

Efficient Riemannian Statistical Shape Analysis with Applications in Disease Assessment

Dissertation

zur Erlangung des Grades eines
Doktors der Naturwissenschaften
(Dr. rer. nat.)

am Fachbereich Mathematik und Informatik
der Freien Universität Berlin

vorgelegt von

Felix Ambellan

Berlin, Juli 2022

Erstgutachter:

Professor Dr. Christof Schütte, Freie Universität Berlin, Deutschland

Zweitgutachter:

Directeur de recherche Dr. Xavier Pennec, INRIA Sophia Antipolis, France

Tag der Disputation: 14.11.2022

Acknowledgements

I am delighted to thank all those who supported and encouraged me during the research for and the writing of this thesis. I wish to thank my supervisor Christof Schütte for providing me with an excellent research environment both at Zuse Institute Berlin and Freie Universität Berlin, as well as my former advisor Stefan Zachow who put trust in me and gave me the opportunity to join his working group and enter science. I am in particular grateful to my current advisor Christoph von Tycowicz for the numerous fruitful conversations, as well as his numerous constructive and critical remarks. Without his inestimable advises and support this thesis would not exist. I furthermore want to thank Xavier Pennec for his work on the methodological foundations in the field of *Geometric Statistics* that my research builds upon and for honoring me with reviewing this thesis.

I have been fortunate to collaborate with Alexander Tack; working with him has been a pleasure. I am indebted to the DFG cluster of excellence MATH+ “Transforming the World through Mathematics” for funding my research. I also wish to thank Helga Baum who was my first teacher in the field of differential geometry and Lie groups at Humboldt-Universität Berlin. I have learned and benefited a lot from the present and former members of the Computational Diagnosis and Therapy Planning, as well as the Geometric Data Analysis and Processing group at Zuse Institute Berlin. I would like to thank all of them for innumerable stimulating discussions and an exemplary working atmosphere.

Finally, this thesis would be unimaginable without my family and friends who accompanied me during the last few years. Words can hardly be found expressing the thankfulness to my wife Caro for supporting me in any moment, as well as to little Marthi, who interrupted me whenever she felt it was necessary.

Contents

Acknowledgements	iii
1. Introduction	1
1.1. Main Achievements	6
1.2. Prior Publications	8
1.3. Outline	8
2. Geometric Foundations and Preliminaries	11
2.1. Riemannian Manifolds and Geodesic Lines	11
2.2. Lie Groups, Lie Algebras and Bi-invariance	14
2.3. Statistics on Riemannian Manifolds and Lie Groups	20
2.4. Statistical Shape Models	25
3. Deformation Gradient based Statistical Shape Modeling	29
3.1. Differential Coordinates	29
3.2. Solving the Inverse Problem	30
3.3. $GL^+(3)$ -based Statistical Shape Model	34
3.4. Differential Coordinates Model	38
3.5. Fundamental Coordinates	42
3.6. Solving the Inverse Problem	44
3.7. Fundamental Coordinates Model	53
4. Experimental Model Evaluation and Comparison	57
4.1. Quantitative Assessment	57
4.2. Qualitative Assessment	62

5. Applications in Disease Assessment	65
5.1. Classification of Knee Osteoarthritis	65
5.2. Stratification of Knee Osteoarthritis	71
5.3. Classification of Alzheimer’s Disease	81
6. Conclusion and Future Work	85
6.1. Conclusion	85
6.2. Future Work	88
List of Abbreviations	91
Bibliography	93
A. Information on Data and Further Experiments	105
A.1. Employed Datasets	105
A.2. Additional Experiments	108
A.3. Technical Acknowledgements	112
B. Zusammenfassung	115
C. Selbstständigkeitserklärung	117

CHAPTER 1.

Introduction

Statistical models of shape have been established as one of the most successful methods for understanding and assessing the geometric variability of anatomical structures [Ambellan et al., 2019a]. 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. Shape modeling is of particular interest in image guided diagnosis where morphological changes of anatomies have been hypothesized to be linked to various disorders, often as early indicators. For example, musculoskeletal disorders affecting large proportions of the adult population such as Osteoarthritis (OA) [Lawrence et al., 2008] are associated with morphological changes. Hence, the overall socio-economic burden [Conaghan et al., 2014] associated with this and other diseases provides a strong impetus to develop novel computational approaches for the support of treatment and prevention strategies. One of those are advanced SSMs. Based on a set of consistently discretized training shapes and a notion of shape space, SSMs 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]. Given a set of samples from an object class under study, SSMs estimate the distribution of the underlying population in terms of a reference shape and a hierarchy

of principle modes encoding the main trends in the variation of the samples around that reference. Moreover, representing samples within the basis of principle modes provides a highly discriminative, yet compact description that is susceptible for analysis, inference by learning algorithms (e.g. classification or clustering) and exploration. For example, descriptors based on statistical shape modeling have proven effective for predicting the onset and progression of OA [Bredbenner et al., 2010; Neogi et al., 2013; Thomson et al., 2015, 2016], in anthropometric studies [Robinette et al., 1999; Hasler et al., 2009] and in detection of Alzheimer’s [Shen et al., 2012]. 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].

Since the pioneering work of Kendall [Kendall, 1989], which introduced a rigorous mathematical foundation for shape spaces and statistics thereon, numerous practical approaches were devised. 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 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. 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 [Kelemen et al., 1999]. This approach was introduced by Cootes et al. [1995], we will refer to it as the Point Distribution Model (PDM) throughout the thesis. However, the quality of such 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). Figure 1.1 shows an example from phytomorphology of a shape imposing difficulties within this setting.

How this technique can be generalized to shape spaces with non-Euclidean structures is still an active area of research. 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




Figure 1.1.: A naturally twisted structure that is difficult to be deformed plausibly in shape spaces based on landmark representations in Euclidean space.

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]. Other exciting contributions accounting for nonlinearity have been presented ranging from the large deformation framework [Miller et al., 2015] in computational anatomy, based on diffeomorphisms of the ambient space, to modeling the variability of surfaces employing 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]. An overview of the various concepts can be found in the chapter by Rumpf and Wirth [Rumpf and Wirth, 2015]. Another 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 correspondences, 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 OA. 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 non-Euclidean methods has been linked to the numerical instability and local minima [Heimann and Meinzer, 2009] that appear as typical drawbacks in nonlinear, high-dimensional optimization problems. Furthermore, while the development of 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. Moreover, even for non-interactive applications the order of magnitude of computational cost becomes a decisive factor when processing thousands or even tens of thousands subjects from large-scale morphological studies.

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 also exists 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). In fact, Fletcher et al. propose to linearly approximate PGA—nowadays known as Linearized Principal Geodesic Analysis (LPGA)/Tangent Principal Component Analysis (TPCA). While intrinsic distances between training shapes and mean are preserved under the logarithmic map this is generally not the case for the distances between shapes. *Exact PGA* [Sommer et al., 2010], *geodesic PCA* [Huckemann et al., 2010] as well as *barycentric subspace analysis* [Pennec, 2018] 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 required operations do not admit closed-form solutions. To alleviate these challenges, one line of work, closely related to this thesis’ approach, models shapes explicitly by a collection of elementary building blocks called primitives (e.g. triangles, M-reps etc.) [Fletcher et al., 2003; Hefny et al., 2015] with a natural, geometric structure. Especially if employing triangles as primitives, local changes in shape can be effectively encoded piecewise as orientation preserving *affine* transformations and be equipped with a Lie group structure. In contrast to the former, utilizing a differential description, namely differential coordinates induced by the deformation gradient, provides a local description of the geometry rather than absolute positions [Hasler et al., 2009; Freifeld and Black, 2012; Gao et al., 2016]. 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 spatial 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, this is similar to simplicial maps underlying the construction in [Freifeld and Black, 2012] as well as our setup. However, the latter employs differential characterizations of such maps that remove translational components and put local geometric variability into focus, i.e. provides invariance under translation. 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. Within this thesis we provide continuous model for the differential characterization that alleviates these issues. Typically differential coordinates that are derived from the (deformation) gradient of the map that encodes the shape relative to a reference naturally belong to the group of orientation preserving linear transformations $GL^+(3)$. However, to the best of our knowledge and in contrast to this thesis, previous work does not account for the rich geometric structure inherent to $GL^+(3)$. Approaches like [Zacur et al., 2014] based on the Riemannian framework are not stable according to group operations (composition and inversion) due to the lack of bi-invariant metrics for $GL^+(3)$. Anyhow, consistency/equivariance with group operations is desirable as it provides invariance w.r.t. changes of reference and data coordinate systems and thus prevents bias due to arbitrary choices thereof. We also propose an alternative $SO(3) \times \text{Sym}^+(3)$ -group structure consisting of rotations and symmetric positive definite matrices (representing stretches) that in turn provides bi-invariance and grants relief w.r.t. existence and uniqueness of logarithmic maps by the price of ignoring the original, canonical structure. Furthermore, while Woods [Woods, 2003] proposes a similar approach for image deformation, he employs a surface representation that is not group-valued. 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. This thesis devises a shape representation based on rotation invariant differential coordinates, that we refer to as fundamental coordinates due to their surface-theoretic interpretation. These fundamental coordinates are invariant under Euclidean motion and, thus, not susceptible to any bias due to misalignment by definition. 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], 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. 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. We will present a physically motivated approach based on differential coordinates for which the inverse problem is well-known and can be solved at (close to) linear cost. Furthermore, we also propose an efficient local/global algorithm to solve the inverse problem for the fundamental coordinates.

1.1. Main Achievements

In this work, we address the challenge of developing SSMs that account for the non-Euclidean nature inherent to (anatomical) shape variation and at the same time offer fast, numerically robust processing and as much invariance as possible regarding translation and rotation, i.e. Euclidean motion. This work makes the following major contributions to the body of research:

- 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. We follow two different tracks endowing the differential representation with a Riemannian structure in order to develop two different SSMs:
 - **GL⁺(3)-Model (GLM)**: We derive a novel SSM based on linear differential coordinates as elements in $GL^+(3)$ that is as-invariant-as-possible and, hence, promises increased consistency and reduced bias. To this end, we adapt the notion of bi-invariant means as proposed in [Pennec and Arsigny, 2013] employing an affine connection structure on $GL^+(3)$. Furthermore, we perform second-order statistics based on a family of Riemannian metrics providing the most possible invariance, viz. $GL^+(3)$ -left-invariance and $O(3)$ -right-invariance. We evaluate the performance of the derived model in terms of shape-based classification of pathological malformations of the human knee demonstrating state-of-the-art accuracy.
 - **Differential Coordinates Model (DCM)**: 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 able to correctly capture nonlinear deformations, e.g. inherent to articulation and pathological morphology. We further derive a statistical shape descriptor that captures systematic differences in shape between normal and diseased subpopulations. We show that these descriptors are very well suited for classification tasks in qualitative as well as quantitative evaluation for OA of the distal femur.

- **Fundamental Coordinates Model (FCM)**: This work further presents a novel shape representation based on discrete fundamental forms that is invariant under Euclidean motion, namely the fundamental coordinates. 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. Experimental evaluation for knee OA and Alzheimer’s demonstrates great performance for shape-based disease classification.

Additionally, this thesis aims on shape-based, continuous disease stratification to provide means that objectify disease assessment over the current clinical practice of ordinal grading systems. This work develops one major contribution to the field:

- We derive the **geodesic B-score**, a generalization of the recently proposed B-score to manifold shape spaces that adheres to the rich geometric structure thereof and at the same time is consistent with its Euclidean counterpart. To this end, we build upon a solid mathematical foundation employing concepts from differential geometry and geometric statistics. We further present an original Newton-type fixed point iteration for projection onto geodesics that is both simple to implement and computationally efficient. To the best of our knowledge, previous algorithms restrict to first-order descent schemes [Sommer et al., 2014] or are tailored to special manifolds [Huckemann and Hotz, 2009; Chakraborty et al., 2016]. On the application side, we show that the derived geodesic B-score features improved predictive performance on assessing the risk of Total Knee Replacement (TKR) surgery within 8 years time using a single time point.

We want to emphasize that all presented approaches to statistical shape modeling as well as the basic building blocks for the geodesic B-score are publicly available as part of the *Morphomatics: Geometric morphometrics in non-Euclidean shape spaces* Python library¹ [Ambellan et al., 2021a]. This library constitutes a joint effort of the author together with M. Hanik and C. von Tycowicz—it should not be considered as a direct part of the thesis at hand.

¹morphomatics.github.io

1.2. Prior Publications

The results on Riemannian statistical shape modeling presented in this dissertation were published in two journal and three conference articles. More precisely, the two journal articles have been published in the high-ranking *Medical Image Analysis*: [von Tyrowicz et al., 2018]² (DCM), [Ambellan et al., 2021c] (FCM and quasi-isometric flattening). The author further presented results at the internationally appreciated, top-tier medical image analysis conferences *Medical Image Computing and Computer Assisted Intervention (MICCAI) 2019* including the main conference [Ambellan et al., 2019d] (FCM) as well as the *International Workshop on Mathematical Foundations of Computational Anatomy (MFCA)* [Ambellan et al., 2019c] (GLM) and the *International Conference on Information Processing in Medical Imaging (IPMI) 2021* [Ambellan et al., 2021b] (geodesic B-score).

1.3. Outline

The thesis at hand comprises six chapters: this introduction, one chapter on the mathematical background, three chapters that form the centerpiece and develop the contributions, as well as one final chapter concluding the manuscript and proposing directions for future research.

The mathematical foundations on Riemannian geometry, Lie groups, bi-invariance and geometric statistics are given in Chapter 2 in a compact, yet self-contained manner. The chapter closes with a formal definition of what we understand as SSM. Chapter 3 details on the methodological development of different approaches on deformation gradient based statistical shape modeling. Differential coordinates are introduced in a continuous as well as a discrete setting. An efficient algorithm to solve the inverse problem of reconstructing a shape from coordinates is given. We propose different Lie group structures defined thereon allowing for different paths to novel SSMs based on mean-variance analysis, i.e. the GLM and the DCM. Then the fundamental coordinates are developed as Euclidean motion invariant extension of the differential coordinates with its foundations in classical surface-theory. An effective local/global algorithm to solve the nonlinear inverse problem is devised as well as an approach to quasi-isometric surface flattening through manipulation of the coordinates. These coordinates admit a bi-invariant Lie group

²C. von Tyrowicz and F. Ambellan share the first authorship.

structure that serves as a basis for the formation of the FCM. The chapter closes with the description of the fundamental coordinates. In Chapter 4 we provide quantitative and qualitative evaluation and comparison between the developed methods as well as other modeling approaches including standard PDM and MeshCNN. This evaluation includes, among others, assessment of model *compactness*, *specificity* and *generalization ability* that are widely accepted as standard measures for SSMs. Finally, Chapter 5 focuses on applications in disease assessment for knee OA and Alzheimer's as well as disease stratification for knee OA featuring the geodesic B-score. Moreover, we provide further insight on the employed data and additional experiments on parameters in Appendix A.

CHAPTER 2.

Geometric Foundations and Preliminaries

The chapter recalls the mathematical foundations this thesis is built upon, in order to keep it as self-contained as possible. Hence, descriptions are kept concise and restricted to the strictly necessary. There are various text books available providing more details and proofs on all touched topics in Riemannian geometry and Lie group theory, e.g. Gallier and Quaintance [2020]; Postnikov [2013]; Baum [2009]; Helgason [2001]; do Carmo [1992]; O’Neill [1983]; Kobayashi and Nomizu [1963]. We refer the interested readers to these references.

2.1. Riemannian Manifolds and Geodesic Lines

Definition 2.1. *A Riemannian manifold (M, g) is a pair consisting of a real, smooth manifold M and a Riemannian metric g , i.e. a positive-definite, symmetric, nondegenerate bilinear form g_p on T_pM for each $p \in M$, smoothly varying with p .*

In order to have a notion of how to differentiate vector fields (along curves) in direction of vector fields (along curves) on M , we need to define the affine connection.

Definition 2.2. An affine connection on a manifold M is a mapping

$$\begin{aligned} \nabla : \Gamma(TM) \times \Gamma(TM) &\rightarrow \Gamma(TM) \\ (X, Y) &\mapsto \nabla_X Y, \text{ with} \end{aligned}$$

$$\begin{aligned} \nabla_X Y &\text{ is } C^\infty(M, \mathbb{R})\text{-linear in } X, \\ \nabla_X Y &\text{ is } \mathbb{R}\text{-linear in } Y, \\ \nabla_X fY &= X(f)Y + f\nabla_X Y \text{ for all } f \in C^\infty(M, \mathbb{R}). \end{aligned}$$

Definition 2.3. A smooth curve $\gamma : [a, b] \rightarrow M$ is called geodesic if $\nabla_{\dot{\gamma}} \dot{\gamma} = 0$.

The above definition signifies that geodesics feature a velocity field $\dot{\gamma} = d\gamma/dt$ that is parallel along the curve, marking these curves as ‘straight lines’.

Theorem 2.4. Given $p \in M$ and $v \in T_p M$ there exists unique maximal (w.r.t. domain) geodesic $\gamma : I \subseteq \mathbb{R} \rightarrow M$, s.t. $\gamma(0) = p$ and $\dot{\gamma}(0) = v$.

Definition 2.5. We call a Riemannian space (M, g) geodesically complete if every maximal geodesic is defined on all \mathbb{R} .

There is a particular relationship between *metric* and *affine* geometry stated by the following theorem.

Theorem 2.6. Among all affine connections on (M, g) , there is one uniquely determined by fulfilling the Koszul formula

$$\begin{aligned} 2g(\nabla_X Y, Z) &= X(g(Y, Z)) + Y(g(X, Z)) - Z(g(X, Y)) \\ &\quad + g([X, Y], Z) - g([X, Z], Y) - g([Y, Z], X), \end{aligned}$$

where $X(\cdot)$ denotes the directional derivative for functions/vector fields in direction of X and $[\cdot, \cdot]$ the Lie bracket for vector fields defined by $[X, Y](f) = X(Y(f)) - Y(X(f))$ with f a scalar function.

This affine connection is called Levi-Civita connection.

As of now we will assume ∇ to be the *Levi-Civita* connection if not stated otherwise.

Definition 2.7. Let (M, g) be a Riemannian manifold and $\gamma : [a, b] \rightarrow M$ a piecewise continuously differentiable (C^1) curve. The length of γ is given by

$$l(\gamma) := \int_a^b \sqrt{g_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t))} dt.$$

Now we can define the distance between arbitrary elements in Riemannian manifolds. We note that if γ is a geodesic, then $g_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t)) \equiv \text{const}$, i.e. γ is parametrized proportionally to arclength.

Definition 2.8. For two points $p, q \in (M, g)$ we define their distance as

$$d(p, q) := \inf \{ l(\gamma) \mid \gamma : [a, b] \rightarrow M, \gamma \text{ piecewise } C^1 \text{ s.t. } \gamma(a) = p \text{ and } \gamma(b) = q \}$$

In case we have found a length-minimizing curve, suitably parametrized, it is already a geodesic as stated by the following theorem.

Theorem 2.9. Let (M, g) be a connected Riemannian manifold, $p, q \in M$ and $\gamma : [a, b] \rightarrow M$, with $\gamma(a) = p, \gamma(b) = q$ length-minimizing and parametrized proportionally to arclength, then γ is a geodesic line.

If we now consider the star-shaped (around $o \in T_p M$) set $D_p := \{v \in T_p M \mid 1 \in I_v\} \subseteq T_p M$, i.e. the set of all vectors serving as initial directions for maximal geodesics γ_v starting at $\gamma_v(0) = p \in M$, s.t. $\gamma_v(1)$ is valid, we can define the Riemannian exponential map.

Definition 2.10. For $p \in M$ and D_p as explained above we call the smooth map

$$\begin{aligned} \text{Exp}_p : D_p \subseteq T_p M &\rightarrow M \\ v &\mapsto \gamma_v(1) \end{aligned}$$

Riemannian exponential at p .

The Riemannian exponential comes with some interesting properties.

Theorem 2.11. *Properties of Exp_p .*

- $\text{Exp}_p(tv) = \gamma_{tv}(1) = \gamma_v(t)$ (line segments in $T_p M$ are mapped to geodesic segments in M).

- Exp_p is a local diffeomorphism around the origin $o \in T_p M$, with $d_o \text{Exp}_p = \text{Id}_{T_p M}$, we name the inverse mapping Riemannian logarithm (Log_p).

A first example is the n -dimensional real vector space with standard metric $(\mathbb{R}^n, \langle \cdot, \cdot \rangle)$, where $\text{Exp}_p(v) = p + v$ and $\text{Log}_p(q) = q - p$, since its geodesics are simply straight lines.

Definition 2.12. Let $D \subseteq T_p M$ be a star domain around $o \in T_p M$, s.t. $\text{Exp}_p : D \rightarrow \text{Exp}_p(D) = U \ni p$ is a diffeomorphism. We call U normal neighbourhood of p .

Theorem 2.13. For every q in a normal neighbourhood U of $p \in M$ exists a unique geodesic connecting $\gamma(0) = p$ and $\gamma(1) = q$ within U , namely $\gamma(t) = \text{Exp}_p(t \text{Log}_p(q))$, with $t \in [0, 1]$. Moreover, there exists an $\varepsilon > 0$, s.t. if $q \in \text{Exp}(B_\varepsilon(o)) \subset U$, where $B_\varepsilon(o)$ denotes the ball of radius ε around $o \in T_p M$, then γ is also length-minimizing and $d(p, q) = \sqrt{g_p(\text{Log}_p(q), \text{Log}_p(q))}$.

Definition 2.14. Given (M, g) a geodesically complete space. For $p \in M$ we call the largest value $i(p, M)$, s.t. $\text{Exp}_p|_{B_{i(p, M)}(o)}$ is a diffeomorphism the injectivity radius of M in p . Furthermore, we call $i(M) = \inf_{p \in M} \{i(p, M)\}$ the injectivity radius of M .

Theorem 2.15 (Hopf-Rinow). Let (M, g) be a connected Riemannian manifold, then the following statements are equivalent:

- (i) M is geodesically complete;
- (ii) M together with its geodesic distance forms a complete metric space;
- (iii) There exists a point $p \in M$, s.t. Exp_p is defined on whole $T_p M$;
- (iv) Every closed and bounded subset of M is compact.

2.2. Lie Groups, Lie Algebras and Bi-invariance

Definition 2.16. A Lie group G is a group with a smooth manifold structure s.t.

$$\begin{aligned} G \times G &\rightarrow G \\ (g, h) &\mapsto g \cdot h^{-1} \end{aligned}$$

is a smooth map, i.e. multiplication and inversion are smooth operations.

A very important class of Lie groups are matrix groups. Among these three examples are of highest interest for this thesis:

- (i) the (positive) general linear group $\text{GL}^+(n) = \{A \in \mathbb{R}^{n \times n} \mid \det(A) > 0\}$,
- (ii) the special orthogonal group $\text{SO}(n) = \{R \in \mathbb{R}^{n \times n} \mid \det(R) = 1, R^{-1} = R^T\}$,
- (iii) the symmetric positive-definite matrices equipped with log-Euclidean multiplication [Arsigny et al., 2006]

$$\text{Sym}^+(n) = \{U \in \mathbb{R}^{n \times n} \mid U = U^T, x^T U x > 0 \text{ for all } o \neq x \in \mathbb{R}^n\}$$

$$U \circ V := \exp(\log(U) + \log(V)), \text{ using matrix exponential and logarithm.}$$

Definition 2.17. For a fixed $g \in G$ we call the diffeomorphisms

$$L_g : G \rightarrow G; L_g(h) := g \cdot h$$

$$R_g : G \rightarrow G; R_g(h) := h \cdot g$$

left and right translation.

Obviously, left and right translation act simply transitive and commute with each other.

Definition 2.18. A Vector field X on Lie group G is called left-invariant (right-invariant) if $d_h L_g(X(h)) = X(gh)$ ($d_h R_g(X(h)) = X(hg)$) for all $g, h \in G$. Moreover, if X is left- and right-invariant we call it bi-invariant.

Definition 2.19. Let V be a vector space and $[\cdot, \cdot] : V \times V \rightarrow V$ an alternating bilinear map fulfilling the Jacobi identity

$$[[u, v], w] + [[v, w], u] + [[w, u], v] = 0 \text{ for all } u, v, w \in V.$$

We name $(V, [\cdot, \cdot])$ Lie algebra and $[\cdot, \cdot]$ Lie bracket.

We can now equip the vector space of left-invariant vector fields \mathfrak{g} with the vector field bracket $[X, Y](f) = X(Y(f)) - Y(X(f))$, for $f \in C^\infty(G, \mathbb{R})$, s.t. it turns into a Lie algebra. Please note that $X(f)$ specifies the derivative of function f in direction of vector field X and not the evaluation of X in a certain group element.

Definition 2.20. The left-invariant vector fields of a Lie group G together with the vector field bracket $(\mathfrak{g}, [\cdot, \cdot])$ are called Lie algebra of Lie group G .

Since left-invariant vector fields are uniquely determined by their value at $e \in G$ through $X(g) = d_e L_g(X(e))$ we have a natural identification $\mathfrak{g} \simeq T_e G$. Keeping this in mind both terms are from now on used interchangeably.

Theorem 2.21. *Given Lie group G and its Lie algebra \mathfrak{g} . Let $X \in \mathfrak{g}$ be a left-invariant vector field and γ_X its maximal integral curve through $e \in G$.*

- (i) γ_X is defined on all \mathbb{R} .
- (ii) $\gamma_X : \mathbb{R} \rightarrow G$ is a Lie group homomorphism (1-parameter subgroup), i.e. $\gamma_X(0) = e$ and $\gamma_X(s+t) = \gamma_X(s) \cdot \gamma_X(t)$ for all $s, t \in \mathbb{R}$.
- (iii) $\gamma_{sX}(t) = \gamma_X(s \cdot t)$ for all $s, t \in \mathbb{R}$.

Definition 2.22. *For \mathfrak{g} being the Lie algebra of G we call the smooth map*

$$\begin{aligned} \exp : \mathfrak{g} &\rightarrow G \\ X &\mapsto \gamma_X(1) \end{aligned}$$

exponential of G .

Theorem 2.23. *Properties of \exp .*

- $\exp(tX) = \gamma_{tX}(1) = \gamma_X(t)$ (lines in \mathfrak{g} are mapped to integral curves of left-invariant vector fields in G).
- \exp is a local diffeomorphism around $o \in \mathfrak{g}$, with $d_o \exp = \text{Id}_{\mathfrak{g}}$, we name the inverse mapping logarithm (\log).
- $\exp(o) = e$ and $\exp(-X) = (\exp(X))^{-1}$.
- $L_g(\exp(tX))$ is the maximal integral curve of X through $g \in G$.

Theorem 2.24. *For Lie groups that are matrix groups the Lie group exponential/logarithm coincides with the matrix exponential/logarithm.*

2.2.1. Bi-invariant Connections

Definition 2.25. An affine connection ∇ on Lie group G is called *left-invariant* if $\nabla_X Y$ is left-invariant for all left-invariant X, Y , i.e. $\nabla_{d_h L_g(X(h))} d_h L_g(Y(h)) = d_h L_g(\nabla_X Y(h)) = \nabla_X Y(gh)$ for all $g, h \in G$. Right-invariant affine connections are defined analogously. Moreover, if ∇ is left- and right-invariant we call it *bi-invariant*.

Theorem 2.26. Any bilinear operator $\alpha(\cdot, \cdot)$ on $\mathfrak{g} \simeq T_e G$ can be extended to a left-invariant connection ∇ on G by defining, for each $h \in G$ and all $X, Y \in \mathfrak{g}$,

$$\nabla_X Y(h) = d_e L_h(\nabla_X Y(e)) = d_e L_h(\alpha(X, Y))$$

In fact, there is a bijective relationship between left-invariant connections and bilinear operators on $\mathfrak{g} \simeq T_e G$. An analogous result holds for right-invariant connections.

Definition 2.27. A left-invariant connection ∇ on a Lie group G is called a *Cartan-Shouten connection* if its geodesics coincide with the integral curves of left-invariant vector fields, i.e. left-translates of 1-parameter subgroups of G .

Corollary 2.28. Lie groups equipped with Cartan-Shouten connection are geodesically complete, i.e. all geodesics are defined on all \mathbb{R} .

Theorem 2.29. A left-invariant connection ∇ on a Lie group G emerging from bilinear operator α on \mathfrak{g} is a Cartan-Shouten connection if, and only if α is skew-symmetric.

Definition 2.30. A connection ∇ is called *torsion-free (symmetric)* if $\nabla_X Y - \nabla_Y X = [X, Y]$, i.e. $\alpha(X, Y) - \alpha(Y, X) = [X, Y]$ for all $X, Y \in \mathfrak{g}$.

Theorem 2.31. Given a Lie group G , there exists one unique torsion-free Cartan-Shouten connection given by:

$$\alpha(X, Y) = \frac{1}{2}[X, Y], \text{ for all } X, Y \in \mathfrak{g}.$$

We refer to this connection as the *canonical Cartan-Shouten connection* on G .

2.2.2. Bi-invariant Metrics

Definition 2.32. Let G be a Lie group. A Riemannian metric g on G is said to be left-invariant, if every left translation is an isometry of g . That is, if for each $g, h \in G$ and every $v, w \in T_h G$,

$$g_h(v, w) = g_{gh}(d_h L_g(v), d_h L_g(w)).$$

Right-invariant metrics are defined analogously. Furthermore, a metric is called bi-invariant if it is both left- and right-invariant.

Theorem 2.33. Any inner product $\langle \cdot, \cdot \rangle$ on $T_e G$ can be extended to a left-invariant metric on G by defining, for each $h \in G$ and all $v, w \in T_h G$,

$$g_h(v, w) = g_e(d_h L_{h^{-1}}(v), d_h L_{h^{-1}}(w)) = \langle d_h L_{h^{-1}}(v), d_h L_{h^{-1}}(w) \rangle.$$

In fact, there is a bijective relationship between left-invariant metrics and inner products on $T_e G$. An analogous result holds for right-invariant metrics.

Bi-invariant metrics do not exist for every Lie group, but there are important special cases where its existence can be assured.

Theorem 2.34. Every compact Lie group G admits a bi-invariant metric.

Corollary 2.35. $\text{SO}(n) = \det^{-1}(1)$ is compact. It features a well known bi-invariant metric emerging from left translation of the Frobenius inner product $\langle X, Y \rangle_F = \text{Tr}(X^T Y)$.

Note, although quite common in the literature, e.g. [Rentmeesters and Absil, 2011; Penec and Arsigny, 2013], we do not employ any additional factor $1/2$ scaling the trace.

Theorem 2.36. If G is Abelian, left and right translations coincide, so every left-invariant metric is bi-invariant.

Corollary 2.37. $\text{Sym}^+(n)$ is Abelian and hence features a bi-invariant metric emerging from left translation of the Frobenius inner product $\langle U, V \rangle_F = \text{Tr}(U^T V)$.

We can now state the main theorem that sheds light on the connection between Lie groups and Riemannian manifolds.

Theorem 2.38. *Given (G, g) a Lie group with bi-invariant metric the geodesics of G coincide with the collection of left-translates of 1-parameter subgroups of G and the exponential map coincides with the usual Riemannian exponential map at the identity element.*

Theorem 2.39. *For a Lie group with bi-invariant metric (G, g) the canonical Cartan-Shouten connection (cf. Theorem 2.31) coincides with the Levi-Civita connection on G .*

In the light of the above theorem we note that the canonical Cartan-Shouten connection on a group is the closest relative to a Levi-Civita connection on that same group carrying a bi-invariant metric. This makes it a very reasonable choice for further consideration in groups that do not admit a bi-invariant metric, as e.g. $\text{GL}^+(n)$.

Corollary 2.40. *Lie groups equipped with bi-invariant metric are geodesically complete, i.e. all geodesics are defined on all \mathbb{R} .*

This essentially means that Lie group structures with bi-invariant metric are compatible with their respective Riemannian structure. Thus, the rather complicated Riemannian maps Exp and Log can be expressed in terms of their comparatively simple Lie group counterparts.

Acknowledging their importance, the following paragraphs will give a brief description of the structure for two Lie groups already introduced, namely, the rotation group (in dimension 3) and the group of symmetric positive matrices (in dimension n , equipped with log-euclidean multiplication).

The Rotation Group $\text{SO}(3)$

The Lie algebra of the rotation group $\text{SO}(3)$ consists of all skew-symmetric matrices $\mathfrak{so}(3) = \{X \in \mathbb{R}^{3 \times 3} \mid X = -X^T\}$. This is substantially the angle-axis representation of rotations, since every $X \in \mathfrak{so}(3)$ describes a rotation by angle $\theta = \|X\|_F / \sqrt{2}$ around axis $v = \theta^{-1} \cdot (X_{32}, X_{13}, X_{21})$. However, formally the exponential is given by the Rodrigues' formula. Let $X, Y \in \mathfrak{so}(3)$, s.t. $\|Y\|_F = \sqrt{2}$ and $X = \theta Y$

$$R = \exp(X) = \exp(\theta Y) = I + \sin(\theta)Y + (1 - \cos(\theta))Y^2.$$

The logarithm of a rotation $R \in \text{SO}(3)$ can also be calculated explicitly by

$$\log(R) = \begin{cases} \frac{\theta}{2 \sin(\theta)}(R - R^T) & , 0 \neq \theta \in (-\pi, \pi) \\ 0 & , \theta = 0 \end{cases},$$

where $\theta = \arccos(\text{Tr}(R)-1/2)$ denotes the angle of rotation R . Note that the logarithm is not globally defined, due to the ambiguity of rotations by $\pm\pi$. The distance for $R, Q \in \text{SO}(3)$ emerging from the bi-invariant metric (Corollary 2.35) also features an explicit formula $d_{\text{SO}(3)}(R, Q) = \|\log(R^T Q)\|_F$.

Symmetric Positive Definite Matrices $\text{Sym}^+(n)$

Since the matrix logarithm is compatible with change of basis and every symmetric positive-definite matrix $U \in \text{Sym}^+(n)$ is diagonalizable with only positive eigenvalues, it is given by $\log(U) = P \log(D) P^{-1}$, where P is a basis of eigenvectors and D a diagonal matrix formed by the respective eigenvalues. This indicates that $\log(U)$ is symmetric but not necessarily positive-definite, leading to the identification $\mathfrak{sym}(n) = \{X \in \mathbb{R}^{n \times n} \mid X = X^T\}$. Conforming to this, the exponential is mapping symmetric matrices to symmetric positive-definite matrices, it respects change of basis and hence is given analogously by $\exp(X) = P \exp(D) P^{-1}$ for every $X \in \mathfrak{sym}(n)$. Note that exponential and logarithm are globally diffeomorphic. If we finally consider the distance of $U, V \in \text{Sym}^+(n)$ associated to the bi-invariant metric (Corollary 2.37) we again have a closed form expression $d_{\text{Sym}^+(n)}(U, V) = \|\log(U) - \log(V)\|_F$.

2.3. Statistics on Riemannian Manifolds and Lie Groups

In the first part of this section we will introduce the notion of mean value for Riemannian manifolds and Lie groups and PGA for Riemannian manifolds, two basic concepts within the field of geometric statistics. In the second part a formal definition of a SSM is proposed, serving as basis for further methodological development in Chapter 3.

2.3.1. Bi-invariant Lie Group Mean

This subsection introduces the notion of mean for the metric-free but affine setting via exponential barycenters following [Pennec and Arsigny, 2013].

Definition 2.41. *Given G a Lie group with canonical Cartan-Shouten connection and $\{g_1, \dots, g_n\} \subset G$ a finite set of elements in G . We call μ bi-invariant sample group mean if it fulfills the Karcher equation*

$$\sum_{i=1}^n \log(\mu^{-1}g_i) = 0.$$

This characterizes μ as the barycenter of the input data w.r.t. geodesics leaving μ towards each input element. Note that bi-invariant refers to that μ behaves equivariant under left-translation, right-translation and inversion of the input samples.

Theorem 2.42. *Given $\{g_1, \dots, g_n\} \subset G$ a finite set of elements in Lie Group G with canonical Cartan-Shouten connection. If the data belongs to a sufficiently small normal convex neighborhood for some $g \in G$, then there exists a unique bi-invariant group mean.*

Especially in view of our applications (Chapter 5) we will necessarily always assume well-localizedness of our data.

2.3.2. Fréchet Mean

In this subsection the Fréchet mean on Riemannian manifolds is defined. Further details on existence and uniqueness are discussed roughly following Chapter 2 of [Pennec et al., 2019].

Definition 2.43. *Given (M, g) a Riemannian space with geodesic distance d and $\{p_1, \dots, p_n\} \subset M$ a finite set of elements in M . We call the minimizers of the sum-of-squared-distances*

$$\mu = \arg \min_{p \in M} \sum_{i=1}^n d^2(p_i, p)$$

the sample Fréchet means.

We note that every Fréchet mean on a Lie group fulfills a Riemannian version of the Karcher equation of Definition 2.41 (Riemannian logarithm w.r.t. Levi-Civita connection replaces group logarithm). In fact, every solution to the Riemannian Karcher equation is a critical point of the Fréchet mean objective. Furthermore, as for a given bi-invariant metric both logarithms coincide, we see that the bi-invariant group mean is a natural generalization of the Fréchet mean to the metric-free affine setting.

Proposition 2.44. *For the \mathbb{R}^m with standard metric the notion of Fréchet mean agrees to the usual arithmetic mean.*

Theorem 2.45. *Let (M, g) be a Riemannian manifold, that is a complete metric space w.r.t. its geodesic distance d . Then the Fréchet mean exists for any a finite set of elements $\{p_1, \dots, p_n\}$ in M .*

Corollary 2.46. *Since Lie groups with bi-invariant metric are geodesically complete it follows with Theorem 2.15 that the Fréchet mean always exists.*

Theorem 2.47. *For a complete Riemannian space (M, g) with sectional curvature bounded from above by κ and with injectivity radius $i(M)$ the Fréchet mean of a finite set $\{p_1, \dots, p_n\}$ in M is unique if all elements are contained in a geodesic ball with radius*

$$\varepsilon = \frac{1}{2} \begin{cases} \min(i(M), \pi/\sqrt{\kappa}) & , \kappa > 0, \\ i(M) & , \kappa \leq 0. \end{cases}$$

Note that the Fréchet mean on manifolds is occasionally also called Karcher mean. This clash of nomenclature is due to Fréchet primarily working on metric spaces and Karcher approaching Riemannian manifolds. Additionally, Fréchet means are considered to be global, whereas Karcher means are considered to be local. A well-arranged overview on this is given in Chapter 8 of [Nielsen and Bhatia, 2013]. Within this thesis we will use both terms interchangeably, since we on the one hand are taking the manifold perspective and on the other hand assume well-localizedness of our data s.t. uniqueness is granted.

Corollary 2.48. *$SO(3)$ with bi-invariant metric emerging from the Frobenius inner product has non-negative curvature bounded by $\kappa = 1/8$ [Cheeger and Ebin, 2008, Chapter 3], [Ge, 2014]. As discussed above, its injectivity radius is $i(M) = \sqrt{2}\pi$. Hence, $\varepsilon = \sqrt{2}\pi/2$, i.e. the respective geodesic ball contains all rotations differing by an angle of less than $\pm\pi/2$ from the identity.*

Corollary 2.49. $\text{Sym}^+(n)$ is an Abelian Lie group and thus flat, i.e. $\kappa = 0$. As discussed above \exp is a global diffeomorphism, i.e. $i(M) = \infty$. Hence, the Fréchet mean is always unique.

Note that this global uniqueness is still granted if we consider the symmetric positive-definite matrices with the affine-invariant metric (cf. [Thanwerdas and Pennec, 2019]), since their sectional curvature is non-positive, i.e. $\kappa \leq 0$ as can be found e.g. in [Pennec et al., 2019].

2.3.3. Principal Geodesic Analysis

In this section we follow the work of [Fletcher et al., 2003, 2004] and at first give a notion of variance based on [Fréchet, 1948] and thereon define PGA as mechanism to find a sequence of nested geodesic submanifolds that maximize the projected variance of the data. Furthermore, we briefly discuss the special case of PCA and conclude with a description of LPGA, a.k.a. TPCA that will be of practical importance throughout the rest of the thesis.

Definition 2.50. Given (M, g) a connected complete Riemannian manifold and denote $U = \text{Exp}_\mu(B_{i(\mu, M)}) \subset M$. Further let μ be the Fréchet mean of a finite set $\{p_1, \dots, p_n\} \subset U$, then we define its sample variance as

$$\sigma^2 = \frac{1}{n} \sum_{i=1}^n d^2(p_i, \mu) = \frac{1}{n} \sum_{i=1}^n \|\text{Log}_\mu(p_i)\|^2. \quad (2.1)$$

In order to relate data samples to geodesic submanifolds we need a notion of distance that is achieved by distance minimizing projection.

Definition 2.51. Let (M, g) be a connected complete Riemannian manifold and $H \subset M$ be a geodesic submanifold in $p^* \in H$ (i.e. geodesics of H through p^* are also geodesics of M). Then we call

$$\begin{aligned} \pi_H : M &\rightarrow H \\ p &\mapsto \pi_H(p) = \arg \min_{q \in H} d^2(p, q), \end{aligned} \quad (2.2)$$

metric projection onto H .

Further information on existence and uniqueness of the metric projection operation can be found e.g. in [Walter, 1974]. However, as [Fletcher et al., 2004] points out existence and uniqueness are granted for sufficiently small neighborhoods around $p^* \in H$. Nevertheless, assessing the projection without further simplification or in-depth examination of the manifold's properties remains a rather intricate task.

Anyway, if we now combine Equations (2.1) and (2.2) we can define sample variance maximizing *principal modes of variation*.

Definition 2.52. *Given (M, g) a connected complete Riemannian manifold and $U = \text{Exp}_\mu(B_{i(\mu, M)}) \subset M$. Further let μ be the Fréchet mean of a finite set $\{p_1, \dots, p_n\} \subset U$, then we define its principal modes of variation by*

$$\begin{aligned} \vartheta_k &= \arg \max_{\substack{\|v\|=1 \\ g_\mu(\vartheta_j, v)=0, j < k}} \sum_{i=1}^n \|\text{Log}_\mu(\pi_H(p_i))\|^2, \quad k = 1, \dots, n-1; \\ H &= \text{Exp}_\mu(\text{span}(\{\vartheta_1, \dots, \vartheta_{k-1}, v\})) \cap U. \end{aligned} \quad (2.3)$$

Note that the number of modes is actually bounded by $\min(n-1, \dim(M))$. However, since we study small sized sample populations in high-dimensional spaces throughout this thesis, we will deliberately assume $n-1 \leq \dim(M)$.

Having determined the modes of variation every input element p_i features a unique low-dimensional representation in the tangent space at μ employing exponential and logarithm. It is given by means of a principal modes linear combination

$$p_i = \text{Exp}_\mu \left(\underbrace{\sum_{k=1}^{n-1} g_\mu(\text{Log}_\mu(p_i), \vartheta_k) \vartheta_k}_{\alpha_i^k} \right),$$

with coefficients α_i^k called *principal scores* or *shape weights* of p_i .

The assessment of principal modes and scores in order to analyze variability within a given set of manifold-valued data and to reduce its dimensionality is called PGA. It can be furthermore understood as generalization of PCA to manifold data, as the following proposition states.

Proposition 2.53. *For $(\mathbb{R}^n, \langle \cdot, \cdot \rangle)$ Equation (2.3) reduces to*

$$\begin{aligned} \vartheta_k &= \arg \max_{\substack{\|v\|=1 \\ \langle \vartheta_j, v \rangle = 0, j < k}} \sum_{i=1}^n \langle p_i - \mu, v \rangle^2, \quad k = 1, \dots, n-1, \end{aligned} \quad (2.4)$$

The given definition of principal modes can further be seen as finding the directions maximizing the sample variance $\|\text{Log}_\mu(\pi_H(p_i))\|^2$ by nested geodesic subspaces. An alternative definition was introduced by [Sommer et al., 2010] in terms of minimization of residual reconstruction errors $\|\text{Log}_{p_i}(\pi_H(p_i))\|^2$. As consequence of the Pythagorean theorem both definitions agree for the Euclidean setting, yet not in the general Riemannian case, i.e. they have to be understood as different generalizations of Euclidean PCA.

Note that the optimization problem given in Definition 2.52 comes with neither existence nor uniqueness assurance for general manifolds. This makes its practical application rather uncomfortable. However, Fletcher et al. [2004] already provided a feasible solution to this in their original work through linearization of the projection operation.

Proposition 2.54. *Assuming $\text{Log}_\mu(\pi_{H_k}(p_i)) \approx \sum_{j=1}^k g_\mu(\text{Log}_\mu(p_i), \vartheta_j) \cdot \vartheta_j$, where $H_k = \text{Exp}_\mu(\text{span}(\{\vartheta_1, \dots, \vartheta_k\})) \cap U$, i.e. the submanifold projection of p_i can be well approximated by linear projection of $\text{Log}_\mu(p_i)$ onto the respective generating subspace, Equation (2.3) reduces to*

$$\vartheta_k = \underset{\substack{\|v\|=1 \\ g_\mu(\vartheta_j, v)=0, j < k}}{\text{arg max}} \sum_{i=1}^n g_\mu^2(\text{Log}_\mu(p_i), v), \quad k = 1, \dots, n-1. \quad (2.5)$$

This equation can be solved uniquely as long as the particular logarithms exist and are unique. Moreover, it obviously agrees for $(\mathbb{R}^m, \langle \cdot, \cdot \rangle)$ with Equation (2.4). Solving this linearized problem in order to approximate PGA is called performing LPGA/TPCA.

2.4. Statistical Shape Models

Within this section we propose a (rough) formalization of the term SSM as we will understand it for the rest of this thesis and how it fits the context provided by the definitions given in the previous section. Throughout this thesis we consider a *shape* as two-dimensional (2D) orientable surface modulo translation and rotation embedded in 3D Euclidean space. However, we keep the data inherent scale untouched as it, in the medical context, potentially correlates with pathological anatomical changes. We additionally consider a shape compact but possibly with boundary, i.e. not necessarily a closed surface. For example, in this thesis we will work with hippocampus shapes (closed) and shapes representing the distal part of the femur bone (not closed).

Definition 2.55. Let (M, g) be a Riemannian manifold and $\{s_1, \dots, s_n\} \subset M$ a finite set, where each element uniquely represents an input shape. We define a data-driven shape model as pair

$$(\mu, \{\vartheta_k\}_1^d),$$

where μ is a fixed element in M and $\{\vartheta_k\}_1^d$ is a minimal set of g -orthonormal vectors in $T_\mu M$ s.t. $s_i = \text{Exp}_\mu\left(\sum_{k=1}^d g_\mu(\text{Log}_\mu(s_i), \vartheta_k)\vartheta_k\right)$ holds for all s_i . We furthermore call $(\mu, \{\vartheta_k\}_1^d)$ SSM if μ is a Fréchet mean (or an exponential barycenter) and $\{\vartheta_k\}_1^d$ is constructed employing PGA or LPGA as detailed in the previous section.

The manifold M is, depending on the context, referred to as *shape space*, *feature space* or *configuration space*.

It is worth mentioning that this definition leaves the inverse problem of shape reconstruction from the manifold-valued representation completely open. This is because reconstruction poses no obstruction for shape statistics. From the practical point of view this means an SSM, as we understand it here, does not necessarily need to be a generative model, in the sense that new instances similar to the inputs can be synthesized from it. However, all models discussed in Chapter 3 are generative and algorithms for shape reconstruction are given.

One, if not the most prominent, SSM-type is the PDM. It was first put into focus of a broad public and applied for medical image segmentation by Cootes et al. [1995]. To express the PDM in our technical setting we assume having given m shapes discretized as point clouds, formed by a finite number of l points, coming with point correspondence, i.e. every two shapes feature a semantic bijection of their discretizations. In a medical context this correspondence usually aligns with anatomical characteristics, i.e. points tend to represent distinct anatomical landmarks. The i -th shape is now uniquely represented via its vectorized coordinates: $s_i = (x_1, y_1, z_1, \dots, x_l, y_l, z_l) \in \mathbb{R}^{3l}$.

Definition 2.56. Let $(\mathbb{R}^{3l}, \langle \cdot, \cdot \rangle)$ be the $3l$ -dimensional Euclidean space with standard metric and let further $\{s_1, \dots, s_n\} \subset \mathbb{R}^{3l}$ be a finite set, where each element contains vectorized coordinates representing an input shape. All input shapes are moreover assumed to feature point correspondence. The SSM $(\mu, \{\vartheta_k\}_1^d)$ consisting of the arithmetic mean μ and $\{\vartheta_k\}_1^d$ constructed by PCA is called PDM.

The PDM is also called linear/Euclidean model, due to the respective properties of $(\mathbb{R}^{3l}, \langle \cdot, \cdot \rangle)$. It is obviously generative by definition.

Note that the landmark-based PDM exhibits an interesting relationship to a similar construction on Kendall's shape space Σ_3^l [Kendall et al., 1999; Dryden and Mardia, 2016; Nava-Yazdani et al., 2020] (for l 3D landmarks) if we consider only two input shapes s_1, s_2 that are superimposed utilizing Procrustes alignment [Gower, 1975] (i.e. including scale removal). In fact, the geodesic joining s_1, s_2 is realized by the respective geodesic in the Sphere S^{3l-1} . Moreover, the geodesic joining both in \mathbb{R}^{3l} (i.e. the straight line) is related to this great-circle arc simply through normalization. This in particular implies that the shape at half way between them, realizing the respective means in Σ_3^l and \mathbb{R}^{3l} , differ only in scale. Thus, exemplification and discussion in this thesis referring to inaccurate means of two input shapes in the Euclidean case (cf. Sections 3.3, 3.4, 3.7) can, without explicit mention, as the inaccuracies obviously do not depend on scaling, be directly extended to Kendall's shape space.

CHAPTER 3.

Deformation Gradient based Statistical Shape Modeling

In this chapter we introduce the notion of deformation gradient between two objects and differential coordinates respectively. We present three different approaches to further process these coordinates and establish Riemannian and/or Lie group structures thereon in order to perform geometric statistics allowing for SSM construction as defined in Chapter 2.

3.1. 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 $\bar{\mathcal{S}} \subset \mathbb{R}^3$, where \mathcal{C} denotes the space of (smooth enough) deformations $\phi : \bar{\mathcal{S}} \rightarrow \mathbb{R}^3$. 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 $\bar{\mathcal{S}}$ denote the entire volume of the object and defer the extension to the 2D case to Section 3.2.2.

In order to derive a notion of shape space, \mathcal{C} is typically projected onto a non-trivial space that is obtained by quotienting out rigid motions or reparametrizations. However, the analysis of the geometry of such spaces and the de-

velopment 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 $p \in \bar{\mathcal{S}}$, it is reasonable to approximate ϕ by its first-order Taylor expansion: $\phi(q) \approx \phi(p) + \nabla\phi|_p(q - p)$. Here, $\nabla\phi|_p$ is the 3×3 matrix of partial derivatives of ϕ at p and is called the deformation gradient (although it is technically no 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 $\nabla\phi|_p$ 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 $\nabla\phi|_p$ 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 $s : \bar{\mathcal{S}} \rightarrow \text{GL}^+(3)$ with $s_i(p) = \nabla\phi_i|_p$ and $\text{GL}^+(3)$ the group of 3×3 matrices with positive determinant.

3.2. Solving the Inverse Problem

Employing gradient fields as differential representation has the additional advantage that the inverse problem of mapping differential coordinates $s \in \mathcal{F}$ to configurations $\pi(s) \in \mathcal{C}$ 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

$$\pi(s) := \arg \min_{\phi \in \mathcal{C}} \frac{1}{2} \int_{\bar{\mathcal{S}}} \|\nabla\phi - s\|^2 dV, \quad (3.1)$$

where dV is the volume form on $\bar{\mathcal{S}}$. The solutions are unique up to additive constants for each connected component of $\bar{\mathcal{S}}$ 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 (3.1) is the well-known Poisson equation

$$\Delta\phi = \nabla \cdot s, \quad (3.2)$$

where Δ and $\nabla \cdot$ denote the Laplacian and divergence operator, respectively. We emphasize that Equation (3.2) 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 tends to distribute errors uniformly such that local gradient field inconsistencies are attenuated (see Figure 4.7 left for mappings from \mathcal{F} to \mathcal{C}). For additional information and description we refer to [Yu et al., 2004]. Throughout the thesis we will additionally identify a deformation $\pi(s)$ with its image, i.e. $\pi(s)(\bar{\mathcal{S}}) = \mathcal{S} = \mathcal{S}(s)$ to emphasize its character as geometric realization of given coordinates s .

3.2.1. 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 (*tangent*) *principal geodesic analysis* [Fletcher et al., 2004] to extract the dominant modes of variation. In this section we will hereby focus on the smooth setting, used discretizations are to be introduced later along with the respective shape modeling approach.

Intrinsic mean

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

$$d_{\mathcal{F}}(s_1, s_2) = \left(\int_{\bar{\mathcal{S}}} d_F(s_1(p), s_2(p))^2 dV \right)^{\frac{1}{2}}. \quad (3.3)$$

The intrinsic or Fréchet mean μ of a collection of n points $s_1, \dots, s_n \in \mathcal{F}$ is then defined as the minimizer of the functional $\sum_{i=1}^n d_{\mathcal{F}}(s, s_i)^2$, i.e. the sum-of-squared distances to each point. This functional can also be written as $\int_{\bar{S}} \mathcal{L}(s(p)) dV$, where $\mathcal{L}(\cdot) = \sum_{i=1}^n d_F(\cdot, s_i)^2$ is the sum-of-squared point-wise distances. Due to its simple form, the Euler-Lagrange equations reduce to $\partial \mathcal{L} / \partial s_{qp} = 0$ (with natural boundary conditions trivially satisfied) and, hence, μ can be computed point-wise, i.e. $\mu(p)$ depends only on the gradients $s_i(p)$ and the choice of d_F . 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.

Tangent Principal Geodesic Analysis

As detailed in Section 2.3, 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

$$g_{\mu}(v, w) = \int_{\bar{S}} \langle v(p), w(p) \rangle_{\mu(p)} dV \quad (3.4)$$

be the inner product on $T_{\mu}\mathcal{F}$, where $\langle \cdot, \cdot \rangle_{\mu(p)}$ denotes an inner product on $T_{\mu(p)}F$.

Using g , we can define a covariance operator for deviations from the intrinsic mean as

$$C(\cdot) = \frac{1}{n} \sum_{i=1}^n g_{\mu}(\cdot, v_i) v_i, \quad (3.5)$$

where $v_i = \text{Log}_{\mu}(s_i)$ 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 $\text{Log}_{\mu}(\cdot)$ can be computed point-wise and leads to efficient, closed-form expressions entailing the matrix logarithm in F .

C is positive-definite on $\text{span}(v_1, \dots, v_n)$ 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 e.g. Figure 4.7) of the ψ_i define the main modes of variation with the λ_i encoding the corresponding variances.

3.2.2. Discretization for Simplicial Manifolds

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

$$\phi(p) = \sum_{j=1}^{n_0} \varphi_j(p) \sigma_j^0. \quad (3.6)$$

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 is thus uniquely determined by its coefficients $\boldsymbol{\phi} = (\sigma_1^0 \dots \sigma_{n_0}^0)^T \in \mathbb{R}^{n_0 \times 3}$. The gradient is then

$$\nabla \phi|_p = \sum_{j=1}^{n_0} \sigma_j^0 \nabla \varphi_j(p)^T \quad (3.7)$$

and yields a constant 3×3 matrix $D_k = \nabla \phi|_{\sigma_k^d}$ for each d -simplex σ_k^d of \bar{S} . However, note 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 $\nabla \phi|_p$ as deformation gradient yields an efficient dimension reduction with a natural surface-based formulation of the Poisson equation.

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

$$\mathbf{D}_{ij} = \left(\nabla \varphi_j|_{\sigma_{[i/3]}^d} \right)_{i \bmod 3} \quad (3.8)$$

such that $\mathbf{D}\boldsymbol{\phi} = (D_1 \dots D_{n_d})^T$. The discrete version of the Poisson equation (3.2) is then given by

$$\mathbf{D}^T \mathbf{M} \mathbf{D} \boldsymbol{\phi} = \mathbf{D}^T \mathbf{M} \mathbf{s}, \quad (3.9)$$

where \mathbf{M} is a diagonal matrix containing the triangle areas or tetrahedron volumes and \mathbf{s} 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 \bar{S} . The problem in Equation (3.9) 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.

3.3. $GL^+(3)$ -based Statistical Shape Model

In this section we will introduce the GLM that treats triangle-wise deformation-gradients simply as elements in $GL^+(3)$ without any additional structure. We therefore consider shapes to be instances of a class of anatomical objects that are topologically consistent, s.t. they can be represented as a left-acting deformation ϕ of a common reference \bar{S} . We further assume that \bar{S} is discretized as a simplicial surface mesh with l vertices and m triangles. In order to perform analysis on local geometric details rather than absolute coordinates of a shape $S = \phi(\bar{S})$, we can employ a differential representation given by the deformation gradient $\nabla\phi$, i.e. the 3×3 matrix of partial derivatives of ϕ (cf. previous section). Let ϕ be orientation-preserving and affine on each triangle \bar{T}_i of \bar{S} , then the derivatives are constant on each triangle with $\nabla\phi|_{\bar{T}_i} \equiv D_i \in GL^+(3)$. Note, the deformation of a triangle fully specifies an affine map of \mathbb{R}^3 if we assume that triangle normals are mapped onto each other (cf. Kirchhoff–Love kinematic assumptions). Accordingly, a representation of a shape S in linear differential coordinates is given by $s = (D_1, \dots, D_m)^T$.

In order to derive information of our geometric data we perform element-wise geometric statistics on it. Let $\{s_j = (D_1^j, \dots, D_m^j)^T\}_{j=1}^n$ be the set of all input shapes represented in differential coordinates. Following our Definition 2.55 the essential components to set up an SSM are mean value and TPCA to analyze the input as deviations thereof.

3.3.1. Bi-invariant Mean

Since $GL^+(3)$ does not admit a bi-invariant metric there can not exist a bi-invariant *Riemannian* mean. Nevertheless, due to the Lie group structure there exists a naturally bi-invariant candidate for the mean in terms of the group exponential barycenter called *bi-invariant mean*. We follow hereby the work of Pennec and Arsigny [Pennec and Arsigny, 2013] who delivered a comprehensive characterization and analysis on this topic (cf. Chapter 2). The bi-invariant mean μ_i is defined through:

$$\sum_{j=1}^n \log(D_i^j \cdot \mu_i^{-1}) = 0, \quad (3.10)$$

where \log denotes the group logarithm. To solve for the unknown μ_i we apply an iterative fixed point scheme:

$$\mu_i^{k+1} = \exp\left(\sum_{j=1}^n \log\left(D_i^j \cdot (\mu_i^k)^{-1}\right)\right) \cdot \mu_i^k, \quad (3.11)$$

where \exp denotes the group exponential. The local existence and uniqueness of the bi-invariant mean have been proven for data with small enough dispersion, i.e. if the data lies within a sufficiently small normal convex neighborhood of some point of the Lie group. Furthermore, the algorithm given by Equation (3.11) always converges to μ_i at least with linear speed provided that the initialization is chosen sufficiently close to the data.

From Equation (3.11) we see that the group logarithm and exponential of $GL^+(3)$ are essential operations required to determine the mean shape as well as for the statistical analysis in its tangent space (Section 3.3.2). However, it should be emphasized that there does not exist a *real* logarithm for every element in $GL^+(3)$. We can classify such elements by investigating the underlying eigenvalue structure. Let D be an arbitrary element in $GL^+(3)$. It is known that there always exists a *real* Jordan-Decomposition [Gallier, 2018] $D = V \cdot E \cdot V^{-1}$ s.t. E belongs (modulo permutation) to one of the following three types:

$$\begin{aligned} \text{A: } & \begin{pmatrix} \lambda_1 & \alpha & 0 \\ 0 & \lambda_2 & \beta \\ 0 & 0 & \lambda_3 \end{pmatrix}, & \text{where } \lambda_i \in \mathbb{R}^+, \alpha, \beta \in \{0, 1\} \\ \text{B: } & \begin{pmatrix} -\lambda_1 & 0 & 0 \\ 0 & -\lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix}, & \text{where } \lambda_i \in \mathbb{R}^+, \lambda_1 \neq \lambda_2, \\ \text{C: } & \begin{pmatrix} \lambda_1 & \nu & 0 \\ -\nu & \lambda_1 & 0 \\ 0 & 0 & \lambda_2 \end{pmatrix}, & \text{where } \lambda_1 + i\nu \in \mathbb{C} \setminus \mathbb{R}_0^+, \lambda_2 \in \mathbb{R}^+. \end{aligned}$$

As the logarithm is compatible with a change of basis it is enough to consider only matrices of the above form. Both cases A and C admit a real logarithm, contrasting case B that does not allow for its existence. This raises the question what deformation gradients could feature such an eigenvalue configuration and whether it is likely to appear. If we take a closer look at case B we see that it encodes an anisotropic scale with two negative weights. Since the respective deformation is orientation preserving it must invert two edges of a triangle

and change their lengths in a non-uniform fashion. This seems to be a rather unlikely deformation if we consider data to be aligned and without artifacts such as local overfolds. In particular, the two real word datasets we performed our experiments on (Section 4.2.1, Chapter 5) did not admit any element in any input shape that came across with a deformation gradient of this structure. Neither during calculation of the mean nor during analysis.

However, in order to do statistics in $GL^+(3)$ that are robust to such extreme cases we require an alternative auxiliary strategy. To this end, we propose to perform a *pseudo* logarithm operation. Let D be an element of $GL^+(3)$ with no real logarithm. We define its pseudo logarithm *plog* employing polar decomposition as follows:

$$\text{plog}(D) = \text{plog}(RU) := \log(R) + \log(U)$$

In case a real $\log(D)$ exists this formula can be seen as first order (commutator free) approximation in terms of the Baker–Campbell–Hausdorff formula and for commuting R, U this formula would be even exact. This can additionally be interpreted as a fallback to the product structure of the DCM (Section 3.4). Contrary to the logarithm, the matrix exponential always exists and can efficiently be calculated using the scaling-and-squaring method together with Padé approximations [Higham, 2005].

3.3.2. Tangent Principal Component Analysis

In the previous section we were able to circumvent the absence of a (bi-invariant) metric but this is no longer possible if we want to perform higher-order analysis employing TPCA. While there is no bi- $GL^+(3)$ -invariant metric, we are interested in metrics that yield at least invariance under orthogonal transformations, i.e. metrics that are invariant with respect to a change of coordinates obtained by rotating or mirroring the data. Indeed, there exists exactly one family of metrics that is $GL^+(3)$ -left-invariant and $O(3)$ -right-invariant and uniquely determined up to three positive real constants [Martin and Neff, 2016]. We define the metric as usual via the inner product on the respective Lie algebra (cf. Theorem 2.33).

Let $X, Y \in \mathfrak{gl}(3) = \mathbb{R}^{3 \times 3}$ and $\tau, \nu, \kappa \in \mathbb{R}^+$:

$$\langle X, Y \rangle_{\tau, \nu, \kappa} := \tau \langle \text{dev sym } X, \text{dev sym } Y \rangle + \nu \langle \text{skew } X, \text{skew } Y \rangle + \frac{\kappa}{3} \text{tr}(X) \text{tr}(Y),$$

where we have used the following notation:

$$\begin{aligned} \langle X, Y \rangle &= \langle X, Y \rangle_F = \text{tr}(X^T Y) && \text{(standard inner product),} \\ \text{sym } X &= \frac{1}{2}(X + X^T) && \text{(symmetric part of } X), \\ \text{skew } X &= \frac{1}{2}(X - X^T) && \text{(skew-symmetric part of } X), \\ \text{dev } X &= X - \frac{\text{tr } X}{3} I_3 && \text{(deviator of } X). \end{aligned}$$

If we consider X as *infinitesimal transformation* the above terms admit certain geometric interpretations: skew X represents the rotational part and sym X the distortion part. While the trace tr quantifies volume changes, the deviator dev represents the trace-free part and, hence, $\text{dev sym } X$ describes the shearing (volume-preserving distortion) part of X . Furthermore, the above inner product features two interesting properties:

$$\begin{aligned} \langle X, Y \rangle_{1,1,1} &= \langle X, Y \rangle \quad \text{for all } X, Y \in \mathfrak{gl}(3), \\ \langle X, Y \rangle_{\tau,\nu,\kappa} &= 0 \quad X \in \mathfrak{so}(3), Y \text{ symmetric.} \end{aligned}$$

Hence, this family of metrics can be seen as natural generalization of the standard metric arising from the standard inner product for matrices. Let us assume to have n input shapes with m triangles each, then we perform TPCA in the tangent space $T_\mu(\text{GL}^+(3))^m$ at the differential coordinates of the mean shape $\mu = (\mu_1, \dots, \mu_m)$. The k -th mode of variation is hereby given as:

$$\vartheta_k = \underset{\substack{\|v\|^{\tau,\nu,\kappa}=1 \\ g_\mu^{\tau,\nu,\kappa}(\vartheta_i, v)=0, i < k}}{\arg \max} \sum_{j=1}^n g_\mu^{\tau,\nu,\kappa}(\log_\mu(s_j), v)^2, \quad k = 1, \dots, n-1. \quad (3.12)$$

where $s_j = (D_1^j, \dots, D_m^j)$, \log is applied component-wise and $g_\mu^{\tau,\nu,\kappa} = \sum \bar{A}_i \cdot g_{\mu_i}^{\tau,\nu,\kappa}$ is the metric emerging from $\langle \cdot, \cdot \rangle_{\tau,\nu,\kappa}$ scaled with the triangle area of the respective element of the reference shape.

Putting it all together, the tuple $(\mu, \{\vartheta_k\}_1^d)$ consisting of mean μ and modes of variation $\{\vartheta_k\}$ rises to an SSM as specified in Definition 2.55. It is furthermore, enabled through Equation (3.9), a generative model.

3.3.3. 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, in other words $\bar{S} = \bar{S}(\mu)$, i.e. the reference shape equals the geometric realization of μ . 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.

In summary: The given $GL^+(3)$ shape model comprises an exponential barycenter based on the affine connection structure on $GL^+(3)^m$, since $GL^+(3)$ does not admit bi-invariant metric. The modes of variation are determined by TPCA carried out utilizing the metric closest as possible to bi-invariant. It is worth mentioning that the shape representation is, due to its simple structure, compatible with piecewise (triangle-wise) affine transformations of the input shapes, i.e. taking the differential representation and applying the transformation commutes. However, although invariant under translation, the barycenter and the modes of variation do depend on the input shapes’ orientation in the ambient space calling for an initial generalized Procrustes alignment of the input to minimize the bias introduced through data misalignment.

3.4. Differential Coordinates Model

As we have seen in Section 3.2.1 calculation of the mean in differential coordinates can be carried out point-wise. However, since we are now assuming triangular meshes as discretizations, we find the deformation gradient triangle-wise constant, allowing to perform calculations at triangle-level.

To compute the local intrinsic mean, we derive a distance d_F that is tailored to our differential coordinates. Let $D = \mathcal{s}(p)$ be the deformation gradient at $p \in \bar{\mathcal{S}}$. By the polar decomposition from linear algebra, D can be uniquely factored as $D = RU$, 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

$$F := \text{SO}(3) \times \text{Sym}^+(3), \quad (3.13)$$

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

While there exists a well-developed theory for $\text{SO}(3)$ with a natural notion of distance, there is no canonical setup for stretch tensors in $\text{Sym}^+(3)$. 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 $\text{Sym}^+(3)$ into a commutative Lie group (with Lie algebra $\mathfrak{sym}^+(3)$) based on the logarithmic multiplication defined as

$$U \circ U' := \exp(\log(U) + \log(U')). \quad (3.14)$$

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

$$\langle u, u' \rangle_{\mathfrak{sym}^+(3)} := \langle u, u' \rangle = \text{tr}(u^T u'), \quad (3.15)$$

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

$$d_{\text{Sym}^+(3)}(U, U') = \|\log(U') - \log(U)\|_F, \quad (3.16)$$

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, F is itself a Lie group and we can define the distance d_F as

$$d_F^2 = \omega d_{\text{SO}(3)}^2 + d_{\text{Sym}^+(3)}^2. \quad (3.17)$$

Here, ω is a positive weighting factor (see also Section 3.2.1) and $d_{\text{SO}(3)}$ denotes the canonical distance on $\text{SO}(3)$

$$d_{\text{SO}(3)}(R, R') = \|\log(R^T R')\|_F \quad (3.18)$$

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

- Geodesic completeness of F , i.e. invalid tensors are at an infinite distance from any element in F ,
- 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. Putting it all together, our discrete feature space can simply be defined as $G = F^m$ as we have one gradient per triangle. This leads us directly to the discretized version of Equation (3.3)

$$d_G(s_1, s_2) = \left(\sum_{j=1}^m \bar{A}_j \cdot d_F(s_1^j, s_2^j)^2 \right)^{\frac{1}{2}}, \quad (3.19)$$

where \bar{A}_j is the triangle area of the j -th triangle of the reference surface. As for bi-invariant Lie groups Riemannian exponential/logarithm and group exponential/logarithm (and hence Riemannian and group mean) coincide we can efficiently solve for the mean μ in differential coordinates employing a Gauss-Newton descent algorithm [Pennec, 2006; Arsigny et al., 2006]:

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

The construction of an anatomical atlas can then be carried out analogously to the description in the previous section on the GLM.

In order to perform TPCA we need a discrete version of the Riemannian metric given in Equation (3.4). We employ the following:

$$g_\mu(v, w) = \sum_{j=1}^m \bar{A}_j \cdot \langle v_j, w_j \rangle_{\mu_j}. \quad (3.20)$$

Moreover, in line with the definition of d_F in Equation (3.17), the rotational component in $\langle \cdot, \cdot \rangle_{\mu(p)}$ 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.

Let us recall the definition of the main modes of variation and assume to have n input shapes with m triangles each, then we perform TPCA in the tangent space $T_\mu G$ at the differential coordinates of the mean shape μ . The k -th mode of variation is hereby given as:

$$\vartheta_k = \arg \max_{\|v\|=1} \sum_{j=1}^n g_\mu(\log_\mu(s_j), v)^2, \quad k = 1, \dots, n-1. \quad (3.21)$$

$g_\mu(\vartheta_i, v) = 0, \quad i < k$

To practically solve for the main modes algorithmically we perform eigen-decomposition of the Gram matrix $(c_{ij})_{ij}$, with $c_{ij} = g_\mu(\log_\mu(s_i), \log_\mu(s_j))$ (cf. Younes [2010] Section E.2.2).

As in the previous section on the GLM, the tuple $(\mu, \{\vartheta_k\}_1^d)$ consisting of mean μ and modes of variation $\{\vartheta_k\}$ rises to an SSM as defined in Definition 2.55. It is furthermore, enabled through Equation 3.9, a generative model. We will refer to it as the DCM.

In summary: The developed shape model based on differential coordinates comprises a Riemannian center of mass due to a bi-invariant metric structure introduced on G and the modes of variation determined by TPCA utilizing this metric. The given shape representation, compared to the previous $GL^+(3)$ -based one, is as well invariant under translation and comes with the additional advantage that a *real* logarithm is always defined. However, mean shape and modes of variation do depend on the input shapes' orientation in the ambient space posing the need for an initial generalized Procrustes alignment of the input to minimize the bias introduced through misalignment. Furthermore, problems may arise due to the ambiguity of rotations by angles of $(\pm\pi)$, i.e. large rotational deformation might induce discontinuous rotations for neighboring triangles (Figure 4.8).

3.5. Fundamental Coordinates

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. [Pennecc and Lorenzi, 2020]).

3.5.1. Fundamental Forms and Surface Theory

As shape statistics should *neither* depend on a shape’s orientation *nor* a shape’s position in the ambient space, we are in need of a shape representation that is invariant under Euclidean motion. In order to derive this representation we will at first take a short side trip into classical 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] Section 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, i.e. per connected component) 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.

3.5.2. 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 \bar{S} given as triangulated surface. Let deformation ϕ be affine on each triangle \bar{T}_i of \bar{S} , then the deformation gradient $\nabla\phi$ is the 3×3 matrix of partial derivatives of ϕ and constant on each triangle




Figure 3.1.: Relations between reference shape \bar{S} (left) and shape $S = \phi(\bar{S})$, a deformation thereof (right), s.t. $D_i := \nabla\phi|_{\bar{T}_i}$. Note that each frame $F_i = R_i\bar{F}_i$ is defined solely on the respective triangle T_i and all neighboring frames are connected across the shared edge of their underlying triangles via $F_i C_{ij} = F_j$.

$D_i := \nabla\phi|_{\bar{T}_i}$ (see e.g. Botsch et al. [2006] for detailed expressions). Note that transition from deformation to deformation gradient 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 $\tilde{U}_i := [\bar{F}_i^T U_i \bar{F}_i]_{3,3} = \mathbb{I}|_{\bar{T}_i}^{1/2}$, 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. $\mathbb{I}_p(v, w) = \mathbb{I}_p(-dN_p(v), w)$ 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 Figure 3.1.

3.6. Solving the Inverse Problem

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] or Section 3.2). However, in our representation the rotations are only given implicitly in terms of the transition rotations. In particular, an immediate computation shows that $R_j = R_i \bar{F}_i C_{ij} \bar{F}_j^T =: R_{i \rightarrow j}$ for an integrable field $\{C_{ij}\}$. Based on this condition, for each triangle T_i we can formulate a residual term

$$\varepsilon_i(\phi, \{R_i\}) = \sum_{j \in \mathcal{N}_i} \frac{1}{|\mathcal{N}_i|} \|D_i - R_{j \rightarrow i} U_i\|_F^2$$

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.

Efficient shape reconstruction 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 Section 3.2.2 we will keep its description at this point rather brief. 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 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 \bar{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.




Figure 3.2.: Initialization procedure propagating an initial rotation R_{i_0} along a pre-computed spanning tree across the data. The red digits indicate the number of propagation steps needed to reach a certain point.

R_{i_0} is propagated along this path employing the local integrability condition, viz. $R_j := R_{i \rightarrow j} = R_i \bar{F}_i C_{ij} \bar{F}_j^T$. 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 $\{C_{ij}\}$ come from an integrable system. In that case the initialization is additionally invariant w.r.t. the choice of the fixed triangle and spanning tree. Figure 3.2 provides a schematic overview on the initialization process.

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_i ,

s.t. the problem decouples into individual low-dimensional optimizations that can be solved in closed-form and allow for massive parallelization. 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 3.1. *The local step within the shape reconstruction algorithm targeting the optimization problem:*

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

where \mathcal{N}_i is the set of indices belonging to edge neighbors of triangle i , can be solved in closed form and the solution is unique.

Proof. For the sake of simplicity let $D_s = \nabla \phi|_{\bar{T}_s}$. We carry out a direct calculation utilizing the definition $\langle A, B \rangle_F := \text{tr}(A^T B)$ and the trace's invariance under cyclic permutations:

$$\begin{aligned} R_i &= \arg \min_{R \in \text{SO}(3)} \sum_{s \in \mathcal{N}_i} \left\| D_s - R \bar{F}_i C_{is} \bar{F}_s^T U_s \right\|_F^2 \\ &= \arg \min_{R \in \text{SO}(3)} \underbrace{\sum_{s \in \mathcal{N}_i} \|D_s\|_F^2}_{\text{const.}} - 2 \underbrace{\langle D_s, R \bar{F}_i C_{is} \bar{F}_s^T U_s \rangle_F}_{\text{const.}} + \underbrace{\|R \bar{F}_i C_{is} \bar{F}_s^T U_s\|_F^2}_{\text{const.}} \\ &= \arg \max_{R \in \text{SO}(3)} \sum_{s \in \mathcal{N}_i} \text{tr} \left(D_s^T R \bar{F}_i C_{is} \bar{F}_s^T U_s \right) \\ &= \arg \max_{R \in \text{SO}(3)} \sum_{s \in \mathcal{N}_i} \langle D_s U_s^T \bar{F}_s C_{is}^T \bar{F}_i^T, R \rangle_F = \arg \max_{R \in \text{SO}(3)} \langle D_{\mathcal{N}_i}, R \rangle_F \end{aligned}$$

Since $D_{\mathcal{N}_i}$ is a nonsingular and orientation-preserving matrix it can be uniquely decomposed via polar decomposition to $R_{\mathcal{N}_i} U_{\mathcal{N}_i}$, where $R_{\mathcal{N}_i} \in \text{SO}(3)$ and $U_{\mathcal{N}_i} \in \text{Sym}^+(3)$ s.t.

$$R_{\mathcal{N}_i} = \arg \max_{R \in \text{SO}(3)} \langle D_{\mathcal{N}_i}, R \rangle_F.$$

□

Figure 3.3 schematically illustrates the underlying neighboring relations framing the local integrability constraints.




Figure 3.3.: Neighboring relations employed in the local step of the integration procedure. Rotations R_i are connected by transition rotations C_{ij} .

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 preprocess allowing for very efficient global solves with close to linear cost.

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.

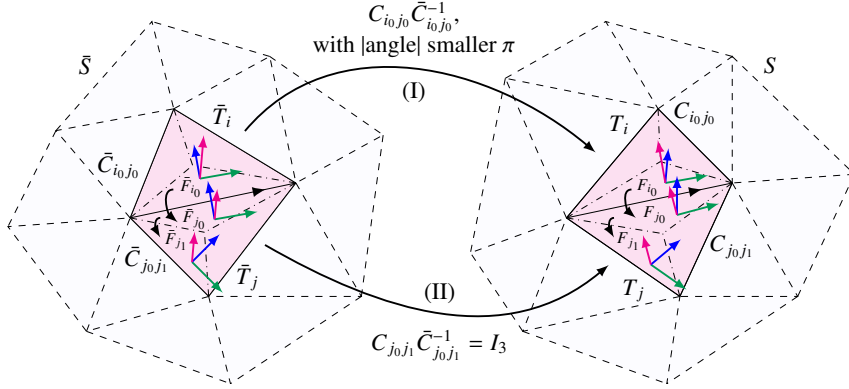


Figure 3.4.: Schematic summary of the argumentation to prove the proposition on relative transition rotations. The given construction allows to separate relative transition rotations into normal (I) and tangential (II) type.

3.6.1. Properties of Relative Transition Rotations

As with other non-Euclidean approaches, existence and uniqueness of the intrinsic mean is only ensured for well-localized data. In particular, for our representation this concerns the rotational components describing the changes in curvature. We would like to remark that this is a rather academic discussion as we did not encounter any example with critical disparity. Indeed, even for the synthetic PIPE dataset (cf. Appendix A.1) representing a severe nonlinear deformation the relative transition rotations are located in a small neighborhood of radius $5\pi/23$, see Figure 3.5. However, the following proposition explains how to (theoretically) control the *relative transition rotations* and thus how to avoid ambiguities regarding the rotational logarithm.

Proposition 3.2. *For any given n shapes S_1, \dots, S_n there exists a common discretization and a frame field, such that all relative transition rotations exhibit angles in $(-\pi, \pi)$. Furthermore, if the angle defect introduced to relative transition rotations through shearing of any two pairs of neighbouring triangles is smaller than $|\pi/2|$ the corridor of possible angles narrows down to $(-\pi/2, \pi/2)$.*

Proof. At the beginning we have a look at the following formula [Altmann, 2005] relating the angle θ^{12} of a composite rotation with those of the two sequential ones (θ^1, θ^2) and their axes of rotation (v^1, v^2) :

$$\cos\left(\frac{\theta^{12}}{2}\right) = \cos\left(\frac{\theta^1}{2}\right)\cos\left(\frac{\theta^2}{2}\right) - \sin\left(\frac{\theta^1}{2}\right)\sin\left(\frac{\theta^2}{2}\right)\langle v^1, v^2 \rangle. \quad (3.22)$$




Figure 3.5.: 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.

Employing standard calculus we can deduce from the above that for minimal/-maximal angles it has to hold that $\theta^{12} = \theta^1 \pm \theta^2$ and $\langle v^1, v^2 \rangle = \pm 1$. We keep this in mind and now poof the result for two shapes \bar{S} , S since the argumentation naturally extends to the whole n shapes. As the vector \bar{v}_{ij} defining the edge connecting two triangles \bar{T}_i, \bar{T}_j of \bar{S} plays an important role in the following, we will detail on how it is deformed to the corresponding counterpart v_{ij} of S . As before we denote $R_i U_i = D_i = \nabla \phi|_{\bar{T}_i}$ determined by polar decomposition and as the gradients have to agree on shared edges $D_i \bar{v}_{ij} = D_j \bar{v}_{ij} = v_{ij}$ holds. However, in general and contrary to this, we have to point out that $v_{ij} \neq R_i \bar{v}_{ij} \neq R_j \bar{v}_{ij} \neq v_{ij}$. In the special situation that U_i, U_j are both realizing isotropic stretch in tangential direction, we remain with the weaker equality $R_i \bar{v}_{ij} / \|\bar{v}_{ij}\| = R_j \bar{v}_{ij} / \|\bar{v}_{ij}\| = v_{ij} / \|v_{ij}\|$. Yet, for the common case of anisotropic stretching even the resulting directions are not equal any longer. This phenomenon is due to the *shearing* components of D_i, D_j , which are essentially the non-isotropic residuals of the respective stretches, and the compensation of the hence introduced angle defect between \bar{v}_{ij} and $U_i \bar{v}_{ij}, U_j \bar{v}_{ij}$ by means of R_i, R_j . This angle defect is bounded, as the shearing angle itself by $\pm \pi/2$ (at-

tion: angel defect \neq shearing angle in general!). These bounds are not sharp as extreme angles correspond to an infinite length of $U_i \bar{v}_{ij}, U_j \bar{v}_{ij}$, i.e. an infinite ratio $\lambda_i^{\max}/\lambda_i^{\min}, \lambda_j^{\max}/\lambda_j^{\min}$ of the largest and smallest stretching eigenvalue. In fact, as we operate in finite realm, we can assume that the angle defect resides in $(-(\pi/2 - \delta), \pi/2 - \delta)$ for some $\delta > 0$ and all triangles of S . Now, 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 $(-\varepsilon, \varepsilon)$ with $\varepsilon < \delta$. This can always be achieved as the shearing is not altered by simple subdivision. From this point onwards the argumentation is additionally summarized in Figure 3.4. 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 \bar{T}_l of \bar{S} (analogously for T_l of S) into three parts $\bar{T}_{l_0}, \bar{T}_{l_1}, \bar{T}_{l_2}$ by means of the incenter and the bisecting lines of the angles. Without loss of generality we can assume that \bar{T}_{j_0} is neighboring \bar{T}_{i_0} (thus \bar{T}_j was already neighboring \bar{T}_i). We fix a frame \bar{F}_{i_0} on \bar{T}_{i_0} ensuring alignment of the first basis vector to the edge shared with \bar{T}_{j_0} . Frame \bar{F}_{j_0} is now defined by rotating \bar{F}_{i_0} around the common edge. This directly implies that $\bar{C}_{i_0j_0}$ realizes an angle with absolute value smaller than ε . For $C_{i_0j_0}$ the situation is slightly more intricate since $F_l = R_l \bar{F}_l$ does not preserve alignment of the frames with the underlying triangles. However, we can decompose $C_{i_0j_0} = Q_{j_0}^{-1} Q_{i_0j_0} Q_{i_0}$, where Q_{i_0} is the rotation s.t. the first basis vector $(F_{i_0})_0 = R_i(\bar{F}_{i_0})_0$ rotates by angle θ_{i_0} around the normal of T_{i_0} ($(F_{i_0})_2$) onto $v_{i_0j_0}$; $Q_{i_0j_0}$ rotates $Q_{i_0}(F_{i_0})_2$ by angle $\theta_{i_0j_0}$ around $v_{i_0j_0}$ and $Q_{j_0}^{-1}$ is the rotation s.t. v_{ij} rotates by angle $-\theta_{j_0}$ around the normal of T_{j_0} onto $(F_{j_0})_0$. In other words: Q_{i_0} and $Q_{j_0}^{-1}$ handle the angle defect due to shearing in D_i/D_j , whereas $Q_{i_0j_0}$ realizes the normal change between T_i and T_j . We analogously define \bar{F}_{i_k} if \bar{T}_{i_k} has neighbors, if not, we simply set $\bar{F}_{i_k} = \bar{F}_{i_0}$. This construction allows to explicitly differentiate two different types of relative transition rotations: type (I) that comes from normal differences (and shearing defects) like $C_{i_0j_0} \bar{C}_{i_0j_0}^{-1}$ and type (II) like $C_{j_0j_1} \bar{C}_{j_0j_1}^{-1}$ that is induced by (non-shearing) tangential change. Since $F_{j_0} = R_j \bar{F}_{j_0}$ and $F_{j_1} = R_j \bar{F}_{j_1}$ we see immediately that $C_{j_0j_1} = \bar{F}_{j_0}^{-1} R_j^{-1} R_j \bar{F}_{j_1} = \bar{C}_{j_0j_1}$ and thus $C_{j_0j_1} \bar{C}_{j_0j_1}^{-1} = I_3$. To clarify (I) let $\phi_{i_0j_0}/-\bar{\phi}_{i_0j_0}$ be the angles of $C_{i_0j_0}/\bar{C}_{i_0j_0}^{-1}$ and $\varphi_{i_0j_0}$ the angle of their composition. We now calculate the bounds of $\varphi_{i_0j_0}$ employing (3.22) together with the decomposition $C_{i_0j_0} = Q_{j_0}^{-1} Q_{i_0j_0} Q_{i_0}$, the domains of the therein appearing

angles and lastly $\varepsilon < \delta$:

$$\begin{aligned} \min \varphi_{i_0 j_0} &\geq \min \phi_{i_0 j_0} + \min -\bar{\phi}_{i_0 j_0} \geq \min \theta_{i_0} + \min \theta_{i_0 j_0} + \min -\theta_{j_0} + \min -\bar{\phi}_{i_0 j_0} \\ &= -(\pi/2 - \delta) - \varepsilon - (\pi/2 - \delta) - \varepsilon > -\pi, \\ \max \varphi_{i_0 j_0} &\leq \max \phi_{i_0 j_0} + \max -\bar{\phi}_{i_0 j_0} \leq \max \theta_{i_0} + \max \theta_{i_0 j_0} + \max -\theta_{j_0} + \max -\bar{\phi}_{i_0 j_0} \\ &= (\pi/2 - \delta) + \varepsilon + (\pi/2 - \delta) + \varepsilon < \pi. \end{aligned}$$

Finally, we have seen that $\varphi_{i_0 j_0}$ is always located in $(-\pi, \pi)$, hence is well-localized. The second statement too follows directly from the above inequalities as a smaller angle defect effectively reduces the bounds of the $\varphi_{i_0 j_0}$. \square

3.6.2. Quasi-Isometric Surface Flattening

In the following section we focus on the quasi-isometric surfaces flattening in the space of fundamental coordinates, i.e. the natural approach within that framework to determine a quasi-isometric chart for a given embedded 2D surfaces in 3D space. We of course assume that the surface under study allows for a single chart, otherwise the presented approach can be applied to a suitable submanifold.

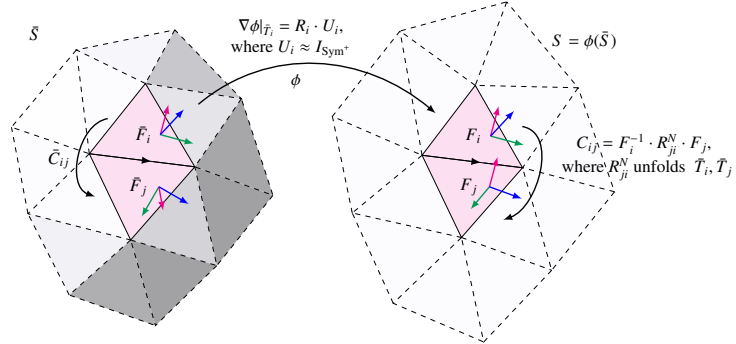


Figure 3.6.: Non-flat surface \bar{S} is employed as reference within the deformation setup. Flat surface S is determined via deformation of \bar{S} by ϕ , s.t. metric distortion, i.e. U_i , is close to identity and R_i is determined by means of C_{ij} that are normal vector fixing modifications of \bar{C}_{ij} .

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




Figure 3.7.: 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.

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 (Figure 3.7). 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 submanifold 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 $\text{SO}(2)^n \times \text{Sym}^+(2)^m$ (cf. next section). See Figure 3.6 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. previous section).

3.7. Fundamental Coordinates Model

In the following section, we will introduce the FCM by first specifying a group structure with notion of distance and second defining mean and modes of variation.

3.7.1. Group Structure and Distance

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. A concise description of that topic is also given in Section 2.2.

Our shape representation consists of transition rotations $C_{ij} \in \text{SO}(3)$ (one per inner edge) and tangential stretches $\tilde{U}_i \in \text{Sym}^+(2)$ (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

$$d_\omega^2(s, t) = \frac{\omega^3}{\bar{A}_\mathcal{E}} \sum_{(i,j) \in \mathcal{E}} \bar{A}_{ij} d_{\text{SO}(3)}^2(C_{ij}^s, C_{ij}^t) + \frac{\omega}{\bar{A}} \sum_{i=1}^m \bar{A}_i d_{\text{Sym}^+(2)}^2(\tilde{U}_i^s, \tilde{U}_i^t), \quad (3.23)$$

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] Section 4.1).

3.7.2. 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 (cf. Chapter 2). 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 can be efficiently computed using the Newton-type 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 half 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 $\pm\pi/2$, what can be brought in relation to too large shearings of pairs of neighbouring triangles. (cf. Proposition. 3.2). As framework for second order statistics we employ LPGA [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 of Equation (3.23). 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] Section 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 3.6. The pair $(\mu, \{\vartheta_k\}_1^d)$ consisting of mean shape and modes of variation constitutes an SSM in the sense of Definition 2.55. It is furthermore a generative model, an algorithm to solve the inverse problem, from shape representation to surface, is given.

In summary: The fundamental coordinates model as introduced above comprises a Riemannian center of mass for a bi-invariant distance and modes of variation determined by TPCA and the respective Riemannian metric. In contrast to the previous two shape representations, the fundamental coordinates are invariant under Euclidean motions and thus allow for alignment-free shape analysis.

CHAPTER 4.

Experimental Model Evaluation and Comparison

In this chapter we want to evaluate the SSMs proposed in the previous chapter both quantitatively and qualitatively w.r.t. certain well-known standard metrics. Furthermore, we perform a comparison to the classic PDM and discuss the outcome.

4.1. Quantitative Assessment

We evaluate the quality of the models under study in terms of specificity, generalization ability, and compactness. A brief description of these measures is given in the following, however, a more detailed discussion of those can be found in Davies et al. [2008].

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:

$$\text{spec}(M) = \frac{1}{N} \sum_{i=1}^N \min_{j=1, \dots, n} \mathcal{W}(M(\alpha_i), S_j).$$

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 Euclidean L^2 or any distance directly related to the previously described SSMs (cf. Sections 3.4, 3.7) could introduce a bias for PDM or any of the proposed models, respectively. Therefore, we employ the independent measure \mathcal{W} that was shown to successfully quantize nonlinear variations in shape [Zhang et al., 2015; Grinspun et al., 2003] known as the *Discrete Shells Energy*. The discrete shells energy is physically plausible for thin shell materials by capturing two important modes of deformation: bending and stretching [Heeren et al., 2018].

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:

$$\text{gen}(M) = \frac{1}{n} \sum_{\substack{j=1 \\ j \neq i}}^n \mathcal{W}(\pi_{M_i}(s_i(S_j)), S_j).$$

where $M_i = M_i(S_1, \dots, \hat{S}_i, \dots, S_n) = (\mu, \{v_k\}_1^d)$ is the model constructed utilizing all input shapes but S_i , $s_i = s_i(S_i)$ denotes the shape space representation of S_i and $\pi_{M_i}(s_i)$ the s_i -distance minimizing element in M_i , i.e. $\pi_{M_i}(s_i) = \arg \min_{\pi \in M_i} d^2(\pi, s_i)$. Lower values correspond to a higher ability to generalize to unseen data. As for specificity, we report generalization ability for the physically-motivated energy \mathcal{W} .

Compactness measures the relative variability within a training set as captured by a SSM. Specifically, it is the ratio of the variance encoded by the k most dominant modes and total variance in the set,

$$\text{com}(M) = \sum_{i=1}^k \frac{\lambda_i}{\sum_{i=1}^d \lambda_i},$$

where the λ_i denote the variances as defined for the respective model (cf. Section 3.2.1). In Figure 4.3 we plot the compactness as a function of k for each of the constructed models.

4.1.1. DCM – Specificity, Generalization Ability, Compactness

We report specificity results for both the Fine Alignment Using Scan Texture (FAUST) and Osteoarthritis Initiative (OAI) datasets (cf. Appendix. A) in Figure 4.2. It can be seen that the proposed method outperforms the other approaches for the physically based measure \mathcal{W} .

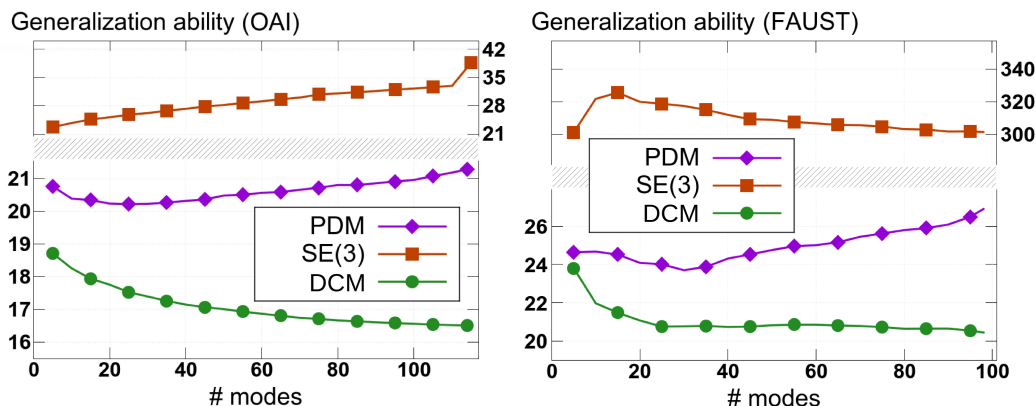


Figure 4.1.: Generalization ability as function of the number of modes (model parameters) for \mathcal{W} of the DCM, PDM, and SE(3)-based model [Hefny et al., 2015].

Figure. 4.1 summarizes the results for generalization ability. In total, the proposed method features a better generalization ability than the other approaches.

The limited performance of the method by Hefny et al. [2015] (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 projecting from SE(3)-based shape space to global vertex positions since each vertex is contained in several triangles that in general

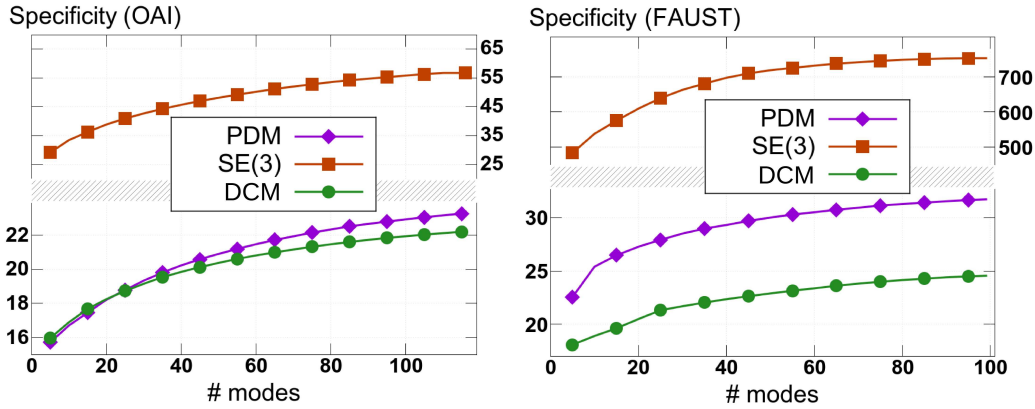


Figure 4.2.: Specificity as function of the number of modes (model parameters) for \mathcal{W} of the DCM, PDM, and SE(3)-based model [Hefny et al., 2015].

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 Figure 4.7, where we provide a visualization of the mean shapes.

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 Figure 4.7), yet account for a substantially different portion (63% and 36% respectively) of the total variance.

4.1.2. FCM – Specificity, Generalization Ability, Compactness

We perform a quantitative comparison with PDM and DCM using the above mentioned standard measures and the physically-based surface energy \mathcal{W} on the OAI dataset. The results show that the FCM is more specific than PDM and DCM (Figure 4.4 center). In terms of generalization ability, the FCM is superior to PDM, yet inferior to DCM (Figure 4.4 left). Finally, the FCM is less compact than PDM and DCM (Figure 4.4 right). Note that compactness

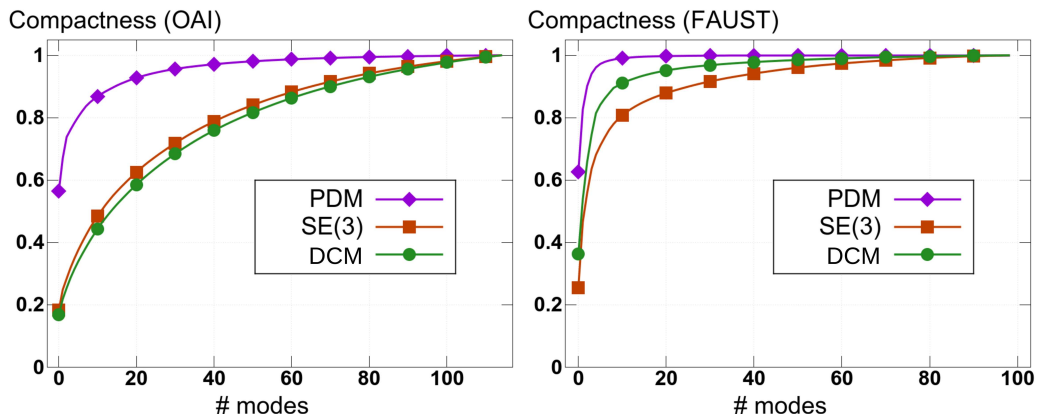


Figure 4.3.: Compactness as function of the number of modes (model parameters) of the DCM PDM, and SE(3)-based model [Hefny et al., 2015].




Figure 4.4.: Generalization ability (left, lower~better), specificity (middle, lower~more specific) and compactness (right, higher~more compact) of the FCM, PDM, and DCM on the OAI dataset.

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 accuracy for disease classification (cf. Appendix. A.2.2).

4.1.3. FCM – 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 nonlinearity in the inverse problem.

4.2. Qualitative Assessment

We qualitatively assessed the model performance of GLM, DCM, and FCM w.r.t. the validity (correctness) of generated shape instances. In a medical or biological context »correctness« can be understood as anatomical plausibility.

4.2.1. GLM – Validity

We perform two qualitative experiments. At first, a comparison of the mean shape of the OAI dataset as determined by the DCM as well as the GLM. To achieve this we align both shapes and calculate the surface distance between them. Both mean shapes are highly similar as can be seen in Figure 4.5. At second, an analysis of the skeletal hand dataset (cf. Appendix A) is performed. We calculate the mean shape of the two input poses, perform TPCA and (visually) investigate the resulting trajectory connecting the two input shapes through the mean w.r.t. plausibility. As shown in Figure 4.6 the principal mode shows natural nonlinear deformation characteristics.

4.2.2. DCM – Validity

In order to get a qualitative impression of the DCM’s ability to generate plausible shapes in presence of large deformations we take two samples from the FAUST dataset, one person with arms up, and the same person with arms down, and calculate the interpolating geodesic between them. Figure 4.7 depicts this experiment for the DCM, the PDM, and the SE(3)-based model [Hefny et al., 2015] (interpreting triangle-wise deformations as elements in SE(3)). It can clearly be seen that the DCM correctly captures the expected interpolating motion between the two inputs, whereas the PDM fails since this linear approach linearizes rotational components of the underlying deformation and hence introducing heavy distortion. The SE(3) model already fails to




Figure 4.5.: Deviations of mean distal femur shape as calculated with the GLM model and the DCM. Absolute values of the surface distance are plotted color-coded on the DCM mean shape.

calculate a mean shape since it does not provide a mechanism to consistently stitch neighboring triangles together.

4.2.3. FCM – Validity

Frequently, datasets feature a high nonlinear variability that are characterized by large rotational components, which are insufficiently captured by linear models like PDM (previous subsection). While DCM treats the rotational components explicitly, it requires them to be well-localized, s.t. the logarithm is




Figure 4.6.: Trajectory as calculated with the GLM connecting the input shapes (left, right) via the exponential mean (center) showing natural deformation characteristics.

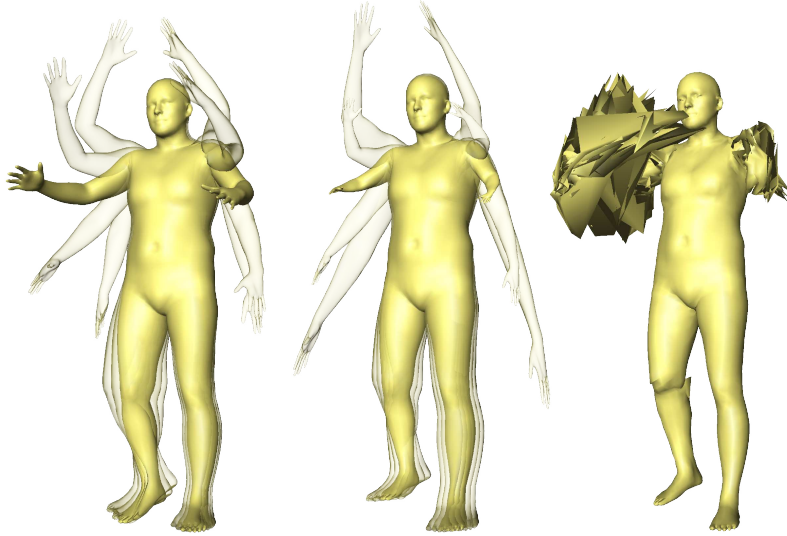


Figure 4.7.: Mean shapes (opaque) and first principal geodesic curve (transparent at ± 0.75 and ± 1.5 standard deviations) for the FAUST dataset obtained using the DCM (left), the PDM (center), and the SE(3)-based model* [Hefny et al., 2015] (right). (* curve not shown due to degeneracies)

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. Proposition 3.2). In Figure 4.8 we illustrate the validity of our model for two extreme examples in comparison to PDM and DCM. Note that, although not depicted in the figure, the GLM features similar problems as the DCM since it is also susceptible for ambiguities with large rotations components.




Figure 4.8.: Interpolating geodesic (mean highlighted) for the pipe surface (left) and FAUST (right) data within (f.l.t.r.) the DCM, the FCM, and PDM, each.

CHAPTER 5.

Applications in Disease Assessment

Main motivation from the application side for the thesis at hand is the field of disease assessment in a purely shape-based manner. The availability of huge databases holding thousands of medical imaging scans and disease related expert labeling (e.g. OAI database and Alzheimer’s Disease Neuroimaging Initiative (ADNI) database) is calling for comprehensive but computational light-weight methods. Hence, the evaluation on real-world data and real-world problems is besides the methodological development the second core component of this thesis.

5.1. Classification of Knee Osteoarthritis

OA is i.a. characterized by changes of the shape of bones composing the knee. Figure 5.1 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. In this experiment we investigate the ability of the proposed models to capture characteristic changes in shape that are incident to radiographic knee OA. To this end, we perform classification of healthy and severely diseased femora according to Kellgren–Lawrence (KL) grades of 0 or 1 and 4, respectively. The KL 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.

GLM – Knee Osteoarthritis Classification

With this experiment we want to investigate the proposed GLM’s sensitivity w.r.t. pathological shape changes and thus its ability to classify knee OA for the OAI dataset of distal femora. To achieve this, we utilize a simple Support Vector Machine (SVM) with linear kernel directly on the 115-dimensional space of shape weights. These weights are the vectors of coefficients w.r.t. the principal modes for each shape. The weights serve as input features to the SVM. The classifier is trained on a balanced set (healthy/diseased) of feature vectors for different shares of randomly chosen data varying from 10% to 90% whereas the testing is performed on the respective complement. Since we have some randomness in our experimental design we carry out the experiment 10000 times for each partition and consider the mean accuracy and the standard deviation. We compare our method to the PDM [Cootes et al., 1995] as well as to the DCM and the FCM, which both achieved highly accurate classification results. To this end, we employ the above outlined classifier setup using the respective model specific shape weights. The following experiments are performed utilizing (rounded) metric parameters $\tau = 0.1$, $\nu = 29.42$, $\kappa = 1.3$ that have been found conducting *hyper parameter optimization* (HPO) w.r.t. best performance in our classification experiment. HPO was carried out within the Scikit-Optimize¹ python framework performing a sequential optimization us-

¹scikit-optimize.github.io




Figure 5.1.: Healthy (left) and osteoarthritic (right) distal femur with delineated pathological changes in shape.

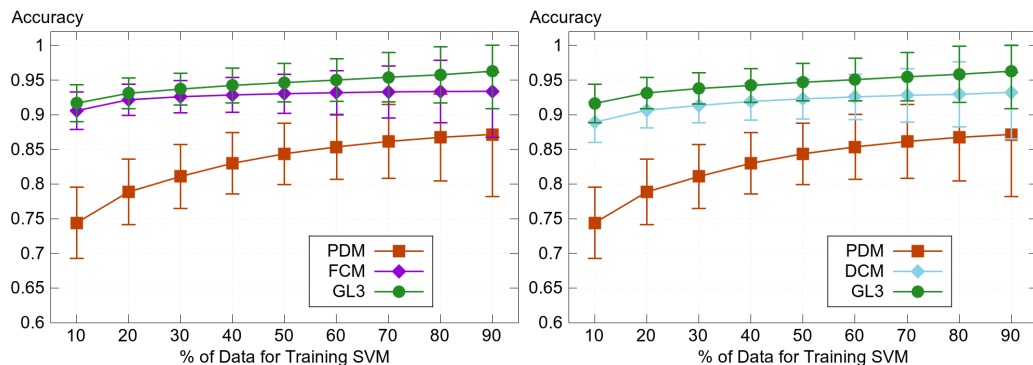


Figure 5.2.: OA classification experiment for the GLM, PDM, as well as the FCM (left) and the DCM (right). The accuracy of the GLM ranges from 91.6% (at 10% training) to 96.3% (at 90% training).

ing decision trees (forest_minimize) on the cubical domain $[0.05, 1000]^3$. Figure 5.2 shows the results in terms of average accuracy and standard deviation. The accuracy of the $GL^+(3)$ model ranges from 91.6% (at 10% training) to 96.3% (at 90% training). Note that solely the proposed GLM method achieves an accuracy of over 91% in case of sparse (10%) training data. However, one should be careful with the resulting accuracy, since HPO was applied directly to the classification experiment. This suggests there might be a possibility that the parameters slightly overfit the data.

DCM – Knee Osteoarthritis Classification

To evaluate the performance of the proposed method, we compare it to the PDM [Cootes et al., 1995] 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 SSM 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 SSM 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 Figure 5.3 we present a visualization of the




Figure 5.3.: Sammon plots visualizing the distribution of statistical shape descriptors obtained by the proposed and the PDM for the dataset of healthy and diseased distal femora from the OAI database.

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 Figure 5.5 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 Figure 5.1)—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 5.4. 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

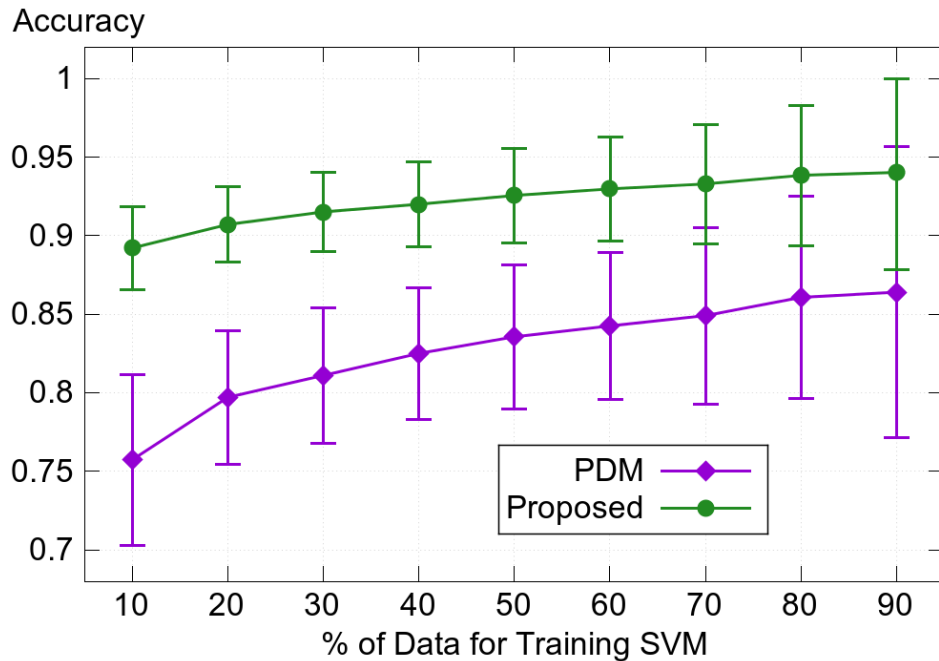


Figure 5.4.: Classification accuracy achieved by the proposed and the PDM as function of the percentage of training data.





Figure 5.5.: Magnitude of displacements under the first principal geodesic curve (± 1 standard deviations) for the DCM (left) and the PDM (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 DCM features more localized and well-delineated shape variations.

training set, we outperform PDM by approximately 15%, which emphasizes further the ability of the proposed method to identify systematic differences between the subpopulations.

FCM – Knee Osteoarthritis Classification

OA is a degenerative disease of the joints that is i.a. characterized by changes of the bone shape (see Figure 5.6). 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 PDM [Cootes et al., 1995] as well as to the DCM, which achieved highly accurate classification results. Figure 5.7 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.

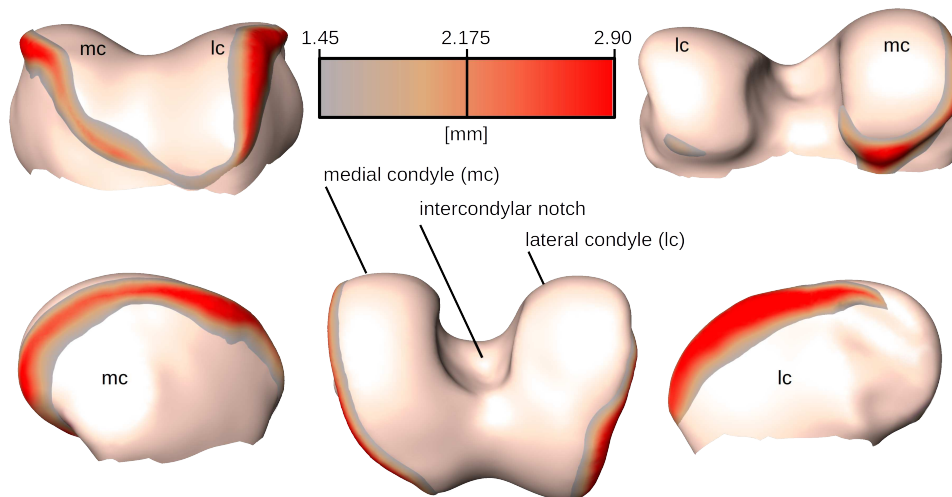


Figure 5.6.: Mean shape of healthy distal femora overlaid with (larger) mean shape of the diseased femora wherever the distance is larger than 1.45mm, colored accordingly.

Transparency

The proposed classifier exposes a high degree of interpretability and explainability due to the generative character of SSMs 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 the OA classification experiment in Figure 5.8 based on SVM instances with average




Figure 5.7.: OA classification experiment for FCM, PDM and DCM.

classification accuracy obtained for 40%/60% training/testing split. In addition to shapes sampled along the discriminating direction, we provide a 2D visualization using orthogonal projection onto the plane spanned by the two principal geodesic modes that retain the highest classification accuracy, viz. ϑ_1, ϑ_2 for OAI.

5.2. Stratification of Knee Osteoarthritis

5.2.1. Introduction

OA is a highly prevalent, degenerative joint disease with a considerable societal and economic impact, in addition to the physical and psychological sequelae it causes in affected individuals. The pathophysiology of OA involves several tissues and is primarily associated with a deterioration of articular cartilage as well as related changes in the underlying bone and at the joint margins. While OA can affect any joint, knee OA accounts for more than 80% of the global disease burden [Vos et al., 2012]. There exist various ways of characterizing OA in the literature ranging from subjective assessment to clinical and radiographic





Figure 5.8.: Visualization of the discriminating direction η and separating hyperplane η^\perp for OA classification showing a 2D 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 $0.15mm$.

ones, albeit with a limited degree of concordance between them. In practice, plain radiography remains a mainstay for the diagnosis of OA with the KL grading system posing the de-facto standard classification scheme. However, due to its sensitivity on acquisition method and rater reliability, which is reflected in the high number of disagreements between the readers (cf. [Bowes et al., 2020]), there is a dire need for accurate and reliable assessment of OA status.

Whereas plain radiography only provides 2D projections, advances in imaging technologies, especially in Magnetic Resonance Imaging (MRI), have enabled the understanding of 3D OA structural pathology. In particular, bone shape derived from MRI has been found to be associated with radiographic structural progression [Hanik et al., 2020], to predict radiographic onset of OAI [Neogi et al., 2013], and to discriminate knees w.r.t. osteophyte formation [von Ty-cowicz, 2020] and OA status (previous section). These findings suggest that bone morphology validly relates to a broader construct of OA pathology. Furthermore, shape-based assessment holds the promise of reduced sensitivity on image appearance and data acquisition set-ups; e.g. systematic changes due to regular technology upgrades with full hardware replacements every 5 to 10 years. In this light, Bowes et al. [Bowes et al., 2020] recently introduced a novel, geometrically derived measure to quantify knee OA from bone morphology termed *B-score*. Contrary to the semi-quantitative KL grade, the B-score is determined fully automatically from femur bone shape and, thus, does not suffer from the subjectivity of the practitioner. Being a continuous score it en-

ables fine-grained stratification of OA-related structural changes and increases discrimination of risk for clinically important outcomes such as TKR surgery.

Despite these recent advances, the formulation of B-score builds upon the popular PDM [Cootes et al., 1995] that treats shapes as elements of Euclidean spaces. However, such linearity assumptions are often inadequate for capturing the high, natural variability in biological shapes (see [Ambellan et al., 2019a] and the references therein). In particular, sizable empirical improvements can be observed when taking the inherent, geometric structure of shape spaces into account [Davis et al., 2010; Zhang et al., 2015; von Tycowicz, 2020].

Shape space

Before we summarize the employed shape representation, we would like to emphasize that the derived concepts and algorithms provided in this work are not tailored towards a particular choice and are indeed applicable to general Riemannian shape spaces.

For experimental evaluation, we opt for the FCM Section 3.7. This model is formulated within the commonly employed deformation-based morphometric framework in which shapes are expressed as deformations of a common reference surface. More precisely, a digital surface \mathcal{S} is encoded via the orientation preserving deformation ϕ of a triangular surface mesh $\bar{\mathcal{S}}$. For simplicial ϕ , the deformation gradient $\nabla\phi$ (also known as Jacobian matrix) is a 3×3 matrix of partial derivatives and constant on each triangle of $\bar{\mathcal{S}}$. In analogy to surface theory, discrete first and second fundamental forms can be derived from $\nabla\phi$ that furnish a complete description of the intrinsic and extrinsic geometry of \mathcal{S} . While the former takes the form of a piece-wise constant (one per triangle) field of 2×2 symmetric positive-definite matrices ($\text{Sym}^+(2)$), the latter is given by 3D rotations ($\text{SO}(3)$) associated with the edges. In particular, let m, n be the number of triangles and inner edges, then the resulting shape space is given as the product $G := \text{SO}(3)^n \times \text{Sym}^+(2)^m$. Remarkably, G can be equipped with a bi-invariant Lie group structure (by virtue of the log-Euclidean framework for $\text{Sym}^+(2)$ [Pennec et al., 2019; Arsigny et al., 2006], cf. Chapter 2, Section 3.7) that lends itself for efficient computations of Riemannian operations. Furthermore, the FCM provides a Euclidean motion invariant—hence alignment-free—shape representation that assures valid shape instances even in presence of strong nonlinear variability.

Geometric statistics

The nonlinear nature of shape spaces implies that there are no such familiar properties as vector space structure or global system of coordinates (that is, linear combinations of shapes do not generally lie in the space again and shape variations w.r.t. to different base shapes are not directly comparable). Consequently, core operations pervasive in machine learning and statistics often have to be generalized based on the geometry and specifics of the data at hand. Approaches that generalize statistical tools to non-Euclidean domains in order to leverage the intrinsic structure belong to the field of geometric statistics and we refer to [Pennec et al., 2019] for an overview.

The simplest—yet also perhaps most fundamentally important—statistic is the sample mean, which estimates the center of a data set. Because a Riemannian manifold \mathcal{M} has a distance $d_{\mathcal{M}}$ (length of the shortest path connecting two points), we can characterize the mean as the point closest to the data points $x_1, \dots, x_N \in \mathcal{M}$. This leads to the notion of (sample) Fréchet mean that is the minimizer of the sum-of-squared geodesic distances to the data:

$$\mu = \arg \min_{x \in \mathcal{M}} \sum_{i=1}^N d_{\mathcal{M}}^2(x, x_i).$$

While closed-form solutions exist in flat spaces, solving this least-squares problem in general requires iterative optimization routines. For geodesic manifolds, solutions always exist and are unique for well-localized data [Pennec et al., 2019] (cf. Chapter 2).

Another fundamental problem is the (statistical) normalization of shape trajectories, i.e. smooth curves in shape space encoding e.g. soft-body motion of anatomical structures. Normalization of such trajectories into a common reference frame is a challenging task in curved spaces (due to holonomy). The aim is to preserve as much as possible of the structural variability, while allowing a precise comparison in a common geometric space. To this end, *parallel transport* [do Carmo, 1992] provides a promising approach with a strong mathematical foundation. Parallel transport allows to propagate a tangent vector (i.e. an infinitesimal shape change) along a path by preserving its properties w.r.t. the space geometry, such as a notion of parallelism.

5.2.2. Geodesic B-score

In this section, we derive a generalization of the recently proposed B-score [Bowes et al., 2020] to Riemannian shape spaces and present a simple,

yet effective computational scheme for the determination thereof. In doing so, our guiding principle is to obtain expressions that take the rich geometric structure of shape space into account (e.g. refraining from linearization) and at the same time are consistent with its Euclidean counterpart (i.e. agree with the original definition for the special case of flat vector spaces). We term the resulting quantity *geodesic B-score* and will refer to the original definition (respectively, its application in linear spaces) as *Euclidean B-score* whenever this distinction is necessary.

Generalization

At the core of the construction in [Bowes et al., 2020] lies the projection to an *OA-vector* that is defined as the line passing through the mean shapes of populations with and without OA as determined by KL grades ≥ 2 and ≤ 1 , respectively. While we can readily rely on the Fréchet mean, differential geometry provides us with a consistent notion of straight lines known as geodesics [do Carmo, 1992]. In particular, we define the OA-geodesic γ as the length minimizing geodesic between the Fréchet means of the two populations (which will be unique under the assumptions for the means and the observed overlap of both distributions [Neogi et al., 2013; Bowes et al., 2020]). A visualization of the OA-geodesic is provided in Figure 5.10 (details on the underlying data are provided in Section 5.2.3). In order to determine the B-score for a shape $\sigma \in G$, we first perform an (intrinsic) projection onto the OA-geodesic:

$$\pi_\gamma(\sigma) := \arg \min_{x \in \gamma} d_G^2(x, \sigma). \quad (5.1)$$

The signed distance of $\pi_\gamma(\sigma)$ along the OA-geodesic w.r.t. the non-OA mean $\mu_{\mathcal{D}}$ (with positive values in direction of the OA mean) then yields the desired notion of geodesic B-score, i.e.

$$B_{\gamma, \lambda}(\sigma) = \lambda g_{\gamma(0)} \left(\dot{\gamma}^{(0)} / \|\dot{\gamma}^{(0)}\|, \text{Log}_{\gamma(0)} \circ \pi_\gamma(\sigma) \right), \quad (5.2)$$

where λ is a positive weighting factor and g , Log denote the Riemannian metric, logarithmic map. A visualization is given in Figure 5.9.

In order to increase interpretability, we take a statistical approach that weights the distances in terms of their distribution within the non-OA population. More precisely, we employ the Mahalanobis distance such that λ is determined as the inverse of the standard deviation $\text{std}(\{B_{\gamma, 1}(\sigma) | \sigma \in \mathcal{H}\})$ for the non-OA group \mathcal{H} . In fact, this statistical re-weighting relates the score to the natural morphological inter-subject variability and renders it unitless and scale-invariant.




Figure 5.9.: Geometric interpretation of geodesic B-score (with $\lambda = 1$) as signed distance along OA-geodesic. Colors indicate different levels of disease severity.

Sex-specific Reference

Females and males have systematically different bone shape [Bowes et al., 2020] introducing a bias during estimation of the B-score. In line with the Euclidean B-score, we correct for this bias using sex-specific OA-geodesics determined by translating γ s.t. it passes through the separately computed non-OA mean shapes for each sex. As a geodesic is uniquely determined by a point and a direction (viz. tangent vector at that point), we perform parallel transport of the defining vector along the geodesic connecting the mixed-sex and the sex-specific mean of the respective non-OA group. Given sex-specific OA geodesics $\gamma^\sigma, \gamma^\varphi$ we also estimate weighting factors $\lambda^\sigma, \lambda^\varphi$ for each sex and define the B-score as

$$B(\sigma) = \begin{cases} B_{\gamma^\varphi, \lambda^\varphi}(\sigma), & \sigma \text{ female} \\ B_{\gamma^\sigma, \lambda^\sigma}(\sigma), & \sigma \text{ male.} \end{cases}$$

Algorithmic Treatment

Determining solutions to the projection problem in Equation (5.1) does not admit closed-form expressions (except for the special case of constant curvature





Figure 5.10.: Signed vertex deviation from mean shape of mixed-sex non-OA group along the OA-geodesic. Color-coding: $-1mm$  $5mm$, with neutral window (i.e. yellowish bone color) from $-0.3mm$ to $0.4mm$.

manifolds [Chakraborty et al., 2016]), thus, requiring iterative optimization. However, this step is an essential ingredient for the computation of the geodesic B-score. In order to derive an efficient numerical scheme we assume (without loss of generality) $\gamma : t \mapsto \gamma(t) \in G$ to be an arc-length parameterized geodesic and express the projection problem as an unconstrained optimization over t with objective function $F(t) = d_G^2(\gamma(t), \sigma)$. A well-established scheme for this type of problem is Newton's method that employs second-order approximations to gain greatly in convergence speed, achieving quadratic convergence rate when close enough to the optimum. Analogously, a quadratic approximation for the objective F is given by

$$F(t + \delta) \approx F(t) + \frac{d}{dt}F(t) \cdot \delta + \frac{1}{2} \frac{d^2}{dt^2}F(t) \cdot \delta^2, \text{ with}$$

$$\frac{d}{dt}F(t) = -2g_{\gamma(t)} \left(\text{Log}_{\gamma(t)}(\sigma), \dot{\gamma}(t) \right), \text{ and}$$

$$\frac{d^2}{dt^2}F(t) = -2g_{\gamma(t)} \left(d_{\gamma(t)} \text{Log}_{\gamma(t)}(\sigma)(\dot{\gamma}(t)), \dot{\gamma}(t) \right).$$

Additionally, employing the first-order approximation for the differential of the logarithm $d_{\gamma(t)} \text{Log}_{\gamma(t)}(\sigma) \approx -Id$ [Pennec, 2017, Equation (5)] (Id denoting the identity) we can obtain an optimal step size δ^* for this quadratic model as

$$\delta^* = g_{\gamma(t)} \left(\text{Log}_{\gamma(t)}(\sigma), \dot{\gamma}(t) \right). \quad (5.3)$$

Indeed, verifiable by direct calculation, this step agrees with the explicit solution for the case of flat spaces (Figure 5.11, right). Eventually, we derive the

Newton-type fixed point iteration

$$\pi_{i+1} = \text{Exp}_{\pi_i}(\delta_i^* \dot{\gamma}_i), \quad (5.4)$$

where Exp denotes the Riemannian exponential map (Figure 5.11, left). In our setting, the choice $\pi_0 = \gamma(0)$ as initial guess is reasonable, since it is the healthy mean.

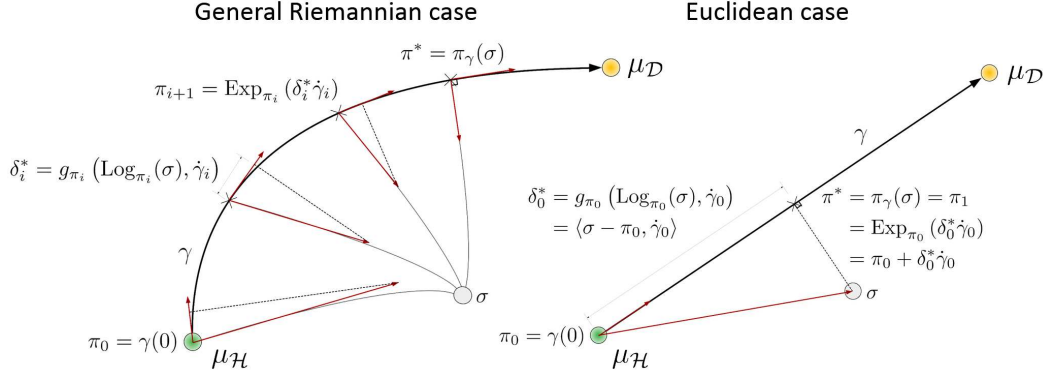


Figure 5.11.: Projection algorithm onto OA-geodesic, general Riemannian case (left) and Euclidean case (right).

5.2.3. Results and Discussion

Data description

Within this practical evaluation we rely on 9290 shapes of left and right distal femora from the OAI database spread over all KL grades of OA severity. Further details and demographic information can be found in Appendix A.1.

Since the shape space we employ is not scale invariant (as well as the PDM) this leaves the option to factor it out. However, since femoral osteoarthritis, among others, leads to flattening and widening of the condyle region that at least partially appears as deviation in scale w.r.t. a healthy configuration we forego scale alignment to preserve sensitivity for scale. Based on the geodesic B-score as derived in Section 5.2.2, we restrict our study population to the B-score percentile range from 0.75 to 99.25 (in terms of B-score: -3.12, 14.65) in order to exclude outliers. The resulting distribution of geodesic B-scores per KL grade is shown in Figure 5.12, visualizing the positive correlation of both grading schemes. Note that the depicted distribution is normalized to account for imbalance within the OAI database of KL grade frequencies, i.e. re-weighted as if KL groups were of equal cardinality.

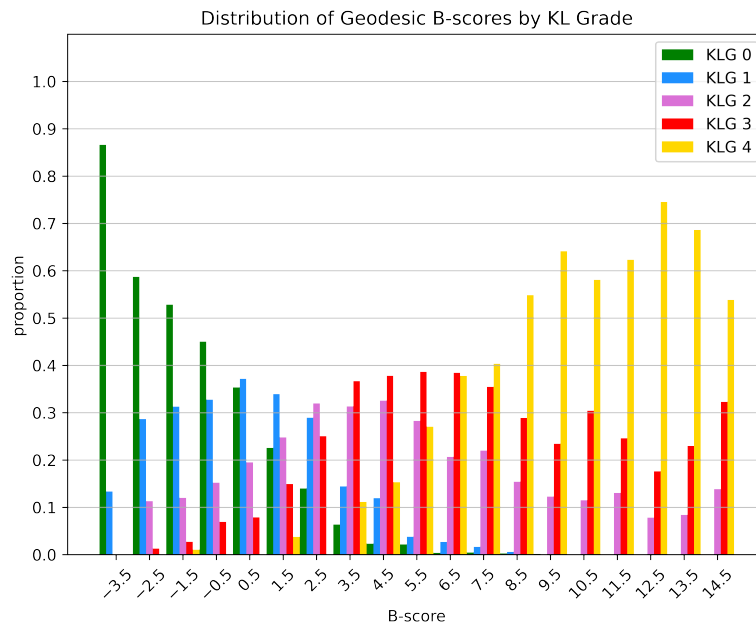


Figure 5.12.: Distribution of B-scores by KL grades normalized w.r.t. KL grade imbalance.

Efficiency of projection algorithm

We empirically evaluate the performance of the derived Newton-type iteration listed in Equation (5.4) using a python-based prototype implementation without parallelization, publicly available as part of the Morphomatics² library. To this end, we computed projections of 100 randomly selected femur shapes. We were able to observe quadratic convergence of the algorithm for all cases with 0.97s and three iterations per case in average.

Predictive validity

We assess the value of the geodesic B-score as a measure of OA status by examining its relationship with risk of TKR surgery—an important clinical outcome. Here and throughout, we refer to the risk of an outcome as the proportion w.r.t. a population. Additionally, we perform a comparison of the predictive performance between the geodesic and Euclidean B-score. To this

²morphomatics.github.io




Figure 5.13.: Comparison of TKR risk assessment for patient groups with TKR and No TKR clinical outcome respectively. Orange plus: mean, red line: median, box range: 25th to 75th percentile, whisker range: 5th to 95th percentile.

end, we follow the proposed setup from [Bowes et al., 2020] by modeling the predictor of TKR (within the follow-up period of 8 years) against B-score using logistic regression.

The determination of Euclidean B-scores is based on the space of vertex coordinates. To reduce confounding effects due to misalignment of the input shapes we employed generalized Procrustes analysis [Cootes et al., 1995] (adding a certain degree of nonlinearity over the approach in [Bowes et al., 2020]). No such considerations apply for the FCM-based geodesic B-score as it inherits the invariance to rigid motions. For both scores, computations were performed on the same input meshes using a modular software design sharing all routines that are not specific to the respective shape space. To compare the predictive performance of the derived models we grouped the study population into a TKR cohort that did receive TKR and a non-TKR cohort that did not. In Figure 5.13 we provide box plots for the resulting risk distributions that show clear differences in median risk between non-TKR and TKR. Furthermore, for the non-TKR cohort the geodesic B-score model validly yields median risks that are half of those for the Euclidean model. All these differences are sta-

tistically significant as determined using Mann–Whitney U tests. While both approaches yield the same median risk for the TKR cohort, the distribution of the geodesic B-score model is skewed towards higher risks. These findings substantiate an improved predictive power for the geodesic B-score.

5.3. Classification of Alzheimer’s Disease

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.

5.3.1. FCM – Alzheimer’s Classification

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.

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 authors³ 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

³github.com/ranahanocka/MeshCNN




Figure 5.14.: Alzheimer’s classification experiment for FCM, MeshCNN [Hanocka et al., 2019] and PDM [Cootes et al., 1995].

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 5.14 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].

Explainability

Analogously to the OA classification we provide a visualization of the SVM discriminating direction for the Alzheimer’s classification experiment in Figure 5.16 based on SVM instances with average classification accuracy obtained

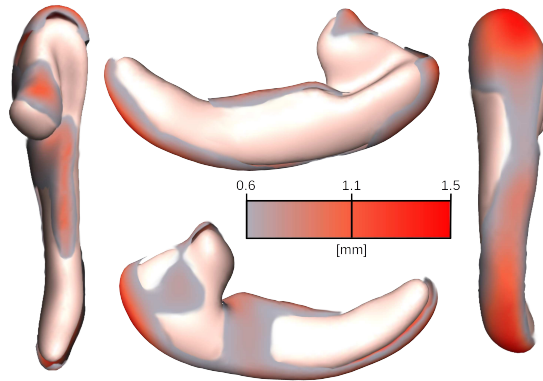


Figure 5.15.: Mean shape of diseased right hippocampi overlaid with mean shape of the healthy hippocampi wherever the distance is larger than 0.6mm, colored accordingly.

for 40%/60% training/testing split. In addition to shapes sampled along the discriminating direction, we further provide a 2D 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 for ADNI, respectively.

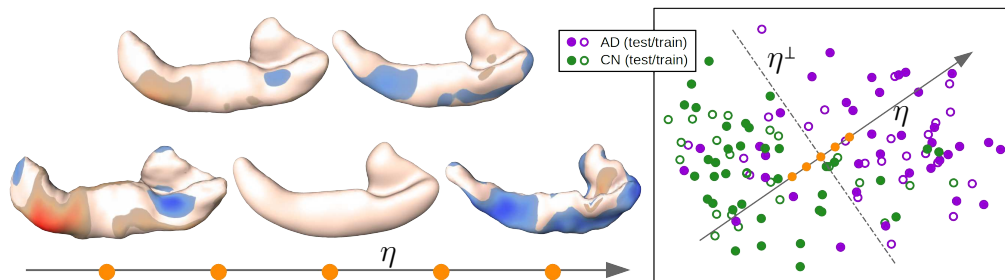



Figure 5.16.: Visualization of the discriminating direction η and separating hyperplane η^\perp for Alzheimer's classification showing a 2D projection (right) and corresponding shapes (left) equidistantly sampled within the interval containing the input data (note that projections onto η and the visualizing 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 $0.5mm$.

CHAPTER 6.

Conclusion and Future Work

6.1. Conclusion

6.1.1. Conclusion on Aspects of statistical shape modeling

In this work, we presented *three* novel approaches for nonlinear statistical analysis of shapes, which we consider discretized as a collection of corresponding primitives (e.g. triangles). Shape variability within a shape population is hereby encoded as deformations of a common reference shape. These deformations are represented through the deformation gradient, i.e. as differential coordinates. On these coordinates we defined different manifold structures allowing for different approaches to shape analysis. In contrast to the standard PDM every shape model shown in this thesis comes at least with invariance under translation, due to the nature of the underlying differential shape representation. An algorithm to solve the inverse problem, devising shapes from given coordinates, employing Poisson-based reconstruction is given.

The **GLM** was based on the canonical $GL^+(3)$ structure of the deformation gradient. It utilized the bi-invariant Lie group mean and a tangent principal component analysis employing a $GL^+(3)$ -left-invariant, $O(3)$ -right-invariant metric in $GL^+(3)$. It can thus be considered as as-invariant-as-possible w.r.t. the canonical $GL^+(3)$ structure. This structure on the one hand is compatible

with piecewise (triangle-wise) affine transformations of the input shapes, i.e. taking the differential representation and applying the transformation commutes. However, the barycenter and the modes of variation do depend on the input shapes' orientation in the ambient space calling for an initial generalized Procrustes alignment. Furthermore, we have shown that the GLM possesses a high descriptiveness w.r.t. natural biological differences in shape. In particular, we conducted experiments on OA classification. In order to determine the parameters of the metric we applied a hyper parameter optimization targeting classification accuracy.

We presented the **DCM** as second 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, consisting of deformation gradients split into rotational and stretch components, with a Riemannian structure that provides our shape space with strong theoretical properties. In particular, physically invalid (stretch-)tensors are at an infinite distance from any element in our space. Despite its non-Euclidean structure our DCM 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 DCM and compared it to alternative approaches. Specifically, in comparison to the state-of-the-art, physically-based Shell PCA model 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 knee OA. Additionally, we showed that the DCM is better able to capture the nonlinear variations present in articulated body pose and disease-specific data by comparing it to the standard PDM and the recent SE(3)-based model [Hefny et al., 2015]. Notably, the DCM 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 feature angles less than π .

Finally, we presented the **FCM** 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 local-global algorithm to solve the inverse problem that does not require any numerical safeguards.

We showed that FCM yields highly differentiating shape descriptors that promote great performance for shape-based disease state classifications for knee OA and Alzheimer’s: (1) In comparisons our FCM descriptors are superior to the DCM ones, as well as the popular linear PDM based descriptors; (2) Remarkably, the FCM-based classifier significantly outperformed 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 thesis 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.

6.1.2. Conclusion OA Stratification

We introduced a consistent generalization of the recently presented (Euclidean) B-score to Riemannian shape spaces. We showed that the obtained formulation features superior predictive power in an experiment on TKR risk assessment, thus, suggesting improved discrimination of morphological status across the range of OA severity. These advances foster the potential of B-score to replace imprecise and insensitive measures for the assessment of OA status based on plain radiography. Moreover, we further presented an original algorithm for the projection of points in a Riemannian manifold onto a geodesic. In particular, the obtained iteration exposes fast, quadratic convergence and is simple to implement. We chose a shape space based on fundamental coordinates, because—due to its deep foundation in differential geometry and link to thin elastic shells—it faithfully captures nonlinear shape variability, while offering fast processing of large-scale shape collections. On the theoretical side, the price to pay is that there is no guarantee that the projection is diffeomorphic.

However, we would like to remark that the estimated OA-geodesic contains only diffeomorphic deformations within the confidence interval, guaranteeing valid instances even if the input shapes are not. Furthermore, contrary to shape spaces based on diffeomorphic metric mapping, ours is invariant under Euclidean motion and, thus, not susceptible to any bias due to misalignment.

Finally, all presented shape modeling approaches, as well as all components of the Geodesic B-score are implemented in the publicly available *Morphomatics*¹ - *Geometric morphometrics in non-Euclidean shape spaces* Python library [Ambellan et al., 2021a]. This, together with the detailed Information on the employed data (A.1) enables reproducibility of results and might stimulate further research in the field.

6.2. Future Work

We consider it valuable and interesting to also investigate the purely Riemannian perspective associated with the metric in the GLM setting and compare it to our present work. However, this comes at a price as, although geodesics can be evaluated in closed form for a given direction and the existence of a shortest geodesic connecting two arbitrary points is theoretically guaranteed, no closed form solution to determine the direction of one (and not necessarily the shortest) connecting geodesic is known [Martin and Neff, 2016].

As another future direction we plan to extend our proposed models to also account for appearances. This in 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 DCM metric. In particular, while the rotational component is affected by local and global rigid motions it also contains curvature information.

One possible and interesting way to proceed with DCM, FCM in the future is to replace the log-Euclidean structure with the affine-invariant one, which can be considered the natural on the symmetric positive-definite matrices. Additionally, investigation of the log-Cholesky structure [Lin, 2019] appears to be promising. 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

¹morphomatics.github.io

three-dimensional problems into the well-known two-dimensional image-based deep learning setup. Finally, the variance analysis by TPCA could be replaced, where feasible, by *exact* PGA to avoid any confounding information via linearization in the models.

Considering of the numerous size of available datasets (e.g. OAI database), another promising line of work is the extension of SSM-based disease classification methods to a broader data base, namely tens of thousands of samples. This is feasible due to the numerical efficiency of the presented methods and might allow for an even more accurate disease assessment. A first step in this direction has already been taken by Tack et al. [2021] working with all subjects from the OAI database.

In this work, we carefully generalized the B-score mimicking the geometric construction of its Euclidean counterpart. However, there are various statistical approaches that allow to estimate submanifolds based on separation or regression considerations, e.g. geodesic discriminant analysis [Louis et al., 2018] or higher-order regression [Hanik et al., 2020], respectively. An interesting direction for future work is to investigate to which extend such geometric statistics can serve as a foundation for advanced notions of an intrinsic B-score. From a medical perspective, it will be most interesting to explore the relationship of the geodesic B-score with further clinically important outcomes such as pain and loss of function. In particular, we will investigate to which degree the geodesic B-score can improve the related risk assessment. Since the presented statistical approach can directly be extended to multiple connected components, another line of work will aim on extension to multi-structure B-scores, e.g. for Femur and Tibia. Moreover, for the future we envision a longitudinal characterization beyond the static B-score that takes subject-specific shape developments into account.

List of Abbreviations

Notation	Description	First
ADNI	Alzheimer’s Disease Neuroimaging Initiative.	65
CNN	Convolutional Neural Network.	9
DCM	Differential Coordinates Model.	6
FAUST	Fine Alignment Using Scan Texture.	59
FCM	Fundamental Coordinates Model.	7
GLM	$GL^+(3)$ -Model.	6
LPGA	Linearized Principal Geodesic Analysis.	4
MRI	Magnetic Resonance Imaging.	72
OA	Osteoarthritis.	1
OAI	Osteoarthritis Initiative.	59
PCA	Principal Component Analysis.	2
PDM	Point Distribution Model.	2
PGA	Principal Geodesic Analysis.	4
SSM	Statistical Shape Model.	1
SVM	Support Vector Machine.	66

List of Abbreviations

Notation	Description	First
TKR	Total Knee Replacement.	7
TPCA	Tangent Principal Component Analysis.	4
weDESS	dual-echo steady state with selective water excitation.	106

Bibliography

- Alexandrino, M. M. and Bettiol, R. G. (2015). *Lie groups and geometric aspects of isometric actions*, volume 8. Springer.
- Altmann, S. L. (2005). *Rotations, quaternions, and double groups*. Courier Corporation.
- Ambellan, F., Hanik, M., and von Tycowicz, C. (2021a). Morphomatics: Geometric morphometrics in non-euclidean shape spaces. One Python Library to Rule Them All.
- Ambellan, F., Lamecker, H., von Tycowicz, C., and Zachow, S. (2019a). Statistical shape models - understanding and mastering variation in anatomy. In Rea, P. M., editor, *Biomedical Visualisation*, volume 3, pages 67 – 84. Springer.
- Ambellan, F., Tack, A., Ehlke, M., and Zachow, S. (2019b). 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., and von Tycowicz, C. (2019c). An as-invariant-as-possible $GL+(3)$ -based statistical shape model. In *Proc. 7th MICCAI workshop on Mathematical Foundations of Computational Anatomy (MFCA)*, volume 11846, pages 219 – 228.
- Ambellan, F., Zachow, S., and von Tycowicz, C. (2019d). A surface-theoretic approach for statistical shape modeling. In *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI), Part IV*, volume 11767, pages 21 – 29.

- Ambellan, F., Zachow, S., and von Tycowicz, C. (2021b). Geodesic b-score for improved assessment of knee osteoarthritis. In *International Conference on Information Processing in Medical Imaging*, pages 177–188. Springer.
- Ambellan, F., Zachow, S., and von Tycowicz, C. (2021c). Rigid motion invariant statistical shape modeling based on discrete fundamental forms. *Medical Image Analysis*, page 102178.
- Arsigny, V., Commowick, O., Ayache, N., and Pennec, X. (2009). A fast and log-euclidean polyaffine framework for locally linear registration. *Journal of Mathematical Imaging and Vision*, 33(2):222–238.
- Arsigny, V., Fillard, P., Pennec, X., and 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., and 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*, pages 829–837. Springer.
- Bauer, M., Bruveris, M., and Michor, P. W. (2014). Overview of the geometries of shape spaces and diffeomorphism groups. *J. Math. Imaging Vis.*, 50(1-2):60–97.
- Baum, H. (2009). *Eichfeldtheorie*. Springer.
- Bhatia, H., Norgard, G., Pascucci, V., and Bremer, P.-T. (2013). The Helmholtz-Hodge decomposition—a survey. *IEEE Trans Vis Comput Graph*, 19(8):1386–1404.
- Bogo, F., Romero, J., Loper, M., and Black, M. J. (2014). FAUST: Dataset and evaluation for 3D mesh registration. In *CVPR*. IEEE.
- Bonner-Jackson, A., Mahmoud, S., Miller, J., and Banks, S. J. (2015). Verbal and non-verbal memory and hippocampal volumes in a memory clinic population. *Alzheimer’s research & therapy*, 7(1):61.
- Botsch, M., Sumner, R., Pauly, M., and Gross, M. (2006). Deformation transfer for detail-preserving surface editing. In *VMV*, pages 357–364.
- Bowes, M. A., Kacena, K., Alabas, O. A., Brett, A. D., Dube, B., Bodick, N., and Conaghan, P. G. (2020). Machine-learning, MRI bone shape and important clinical outcomes in osteoarthritis: data from the osteoarthritis initiative. *Annals of the Rheumatic Diseases*. published online first: 13 November 2020.

- Brandt, C., von Tycowicz, C., and Hildebrandt, K. (2016). Geometric flows of curves in shape space for processing motion of deformable objects. *Comput Graph Forum*, 35(2).
- Bredbenner, T. L., Eliason, T. D., Potter, R. S., Mason, R. L., Havill, L. M., and 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(9):1780–1786.
- Brunton, A., Salazar, A., Bolkart, T., and Wuhrer, S. (2014). Review of statistical shape spaces for 3D data with comparative analysis for human faces. *Comput Vis Image Underst*, 128:1–17.
- Chakraborty, R., Seo, D., and Vemuri, B. C. (2016). An efficient exact-PGA algorithm for constant curvature manifolds. In *Proc. Computer Vision and Pattern Recognition*, pages 3976–3984.
- Charpiat, G., Faugeras, O., Keriven, R., and Maurel, P. (2006). Distance-based shape statistics. In *ICASSP*, pages V925–V928. IEEE.
- Cheeger, J. and Ebin, D. G. (2008). *Comparison theorems in Riemannian geometry*, volume 365. American Mathematical Soc.
- Ciarlet, P. G. (2005). An introduction to differential geometry with applications to elasticity. *Journal of Elasticity*, 78(1):1–215.
- Conaghan, P. G., Kloppenburg, M., Schett, G., Bijlsma, J. W., et al. (2014). Osteoarthritis research priorities: a report from a eular ad hoc expert committee. *Annals of the rheumatic diseases*, 73(8):1442–1445.
- Cootes, T., Taylor, C. J., Cooper, D. H., and 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., and Ovsjanikov, M. (2017). Functional characterization of intrinsic and extrinsic geometry. *ACM Trans Graph*, 36(2):1–17.
- Davies, R., Twining, C., and Taylor, C. (2008). *Statistical Models of Shape: Optimisation and Evaluation*. Springer.
- Davis, B., Fletcher, P. T., Bullitt, E., and Joshi, S. (2010). Population shape regression from random design data. *Int J Comput Vis*, 90(2):255–266.

- de Toledo-Morrell, L., Dickerson, B., Sullivan, M., Spanovic, C., Wilson, R., and 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.
- do Carmo, M. P. (1976). *Differential geometry of curves and surfaces*. Prentice-Hall.
- do Carmo, M. P. (1992). *Riemannian Geometry; 2nd ed.* Mathematics : Theory and Applications. Birkhäuser, Boston, MA.
- Dryden, I. L. and Mardia, K. V. (2016). *Statistical shape analysis: with applications in R*, volume 995. John Wiley & Sons.
- Durrleman, S., Prastawa, M., Charon, N., Korenberg, J. R., Joshi, S., Gerig, G., and Trounev, A. (2014). Morphometry of anatomical shape complexes with dense deformations and sparse parameters. *NeuroImage*, 101(0):35 – 49.
- Ezuz, D. and Ben-Chen, M. (2017). Deblurring and denoising of maps between shapes. In *Computer Graphics Forum*, volume 36, pages 165–174. Wiley Online Library.
- Fletcher, P., Lu, C., Pizer, S., and Joshi, S. (2004). Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Trans Med Imaging*, 23(8):995–1005.
- Fletcher, P. T., Lu, C., and Joshi, S. (2003). Statistics of shape via principal geodesic analysis on lie groups. In *CVPR*, volume 1, pages I–95. IEEE.
- Fréchet, M. (1948). Les éléments aléatoires de nature quelconque dans un espace distancié. *Annales de l’institut Henri Poincaré*, 10(4):215–310.
- Freifeld, O. and Black, M. J. (2012). Lie bodies: A manifold representation of 3d human shape. In *ECCV*, pages 1–14. Springer.
- Fuchs, M., Jüttler, B., Scherzer, O., and Yang, H. (2009). Shape metrics based on elastic deformations. *J Math Imaging Vis*, 35(1):86–102.
- Gallier, J. (2018). Logarithms and square roots of real matrices existence, uniqueness and applications in medical imaging. *arXiv preprint arXiv:0805.0245*.
- Gallier, J. and Quaintance, J. (2020). *Differential geometry and Lie groups: a computational perspective*, volume 12. Springer Nature.

- Gao, L., Lai, Y.-K., Liang, D., Chen, S.-Y., and Xia, S. (2016). Efficient and flexible deformation representation for data-driven surface modeling. *ACM Trans Graph*, 35(5):158.
- Ge, J. (2014). Ddvv-type inequality for skew-symmetric matrices and simons-type inequality for riemannian submersions. *Advances in Mathematics*, 251:62–86.
- Gilles, B., Bousquet, G., Faure, F., and Pai, D. K. (2011). Frame-based elastic models. *ACM transactions on graphics (TOG)*, 30(2):1–12.
- Goh, A., Lenglet, C., Thompson, P. M., and 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.
- Goodall, C. (1991). Procrustes methods in the statistical analysis of shape. *J. Roy. Statist. Soc. Series B (Methodological)*, 53(2):285–339.
- Gower, J. C. (1975). Generalized procrustes analysis. *Psychometrika*, 40(1):33–51.
- Grinspun, E., Hirani, A. N., Desbrun, M., and Schröder, P. (2003). Discrete Shells. In *ACM SIGGRAPH/Eurographics SCA*, pages 62–67.
- Hanik, M., Hege, H.-C., Hennemuth, A., and von Tycowicz, C. (2020). Nonlinear regression on manifolds for shape analysis using intrinsic Bézier splines. In *Proc. Medical Image Computing and Computer-Assisted Intervention*, pages 617–626.
- Hanocka, R., Hertz, A., Fish, N., Giryes, R., Fleishman, S., and 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., and 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., and Smith, W. (2018). Principal geodesic analysis in the space of discrete shells. *Comput Graph Forum*, 37(5):173–184.
- Hefny, M. S., Okada, T., Hori, M., Sato, Y., and Ellis, R. E. (2015). A liver atlas using the special Euclidean group. In *MICCAI*, pages 238–245. Springer.
- Heimann, T. and Meinzer, H.-P. (2009). Statistical shape models for 3d medical image segmentation: A review. *Med Image Anal*, 13(4):543–563.

- Helgason, S. (2001). *Differential geometry and symmetric spaces*, volume 341. American Mathematical Soc.
- Higham, N. J. (2005). The scaling and squaring method for the matrix exponential revisited. *SIAM Journal on Matrix Analysis and Applications*, 26(4):1179–1193.
- Huckemann, S. and Hotz, T. (2009). Principal component geodesics for planar shape spaces. *Journal of Multivariate Analysis*, 100(4):699–714.
- Huckemann, S., Hotz, T., and Munk, A. (2010). Intrinsic shape analysis: Geodesic PCA for Riemannian manifolds modulo isometric Lie group actions. *Stat Sin*, pages 1–58.
- Joshi, S., Davis, B., Jomier, M., and 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., and Zachow, S. (2013). Omnidirectional displacements for deformable surfaces. *Med Image Anal*, 17(4):429 – 441.
- Karcher, H. (1977). Riemannian center of mass and mollifier smoothing. *Commun Pure Appl Math*, 30(5):509–541.
- Kelemen, A., Székely, G., and Gerig, G. (1999). Elastic model-based segmentation of 3-d neuroradiological data sets. *IEEE Trans Med Imaging*, 18(10):828–839.
- Kellgren, J. and Lawrence, J. (1957). Radiological assessment of osteoarthritis. *Ann Rheum Dis*, 16(4):494.
- Kendall, D., Barden, D., Carne, T. K., and Le, H. (1999). *Shape and shape theory*. John Wiley & Sons.
- Kendall, D. G. (1989). A survey of the statistical theory of shape. *Statistical Science*, pages 87–99.
- Kilian, M., Mitra, N. J., and Pottmann, H. (2007). Geometric modeling in shape space. *ACM Transactions on Graphics (SIGGRAPH)*, 26(3):#64, 1–8.
- Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kircher, S. and Garland, M. (2008). Free-form motion processing. *ACM Trans Graph*, 27(2):12.

- Kirschner, M., Becker, M., and Wesarg, S. (2011). 3d active shape model segmentation with nonlinear shape priors. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 492–499. Springer Berlin Heidelberg.
- Kobayashi, S. and Nomizu, K. (1963). *Foundations of differential geometry*, volume 1. New York, London.
- 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., and Ropinski, T. (2018). A survey of flattening-based medical visualization techniques. In *Computer Graphics Forum*, volume 37, pages 597–624. Wiley Online Library.
- Lamecker, H. and Zachow, S. (2016). Statistical shape modeling of musculoskeletal structures and its applications. In *Computational Radiology for Orthopaedic Interventions*, volume 23, pages 1 – 23. Springer.
- Lawrence, R. C., Felson, D. T., Helmick, C. G., Arnold, L. M., Choi, H., Deyo, R. A., Gabriel, S., Hirsch, R., Hochberg, M. C., Hunder, G. G., et al. (2008). Estimates of the prevalence of arthritis and other rheumatic conditions in the united states: Part ii. *Arthritis & Rheumatology*, 58(1):26–35.
- Lin, Z. (2019). Riemannian geometry of symmetric positive definite matrices via cholesky decomposition. *SIAM Journal on Matrix Analysis and Applications*, 40(4):1353–1370.
- Louis, M., Charlier, B., and Durrleman, S. (2018). Geodesic discriminant analysis for manifold-valued data. In *Proc. Workshops at Computer Vision and Pattern Recognition*, pages 332–340.
- Martin, R. J. and Neff, P. (2016). Minimal geodesics on $GL(n)$ for left-invariant, right- $O(n)$ -invariant Riemannian metrics. *Journal of Geometric Mechanics*, 8(3):323–357.
- McLeod, K., Sermesant, M., Beerbaum, P., and Pennec, X. (2015). Spatio-temporal tensor decomposition of a polyaffine motion model for a better analysis of pathological left ventricular dynamics. *IEEE transactions on medical imaging*, 34(7):1562–1575.

- Mika, S., Schölkopf, B., Smola, A., Müller, K.-R., Scholz, M., and Rätsch, G. (1999). Kernel PCA and de-noising in feature spaces. In *NIPS*, pages 536–542, Cambridge, MA, USA. MIT Press.
- Miller, M. I., Trounev, A., and Younes, L. (2015). Hamiltonian systems and optimal control in computational anatomy: 100 years since d’arcy thompson. *Annual review of biomedical engineering*, 17:447–509.
- Nava-Yazdani, E., Hege, H.-C., Sullivan, T. J., and von Tycowicz, C. (2020). Geodesic analysis in kendall’s shape space with epidemiological applications. *Journal of Mathematical Imaging and Vision*, 62(4):549–559.
- Neogi, T., Bowes, M. A., Niu, J., Souza, K. M., Vincent, G. R., Goggins, J., Zhang, Y., and Felson, D. T. (2013). Magnetic resonance imaging-based three-dimensional bone shape of the knee predicts onset of knee osteoarthritis. *Arthritis Rheum*, 65(8):2048–2058.
- Nielsen, F. and Bhatia, R. (2013). *Matrix information geometry*. Springer.
- O’Neill, B. (1983). *Semi-Riemannian geometry with applications to relativity*, volume 103. Academic press.
- Pennecc, X. (2006). Intrinsic statistics on riemannian manifolds: Basic tools for geometric measurements. *J Math Imaging Vis*, 25(1):127.
- Pennecc, X. (2017). Hessian of the Riemannian squared distance. *Technical report, Université Côte d’Azur and Inria Sophia-Antipolis Méditerranée*.
- Pennecc, X. (2018). Barycentric subspace analysis on manifolds. *The Annals of Statistics*, 46(6A):2711–2746.
- Pennecc, X. (2020). Advances in geometric statistics for manifold dimension reduction. In *Handbook of Variational Methods for Nonlinear Geometric Data*, pages 339–359. Springer.
- Pennecc, X. and Arsigny, V. (2013). Exponential barycenters of the canonical Cartan connection and invariant means on Lie groups. In *Matrix Information Geometry*, pages 123–166. Springer.
- Pennecc, X. and Lorenzi, M. (2020). Beyond riemannian geometry: The affine connection setting for transformation groups. In *Riemannian Geometric Statistics in Medical Image Analysis*, pages 169–229. Elsevier.
- Pennecc, X., Sommer, S., and Fletcher, T. (2019). *Riemannian Geometric Statistics in Medical Image Analysis*. Academic Press.

- Peterfy, C., Gold, G., Eckstein, F., Cicuttini, F., Dardzinski, B., and Stevens, R. (2006). MRI protocols for whole-organ assessment of the knee in osteoarthritis. *Osteoarthritis and cartilage*, 14:95–111.
- Postnikov, M. (2013). *Geometry VI: Riemannian Geometry*, volume 91. Springer.
- Rathi, Y., Dambreville, S., and Tannenbaum, A. (2006). Comparative analysis of kernel methods for statistical shape learning. In *CVAMIA*, pages 96–107. Springer.
- Rentmeesters, Q. and Absil, P.-A. (2011). Algorithm comparison for karcher mean computation of rotation matrices and diffusion tensors. In *2011 19th European Signal Processing Conference*, pages 2229–2233. IEEE.
- Robinette, K. M., Daanen, H., and Paquet, E. (1999). The CAESAR project: A 3-D surface anthropometry survey. In *3DIM*, pages 380–386.
- Rumpf, M. and Wirth, B. (2011). An elasticity-based covariance analysis of shapes. *Int J Comput Vis*, 92(3):281–295.
- Rumpf, M. and Wirth, B. (2015). Variational methods in shape analysis. *Handbook of Mathematical Methods in Imaging*, 2:1819–1858.
- Rustamov, R. M., Ovsjanikov, M., Azencot, O., Ben-Chen, M., Chazal, F., and Guibas, L. (2013). Map-based exploration of intrinsic shape differences and variability. *ACM Trans Graph*, 32(4):1–12.
- Saito, A., Nawano, S., and 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. *Med Image Anal*, 28:46–65.
- Sammon, J. W. (1969). A nonlinear mapping for data structure analysis. *IEEE Trans Comput*, C-18(5):401–409.
- Sarkalkan, N., Weinans, H., and Zadpoor, A. A. (2014). Statistical shape and appearance models of bones. *Bone*, 60:129–140.
- Seim, H., Kainmüller, D., Lamecker, H., Bindernagel, M., Malinowski, J., and Zachow, S. (2010). Model-based auto-segmentation of knee bones and cartilage in MRI data. In v. Ginneken, B., editor, *MICCAI Workshop Med Image Anal for the Clinic*, pages 215 – 223.

- Seo, D., Ho, J., and Vemuri, B. C. (2016). Covariant image representation with applications to classification problems in medical imaging. *Int J Comput Vis*, 116(2):190–209.
- Shen, K.-k., Fripp, J., Mériaudeau, F., Chételat, G., Salvado, O., Bourgeat, P., Initiative, A. D. N., et al. (2012). Detecting global and local hippocampal shape changes in alzheimer’s disease using statistical shape models. *Neuroimage*, 59(3):2155–2166.
- Sommer, S., Lauze, F., Hauberg, S., and Nielsen, M. (2010). Manifold valued statistics, exact principal geodesic analysis and the effect of linear approximations. In *ECCV*, pages 43–56. Springer.
- Sommer, S., Lauze, F., and Nielsen, M. (2014). Optimization over geodesics for exact principal geodesic analysis. *Advances in Computational Mathematics*, 40(2):283–313.
- Tack, A., Ambellan, F., and Zachow, S. (2021). Towards novel osteoarthritis biomarkers: Multi-criteria evaluation of 46,996 segmented knee mri data from the osteoarthritis initiative. *PloS one*, 16(10):e0258855.
- Thanwerdas, Y. and Pennec, X. (2019). Is affine-invariance well defined on spd matrices? a principled continuum of metrics. In *International Conference on Geometric Science of Information*, pages 502–510. Springer.
- Thomson, J., O’Neill, T., Felson, D., and Cootes, T. (2016). *Detecting Osteophytes in Radiographs of the Knee to Diagnose Osteoarthritis*, pages 45–52. Springer International Publishing.
- Thomson, J., O’Neill, T., Felson, D., and Cootes, T. (2015). Automated shape and texture analysis for detection of osteoarthritis from radiographs of the knee. In *MICCAI*, pages 127–134. Springer.
- Twining, C. J. and Taylor, C. J. (2011). Specificity: A graph-based estimator of divergence. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(12):2492–2505.
- von Tycowicz, C. (2020). Towards shape-based knee osteoarthritis classification using graph convolutional networks. In *International Symposium on Biomedical Imaging*, pages 750–753.
- von Tycowicz, C., Ambellan, F., Mukhopadhyay, A., and 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., and Hildebrandt, K. (2015). Real-time nonlinear shape interpolation. *ACM Trans Graph*, 34(3):34:1–34:10.
- Vos, T., Flaxman, A., Naghavi, M., Lozano, R., Michaud, C., Ezzati, M., Shibuya, K., Salomon, J., Abdalla, S., Aboyans, V., et al. (2012). Years lived with disability (YLDs) for 1160 sequelae of 289 diseases and injuries 1990–2010: a systematic analysis for the global burden of disease study 2010. *The lancet*, 380(9859):2163–2196.
- Walter, R. (1974). On the metric projection onto convex sets in riemannian spaces. *Archiv der Mathematik*, 25(1):91–98.
- Wang, Y., Ben-Chen, M., Polterovich, I., and Solomon, J. (2018). Steklov spectral geometry for extrinsic shape analysis. *ACM Trans Graph*, 38(1):1–21.
- Woods, R. P. (2003). Characterizing volume and surface deformations in an atlas framework: theory, applications, and implementation. *NeuroImage*, 18(3):769–788.
- Younes, L. (2010). *Shapes and diffeomorphisms*, volume 171. Springer.
- Yu, Y., Zhou, K., Xu, D., Shi, X., Bao, H., Guo, B., and Shum, H.-Y. (2004). Mesh editing with Poisson-based gradient field manipulation. *ACM Trans Graph*, 23(3):644–651.
- Zacur, E., Bossa, M., and Olmos, S. (2014). Multivariate tensor-based morphometry with a right-invariant riemannian distance on $gl+(n)$. *Journal of mathematical imaging and vision*, 50(1-2):18–31.
- Zhang, C., Heeren, B., Rumpf, M., and Smith, W. A. (2015). Shell PCA: Statistical shape modelling in shell space. In *ICCV*, pages 1671–1679. IEEE.

APPENDIX A.

Information on Data and Further Experiments

The following chapter of the appendix will give detailed information on the employed datasets as well as on additional experiments that were carried out in order to substantiate the presented SSMs.

A.1. Employed Datasets

In order to experimentally validate the performance of the presented SSMs within this thesis we did employ six datasets in total. Furthermore, we globally aligned the objects within each dataset w.r.t. position and orientation using the standard generalized Procrustes method [Goodall, 1991]. This is, in view of the FCM, not necessary, but to allow for a fair comparison between different models with different invariances simply inevitable.

(i) **OAI - Right distal femora** (see e.g. Figure 4.5) from the 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 nda.nih.gov/oai [Peterfy et al., 2006]. From the baseline dataset, we chose 58 severely diseased subjects and 58 healthy subjects according to their KL score [Kellgren and Lawrence, 1957] (diseased: grade 4 healthy: grade 0,1). For the 116 subjects we extracted

surfaces of the distal femora from the respective 3D dual-echo steady state with selective water excitation (weDESS) MRI (0.37×0.37 mm matrix, 0.7 mm slice thickness) using [Seim et al., 2010]. The used segmentation masks have been published along with [Ambellan et al., 2019b] and are freely available as part of the OAI-ZIB dataset¹. In a supervised post-process, the quality of segmentations as well as the correspondence of the resulting meshes (8988 vertices, 17829 triangles) were ensured. A detailed list of the exact subjects that are included in the experiment OA classification experiment, as well as their disease state are given in the following Table A.1.

Table A.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	

(ii) **ADNI - Right hippocampi** (see e.g. Figure 5.15) from the ADNI² 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, featuring three different Alzheimer’s states: Cognitive normal (CN), mild cognitive impairment and Alzheimer’s

¹doi.org/10.12752/4.ATEZ.1.0 or at pubdata.zib.de

²adni.loni.usc.edu

disease (AD). 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]³ 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 scan ids in Table A.2 since the used hippocampus segmentations are publicly available as part of the ADNI database.

Table A.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

(iii) **FAUST** - An anthropological, open-access dataset of human body scans that have been acquired with a high-accuracy 3D multi-stereo system Bogo et al. [2014]. 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 (6890 vertices, 13776 faces).

³code online available: cs.technion.ac.il/~mirela/code/fmap2p2p.zip

Note that the male human body in two poses used e.g. in Figure 4.8 (right) is a subset of the whole FAUST set.

(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 Figure 4.8 (left).

(v) **Skeletal human hand** - A pair of skeletal human hands (see Figure 4.6) taken from the publicly available data⁴ of [Kilian et al., 2007] that is based on data of the Large Geometric Models Archive from the Georgia Institute of Technology. It features 13072 vertices and 26160 triangles.

(vi) **OAI-B-score** - Within B-score evaluation we rely on 3D sagittal weDESS MRI acquired at baseline as part of the OAI database (similar to dataset (i)). We segmented the distal femur bone for 9290 (of 9345) scans and established pointwise correspondence employing a fully automatic method combining CNNs with statistical shape knowledge achieving state-of-the-art accuracy [Ambellan et al., 2019b]. The 55 cases not taken into consideration were omitted due to imaging artifacts or failure of the segmentation pipeline. All reconstructed distal femur shapes are taken as left femurs, i.e. all right shapes were mirrored suitably and every shape consists of 8988 vertices and 17829 triangles (cf. Figure 5.10). Apart from image data the OAI database also provides clinical scores as KL and information about clinical outcomes such as TKR surgery. An overview on the employed data is given in Table A.3. Note that the list of unique MRI scan IDs defining the study population cannot be attached to the thesis at hand, because of its extend, but is publicly available as ancillary file (20 pages) accompanying the preprint⁵ to [Ambellan et al., 2021b].

A.2. Additional Experiments

In order to perform parameter studies and to access the influence of parameters on classification accuracy and model compactness we perform some additional experiments for the DCM and FCM.

⁴graphics.stanford.edu/~niloy/research/shape_space/shape_space_sig_07.html

⁵arxiv.org/abs/2104.01107

Table A.3.: Demographic information for the data involved to the B-score experiments.

No. of Shapes	9290
Laterality (left, right)	4678, 4612
Sex (male, female)	3825, 5465
Age [years]	61.1 ± 9.2
BMI [k/m^2]	$28,6 \pm 4.8$
KL 0, 1, 2, 3, 4	3382, 1557, 2313, 1198, 282
TKR within 8 years	508

A.2.1. DCM – Additional Experiments

Commensuration strategies

We evaluated the impact of the different commensuration strategies discussed in Section 3.4 by assessing the ability of the resulting model to capture characteristic changes in shape that are incident to radiographic knee OA. To this end, we performed the classification experiment presented in Section 5.1 using the proposed unweighted strategy, stretch-dominant weighting (ω close to zero) as motivated from 3D elasticity, and the data-driven approach.

The results are shown in Figure A.1. While all strategies yield statistical shape descriptors that outperform the PDM, the unweighted version ($\omega = 1$) shows the best performance. However, the accuracy of the data-driven approach improves with the amount of available training data suggesting a superior performance for large datasets.

Specificity

Low values for specificity (cf. Section 4.1.1) indicate that the model under consideration generates instances similar to the training data. However, it does not allow to differentiate whether the model is specific to the nonlinear distribution underlying the training data or just a few of its training instances. Therefore, we also evaluate how often each training instance was most similar (i.e. closest) to one of the sampled instances. The histograms presented in Figures A.3 and A.2 show that the proposed model is indeed specific to the distribution from which the training data was drawn.




Figure A.1.: Classification accuracy achieved by the DCM approach (employing different commensuration strategies) and the PDM as function of the percentage of training data.

A.2.2. FCM – Additional Experiments

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 3.7) 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 Figure A.4 the disease classification accuracy increases as ω increases. Looking at Equation (3.23) 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 Figure A.4. 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 Figure A.5. For values $\omega \gtrsim 10$ the classification accuracy lies below

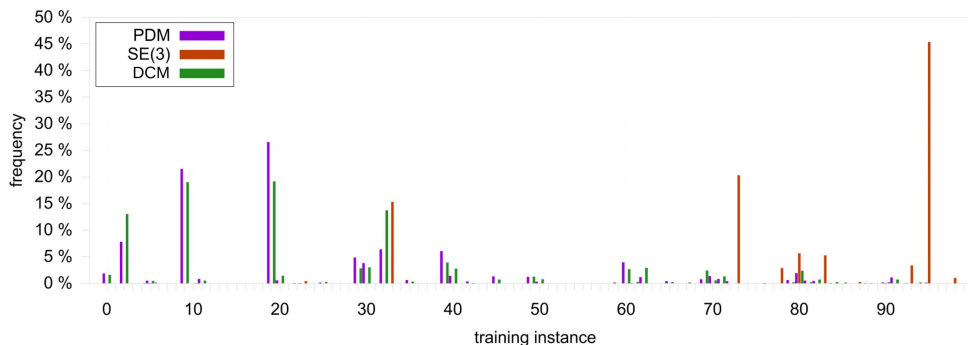


Figure A.2.: Frequency of being the closest shape (w.r.t. \mathcal{W}) during specificity evaluation (full model) for each training shape within FAUST dataset.

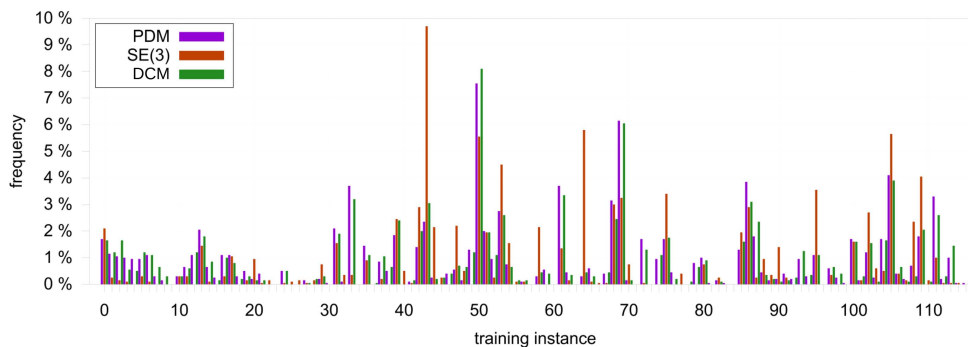


Figure A.3.: Frequency of being the closest shape (w.r.t. \mathcal{W}) during specificity evaluation (full model) for each training shape within OAI dataset.

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 Figure A.5, 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 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.

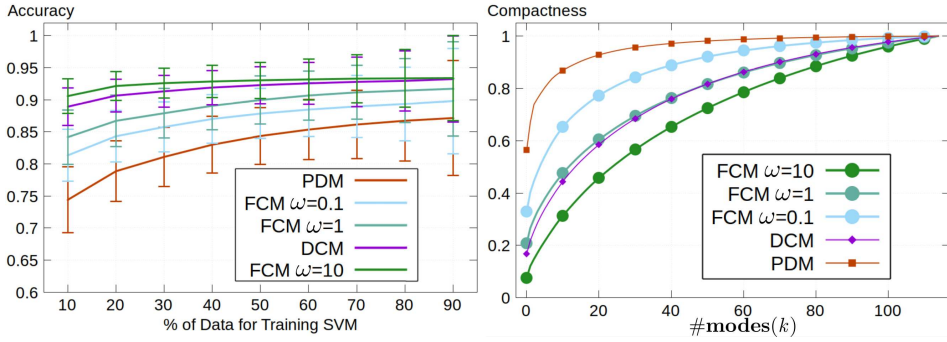


Figure A.4.: Varying commensuration parameter ω , as it increases the classification accuracy does and compactness decreases.

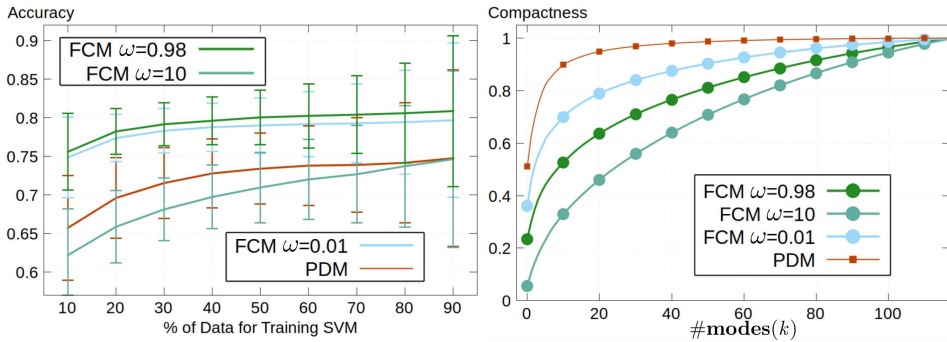


Figure A.5.: Varying commensuration parameter ω , with peak classification accuracy for $\omega \approx 0.98$. As it increases the model compactness decreases.

A.3. Technical Acknowledgements

During the research period the author was 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) and the Bundesministerium für Bildung und Forschung (BMBF) through project *TOKMIS: Treating Osteoarthritis in Knee with Mimicked Interpositional Spacer* (01EC1406E). This work partly relies on data of the Osteoarthritis Initiative. 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 fund-

ing 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. It moreover partly relies on data of the Alzheimer's Disease Neuroimaging Initiative (ADNI). 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. 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].

APPENDIX B.

Zusammenfassung

Die vorliegende Arbeit beschäftigt sich mit der Herausforderung der Entwicklung statistischer Formmodelle, welche die inhärente nichteuklidische Struktur (anatomischer) Formvariation berücksichtigen, gleichzeitig eine effiziente, numerisch robuste Verarbeitung erlauben und zusätzlich möglichst viel Invarianz unter euklidischen Bewegungen der verwendeten Daten bieten.

Dazu schlagen wir einen kontinuierlichen und physikalisch motivierten Formenraum basierend auf Deformationsgradienten vor. Wir verfolgen zwei verschiedene Ansätze, um auf diesem eine Riemannsche Struktur und damit ein statistisches Formmodell zu etablieren. (1) Wir entwickeln ein Modell für die differentiellen Koordinaten als Elemente in $GL^+(3)$. Zu diesem Zweck adaptieren wir den Begriff des biinvarianten Mittelwerts bezüglich eines affinen Zusammenhanges auf $GL^+(3)$ und führen Statistik zweiter Ordnung basierend auf einer Familie maximal invarianter, d.h. $GL^+(3)$ -links- und $O(3)$ -rechts-invarianter, Riemannscher Metriken durch. (2) Wir versehen die differentiellen Koordinaten mit einer nichteuklidischen Struktur, die angelehnt an die Lie-Produktgruppe aus Streckungen und Rotationen ist. Diese lässt eine biinvariante Metrik und damit eine konsistente Analyse mittels mannigfaltigkeitswertiger Statistik im Riemannschen Rahmen zu.

Die vorliegende Arbeit präsentiert überdies die *fundamentalen Koordinaten*, eine neue Formrepräsentation basierend auf diskreten Flächenfundamentalformen, welche auf natürliche Weise invariant unter euklidischen Bewegungen ist. Wir versehen diese Repräsentation mit einer Lie-Gruppenstruktur, die eine biinvariante Metrik und somit Riemannsche Statistik erlaubt. Darüber hinaus entwickeln wir einen einfachen, effizienten, robusten, sowie akkuraten (d.h. ohne Rückgriff auf Modellapproximationen) Löser für die Rückabbildung von den Koordinaten zur Form im Raum. Neben der statistischen Formanalyse erlaubt der beschriebene Ansatz auch Anwendungen in der Geometrieverarbeitung, insbesondere zur quasi-isometrischen Oberflächenverflachung.

Der letzte Abschnitt der Arbeit befasst sich mit der Entwicklung kontinuierlicher formbasierter Erkrankungsstratifikationen, um die Krankheitsbewertung über die aktuelle klinische Praxis ordinaler Bewertungssysteme hinaus zu objektivieren. Hierzu entwickeln wir den *geodesic B-score* in gekrümmten Formenräumen zur Bewertung von Kniegelenksarthrose als Generalisierung des euklidischen B-scores. In diesem Rahmen beschreiben wir eine Newton-Typ-Fixpunktiteration zur Bestimmung der Projektion auf Geodätische im Formenraum. Als Anwendung zeigen wir, dass der geodätische B-score gegenüber seinem euklidischen Gegenstück eine verbesserte Vorhersageleistung hinsichtlich der Risikobewertung bezüglich einer totalen Kniearthroplastie besitzt.

APPENDIX C.

Selbstständigkeitserklärung

Ich erkläre gegenüber der Freien Universität Berlin, dass ich die vorliegende Dissertation selbstständig und ohne Benutzung anderer als der angegebenen Quellen und Hilfsmittel angefertigt habe. Die vorliegende Arbeit ist frei von Plagiaten. Alle Ausführungen, die wörtlich oder inhaltlich aus anderen Schriften entnommen sind, habe ich als solche kenntlich gemacht. Diese Dissertation wurde in gleicher oder ähnlicher Form noch in keinem früheren Promotionsverfahren eingereicht.

Mit einer Prüfung meiner Arbeit durch ein Plagiatsprüfungsprogramm erkläre ich mich einverstanden.

Berlin,

Felix Ambellan