennec, 2006). In contrast, $\text{SO}(3)$ features a less trivial structure exhibiting, e.g. a non-empty *cut locus*. However, as long as the data is located within some ε -ball, with ε smaller than the injectivity radius of the exponential map, the existence and (local) uniqueness of the mean can be guaranteed (Pennec, 2020) and thus convergence of the algorithm. Note that this assumption is only violated for transition rotations differing by more than $\pm\pi$, what can be practically ruled out. (cf. Appendix B). As framework for second order statistics we employ (linearized) Principal Geodesic Analysis (Fletcher et al., 2004) at μ that is an extension of the common Principal Component Analysis to Riemannian manifolds allowing for covariance analysis. In particular, we solve

$$\begin{aligned} \vartheta_p &= \arg \max_{\vartheta \in T_\mu} G \sum_i g_\omega^\mu(\vartheta, \log_\mu(s_i))^2, \\ \text{s.t. } g_\omega^\mu(\vartheta_p, \vartheta_l) &= \delta_{pl}, \text{ for } 1 \leq l \leq p \end{aligned}$$

for the main modes of variation ϑ_p , where g_ω is the metric associated to distance d_ω (Eq. 1). The solution is found algorithmically by eigendecomposition of the Gram matrix $C = (c_{ij})_{ij}$, with $c_{ij} = g_\omega^\mu(\log_\mu(s_i), \log_\mu(s_j))$ (cf. Younes (2010) E.2.2). In order to avoid a systematic bias due to the choice of reference shape $\bar{S}$, we require it to agree with the mean of the training data ($\bar{S} = \bar{S}(\mu)$) as proposed in Joshi et al. (2004). Details on how to determine a shape for given representation parameters are given in Section 4.

3.2. Quasi-isometric surface flattening

Apart from shape analysis, the proposed representation provides an effective framework for processing operations. In this section, we derive an approach for the calculation of a quasi-isometric surface chart, i.e. a low-distortion immersion of a given surface into the two dimensional Euclidean space. Since flattening techniques provide a way to access problems of three dimensional context in a two dimensional fashion, such an approach facilitates practically relevant applications like visualization and deep learning based assessment of knee cartilage thickness (Fig. 3). For a broader overview on application examples we refer to Kreiser et al. (2018), who published a survey on flattening-based medical visualization techniques. The key idea behind our flattening approach is to consider the set of flat immersions of the reference shape $\bar{S}$ as a submanifold in shape space. This sub-

manifold has a particularly convenient characterization in our representation space G allowing for a simple, isometric projection: We fix the metric part $\{U_i = I_{\text{Sym}^+}\}$ as identity (no metric distortion) and choose transition rotations s.t. they act as identity on the normals (zero curvature). In particular, the latter are given by $\{C_{ij} = F_i^{-1} \cdot R_{ji}^N \cdot F_j\}$, where R_{ji}^N unfolds triangle $\bar{T}_j$ to the plane of triangle $\bar{T}_i$. Phrasing it in the group setting this means we project the transition rotations to $\text{SO}(2)$ (embedded in $\text{SO}(3)$) since all feasible flat shape representations necessarily have to be elements of $\text{SO}(2)^n \times \text{Sym}^+(2)^m$. See Fig. 2 for a schematic overview. Note that the obtained projection corresponds to a realizable deformation, iff the input shape $\bar{S}$ is isometric to the plane. In general, a near-isometric flattening can be efficiently computed using our reconstruction (cf. next section).

4. Efficient numerics

In this section, we propose an efficient numerical algorithm to solve for the inverse problem of mapping a point in representation space G to a corresponding shape $S = \phi(\bar{S})$. If the corresponding rotations $\{R_i\}$ were known, ϕ could be obtained as the minimizer of $\sum_{i=1}^m \bar{A}_i \|D_i - R_i U_i\|_F^2$ by solving the well-known Poisson equation (see e.g. Botsch et al. (2006); von Tycowicz et al. (2018)). However, in our representation the rotations are only given implicitly in terms of the transition rotations. In particular, an immediate computation shows that $R_j = R_i \bar{F}_{ij} C_{ij} \bar{F}_j^T =: R_{i \rightarrow j}$ for an integrable field $\{C_{ij}\}$. Based on this condition, for each triangle T_i we can formulate a residual term $\varepsilon_i(\phi, \{R_i\}) = \sum_{j \in \mathcal{N}_i} 1/|\mathcal{N}_i| \|D_i - R_{j \rightarrow i} U_i\|_F^2$ in terms of the rotations of neighboring triangles (indexed by $\mathcal{N}_i$). Then, the objective for the inverse problem is given as $E(\phi) = \min_{\{R_i \in \text{SO}(3)\}} E(\phi, \{R_i\})$, where $E(\phi, \{R_i\}) = \sum_{i=1}^m \bar{A}_i \varepsilon_i(\phi, \{R_i\})$. Although $E(\phi)$ is a nonlinear function calling for iterative optimization routines, it exhibits a special structure amenable to an efficient alternating minimization technique. Specifically, we employ a block coordinate descent strategy that alternates between a local and a global step:

Local step: First, we minimize $E(\phi, \{R_i\})$ over the rotations $\{R_i\}$ keeping ϕ (hence D_i) fixed. Each summand in ε_i depends on a single rotation R_j , s.t. the problem decouples into individual low-dimensional optimizations that can be solved in closed-form and allow for massive parallelization. Note that the problem reduces to the well-known orthogonal Procrustes problem. Further details are given visually and formally in Appendix A.

Global step: Second, we minimize $E(\phi, \{R_i\})$ over ϕ with rotations $\{R_i\}$ fixed leading to a quadratic optimization problem for which the optimality conditions are determined by a Poisson equation. As the system matrix is sparse and depends only on the reference shape, it can be factorized once during the preprocessing allowing for very efficient global solves with close to linear cost.




Fig. 4. Mean shape of healthy distal femora overlaid with (larger) mean shape of the diseased femora wherever the distance is larger than 1.45mm, colored accordingly.




Fig. 5. Mean shape of diseased right hippocampi overlaid with mean shape of the healthy hippocampi wherever the distance is larger than 0.6mm, colored accordingly.

Note that the objective is bounded from below and that both local and global steps feature unique solutions that are guaranteed to weakly decrease the objective making any numerical safeguards unnecessary. This contrasts with classical approaches that require precautions, such as line search strategies and modification schemes for singular or indefinite Hessians, to guarantee robustness.

Initialization To provide the solver with a warm start, we compute an initial guess for the rotation field $\{R_i\}$. To this end, we employ the local integrability condition $R_j = R_{i \rightarrow j}$ to propagate an initial rotation matrix from an arbitrary seed along a precomputed spanning tree of the dual graph of $\tilde{S}$. Note, that this strategy recovers the rotation field exactly for integrable $\{C_{ij}\}$. In case of non-integrable fields, one advantage of the Poisson-based reconstruction (global step) is that it distributes errors uniformly s.t. local inconsistencies are attenuated. More details on the initialization procedure can be found in [Appendix A](#).

5. Experiments and results

Except where stated otherwise all experiments are performed employing a fixed metric commensuration weight $\omega = 10$ that empirically shows the best performance in our knee osteoarthritis classification experiment (cf. [Appendix C](#)).

5.1. Data

We employ four different datasets to ensure a qualitative and quantitative as well as a technical and application-oriented assessment of the proposed FCM.




Fig. 6. OA classification experiment for the proposed FCM, PDM (Cootes et al., 1995) and DCM (von Tycowicz et al., 2018).




Fig. 7. Alzheimer's classification experiment for the proposed FCM, MeshCNN (Hanocka et al., 2019) and PDM (Cootes et al., 1995).

(i) **OAI - Right distal femora** (see Fig. 4) from the Osteoarthritis Initiative⁵ (OAI) database. All subjects are rated w.r.t. knee osteoarthritis using the Kellgren and Lawrence score (0, healthy $\rightarrow$ 4, severely diseased) (Kellgren and Lawrence, 1957). The dataset consists of 58 severely diseased (grade 4) and 58 healthy subjects (grade 0,1) that were also used for evaluation in von Tycowicz et al. (2018) to which we refer for further details on the data, especially with regards to the arrangement of correspondence (8988 vertices, 13,776 triangles). We added a list of patient ids to [Appendix D](#) since the underlying segmentation masks are publicly available as part of publication (Ambellan et al., 2019a).

(ii) **ADNI - Right hippocampi** (see Fig. 5) from the Alzheimer's Disease Neuroimaging Initiative⁶ (ADNI) consisting of 60 subjects showing Alzheimer's disease and 60 cognitive normal controls. We prepared this dataset using imaging data from the ADNI database that contains, among others, 1632 brain MRI scans collected on four different time points with segmented hippocampi. We established surface correspondence (2280 vertices, 4556 triangles) in a fully automatic manner employing the deblurring and denoising of functional maps approach (Ezuz and Ben-Chen, 2017)⁷ for isosurfaces extracted from the given segmentations. The dataset was randomly assembled from the baseline shapes for which segmentations were simply connected and remeshed surfaces were well-approximating ($\leq 10^{-5}$ mm root mean square surface distance to the isosurface). Similar as for the OAI dataset we added a list of

⁵ nda.nih.gov/oai.

⁶ adni.loni.usc.edu.

⁷ code online available: cs.technion.ac.il/~mirela/code/fmap2p2p.zip.




Fig. 8. Visualization of the discriminating direction η and separating hyperplane $\eta^\perp$ for OA classification showing a 2-dimensional projection (right) and corresponding shapes (left) equidistantly sampled within the interval containing the input data (note that projections onto η and the visualizing plane do not commute causing the interval to appear smaller). Point-wise distance to the middle shape colored using -0.5mm 0.5mm with neutral window (i.e. rosy color) from -0.15mm to 0.15mm .




Fig. 9. Visualization of the discriminating direction η and separating hyperplane $\eta^\perp$ for Alzheimer's classification showing a 2-dimensional projection (right) and corresponding shapes (left) equidistantly sampled within the interval containing the input data (note that projections onto η and the visualizing plane do not commute causing the interval to appear smaller). Point-wise distance to the middle shape colored using -2.5mm 2.5mm with neutral window (i.e. rosy color) from -0.5mm to 0.5mm .

scan ids to [Appendix D](#) since the used hippocampus segmentations are publicly available as part of the ADNI database.

(iii) **FAUST - A male human body in two poses** being part of the anthropological, open-access Fine Alignment Using Scan Texture (FAUST) (Bogo et al., 2014)) dataset of whole body scans featuring high-quality, dense correspondences (6890 vertices, 13,776 faces). We chose two poses of the same (male) person, lifting the arms up and down alongside its body (Fig. 10, right).

(iv) **PIPE - A pair of synthetic pipe surfaces**, one in a cylindrical and one in a helical configuration consisting of 1220 triangles and 612 vertices as can be seen in Fig. 10 (left).

Throughout the manuscript we will refer to the datasets using the above acronyms relating to the data origin or content.

The datasets OAI and ADNI will be used for quantitative analysis and comparison to other shape models, whereas FAUST and PIPE serve for qualitative assessment of the proposed model.

5.2. Disease classification

To assess the sensitivity of the proposed FCM for degenerative shape changes and compare to different other explicit shape modeling approaches, we will perform two binary disease classification experiments, one regarding knee osteoarthritis on the distal femur and one concerning Alzheimer's disease w.r.t. the right hippocampus. Although these diseases are very different and not comparable in a direct way we will make use of the same classification pipeline for both of them assessing the generalization potential of the proposed FCM regarding classification tasks. To this end, we train a support vector machine (SVM) with linear kernel on feature vectors comprising shape weights, i.e. coefficients of the basis representation in terms of the principal modes for every input shape. By construction this representation is exact for all input shapes and every shape model type. The coefficients hence forthrightly reflect the underlying model approach to gauge variation. The classifier is trained on a balanced set of feature vectors for different shares of

data varying from 10% to 90% with testing on the respective complement. To address the randomness in our experimental design, we perform a Monte Carlo cross-validation drawing 10,000 times per partitioning.

5.2.1. Knee osteoarthritis classification

Osteoarthritis (OA) is a degenerative disease of the joints that is i.a. characterized by changes of the bone shape (see Fig. 4). Here, we investigate the proposed FCM's ability to classify knee OA for the OAI dataset of distal femora. Since our test set contains 58 healthy and 58 diseased cases the SVM classifier is trained on 115-dimensional feature vectors. We compare to the popular *point distribution model* (Cootes et al., 1995) (PDM) as well as to the differential coordinates model (von Tycowicz et al., 2018) (DCM), which recently achieved highly accurate classification results. Fig. 6 shows the results in terms of average accuracy and standard deviation. Note that solely the FCM achieves an accuracy of over 90% in case of sparse (10%) training data.

5.2.2. Alzheimer's classification

There is a substantial body of work confirming the well-known connection between hippocampal volume loss and Alzheimer's progression (Köhler et al., 1998; de Toledo-Morrell et al., 2000; Bonner-Jackson et al., 2015). In line with these findings, we observe $\approx 1/4$ volume loss between the FCM-based mean shapes of the diseased subjects to the one of the healthy controls. This motivates a classification experiment regarding the shape of right hippocampi and the disease state to further evaluate the descriptiveness of our shape representation. For this experiment we employ the commensuration parameter $\omega = 0.98$ that empirically performs best w.r.t. classification accuracy, i.e. metric and curvature related differences are weighted almost equally within the shape analysis. Since our test set contains 60 cognitive normal and 60 diagnosed Alzheimer's cases the SVM is trained on 119-dimensional feature vectors.




Fig. 10. Interpolating geodesic (mean highlighted) for the pipe surface (left) and FAUST (right) data within (f.l.t.r.) the DCM (von Tycowicz et al., 2018), the proposed FCM, and PDM (Cootes et al., 1995), each.

Given the coarse discretization of the hippocampal surface (other than the OAI data) and, thus, moderate hardware requirements, we can perform a direct comparison to MeshCNN (Hanocka et al., 2019, i.e. a state-of-the-art surface-based classifier from the field of geometric deep learning. Specifically, we employed the implementation of the authors⁸ performing training on an Nvidia GTX 980 TI graphics card (6GB memory). Due to the lack of proper stopping criteria (no option for a validation set), we report the best test accuracy attained in the first 100 epochs, which is rather an upper bound for the classification performance. Due to the high computational cost (> 2 hours for training) we restrict to 10 samples per partitioning during Monte Carlo cross-validation, which increases variability as evident by the lack of monotonicity of the estimated dependency of the accuracy on the training size. Partitioning is carried out analogously to the SVM classifier and training employs the Adam optimizer (Kingma and Ba, 2014) with weight decay $\beta_1 = 0.9$ and $\beta_2 = 0.999$.

Figure 7 shows the obtained classification accuracies for MeshCNN as well as our FCM-based and (for reference) PDM-based SVM. Note that the FCM reaches average accuracies ranging from 75.6% (10% training) up to 80.8% (90% training) with values above 80% for all data shares $\geq 50\%$. Remarkably, the FCM-based classifier not only outperforms the PDM one but is also superior to MeshCNN especially in presence of sparse training data. Note, these results have to be understood in the context of data used, namely the shape of *one* single anatomy. Higher classification accuracy is possible if more data is utilized, as e.g. 3D MRI scans of the *whole* brain in Seo et al. (2016).

5.2.3. Transparency

The proposed classifier exposes a high degree of interpretability and explainability due to the generative character of statistical shape models and the linearity of the employed SVM. In particular, the discriminating direction underlying the SVM corresponds to a geodesic in representation space that directly encodes the *single* type of morphological variation that determines the classifier's prediction. We provide a visualization of the discriminating direction for both experiments (OAI and ADNI) in Figs. 8 and 9 based on SVM instances with average classification accuracy obtained for 40%/60% training/testing split. In addition to shapes sampled along the discriminating direction, we provide a 2-dimensional visualization using orthogonal projection onto the plane spanned by the two principal geodesic modes that retain the highest classification accuracy, viz. ϑ_1, ϑ_3 and ϑ_1, ϑ_2 for ADNI and OAI, respectively.

5.3. Validity

Frequently, datasets feature a high nonlinear variability that are characterized by large rotational components, which are insufficiently captured by linear models like PDM. While DCM treats the rotational components explicitly, it requires them to be well-localized, s.t. the logarithm is unambiguous. This assumption may not be satisfied for data with large spread in shape space. Contrary, our model overcomes this limitation by utilizing a relative encoding via transition rotations, which will never exceed $\pm\pi$ in practical scenarios (cf. Appendix B). In Fig. 10 we illustrate the validity of our model for two extreme examples in comparison to PDM and DCM.

5.4. Computational performance

We compare our framework in terms of computational efficiency to two state-of-the-art approaches: The *large deformation diffeomorphism metric mapping* (LDDMM) using the open-source Deformetrica (Durrleman et al., 2014) software, and the recent DCM. To this end, we compute the mean shape on 100 randomly sampled pairs from the OAI dataset. Overall, the LDDMM approach requires 172.8s (± 44.8 s) in average whereas the proposed FCM features an average runtime of only 2.3s (± 1.9 s), hence a two orders of magnitude speedup. In comparison to the highly efficient DCM—requiring 1.1s (± 0.3 s) in average—our model achieves runtimes within the same order of magnitude, despite the added non-linearity in the inverse problem.

5.5. Specificity, generalization ability, compactness

We perform a quantitative comparison with PDM and DCM using standard measures (detailed in Davies et al. (2008)) w.r.t. a physically-based surface distance $\mathcal{W}$ (Heeren et al., 2018) as proposed in von Tycowicz et al. (2018). *Specificity* (Fig. 11 middle) evaluates the validity of the model generated instances in terms of their distance to the training shapes. We estimate it using 1000 randomly generated instances according to the discrete distribution of the respective model. *Generalization ability* (Fig. 11 left) assesses how well a model represents unseen instances. It is calculated in a leave-one-out study. *Compactness* (Fig. 11 right) measures the relative amount of variability of the training set captured by every mode in an accumulated manner. The results show that the FCM is more specific than PDM and DCM. In terms of generalization ability, the FCM is superior to PDM, yet inferior to DCM. Finally, the FCM is less compact than PDM and DCM. Note that compactness is calculated for each model w.r.t its own metric, hence not directly comparable. In particular, we found that decreasing ω leads to increased compactness, albeit at the possible expense of classification accuracy (cf. Appendix C).

⁸ github.com/ranahanocka/MeshCNN.




Fig. 11. Generalization ability (left, lower~better), specificity (middle, lower~more specific) and compactness (right, higher~more compact) of the proposed FCM, PDM (Cootes et al., 1995), and DCM (von Tycowicz et al., 2018) on the OAI dataset.

6. Conclusion and future work

In this work, we presented a novel nonlinear SSM based on a Euclidean motion invariant—hence alignment-free—shape representation with deep foundations in surface theory. The rich structure of the derived shape space assures valid shape instances even in presence of strong nonlinear variability. Moreover, we demonstrated that the proposed shape representation can be used to effectively calculate quasi-isometric flat immersions to the plane. We performed manifold-valued statistics in a consistent Lie group setup allowing for closed-form evaluation of Riemannian operations. Furthermore, we devised an efficient and robust algorithm to solve the inverse problem that does not require any numerical safeguards.

We showed that FCM yields highly differentiating shape descriptors that promote state-of-the-art performance for shape-based disease state classification: (1) In comparisons based on a simple classifier (viz. linear SVM) our descriptors are superior to the recent Riemannian DCM (von Tycowicz et al., 2018) as well as the popular linear PDM (Cootes et al., 1995) based descriptors; (2) Remarkably, the FCM-based SVM significantly outperforms the state-of-the-art, geometric deep learning approach MeshCNN (Hanocka et al., 2019).

We would like to remark that our approach guarantees deformations to be only locally diffeomorphic (i.e. immersions) but not globally. However, we did not observe any non-diffeomorphic instances in our experiments (e.g. all FCM-derived shapes shown in this article are embeddings). Indeed, the FCM correctly captures nonlinear deformations with large rotational components that violate well-localizedness assumptions of previous approaches. On the other hand, in comparison to shape spaces based on diffeomorphic mapping, FCM allows for fast processing of large-scale shape collections and is invariant under Euclidean motion, hence, not susceptible to any bias due to misalignment.

One possible and interesting way to proceed in the future is to replace the log-Euclidean metric with the affine-invariant one, which can be considered the natural metric on the symmetric positive-definite matrices. Another interesting line of future work is to explore the feasibility of fully automatic computer-aided diagnostics based on advanced machine learning, e.g. by combining our shape representation with the approach in von Tycowicz (2020) or utilizing the proposed flattening approach to transform three-dimensional problems into the well-known two-dimensional image-based deep learning setup.

Declaration of Competing Interest

None.

CRediT authorship contribution statement

Felix Ambellan: Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization. **Stefan Zachow:** Conceptualization, Resources, Writing – review & editing, Supervision, Project administration, Funding acquisition. **Christoph von Tycowicz:** Conceptualization, Methodology, Software, Validation, Writing – original draft, Writing – review & editing, Visualization, Supervision.

Acknowledgments

The authors are funded by the [Deutsche Forschungsgemeinschaft](#) (DFG, German Research Foundation) under Germany's Excellence Strategy - The Berlin Mathematics Research Center MATH+ (EXC-2046/1, project ID: 390685689). This work partly relies on data of the Osteoarthritis Initiative⁹ and on data of the Alzheimer's Disease Neuroimaging Initiative (ADNI).¹⁰ Furthermore we are grateful for the open-access dataset FAUST (Bogo et al., 2014) as well as for the open-source software Deformetrica (Durrleman et al., 2014). Finally, we want to thank A. Tack for providing us with cartilage thickness measure-

⁹ The Osteoarthritis Initiative is a public-private partnership comprised of five contracts (N01-AR-2-2258; N01-AR-2-2259; N01-AR-2-2260; N01-AR-2-2261; N01-AR-2-2262) funded by the National Institutes of Health, a branch of the Department of Health and Human Services, and conducted by the OAI Study Investigators. Private funding partners include Merck Research Laboratories; Novartis Pharmaceuticals Corporation, GlaxoSmithKline; and Pfizer, Inc. Private sector funding for the OAI is managed by the Foundation for the National Institutes of Health. This manuscript was prepared using an OAI public use dataset and does not necessarily reflect the opinions or views of the OAI investigators, the NIH, or the private funding partners.

¹⁰ Data collection and sharing for this project was funded by the ADNI (National Institutes of Health Grant U01 AG024904) and DOD ADNI (Department of Defense award number W81XWH-12-2-0012). ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and Bioengineering, and through generous contributions from the following: AbbVie, Alzheimer's Association; Alzheimer's Drug Discovery Foundation; Araclon Biotech; BioClinica, Inc.; Biogen; Bristol-Myers Squibb Company; CereSpir, Inc.; Cogstate; Eisai Inc.; Elan Pharmaceuticals, Inc.; Eli Lilly and Company; EuroImmun; F. Hoffmann-La Roche Ltd and its affiliated company Genentech, Inc.; Fujirebio; GE Healthcare; IXICO Ltd.; Janssen Alzheimer Immunotherapy Research & Development, LLC.; Johnson & Johnson Pharmaceutical Research & Development LLC.; Lumosity; Lundbeck; Merck & Co., Inc.; Meso Scale Diagnostics, LLC.; NeuroRx Research; Neurotrack Technologies; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Piramal Imaging; Servier; Takeda Pharmaceutical Company; and Transition Therapeutics. The Canadian Institutes of Health Research is providing funds to support ADNI clinical sites in Canada. Private sector contributions 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 Therapeutic Research Institute at the University of Southern California. ADNI data are disseminated by the Laboratory for Neuro Imaging at the University of Southern California.

ments for Fig. 3. The present article is an extended version of Ambellan et al. (2019c) presented at MICCAI 2019.

Appendix A. Shape reconstruction - Initialization and Local step

Efficient shape reconstruction as outlined in Section 4 is an essential part of this work and consists of an initialization and an iteration of global and local step until a solution is reached. Since the global step is basically solving a Poisson problem and also exemplified in von Tycowicz et al. (2018) we will omit it here. To ease deeper inside on how to initialize the reconstruction algorithm and how to do the local step we provide two schematic visuals and explicitly determine the solution of the local step through direct calculation.


Initialization. To initialize the algorithm we fix the pose of an arbitrarily chosen triangle i_0 by fixing its rotation R_{i_0} relative to reference shape $\bar{S}$. Starting from triangle i_0 a spanning tree is determined to define a path through the dual graph, passing every triangle exactly once.

R_{i_0} is propagated along this path employing the local integrability condition, viz. $R_j := R_{i \rightarrow j} = R_i \bar{F}_{ij} C_{ij} \bar{F}_{ij}^T$. Finally this procedure provides a field of extrinsic rotations $\{R_i\}$ to initialize the local/global solver. Note that the algorithm will stop right after the first iteration, if the $\{C_{ij}\}$ form an integrable system. In that case the initialization is additionally invariant w.r.t. the choice of the fixed triangle and spanning tree. Fig. A.1 provides a schematic overview on the initialization process.

Local step. Within the local step of the reconstruction algorithm the deformation ϕ of $\bar{S}$ and hence the deformation gradient is fixed. We aim to find R_i for every triangle i , s.t. it, mediated through $\{C_{ij}\}$, locally optimally accompanies ϕ . »Local« hereby has to be understood as the one-ring triangle-neighborhood. This problem has a closed-form solution that we will work out within the following proposition via direct calculation on the optimization target.

Proposition. The local step within the shape reconstruction algorithm targeting the optimization problem:

$$R_i = \arg \min_{R \in \text{SO}(3)} \sum_{s \in \mathcal{N}_i} \|\nabla \phi|_{\bar{T}_s} - R \bar{F}_{is} C_{is} \bar{F}_{is}^T U_s\|_F^2,$$






Fig. B.1. Histogram of angles between transition rotations of the PIPE shapes. The angles are all relatively small and far away from the critical region of angles larger than $\pm\pi$.




Fig. B.2. Schematic summary of the argumentation to prove the proposition on relative transition rotations. The given construction allows to separate relative transition rotations into normal (I) and tangential (II) type.

Proof. We poof the result for two shapes $\tilde{S}$ and S since the argumentation naturally extends to the whole n shapes. At first we have to control change in normal direction in regions of high curvature. We therefore simultaneously refine the triangulations until the angle between normals of any two neighboring triangles lies in $(-\pi/2, \pi/2)$. From this point onwards the argumentation is additionally summarized in Fig. B.2. The basic idea of the following construction is to separate normal from tangential difference and to argue independently on each. We now subdivide every triangle $\tilde{T}_i$ of $\tilde{S}$ (analogously for T_i of S) into three parts $\tilde{T}_{i_0}, \tilde{T}_{i_1}, \tilde{T}_{i_2}$ by means of the incenter and the bisecting lines of the angles. Without loss of generality we can assume that $\tilde{T}_{j_0}$ is neighboring $\tilde{T}_{i_0}$ (thus $\tilde{T}_j$ was already neighboring $\tilde{T}_i$). We fix a frame $\tilde{F}_{i_0}$ on $\tilde{T}_{i_0}$ ensuring alignment of the first basis vector to the edge shared with $\tilde{T}_{j_0}$. Frame $\tilde{F}_{j_0}$ is now defined by rotating $\tilde{F}_{i_0}$ around the common edge. This directly implies that $\tilde{C}_{i_0 j_0}$ realizes an angle with absolute value smaller than $\pi/2$. The same holds for $C_{i_0 j_0}$

since $\tilde{F}_i = R_i \tilde{F}_i$ preserves alignment of the frames with the underlying triangles and through the initial refinement we already ensured normal differing of less then $\pm\pi$. We analogously define $\tilde{F}_{i_k}$ if $\tilde{T}_{i_k}$ has neighbors, if not, we simply set $\tilde{F}_{i_k} = \tilde{F}_{i_0}$. This construction allows to explicitly differentiate two different types of relative transition rotations: type (I) that comes from normal differences like $C_{i_0 j_0} \tilde{C}_{i_0 j_0}^{-1}$ and type (II) like $C_{j_0 j_1} \tilde{C}_{j_0 j_1}^{-1}$ that is induced by tangential change. Since $F_{j_0} = R_j \tilde{F}_{j_0}$ and $F_{j_1} = R_j \tilde{F}_{j_1}$ we see immediately that $C_{j_0 j_1} = \tilde{F}_{j_0}^{-1} R_j^{-1} R_j \tilde{F}_{j_1} = \tilde{C}_{j_0 j_1}$ and thus $C_{j_0 j_1} \tilde{C}_{j_0 j_1}^{-1} = I_3$. To clarify (I) we strip some notation, more precisely, let C^1, C^2 be two transition rotations with angles $-\theta^1, \theta^2$ and axes $-v^1, v^2$, respectively realizing normal change only. Then, the relative transition rotation is given by $C^{12} = C^2 \cdot (C^1)^{-1}$ and of type (I). Now, assuming $\theta^1, \theta^2 \in (-\pi/2, \pi/2)$ and in light of

$$\cos\left(\frac{\theta^{12}}{2}\right) = \cos\left(\frac{\theta^1}{2}\right) \cos\left(\frac{\theta^2}{2}\right) - \sin\left(\frac{\theta^1}{2}\right) \sin\left(\frac{\theta^2}{2}\right) \langle v^1, v^2 \rangle$$


(cf. e.g. Altmann (2005)), it follows that the angle θ^{12} of the composite rotation does not exceed $(-\pi, \pi)$, hence is well-localized. $\square$

Appendix C. Classification with varying commensuration parameter


As we are on the one hand applying shape models for disease classification purposes and on the other hand are in general interested in rather compact models, we did vary the metric commensuration parameter ω (cf. Section 2.2) since it directly affects both. We studied the connection between choice of ω , classification accuracy and model compactness.

OAI - OA Classification. As can be seen in Fig. C.1 the disease classification accuracy increases as ω increases. Looking at Eq. (1) this means putting higher weight on the rotational, thus curvature related term leads to higher classification accuracy. Additionally all examined choices of ω give FCM classification results with a superior performance compared to PDM. However, the development of model compactness is contrary to the classification accuracy as shown in Fig. C.1. The larger ω gets, the less compact is the shape model and none of the examined commensuration parameter choices leads to a compactness as high as for the PDM.

ADNI - Alzheimer's Classification. A similar experiment for Alzheimer's classification reveals a rather different dependency on the commensuration parameter, see Fig. C.1. For values $\omega \gtrsim 10$ the classification accuracy lies below the one achieved by the PDM. For $\omega \approx 0.98$ the peak performance is reached and for values below we note again slight decrease in performance. The compactness instead, as can be seen in Fig. C.1, develops very similar as for the OAI dataset and is still for all ω below that of the PDM. As conclusion to this section we conjecture, that comparison of compactness might be interesting for models that are build on the same shape representation but it becomes less meaningful if different



(C.1.1) OAI data: As ω increases accuracy does and compactness decreases.



(C.1.2) ADNI data: As ω increases compactness decreases, peak accuracy is obtained for $\omega \approx 0.98$.

Fig. C.1. Classification accuracy and compactness for varying commensuration parameter ω .

representations are compared. Furthermore, we find that the most compact models do not (necessarily) give the best classification accuracy. It appears that complex shape variation as it emerges from certain diseases tends to require a less compact encoding for an expressive but specific description.

Appendix D. Data identifiers

The data used within the given classification experiments is publicly available, we thus aim to facilitate reproduction of and comparison to our results. To this end we compiled labeled identifier lists of the data used in our experiments (Table D.1 and D.2).

Table D.1

List of unique patient ids from the OAI database used in the OA classification experiment.

Healthy (KL 0/1)			Diseased (KL 4)		
9008561	9258563	9510418	9246518	9391984	9631713
9013798	9304351	9517914	9256759	9393987	9638953
9017909	9331053	9582487	9263504	9413071	9642550
9036770	9333574	9601162	9266394	9414291	9660708
9036948	9341699	9617689	9267719	9421492	9672573
9039744	9341903	9645577	9271965	9422381	9680800
9089627	9355112	9655592	9284505	9430102	9689922
9108461	9383004	9718992	9287216	9439428	9691663
9116298	9391372	9750072	9301332	9457359	9695135
9120941	9394136	9854269	9326657	9467278	9700341
9132486	9397088	9876530	9331465	9469318	9710479
9141244	9397976	9878765	9340139	9470313	9745458
9153509	9433408	9879069	9349261	9475286	9750090
9171766	9440417	9907090	9364366	9477175	9760079
9184495	9460287	9916140	9365968	9477358	9781749
9189553	9474901	9967815	9375317	9508335	9858216
9207016	9486748	9973322	9379276	9517311	9895555
9211049	9488834	9978579	9389580	9557454	9933836
9245519	9501871	9988421	9391061	9568504	9943227
	9504627			9604541	

Table D.2

List of unique scan ids from the ADNI database used in the Alzheimer's classification experiment.

Cognitive Normal			Alzheimer's Diagnosed		
10312	13681	17207	10064	14974	22310
10605	13717	17232	10468	15001	22938
10813	13737	17487	10568	15145	23375
10835	13893	17527	10764	15287	23446
10883	14104	18109	11633	15315	24659
10960	14488	18236	12000	15935	24672
11006	14513	18321	12365	16313	25082
11161	14559	18450	12375	16924	25357
11314	14818	18827	12381	17191	25455
11584	14959	18909	12402	17337	25763
11594	14991	18917	12468	18077	26038
11928	15079	19971	12583	18094	26136
11974	15527	20352	12836	18151	26143
12081	15727	20753	12952	18189	26314
12419	15789	21817	13839	18373	26431
12485	16048	22439	13976	18390	27061
12563	16099	24338	13990	19296	27414
12992	16553	25680	14199	19386	27584
13191	16759	25829	14629	19395	27673
13556	17131	26899	14699	21207	28133

Regarding the OA classification all cases can be found as segmentation masks accompanying the publicly available OAI-ZIB dataset¹¹, whereas the Alzheimer's classification experiment relies on hippocampus segmentation masks that can be accessed as part of ADNI database.¹²

References

- Alexandrino, M.M., Bettiol, R.G., 2015. Lie Groups and Geometric Aspects of Isometric Actions, Vol. 8. Springer.
- Altmann, S.L., 2005. Rotations, Quaternions, and Double Groups. Courier Corporation.
- Ambellan, F., Tack, A., Ehlke, M., Zachow, S., 2019. Automated segmentation of knee bone and cartilage combining statistical shape knowledge and convolutional neural networks. *Med. Image Anal.* 52, 109–118.
- Ambellan, F., Zachow, S., von Tycowicz, C., 2019. An as-invariant-as-possible GL+(3)-based statistical shape model. In: *Proc. 7th MICCAI Workshop on Mathematical Foundations of Computational Anatomy*, vol. 11846, pp. 219–228.
- Ambellan, F., Zachow, S., von Tycowicz, C., 2019. A surface-theoretic approach for statistical shape modeling. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 21–29.
- Arsigny, V., Commowick, O., Ayache, N., Pennec, X., 2009. A fast and log-euclidean polyaffine framework for locally linear registration. *J. Math. Imaging Vis.* 33 (2), 222–238.
- Arsigny, V., Fillard, P., Pennec, X., Ayache, N., 2006. Log-Euclidean metrics for fast and simple calculus on diffusion tensors. *Magn. Reson. Med.* 56 (2), 411–421.
- Arsigny, V., Pennec, X., Ayache, N., 2003. Polyrigid and polyaffine transformations: a new class of diffeomorphisms for locally rigid or affine registration. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 829–837.
- Bogo, F., Romero, J., Loper, M., Black, M.J., 2014. FAUST: dataset and evaluation for 3D mesh registration. *CVPR. IEEE*.
- Bonner-Jackson, A., Mahmoud, S., Miller, J., Banks, S.J., 2015. Verbal and non-verbal memory and hippocampal volumes in a memory clinic population. *Alzheimer's Res. Ther.* 7 (1), 61.
- Botsch, M., Sumner, R., Pauly, M., Gross, M., 2006. Deformation transfer for detail-preserving surface editing. In: *VMV*, pp. 357–364.
- Brandt, C., von Tycowicz, C., Hildebrandt, K., 2016. Geometric flows of curves in shape space for processing motion of deformable objects. *Comput. Graph. Forum* 35 (2).
- do Carmo, M.P., 1976. *Differential Geometry of Curves and Surfaces*. Prentice-Hall.
- Charpiat, G., Faugeras, O., Keriven, R., Maurel, P., 2006. Distance-based shape statistics. In: *ICASSP. IEEE*, pp. V925–V928.
- Ciarlet, P.G., 2005. An introduction to differential geometry with applications to elasticity. *J. Elast.* 78 (1), 1–215.
- Coates, T.F., Taylor, C.J., Cooper, D.H., Graham, J., 1995. Active shape models-their training and application. *Comput. Vis. Image Underst.* 61 (1), 38–59.
- Corman, E., Solomon, J., Ben-Chen, M., Guibas, L., Ovsjanikov, M., 2017. Functional characterization of intrinsic and extrinsic geometry. *ACM Trans. Graph.* 36 (2), 1–17.
- Davies, R., Twining, C., Taylor, C., 2008. *Statistical Models of Shape: Optimisation and Evaluation*. Springer.
- Davis, B.C., Fletcher, P.T., Bullitt, E., Joshi, S., 2010. Population shape regression from random design data. *Int. J. Comput. Vis.* 90 (2), 255–266.
- Durrleman, S., Prastawa, M., Charon, N., Korenberg, J.R., Joshi, S., Gerig, G., Trounev, A., 2014. Morphometry of anatomical shape complexes with dense deformations and sparse parameters. *Neuroimage* 101 (0), 35–49.
- Fletcher, P., Lu, C., Pizer, S., Joshi, S., 2004. Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Trans. Med. Imaging* 23 (8), 995–1005.
- Ezuz, D., Ben-Chen, M., 2017. Deblurring and denoising of maps between shapes. In: *Comput. Graph. Forum*, vol. 36. Wiley Online Library, pp. 165–174.
- Fletcher, P.T., Lu, C., Joshi, S., 2003. Statistics of shape via principal geodesic analysis on lie groups. In: *CVPR*, vol. 1. IEEE, pp. 1–95.
- Freifeld, O., Black, M.J., 2012. Lie bodies: a manifold representation of 3d human shape. In: *ECCV*. Springer, pp. 1–14.
- Fuchs, M., Jüttler, B., Scherzer, O., Yang, H., 2009. Shape metrics based on elastic deformations. *J. Math. Imaging Vis.* 35 (1), 86–102.
- Gao, L., Lai, Y.-K., Liang, D., Chen, S.-Y., Xia, S., 2016. Efficient and flexible deformation representation for data-driven surface modeling. *ACM Trans. Graph.* 35 (5), 158.
- Gilles, B., Bousquet, G., Faure, F., Pai, D.K., 2011. Frame-based elastic models. *ACM Trans. Graph.* 30 (2), 1–12.
- Goh, A., Lenglet, C., Thompson, P.M., Vidal, R., 2011. A nonparametric Riemannian framework for processing high angular resolution diffusion images and its applications to ODF-based morphometry. *Neuroimage* 56 (3), 1181–1201.
- Hanocka, R., Hertz, A., Fish, N., Giryas, R., Fleishman, S., Cohen-Or, D., 2019. MeshCNN: a network with an edge. *ACM Trans. Graph.* 38 (4), 1–12.
- Hasler, N., Stoll, C., Sunkel, M., Rosenhahn, B., Seidel, H.-P., 2009. A statistical model of human pose and body shape. *Comput. Graph. Forum* 28 (2), 337–346.
- Heeren, B., Zhang, C., Rumpf, M., Smith, W., 2018. Principal geodesic analysis in the space of discrete shells. *Comput. Graph. Forum* 37 (5), 173–184.
- Heimann, T., Meinzer, H.-P., 2009. Statistical shape models for 3D medical image segmentation: a review. *Med. Image Anal.* 13 (4), 543–563.
- Joshi, S., Davis, B., Jomier, M., Gerig, G., 2004. Unbiased diffeomorphic atlas construction for computational anatomy. *Neuroimage* 23, S151–S160.
- Kellgren, J., Lawrence, J., 1957. Radiological assessment of osteo-arthritis. *Ann. Rheum. Dis.* 16 (4), 494.
- Kingma, D. P., Ba, J., 2014. Adam: a method for stochastic optimization. *arXiv:1412.6980*.
- Kircher, S., Garland, M., 2008. Free-form motion processing. *ACM Trans. Graph.* 27 (2), 12.

¹¹ <https://doi.org/10.12752/4.ATEZ.1.0>.

¹² adni.loni.usc.edu.

- Köhler, S., Black, S., Sinden, M., Szekely, C., Kidron, D., Parker, J., Foster, J., Moscovitch, M., Wincour, G., Szalai, J., et al., 1998. Memory impairments associated with hippocampal versus parahippocampal-gyrus atrophy: an MR volumetry study in Alzheimer's disease. *Neuropsychologia* 36 (9), 901–914.
- Kreiser, J., Meuschke, M., Mistelbauer, G., Preim, B., Ropinski, T., 2018. A survey of flattening-based medical visualization techniques. In: *Comput. Graph. Forum*, vol. 37. Wiley Online Library, pp. 597–624.
- McLeod, K., Sermesant, M., Beerbaum, P., Pennec, X., 2015. Spatio-temporal tensor decomposition of a polyaffine motion model for a better analysis of pathological left ventricular dynamics. *IEEE Trans. Med. Imaging* 34 (7), 1562–1575.
- Pennec, X., 2006. Intrinsic statistics on Riemannian manifolds: basic tools for geometric measurements. *J. Math. Imaging Vis.* 25 (1), 127.
- Pennec, X., 2020. Advances in geometric statistics for manifold dimension reduction. In: *Handbook of Variational Methods for Nonlinear Geometric Data*. Springer, pp. 339–359.
- Pennec, X., Lorenzi, M., 2020. Beyond Riemannian geometry: the affine connection setting for transformation groups. In: *Riemannian Geometric Statistics in Medical Image Analysis*. Elsevier, pp. 169–229.
- Rumpf, M., Wirth, B., 2011. An elasticity-based covariance analysis of shapes. *Int. J. Comput. Vis.* 92 (3), 281–295.
- Rustamov, R.M., Ovsjanikov, M., Azencot, O., Ben-Chen, M., Chazal, F., Guibas, L., 2013. Map-based exploration of intrinsic shape differences and variability. *ACM Trans. Graph.* 32 (4), 1–12.
- Seo, D., Ho, J., Vemuri, B.C., 2016. Covariant image representation with applications to classification problems in medical imaging. *Int. J. Comput. Vis.* 116 (2), 190–209.
- de Toledo-Morrell, L., Dickerson, B., Sullivan, M., Spanovic, C., Wilson, R., Bennett, D., 2000. Hemispheric differences in hippocampal volume predict verbal and spatial memory performance in patients with Alzheimer's disease. *Hippocampus* 10 (2), 136–142.
- von Tycowicz, C., 2020. Towards shape-based knee osteoarthritis classification using graph convolutional networks. In: *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*. IEEE, pp. 750–753.
- von Tycowicz, C., Ambellan, F., Mukhopadhyay, A., Zachow, S., 2018. An efficient Riemannian statistical shape model using differential coordinates. *Med. Image Anal.* 43, 1–9.
- von Tycowicz, C., Schulz, C., Seidel, H.-P., Hildebrandt, K., 2015. Real-time nonlinear shape interpolation. *ACM Trans. Graph.* 34 (3), 34:1–34:10.
- Wang, Y., Ben-Chen, M., Polterovich, I., Solomon, J., 2018. Steklov spectral geometry for extrinsic shape analysis. *ACM Trans. Graph.* 38 (1), 1–21.
- Younes, L., 2010. *Shapes and Diffeomorphisms*, vol. 171. Springer.
- Zhang, C., Heeren, B., Rumpf, M., Smith, W.A., 2015. Shell PCA: statistical shape modelling in shell space. In: *ICCV*. IEEE, pp. 1671–1679.