Zuse Institute Berlin

Takustr. 7
14195 Berlin
Germany

CHRISTOPH VON TYCOWICZ[1], FELIX AMBELLAN[2],
ANIRBAN MUKHOPADHYAY[3] STEFAN ZACHOW[4]

An Efficient Riemannian Statistical Shape Model using Differential Coordinates[5]

[1] 0000-0002-1447-4069, equal contribution
[2] 0000-0001-9415-0859, equal contribution, corresponding author
[3] 0000-0003-0669-4018
[4] 0000-0001-7964-3049
[5] to appear in: Medical Image Analysis

{vontycowicz, ambellan, mukhopadhyay, zachow}@zib.de

ZIB Report 16-69 (November 2016)

Zuse Institute Berlin
Takustr. 7
14195 Berlin
Germany

Telephone: +49 30-84185-0
Telefax: +49 30-84185-125

E-mail: bibliothek@zib.de
URL: http://www.zib.de

ZIB-Report (Print) ISSN 1438-0064
ZIB-Report (Internet) ISSN 2192-7782

An Efficient Riemannian Statistical Shape Model using Differential Coordinates

With Application to the Classification of Data from the Osteoarthritis Initiative

Christoph von Tycowicz[a,1], Felix Ambellan[a,1,*], Anirban Mukhopadhyay[b], Stefan Zachow[a]

[a] *Therapy Planning Group, Zuse Institute Berlin, Berlin, Germany*
[b] *Interactive Graphics Systems Group, Technische Universität Darmstadt, Darmstadt, Germany*

Abstract

We propose a novel Riemannian framework for statistical analysis of shapes that is able to account for the nonlinearity in shape variation. By adopting a physical perspective, we introduce a differential representation that puts the local geometric variability into focus. We model these differential coordinates as elements of a Lie group thereby endowing our shape space with a non-Euclidean structure. A key advantage of our framework is that statistics in a manifold shape space becomes numerically tractable improving performance by several orders of magnitude over state-of-the-art. We show that our Riemannian model is well suited for the identification of intra-population variability as well as inter-population differences. In particular, we demonstrate the superiority of the proposed model in experiments on specificity and generalization ability. We further derive a statistical shape descriptor that outperforms the standard Euclidean approach in terms of shape-based classification of morphological disorders.

Keywords: Statistical shape analysis, Classification, Riemannian metrics, Principal geodesic analysis, Manifold valued statistics

1. Introduction

Statistical models of shape have been established as one of the most successful methods for understanding the geometric variability of anatomical structures. Shape modeling is of particular interest in image guided diagnosis where morphological changes of anatomies have been hypothesized to be linked to various disorders. Based on a set of training shapes, statistical shape models efficiently parametrize the geometric variability of the biological objects under study. This in turn is not only useful in imposing shape constraints in synthesis and analysis problems but also in understanding the processes behind growth and disease. Hence, statistical shape modeling has moved well beyond its de-facto application of automatic image segmentation (Cootes et al., 1995; Seim et al., 2010; Kainmüller et al., 2013; Saito et al., 2016). For example, descriptors based on statistical shape modeling have proven effective for predicting the onset and progression of osteoarthritis (Bredbenner et al., 2010; Neogi et al., 2013; Thomson et al., 2015, 2016) and in anthropometric studies (Robinette et al., 1999; Hasler et al., 2009). An overview of the numerous contributions as well as the wide range of applications can be found in (Heimann and Meinzer, 2009; Sarkalkan et al., 2014; Brunton et al., 2014; Lamecker and Zachow, 2016).

In the simplest case, the space of shapes is represented by an Euclidean vector space, e.g. by describing a shape ranging from a sparse set of landmarks up to a dense collection of boundary points. This structure then allows linear statistical tools like principal component analysis (PCA) to be applied to the training data resulting in fast and simple algorithms for the analysis and synthesis of shapes. However, the quality of the model depends on the validity of the assumption that instances of the object class being modeled lie in (or are well-approximated by) a flat Euclidean space. This assumption is a poor choice for data with a large spread or within regions of high curvature in shape space (Huckemann et al., 2010; Zhang et al., 2015) and, thus, is considered a limiting factor for the ability to represent natural biological variability in populations (cf. for example (Davis et al., 2010) and the references therein).

More recently, many exciting ideas and new theoretical insights for treating manifold valued data have been presented addressing the growing demand for nonlinear multivariate statistics accounting for the spaces' intrinsic structure (see Sec. 1.1). In spite of these developments, linear methods are still the most widely used approaches to 3D statistical shape modeling (Sarkalkan et al., 2014). This hesitant adoption of nonlinear methods has been linked to the lack of numerical robustness (Heimann and Meinzer, 2009). Moreover, while the development of

*Corresponding author
Email addresses: vontycowicz@zib.de (Christoph von Tycowicz), ambellan@zib.de (Felix Ambellan), anirban.mukhopadhyay@gris.tu-darmstadt.de (Anirban Mukhopadhyay), zachow@zib.de (Stefan Zachow)
[1] Equal contribution

specific fast-converging algorithms for the underlying optimization problems is still an ongoing field of research, many medical applications such as computer-assisted intervention or interactive segmentation require near-realtime response rates which prohibits the use of computationally expensive, nonlinear models.

In this work, we address the challenge of developing a statistical shape model that accounts for the non-Euclidean nature inherent to (anatomical) shapes and at the same time offers fast and numerically robust processing. This work makes three major contributions:

First, we formulate a continuous and physically motivated notion of shape space based on deformation gradients. Other than Euclidean structures that pronounce shape variations with large displacements even if they are near-isometries, our differential representation puts the (local) geometric changes into focus, which is mechanically sound. Furthermore, the continuous formulation allows for discretization that is consistent w.r.t. resampling and is generalizable to other representations of shape.

Second, we endow the differential coordinates with a non-Euclidean structure that comes with both: Excellent theoretical properties and closed-form expressions yielding simple and efficient algorithms. In particular, the structure stems from a product Lie group of stretches and rotations for which we present a bi-invariant metric. We demonstrate that our model is better able to capture nonlinear deformations, e.g. inherent to articulation and pathological morphology. Specifically, we perform a quantitative comparison to state-of-the-art approaches on datasets of distal femora and human body shapes showing superior performance in terms of specificity and generalization ability.

Third, we derive a statistical shape descriptor that captures systematic differences in shape between normal and diseased subpopulations. We show that these descriptors are better suited for classification tasks in qualitative as well as quantitative evaluation for osteoarthritis of the distal femur.

1.1. Related work

In medical image processing, the construction and application of statistical models of shape in a linear vector space is a standard technique with long tradition (Cootes et al., 1995; Kelemen et al., 1999). It is still an active area of research how this technique can be generalized to shape spaces with non-Euclidean structures. One line of work to introduce nonlinearity are so called kernel methods that embed data into a potentially infinite-dimensional feature space and perform linear statistic analysis there (Kirschner et al., 2011). While this approach is widely applicable, the choice of a kernel is however non-trivial and does usually depend on the specific application. Furthermore, there exists in general no (exact) pre-image in shape space so that one can only settle for approximate reconstructions thereof (Mika et al., 1999). This in turn can impede the synthesis of new shapes considerably. A comparison between different kernel methods is given in (Rathi et al., 2006).

More sophisticated methods to account for the non-Euclidean structure adopt wider geometrical as well as physical perspectives on shape spaces, e.g. based on the notion of the Hausdorff distance (Charpiat et al., 2006), elastic deformations (Rumpf and Wirth, 2011a; von Tycowicz et al., 2015; Zhang et al., 2015), and viscous flows (Fuchs et al., 2009; Heeren et al., 2012; Brandt et al., 2016). For a discussion of the various concepts we refer to the chapter by Rumpf and Wirth (2011b). The price to pay for such variational approaches is that one has to solve nonlinear, high-dimensional optimization problems that suffer from numerical instability and local minima.

For the case of Riemannian shape spaces the strength of shape variations can be measured in terms of geodesic distances. Averages in these spaces have been presented by Fréchet (1948) and were further analysed by Karcher (1977). There exists also a natural linear representation of shapes in the tangent space of the mean via the geodesic logarithmic map. Fletcher et al. (2004) employ this structure to extend the PCA to the manifold setting which is thus referred to as principal geodesic analysis (PGA). While intrinsic distances between training shapes and the mean are preserved under the logarithmic map this is generally not the case for the distances between shapes. *Exact PGA* (Sommer et al., 2010) as well as *geodesic PCA* (Huckemann et al., 2010) both attempt to capture this additional information by employing the true intrinsic distances. In general, geodesic calculus on manifold shape spaces is hard to carry out in practice because most operations do not admit closed-form solutions.

Another line of work, closely related to our framework, alleviates these challenges by explicitly modeling shapes as a collection of primitives equipped with a Lie group structure (Fletcher et al., 2003; Freifeld and Black, 2012; Hefny et al., 2015). However, contrary to our continuous model, these methods are posed in a purely discrete setting that is not readily applicable to different shape representations. Furthermore, the lack of a continuous counterpart leaves questions as for the consistency and convergence of the discrete model unanswered. These aspects are not only of concern in the refinement limit: Results will depend on the resolution and anisotropy of the sampling.

2. Differential Coordinates

In the following, we regard shapes as boundaries of physical objects belonging to a particular class of anatomical structures so that they can be represented as a collection of (orientation-preserving) embeddings $\phi_1, \dots, \phi_n \in \mathcal{C}$ of a common reference $\mathcal{M} \subset \mathbb{R}^3$, where $\mathcal{C}$ denotes the space of (smooth enough) deformations $\phi : \mathcal{M} \to \mathbb{R}^3$. While some applications require the interior of the object, e.g. joint shape and appearance analysis, typically only the object's surface is of importance. For now, we let $\mathcal{M}$ denote

the entire volume of the object and defer the extension to the 2-dimensional case to Sec. 4.1.

In order to derive a notion of shape space, $\mathcal{C}$ is typically projected onto a non-trivial space that is obtained by quotening out rigid motions or reparametrizations. However, the analysis of the geometry of such spaces and the development of practically applicable algorithms is still an active area of research (Bauer et al., 2014). An alternative approach, that we employ here, is to shift the alignment and correspondence problem to a pre-process. This also allows for auxiliary information like appearances to be exploited, improving registration accuracy and robustness (see e.g. Bogo et al., 2014).

In this light, we derive a physically motivated shape representation. Working directly with the Euclidean coordinates given by the ϕ_i introduces confounding information into the analysis as they are subject to overall changes in shape. For example, parts of an object might move rigidly leading to large displacements, albeit without local change in shape. To address this aspect, we propose a representation that puts the local differential properties into focus. In particular, as we are interested in a representation that describes a deformation ϕ within a local neighborhood of a point $x \in \mathcal{M}$, it is reasonable to approximate ϕ by its first-order Taylor expansion: $\phi(x') \approx \phi(x) + \nabla\phi|_x (x' - x)$. Here, $\nabla\phi|_x$ is the 3×3 matrix of partial derivatives of ϕ at x and is called the deformation gradient. As the constant term (i.e. the Euclidean coordinate) in the expansion corresponds to translations and hence causes no changes in shape, the gradient $\nabla\phi|_x$ furnishes a complete description of the shape changes within an infinitesimal neighborhood. This property is utilized, e.g., in continuum mechanics where the strain energy density of a homogenous hyperelastic material is a function of $\nabla\phi|_x$ alone. For these reasons, we propose to perform the analysis on the deformation gradients $\nabla\phi$. We therefore represent the shapes given by the ϕ_i in the space $\mathcal{F}$ of (regular enough) functions $\xi : \mathcal{M} \to \mathrm{GL}^+(3)$ with $\xi_i(x) = \nabla\phi_i|_x$ and $\mathrm{GL}^+(3)$ the group of 3×3 matrices with positive determinant.

Employing gradient fields as differential representation has the additional advantage that the inverse problem of mapping differential coordinates $\xi \in \mathcal{F}$ to configurations $\pi(\xi) \in \mathcal{C}$ is related to the Helmholtz-Hodge decomposition for which existence and uniqueness of solutions is well analyzed (see e.g. Bhatia et al., 2013). In particular, the map π can be formulated as a variational problem

$$\pi(\xi) := \operatorname*{arg\,min}_{\phi \in \mathcal{C}} \frac{1}{2} \int_{\mathcal{M}} ||\nabla\phi - \xi||^2 \, \mathrm{d}V, \tag{1}$$

where $\mathrm{d}V$ is the volume form on $\mathcal{M}$. The solutions are unique up to additive constants for each connected component of $\mathcal{M}$ that correspond to translations. We remove the additive constants by fixing the center of gravity of each component. Note that, due to these constraints, shapes will be recovered exactly, i.e. $\pi(\nabla\phi) = \phi$. The Euler-Lagrange equation of the variational problem (1) is the well-known Poisson equation

$$\Delta\phi = \nabla \cdot \xi, \tag{2}$$

where Δ and $\nabla\cdot$ denote the Laplacian and divergence operator, respectively. We emphasize that (2) is a linear differential equation that can be solved efficiently using e.g. multigrid methods for which computational cost scales linearly with the number of variables. This facilitates fast processing even for large-scale problems. Furthermore, as a global variational approach, the minimizer given by the Poisson equation (2) tends to distribute errors uniformly such that local gradient field inconsistencies are attenuated (see Fig. 1 left for mappings from $\mathcal{F}$ to $\mathcal{C}$). For additional information and description we refer to (Yu et al., 2004).

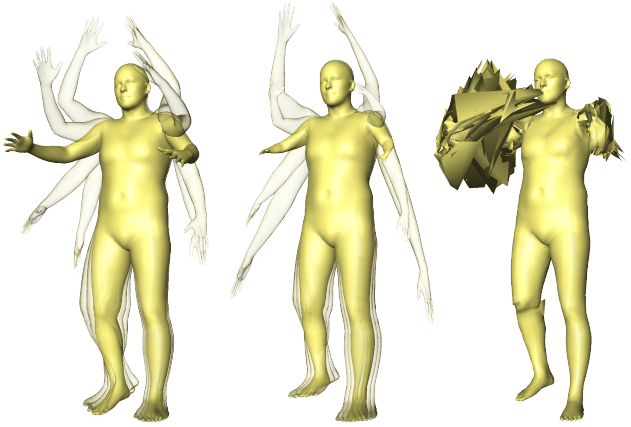

Figure 1: Mean shapes (opaque) and first principal geodesic curve (transparent at ±0.75 and ±1.5 standard deviations) for the FAUST dataset obtained using the proposed method (left), the point distribution model (center), and the SE(3)-based model* (Hefny et al., 2015) (right). ($*$ curve not shown due to degeneracies)

3. Statistical framework

We now develop the methodology for statistical analysis of the differential representation presented in the previous section. To this end, we derive the *intrinsic mean* (Fréchet, 1948) of a set of points in the space of differential coordinates $\mathcal{F}$ and then apply *principal geodesic analysis* (Fletcher et al., 2004) to extract the dominant modes of variation. This section assumes familiarity with Lie groups and we refer to (O'Neill, 1983) for an introduction. Additionally, we provide detailed C++ code in the supplemental material.

3.1. Intrinsic mean

Let $d_G(\cdot,\cdot)$ be a distance for 3×3 matrices with positive determinant, we define the distance in $\mathcal{F}$ as

$$d_{\mathcal{F}}(\xi, \xi') = \left(\int_{\mathcal{M}} d_G(\xi(x), \xi'(x))^2 \, \mathrm{d}V \right)^{\frac{1}{2}}. \tag{3}$$

The intrinsic or Fréchet mean μ of a collection of n points $\xi_1, \ldots, \xi_n \in \mathcal{F}$ is then defined as the minimizer of the functional $\sum_{i=1}^{n} d_{\mathcal{F}}(\xi, \xi_i)^2$, i.e. the sum-of-squared distances to each point. This functional can also be written as $\int_{\mathcal{M}} \mathcal{L}(\xi(x))\,\mathrm{d}V$, where $\mathcal{L}(\cdot) = \sum_{i=1}^{n} d_G(\cdot, \xi_i)^2$ is the sum-of-squared point-wise distances. Due to its simple form, the Euler-Lagrange equations reduce to $\partial\mathcal{L}/\partial\xi_{qp} = 0$ (with natural boundary conditions trivially satisfied) and, hence, μ can be computed point-wise, i.e. $\mu(x)$ depends only on the gradients $\xi_i(x)$ and the choice of d_G. This local characterization of μ is an important property as it reduces the discrete high-dimensional problem to a system of independent, low-dimensional ones that can be handled much more efficiently.

To compute the local intrinsic mean, we derive a distance d_G that is tailored to our differential coordinates. Let $F = \xi(x)$ be the deformation gradient at $x \in \mathcal{M}$. By the polar decomposition from linear algebra, F can be uniquely factored as $F = RU$, where R is a special orthogonal matrix and U—called the *right Biot stretch tensor*—is symmetric positive-definite. Thus, a deformation is locally given to first order by a stretching with U followed by a rotation R. With this in mind, we interpret the gradients as elements in set

$$\mathrm{G} := \mathrm{SO}(3) \times \mathrm{Sym}^+(3), \tag{4}$$

where SO(3) is the Lie group of rotations in $\mathbb{R}^3$ (with Lie algebra $\mathfrak{so}(3)$) and $\mathrm{Sym}^+(3)$ is the space of 3×3 symmetric and positive-definite matrices.

While there exists a well-developed theory for SO(3) with a natural notion of distance, there is no canonical setup for stretch tensors in $\mathrm{Sym}^+(3)$. Resorting to Euclidean operations on symmetric matrices without taking positive-definiteness into account causes defects: The determinant of the Euclidean mean (and thus the volume form) can be larger than those of the original stretch tensors. Furthermore, nonconvex combinations typically yield indefinite stretch tensors corresponding to physically invalid configurations.

To circumvent these shortcomings, we employ the Log-Euclidean approach (Arsigny et al., 2006) that turns $\mathrm{Sym}^+(3)$ into a commutative Lie group (with Lie algebra $\mathfrak{sym}^+(3)$) based on the logarithmic multiplication defined as

$$U \circ U' := \exp(\log(U) + \log(U')). \tag{5}$$

Here, exp and log are the standard matrix exponential and logarithmic map, respectively. Remarkably, bi-invariant metrics exist for this structure and are particularly simple. In particular, we endow $\mathfrak{sym}^+(3)$ with the standard inner product for quadratic matrices

$$\langle u, u' \rangle_{\mathfrak{sym}^+(3)} := \langle u, u' \rangle = \mathrm{trace}(u^T u'), \tag{6}$$

inducing the distance on $\mathrm{Sym}^+(3)$

$$d_{\mathrm{Sym}^+(3)}(U, U') = \left\| \log(U') - \log(U) \right\|_F, \tag{7}$$

where $||\cdot||_F$ denotes the Frobenius norm. This distance is invariant with respect to a change of coordinates obtained by orthogonal transformation and scaling (similarity transforms).

As a product of Lie groups, G is itself a Lie group and we can define the distance d_G as the sum

$$d_G = \omega d_{\mathrm{SO}(3)} + d_{\mathrm{Sym}^+(3)}. \tag{8}$$

Here, ω is a positive weighting factor (see also Sec. 3.2) and $d_{\mathrm{SO}(3)}$ denotes the canonical distance on SO(3)

$$d_{\mathrm{SO}(3)}(R, R') = \left\| \log(R^T R') \right\|_F \tag{9}$$

that is, too, induced by the standard inner product $\langle \cdot, \cdot \rangle$ on $\mathfrak{so}(3)$. Indeed, d_G is also bi-invariant and thus comes with strong theoretical properties:

- Geodesic completeness of G, i.e. invalid tensors are at an infinite distance from any element in G,
- the mean is completely consistent with algebraic operations (left/right composition and inversion), and
- the Lie group and Riemannian exponential coincide yielding particularly simple expressions for geodesics.

Furthermore, the group exponential can be efficiently calculated due to the nilpotency and diagonalizability of elements in $\mathfrak{so}(3)$ and $\mathfrak{sym}^+(3)$, respectively. In contrast to numerical approximation schemes, this provides improved performance and robustness necessary in practical applications, e.g., in medical image processing.

3.2. Principal geodesic analysis

As in (Fletcher et al., 2004), we perform a principal component analysis in the tangent space $T_\mu\mathcal{F}$ at the intrinsic mean to approximate the principal geodesic submanifold that best represents the variability of the shapes. Let

$$g(v, v') = \int_{\mathcal{M}} \langle v(x), v'(x) \rangle_{\mu(x)}\, \mathrm{d}V \tag{10}$$

be the inner product on $T_\mu\mathcal{F}$, where $\langle \cdot, \cdot \rangle_{\mu(x)}$ denotes the inner product on $T_{\mu(x)}G$ obtained by left translation of $\langle \cdot, \cdot \rangle$ on $\mathfrak{so}(3)$ and $\mathfrak{sym}^+(3)$. Moreover, in line with the definition of d_G in (8), the rotational component in $\langle \cdot, \cdot \rangle_{\mu(x)}$ is weighted by ω. While the mean is invariant w.r.t. ω, this weighting does effect the modes of variation. In particular, as the rotational and stretching components are not in the same units, ω provides a mean to commensurate both terms. A data-driven approach is to determine ω s.t. the range of both terms agrees within a population under study. Alternatively, following the paradigms from nonlinear elasticity all change in shape is encoded in the stretch tensor, which suggests values of ω close to zero. However, for surfaces the rotations describe how normals

transform and thus encode curvature information. For implementation details see the supplemental material. Using g, we can define a covariance operator for deviations from the intrinsic mean as

$$C(\cdot) = \tfrac{1}{n} \sum_{i=1}^{n} g(\cdot, \upsilon_i)\upsilon_i, \tag{11}$$

where $\upsilon_i = \mathrm{Log}_\mu(\xi_i)$ is the representation of the i-th shape in $T_\mu\mathcal{F}$. Note that (by the same variational approach as for the mean) the geodesic logarithmic mapping $\mathrm{Log}_\xi(\cdot)$ can be computed point-wise and leads to efficient, closed-form expressions entailing the matrix logarithm in G.

C is positive-definite on $\mathrm{span}(\upsilon_1, \ldots, \upsilon_n)$ and diagonalizing it on this finite-dimensional space yields a set of g-orthogonal eigenfunctions ψ_i and -values λ_i with $C(\psi_i) = \lambda_i\psi_i$. Accordingly, the *principal geodesic curves* given by the exponential mappings (see Fig. 1) of the ψ_i approximate the main modes of variation with the λ_i encoding the corresponding variances.

4. Numerical Treatment

In this section, we present a consistent and convergent finite element based discretization of our framework for volumetric objects together with a dimension reduction to surfaces. In addition, we provide details on the algorithmic aspects involved in the construction of our shape model.

4.1. Discretization

In the discrete setting, we consider a d-dimensional simplicial manifold $\mathcal{M}^h$ as reference shape, i.e. a triangular ($d = 2$) or tetrahedral ($d = 3$) mesh with vertices $\{\sigma_1^0, \ldots, \sigma_{n_0}^0\}$. Deformations ϕ^h of $\mathcal{M}^h$ are then represented as piecewise affine coordinate functions defined by barycentric interpolation of vertices $\sigma_j^0 \in \mathbb{R}^3$, i.e.

$$\phi^h(x) = \sum_{j=1}^{n_0} \varphi_j(x)\, \sigma_j^0. \tag{12}$$

Here, the φ_j denote the Lagrange basis functions that are linear within each simplex and satisfy $\varphi_j(\sigma_k^0) = \delta_{jk}$. Each deformation ϕ^h is thus uniquely determined by its coefficients $\boldsymbol{\phi} = (\sigma_1^0 \ldots \sigma_{n_0}^0)^T \in \mathbb{R}^{n_0 \times 3}$. The gradient is then

$$\nabla\phi^h|_x = \sum_{j=1}^{n_0} \sigma_j^0\, \nabla\varphi_j(x)^T \tag{13}$$

and yields a constant 3×3 matrix $F_k = \nabla\phi^h|_{\sigma_k^d}$ for each d-simplex $\sigma_k^d \in \mathcal{M}^h$. Note, however, that the deformation of a triangle does not fully specify an affine map of $\mathbb{R}^3$ so that the deformation gradient is only defined within its tangent plane. Nevertheless, Botsch et al. (2006) showed that employing $\nabla\phi^h|_x$ as deformation gradient yields an efficient dimension reduction with a natural surface-based formulation of the Poisson equation. This approach is easily combined with our statistical framework as rotations can still be uniquely determined (taking normals into account) and the Lagrangian view of the stretches allows for a simple restriction of the analysis to the tangent plane.

Now, the computation of the gradients F_k can be written with a $3n_d \times n_0$ matrix D, where n_d is either the number of triangles ($d = 2$) or tets ($d = 3$) , and

$$D_{ij} = \left(\nabla\varphi_j|_{\sigma^d_{\lfloor i/3 \rfloor}}\right)_{i \bmod 3} \tag{14}$$

such that $D\boldsymbol{\phi} = (F_1 \ldots F_{n_d})^T$. The discrete version of the Poisson equation (2) is then

$$D^T M D\boldsymbol{\phi} = D^T M\boldsymbol{\xi}, \tag{15}$$

where M is a diagonal matrix containing the triangle areas or tet volumes and $\boldsymbol{\xi}$ is the $3n_d \times 3$ matrix of coefficients determining the piecewise constant matrix-field of differential coordinates. The discrete problem results in a sparse, linear system of equations with positive semi-definite system matrix. We want to emphasize that the matrices depend only on the reference shape $\mathcal{M}^h$. The problem in Eq. (15) can thus be solved efficiently using a direct solver: We compute the factorization once during pre-processing and only perform forward- and backward substitutions with close to linear costs at run-time.

4.2. Atlas construction

In the preceding sections we derived a framework for the statistical analysis of geometric variability of a population under study. To this end, we considered the variability to be represented via embeddings of a reference shape. The reference shape therefore serves as an anatomical atlas that provides a common coordinate system. The construction of such atlases is a key task in population based medical image analysis and has, in turn, been a driving application of statistical shape modeling. A common choice of atlas is to employ either a template or a single subject's anatomy. However, in general, a single anatomy cannot faithfully represent the complex structural variability evident in a population and thus inherently introduces a bias. To avoid this systematic bias we opt for an atlas construction according to Joshi et al. (2004) that requires the anatomical atlas to agree with the mean of the training data. This approach results in an alternating optimization that updates the reference shape with the current mean shape in each iteration. This strategy typically converges rapidly s.t. we limit the number of iterations to ≤ 10 in our implementation.

5. Experiments and Results

All experiments are done employing a fixed balancing weight $\omega \equiv 1$ (see Sec. 3.2) which empirically shows the best performance in our classification experiments.

5.1. Data description

In order to validate the performance of the proposed statistical shape model we conduct experiments on two different open-access datasets.

The first dataset is derived from the Osteoarthritis Initiative (OAI), which is a longitudinal study of knee osteoarthritis maintaining (among others) clinical evaluation data and radiological images from 4,796 men and women of age 45-79 that are available for public access at `http://www.oai.ucsf.edu/`. From the baseline dataset, we chose 58 severely diseased subjects and 58 healthy subjects according to their Kellgren–Lawrence score (Kellgren and Lawrence, 1957) (see also Sec. 5.3). For the 116 subjects we extracted surfaces of the distal femora from the respective 3D weDESS MR images (0.37×0.37 mm matrix, 0.7 mm slice thickness) using (Seim et al., 2010). In a supervised post-process, the quality of segmentations as well as the correspondence of the resulting meshes (8,988 vertices) were ensured.

As second dataset we employ the FAUST (Bogo et al., 2014) dataset of human body scans that have been acquired with a high-accuracy 3D multi-stereo system. Specifically, we use the training data of FAUST that comprises 100 scans of 10 subjects of different shapes and in 10 different wide ranging poses. FAUST features full-body, high-quality correspondences established via registration that combines 3D shape and appearance. The provided FAUST meshes are watertight, genus zero surfaces with 6,890 vertices.

For both datasets, we consider the pointwise correspondences obtained during the acquisition/reconstruction process as ground truth. In particular, all experiments are performed using these original correspondences. Furthermore, we globally aligned the objects within each dataset w.r.t. position and orientation using the standard generalized Procrustes method (Goodall, 1991).

5.2. Efficiency and robustness

The goal of this experiment is to demonstrate the efficiency and numerical robustness of our proposed method in comparison to state-of-the-art nonlinear approaches. To this end, we compare our framework to the established *large deformation diffeomorphism metric mapping* (LDDMM) approach using the open-source *Deformetrica* software (Durrleman et al., 2014). Especially, we configure Deformetrica to employ correspondences via the *Landmark* object type and enable multi-threading and CUDA acceleration (see supplemental material for parameter files). Additionally, we compare to the recently presented *Shell PCA* (Zhang et al., 2015) that employs a physically-based distance $\mathcal{W}$ induced by an elastic shell deformation energy. By virtue of the shell energy, the corresponding Fréchet mean is the solution of a high-dimensional, non-convex optimization problem. In our implementation, we solve these problems using the IPOPT (Wächter and Biegler, 2006) software package, which is an established library for large-scale, nonlinear optimization. As in the Shell PCA article, we employ the *Discrete Shells* (Grinspun et al., 2003) energy.

In this evaluation, we perform first moment statistical analysis on 100 randomly sampled pairs for each of the datasets. Note that the evaluation of the mean shape is computationally the most demanding part. All experiments are run on an 8-core Intel Xeon 2.6GHz.

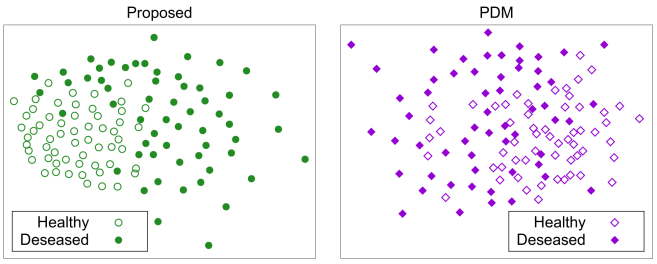

Figure 2: Sammon plots visualizing the distribution of statistical shape descriptors obtained by the proposed and the point distribution model (PDM) for the dataset of healthy and diseased distal femora from the OAI database.

For the OAI dataset, the LDDMM-based approach requires 173 seconds in average whereas our proposed framework achieves an average computation time of one second and hence a two orders of magnitude speedup. Due to numerical stiffness, only 84% of the test cases are successfully solved for the Shell PCA framework. Considering only the successful cases, our method achieves an average speedup factor of 360.

We also perform this experiment for the FAUST dataset containing large (non-linear) deformations. While our framework successfully determines all mean shapes (taking 0.5 seconds in average), the other approaches suffer from numerical instability. In particular, the LDDMM-based approach succeeds only for 94% of the cases and the Shell PCA implementation fails for all cases. Restricting the comparison to the solved cases, computing the differential coordinates based mean is 16,490 times faster than computing the LDDMM mean (taking 8,312 seconds in average).

Since IPOPT is not able to solve for the Shell PCA means, we cannot compare the timings directly. However, Zhang et al. (2015) present results for five poses from a single individual in the FAUST dataset. For remeshed (simplified) versions they report 5 minutes (on an Intel Core 3.4GHz) for computing the Shell PCA. These results suggest a speedup of similar order of magnitude for our method as observed with the OAI dataset.

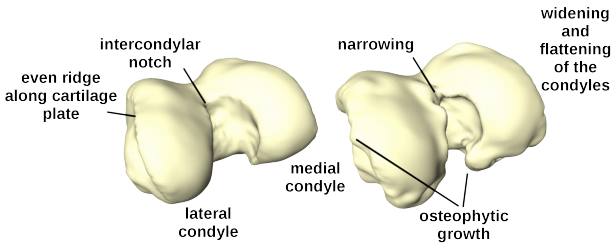

Figure 3: Healthy (left) and osteoarthritic (right) distal femur with delineated pathological changes in shape.

5.3. Classification

In this experiment we investigate the ability of the proposed model to capture characteristic changes in shape that are incident to radiographic knee osteoarthritis (OA). To this end, we perform classification of healthy and severely diseased femora according to Kellgren–Lawrence grades of 0 or 1 and 4, respectively. The Kellgren–Lawrence system (Kellgren and Lawrence, 1957) is a categorial scale of radiographic knee OA that assesses the severity of OA based on a number of radiographic features, e.g. osteophyte formation, joint space narrowing, and sclerosis. Figure 3 shows an example of a healthy as well as a diseased femur within the OAI-based dataset and illustrates a set of geometric features incident to femoral OA.

To evaluate the performance of the proposed method, we compare it to the *point distribution model* (Cootes et al., 1995) (PDM) that is frequently employed to derive systems for the prediction of the onset and progression of OA, see e.g. (Bredbenner et al., 2010; Neogi et al., 2013; Thomson et al., 2015, 2016) as well as (Sarkalkan et al., 2014) and the references therein. In literature the terms statistical shape model and PDM are often used synonymously. We differentiate here and use only PDM to explicitly refer to the particular model since it on one hand emphasizes the underlying character of working with points and on the other hand leaves statistical shape model to be used more general for *any* kind of model. As in previous works, we consider the principal component scores of the training data, i.e. the coefficients encoding the shapes within the basis of principal modes, as representative features. This results, for both models, in a 115-dimensional feature vector of shape mode coefficients for each subject. In Fig. 2 we present a visualization of the distribution of feature vectors using Sammon projection (Sammon, 1969)—a type of dimension reduction that (nonlinearly) maps the high-dimensional features onto the plane such that pairwise distances are preserved as good as possible. This first qualitative comparison already reveals an improved degree of separation for the proposed features between the non-OA and OA shapes promising superior performance as statistical shape descriptor.

The visualization of the most dominant mode shown in Fig. 5 further suggests an improved ability to capture the inter-population differences. While both models exhibit displacements around the cartilage plate—modulating the characteristic increased ridge of osteophytic growth (see also Fig. 3)—the proposed model features more localized and well-delineated shape variations in comparison to PDM.

We additionally perform a quantitative evaluation of the features in terms of classification accuracy. To this end, an SVM with linear kernel is used to classify healthy and diseased femora, where the classifier is trained from only a balanced subset of feature vectors of the training data. The remaining feature vectors are used during testing to measure the classification accuracy. In particular, the percentage of randomly sampled balanced training data w.r.t. the complete dataset is varied from 10% to 90%, and the rest is used for testing as shown in Figure 4. To address the randomness of our experimental design, we run the experiment 1000 times per data partition and plot the mean accuracy together with standard deviations. Note that, not only we are outperforming the PDM in all levels, but by only taking 10% of the data as training set, we outperform PDM by approximately 15%, which emphasizes further the ability of the proposed method to identify systematic differences between the subpopulations.

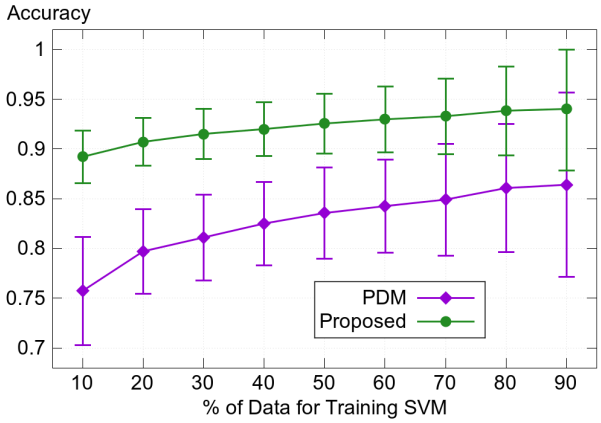

Figure 4: Classification accuracy achieved by the proposed and the point distribution model (PDM) as function of the percentage of training data.

5.4. Model evaluation

We further perform a quantitative comparison between our proposed method, the popular PDM and the recent method by Hefny et al. (2015). The latter is a non-Euclidean model for triangular meshes that is based on SE(3), i.e. the Lie group of rigid motions in $\mathbb{R}^3$. We evaluate the quality of the models in terms of specificity, generalization ability, and compactness (cf. Davies et al., 2008, for a detailed discussion of these measures). For pairwise comparisons we employ the Mann–Whitney U test and consider the results to be statistically significant for p-values less than 0.01.

Specificity is the property that a model generates only valid instances of the class(es) of objects presented in the training set, i.e. instances that are similar to the training shapes. First, we consider the model as a Gaussian distribution in shape space and randomly draw 2000 shapes according to that distribution. Specificity is then measured as the degree to which these sample shapes cluster in the vicinity of the training data by averaging the distances of each sample to its closest training shape. The lower the average the more specific is the model w.r.t. its training set. Twining and Taylor (2011) showed that this measure estimates the divergence between the unknown probability density function from which the training data was drawn and the one underlying the model. As specificity depends directly on the distance used to evaluate the closeness between the samples and the training shapes, employing the

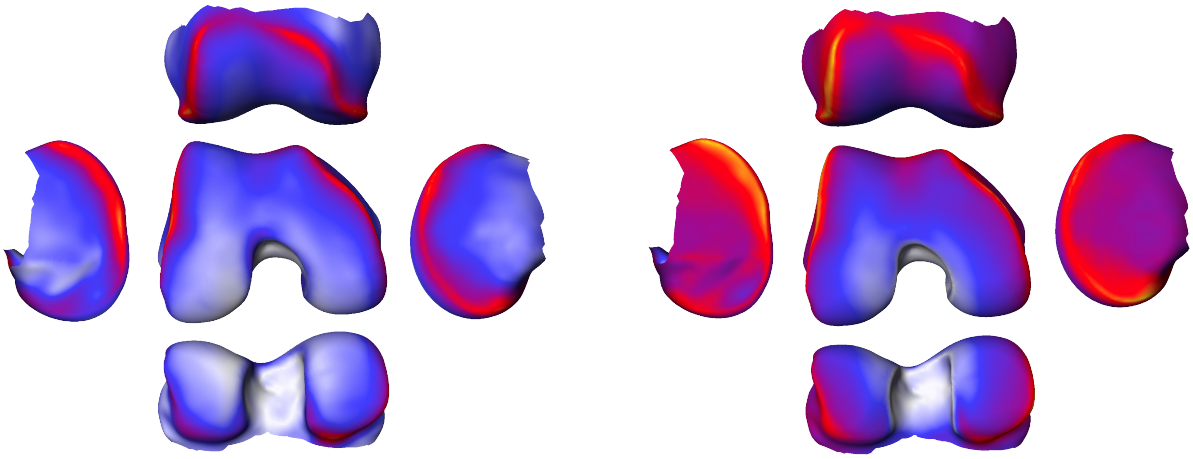

Figure 5: Magnitude of displacements under the first principal geodesic curve (±1 standard deviations) for the proposed (left) and the point distribution model (right) color-coded on the respective mean shape (0.0 mm 5.0 mm). While both curves model the characteristic increased ridge of osteophytic growth, the proposed model features more localized and well-delineated shape variations.

Table 1: Specificity

Method	FAUST			OAI		
	L^2	$d_{\mathcal{F}}$	$\mathcal{W}$	L^2	$d_{\mathcal{F}}$	$\mathcal{W}$
Proposed	0.218	**0.019**	**24.571**	**1.443**	**0.520**	**22.162**
PDM	**0.212**	0.020	31.849	1.466	0.529	23.309
SE(3)	0.272	0.049	746.056	1.506	0.714	56.575

Best results highlighted; all improvements statistically significant ($p < 0.01$).

Euclidean L^2 and the proposed $d_{\mathcal{F}}$ (cf. Eq. (3)) distance could introduce a bias for PDM and the proposed model, respectively. Therefore, we also employ the independent distance $\mathcal{W}$ (see Sec. 5.2) that was shown to successfully quantize nonlinear variations in shape (Zhang et al., 2015).

We report specificity results for both the FAUST and OAI datasets in Fig. 7 and Table 1. It can be seen that the proposed method outperforms the other approaches for distance $\mathcal{W}$ as well as $d_{\mathcal{F}}$. The L^2-based results do not allow for a clear statement: For the OAI data the proposed method is most specific while PDM performs better for the FAUST data. However, as shown in Fig. 1, the Euclidean distance promotes degenerated configurations for the FAUST data (note the short arms and small hands of the PDM mean shape).

Generalization ability measures how well a model is able to represent unseen instances of the class of objects that is modeled. We evaluate the generalization ability from the training shapes by performing leave-one-out cross validation. To this end, we construct a model from all but one training shapes and then fit it to the excluded instance. The accuracy to which the model can represent the unseen instance is measured as the distance between the instance and the reconstruction thereof. This process is repeated, excluding each of the training shapes in turn, and the accuracy is averaged over the complete set of trials. Lower values correspond to a higher ability to generalize to unseen data. As for specificity, we report generalization ability for the physically-motivated distance $\mathcal{W}$ as well as for $d_{\mathcal{F}}$ and L_2. Table 2 and Fig. 6 summarize the results. In total, the proposed method features a better generalization ability than the other approaches. More specifically, for both non-Euclidean measures the proposed method performs best, whereas the results for the Euclidean case vary.

Table 2: Generalization ability

Method	FAUST			OAI		
	L^2	$d_{\mathcal{F}}$	$\mathcal{W}$	L^2	$d_{\mathcal{F}}$	$\mathcal{W}$
Proposed	0.045	**0.005**	**20.447**	0.862	**0.338**	**16.516**
PDM	0.032	0.007	26.897	**0.468**	0.381	21.290
SE(3)	0.096	0.027	308.967	0.904	0.490	38.707

Best results highlighted, if improvement statistically significant ($p < 0.01$).

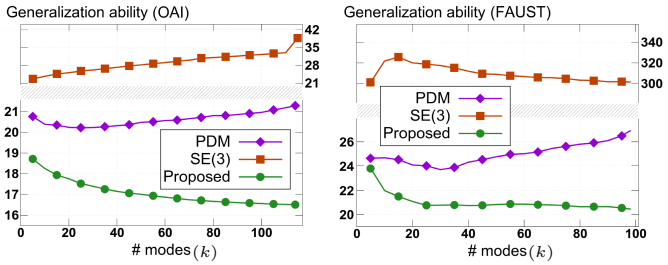

Figure 6: Generalization ability as function of the number of modes (model parameters) for $\mathcal{W}$ of the proposed shape model, point distribution model (PDM), and SE(3)-based model (Hefny et al., 2015).

The limited performance of the method by Hefny et al. (SE(3)) for both datasets in terms of specificity and generalization ability can be attributed to their purely local treatment of shape that models triangles as elements in SE(3). Unlike our formulation, this approach does not model stretches potentially leading to physically invalid configurations. Furthermore, there is no canonical way of

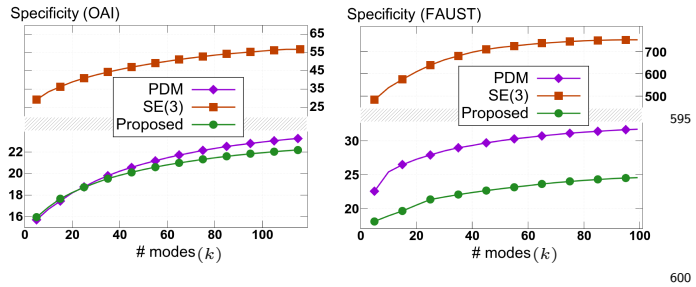

Figure 7: Specificity as function of the number of modes (model parameters) for $\mathcal{W}$ of the proposed shape model, point distribution model (PDM), and SE(3)-based model (Hefny et al., 2015).

projecting from SE(3)-based shape space to global vertex positions since each vertex is contained in several triangles that in general take different paths in SE(3). Contrary, our Poisson-based reconstruction consistently bridges among local differential properties and global effects. The inherent defects of the SE(3) model are also apparent in Fig. 1, where we provide a visualization of the mean shapes.

Compactness measures the relative variability within a training set as captured by a statistical shape model. Specifically, it is the ratio of the variance encoded by the k most dominant modes and total variance in the set, i.e. $\sum_{i=1}^{k} \lambda_i / \sum_{i=1}^{n} \lambda_i$, where the λ_i denote the variances as defined for the respective model (cf. Sec. 3.2). In Fig. 8 we plot the compactness as a function of k for each of the constructed models. For both datasets, PDM concentrates more variance on the first modes as compared to the other models. The behavior implies that variance modeled in Euclidean space is more easily captured than in non-Euclidean spaces (this is consistent with the observations for elastic variance in (Zhang et al., 2015)). However, since the variances are computed in different shape spaces, we should be cautious in how to interpret these results. For example, the most-dominant mode of the PDM and the proposed model for the FAUST data describes similar change in posture (see Fig. 1), yet account for a substantially different portion (63% and 36% respectively) of the total variance.

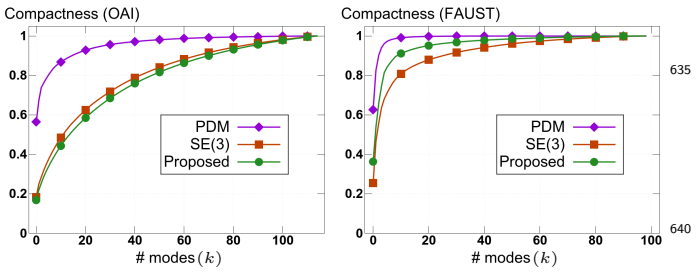

Figure 8: Compactness as function of the number of modes (model parameters) of the proposed shape model, point distribution model (PDM), and SE(3)-based model (Hefny et al., 2015).

6. Conclusion

In this work, we presented a novel approach for nonlinear statistical analysis of shapes. Our formulation employs a differential representation of shape that incorporates a natural (local) measure of deformation in terms of stretch tensors. We furthermore endow this representation with a nonlinear Riemannian structure that provides our shape space with strong theoretical properties. In particular, physically invalid tensors are at an infinite distance from any element in our space. Despite its non-Euclidean structure our shape model admits closed-form expressions for most operations and therefore facilitates fast and numerically robust processing. Indeed, this addresses one of the major shortcomings usually found in nonlinear statistical tools limiting their practical applicability.

We experimentally evaluated the performance of the proposed method and compared it to alternative approaches. Specifically, in comparison to the state-of-the-art, physically-based Shell PCA model (Zhang et al., 2015) we demonstrated a speedup of two orders of magnitude while, at the same time, being more numerically robust. We further derived a shape descriptor based on the principal component scores and showed the superiority of these features for the characterization of inter-population differences in a classifier system for radiographic knee OA. Additionally, we showed that our model is better able to capture the nonlinear variations present in articulated body pose and disease-specific data by comparing it to the standard point distribution model and the recent SE(3)-based model (Hefny et al., 2015). In particular, our proposed model showed significant improvements in specificity and generalization ability.

As with other non-Euclidean approaches, existence and uniqueness of the intrinsic mean is only ensured for 'well-localized' data. In particular, the logarithm is only well defined if the rotational part is also well defined, i.e. if the rotations have angles less than π. However, one advantage of our Poisson-based reconstruction is that it can account for global effects and thus regularizes local inconsistencies.

In the future, we plan to extend our proposed model to also account for appearances. In particular, physically invalid configurations like degenerated or inverted tetrahedra can cause appearance artifacts that will be alleviated by our approach. This is turn can lead to improved accuracy and robustness in applications like 3D reconstruction via analysis-by-synthesis. Furthermore, we want to investigate different strategies for commensuration of the rotational and stretch components in our shape metric. In particular, while the rotational component is affected by local and global rigid motions it also contains curvature information. While we expect a more explicit representation of this information to increase the performance of the shape descriptors it would also render the inverse problem (mapping to configurations in $\mathcal{C}$) numerically less tractable. Another interesting line of future work is to explore the feasibility of fully automatic computer-aided

diagnostics for knee OA based on advanced machine learning.

Acknowledgements

This work was supported by the DFG project *Knee-Laxity: Dynamic multi-modal knee joint registration* (EH 422/1-1), BMBF project *TOKMIS: Treating Osteoarthritis in Knee with Mimicked Interpositional Spacer* (01EC1406E), and BMBF MODAL - MedLab. Furthermore we are grateful for the open-access datasets OAI [2] and FAUST (Bogo et al., 2014) as well as for the open-source software Deformetrica (Durrleman et al., 2014) and IPOPT (Wächter and Biegler, 2006).

References

Arsigny, V., Fillard, P., Pennec, X., Ayache, N., 2006. Log-Euclidean metrics for fast and simple calculus on diffusion tensors. Magn Reson Med 56, 411–421.

Bauer, M., Bruveris, M., Michor, P.W., 2014. Overview of the geometries of shape spaces and diffeomorphism groups. J. Math. Imaging Vis. 50, 60–97.

Bhatia, H., Norgard, G., Pascucci, V., Bremer, P.T., 2013. The Helmholtz-Hodge decomposition—a survey. IEEE Trans Vis Comput Graph 19, 1386–1404.

Bogo, F., Romero, J., Loper, M., Black, M.J., 2014. FAUST: Dataset and evaluation for 3D mesh registration, in: CVPR, IEEE.

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.

Bredbenner, T.L., Eliason, T.D., Potter, R.S., Mason, R.L., Havill, L.M., Nicolella, D.P., 2010. Statistical shape modeling describes variation in tibia and femur surface geometry between control and incidence groups from the osteoarthritis initiative database. J Biomech 43, 1780–1786.

Brunton, A., Salazar, A., Bolkart, T., Wuhrer, S., 2014. Review of statistical shape spaces for 3D data with comparative analysis for human faces. Comput Vis Image Underst 128, 1–17.

Charpiat, G., Faugeras, O., Keriven, R., Maurel, P., 2006. Distance-based shape statistics, in: ICASSP, IEEE. pp. V925–V928.

Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J., 1995. Active shape models-their training and application. Comput Vis Image Underst 61, 38–59.

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, 255–266.

Durrleman, S., Prastawa, M., Charon, N., Korenberg, J.R., Joshi, S., Gerig, G., Trouv, A., 2014. Morphometry of anatomical shape complexes with dense deformations and sparse parameters. NeuroImage 101, 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, 995–1005.

Fletcher, P.T., Lu, C., Joshi, S., 2003. Statistics of shape via principal geodesic analysis on lie groups, in: CVPR, IEEE. pp. I–95.

Fréchet, M., 1948. Les éléments aléatoires de nature quelconque dans un espace distancié, in: Annales de l'institut Henri Poincaré, pp. 215–310.

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, 86–102.

Goodall, C., 1991. Procrustes methods in the statistical analysis of shape. J. Roy. Statist. Soc. Series B (Methodological) 53, 285–339.

Grinspun, E., Hirani, A.N., Desbrun, M., Schröder, P., 2003. Discrete Shells, in: ACM SIGGRAPH/Eurographics SCA, pp. 62–67.

Hasler, N., Stoll, C., Sunkel, M., Rosenhahn, B., Seidel, H.P., 2009. A statistical model of human pose and body shape 28, 337–346.

Heeren, B., Rumpf, M., Wardetzky, M., Wirth, B., 2012. Time-discrete geodesics in the space of shells 31, 1755–1764.

Hefny, M.S., Okada, T., Hori, M., Sato, Y., Ellis, R.E., 2015. A liver atlas using the special Euclidean group, in: MICCAI, Springer. pp. 238–245.

Heimann, T., Meinzer, H.P., 2009. Statistical shape models for 3d medical image segmentation: A review. Med Image Anal 13, 543–563.

Huckemann, S., Hotz, T., Munk, A., 2010. Intrinsic shape analysis: Geodesic PCA for Riemannian manifolds modulo isometric Lie group actions. Stat Sin , 1–58.

Joshi, S., Davis, B., Jomier, M., Gerig, G., 2004. Unbiased diffeomorphic atlas construction for computational anatomy. NeuroImage 23, S151–S160.

Kainmüller, D., Lamecker, H., Heller, M., Weber, B., Hege, H.C., Zachow, S., 2013. Omnidirectional displacements for deformable surfaces. Med Image Anal 17, 429 – 441.

Karcher, H., 1977. Riemannian center of mass and mollifier smoothing. Commun Pure Appl Math 30, 509–541.

Kelemen, A., Székely, G., Gerig, G., 1999. Elastic model-based segmentation of 3-d neuroradiological data sets. IEEE Trans Med Imaging 18, 828–839.

Kellgren, J., Lawrence, J., 1957. Radiological assessment of osteo-arthrosis. Ann Rheum Dis 16, 494.

Kirschner, M., Becker, M., Wesarg, S., 2011. 3d active shape model segmentation with nonlinear shape priors, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer Berlin Heidelberg. pp. 492–499.

Lamecker, H., Zachow, S., 2016. Statistical shape modeling of musculoskeletal structures and its applications, in: Computational Radiology for Orthopaedic Interventions. volume 23, pp. 1 – 23.

Mika, S., Schölkopf, B., Smola, A., Müller, K.R., Scholz, M., Rätsch, G., 1999. Kernel PCA and de-noising in feature spaces, in: NIPS, MIT Press, Cambridge, MA, USA. pp. 536–542.

Neogi, T., Bowes, M.A., Niu, J., Souza, K.M., Vincent, G.R., Goggins, J., Zhang, Y., Felson, D.T., 2013. Magnetic resonance imaging–based three-dimensional bone shape of the knee predicts onset of knee osteoarthritis: data from the osteoarthritis initiative. Arthritis Rheum 65, 2048–2058.

O'Neill, B., 1983. Semi-Riemannian geometry with applications to relativity. volume 103. Academic press.

Rathi, Y., Dambreville, S., Tannenbaum, A., 2006. Comparative analysis of kernel methods for statistical shape learning, in: CVAMIA, Springer. pp. 96–107.

Robinette, K.M., Daanen, H., Paquet, E., 1999. The CAESAR project: A 3-D surface anthropometry survey, in: 3DIM, pp. 380–386.

Rumpf, M., Wirth, B., 2011a. An elasticity-based covariance analysis of shapes. Int J Comput Vis 92, 281–295.

[2] 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 data set and does not necessarily reflect the opinions or views of the OAI investigators, the NIH, or the private funding partners.

Rumpf, M., Wirth, B., 2011b. Variational methods in shape analysis, in: Scherzer, O. (Ed.), Handbook of Mathematical Methods in Imaging. Springer, pp. 1363–1401.

Saito, A., Nawano, S., Shimizu, A., 2016. Joint optimization of segmentation and shape prior from level-set-based statistical shape model, and its application to the automated segmentation of abdominal organs. Medical image analysis 28, 46–65.

Sammon, J.W., 1969. A nonlinear mapping for data structure analysis. IEEE Trans Comput C-18, 401–409.

Sarkalkan, N., Weinans, H., Zadpoor, A.A., 2014. Statistical shape and appearance models of bones. Bone 60, 129–140.

Seim, H., Kainmueller, D., Lamecker, H., Bindernagel, M., Malinowski, J., Zachow, S., 2010. Model-based auto-segmentation of knee bones and cartilage in MRI data, in: v. Ginneken, B. (Ed.), MICCAI Workshop Med Image Anal for the Clinic, pp. 215 – 223.

Sommer, S., Lauze, F., Hauberg, S., Nielsen, M., 2010. Manifold valued statistics, exact principal geodesic analysis and the effect of linear approximations, in: ECCV. Springer, pp. 43–56.

Thomson, J., O'Neill, T., Felson, D., Cootes, T., 2016. Detecting Osteophytes in Radiographs of the Knee to Diagnose Osteoarthritis. Springer International Publishing. pp. 45–52.

Thomson, J., ONeill, T., Felson, D., Cootes, T., 2015. Automated shape and texture analysis for detection of osteoarthritis from radiographs of the knee, in: MICCAI, Springer. pp. 127–134.

Twining, C.J., Taylor, C.J., 2011. Specificity: A graph-based estimator of divergence. IEEE Transactions on Pattern Analysis and Machine Intelligence 33, 2492–2505.

von Tycowicz, C., Schulz, C., Seidel, H.P., Hildebrandt, K., 2015. Real-time nonlinear shape interpolation. ACM Trans Graph 34, 34:1–34:10.

Wächter, A., Biegler, T.L., 2006. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. Math Prog 106, 25–57.

Yu, Y., Zhou, K., Xu, D., Shi, X., Bao, H., Guo, B., Shum, H.Y., 2004. Mesh editing with Poisson-based gradient field manipulation. ACM Trans Graph 23, 644–651.

Zhang, C., Heeren, B., Rumpf, M., Smith, W.A., 2015. Shell PCA: Statistical shape modelling in shell space, in: ICCV, IEEE. pp. 1671–1679.