

Visual analytics methods for
shape analysis of biomedical images
exemplified on rodent skull morphology

Simon Maximilian Hermann

Visual analytics methods for
shape analysis of biomedical images
exemplified on rodent skull morphology

DISSERTATION

zur Erlangung des Doktorgrades (Dr. rer. nat.)
der

Mathematisch-Naturwissenschaftlichen Fakultät
der Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von

Dipl.-Inf. Simon Maximilian Hermann

aus Neuwied am Rhein

Bonn, August 2016

Angefertigt mit Genehmigung der Mathematisch-Naturwissenschaftlichen
Fakultät der Rheinischen Friedrich-Wilhelms-Universität Bonn

Dekan: Prof. Dr. Johannes Beck

1. Referent: Prof. Dr. Reinhard Klein
2. Referent: Jun.-Prof. Dr. Thomas Schultz

Tag der mündlichen Prüfung: 20. April 2017
Erscheinungsjahr: 2017

Universität Bonn
Institut für Informatik II
Friedrich-Ebert-Allee 144, D-53113 Bonn

Contents

Abstract *v*

Acknowledgements *vii*

I Introduction and background

1 Introduction *3*

- 1.1 Motivation *3*
- 1.2 A visual analytics approach to shape analysis *5*
- 1.3 Contributions *9*
- 1.4 Background on concepts of shape analysis *10*
- 1.5 Applications to rodent skull morphology *13*
- 1.6 Outline *15*

2 State-of-the-art *21*

- 2.1 Navigation in shape space *21*
- 2.2 Visualization of shape variability *24*

3 Modeling of form and its variability from images *29*

- 3.1 Introduction *30*
- 3.2 Biomedical images and their deformation *33*
- 3.3 Image registration *35*
- 3.4 Estimating an average template *39*
- 3.5 Statistical deformation model *44*
- 3.6 Summary *49*

II Visual shape analytics

4 Accurate interactive visualization of large deformations 53

- 4.1 Introduction 55
- 4.2 Related work 57
- 4.3 Visual analysis methods 60
- 4.4 Evaluation 70
- 4.5 Visual analysis of rodent anatomy 74
- 4.6 Discussion 78
- 4.7 Conclusion 79

5 Visual analysis based on prior knowledge 81

- 5.1 Introduction 82
- 5.2 Related work 84
- 5.3 Weighted analysis of a region of interest 86
- 5.4 Robust classification using machine learning 87
- 5.5 Semantic exploration 88
- 5.6 Visual analysis of traits in Gerbillinae and Murinae 89
- 5.7 Conclusion 95

6 Visual analysis of anatomic covariation 97

- 6.1 Introduction 98
- 6.2 Related work 100
- 6.3 Inter-point covariance 102
- 6.4 Visualization system 109
- 6.5 Experiment on a toy dataset 111
- 6.6 Visual analysis of modularity in *Mus* 114
- 6.7 Visual analysis of group differences in *Cricetinae* 121
- 6.8 Conclusion 123

III Closing

7 Conclusion 127

- 7.1 Summary 127
- 7.2 Future prospects 128

A Data sets and preprocessing 131

- A.1 Acknowledgements 131

- A.2 Acquisition *131*
- A.3 Data sets used in this thesis *132*
- A.4 Preprocessing *132*

- B Cross-validation for model-based deformation** *135*
 - B.1 Cross-validation algorithm *135*
 - B.2 Smoothness of cross-validation function *136*

- C Proof of error estimate for heuristic inverse** *137*
 - C.1 Linear approximation of inverse *137*

- D Mathematical notation & acronyms** *139*
 - D.1 Mathematical notation *139*
 - D.2 List of acronyms *142*

Abstract

In morphometrics and its application fields like medicine and biology experts are interested in causal relations of variation in organismic shape to phylogenetic, ecological, geographical, epidemiological or disease factors – or put more succinctly by Fred L. Bookstein, morphometrics is *the study of covariances of biological form*.

In order to reveal causes for shape variability, targeted statistical analysis correlating shape features against external and internal factors is necessary but due to the complexity of the problem often not feasible in an automated way. Therefore, a visual analytics approach is proposed in this thesis that couples interactive visualizations with automated statistical analyses in order to stimulate generation and qualitative assessment of hypotheses on relevant shape features and their potentially affecting factors. To this end long established morphometric techniques are combined with recent shape modeling approaches from geometry processing and medical imaging, leading to novel visual analytics methods for shape analysis.

When used in concert these methods facilitate targeted analysis of characteristic shape differences between groups, co-variation between different structures on the same anatomy and correlation of shape to extrinsic attributes. Here a special focus is put on accurate modeling and interactive rendering of image deformations at high spatial resolution, because that allows for faithful representation and communication of diminutive shape features, large shape differences and volumetric structures. The utility of the presented methods is demonstrated in case studies conducted together with a collaborating morphometrics expert. As exemplary model structure serves the rodent skull and its mandible that are assessed via computed tomography scans.

Acknowledgements

I couldn't have accomplished this thesis without the help from many people I want to thank hereby. First of all I have to thank my supervisor Prof. Dr. Reinhard Klein who gave me the opportunity to work on this dissertation project. And I am especially grateful to my second reviewer Jun.-Prof. Dr. Thomas Schultz for his support, the many interesting discussions we had and for the projects I had the opportunity to work on with him. I am also deeply indebted to Dr. Anja C. Schunke who functioned not only as collaborator and morphometric expert in the project and provided a very exquisite dataset of 3D μ CT images but was a continuous source of motivation with her infectious enthusiasm about the project.

Second, I want to thank the people at the department of computer science II that made working there most of the time fun and interesting. In particular I'd like to thank Stefan Hartmann and Raoul Wessel for bearing with me as their room-mate, Vitalis Wiens for bearing with me as his part-time supervisor, Björn Krüger for the interesting collaboration and Roland Ruiters for the always inspiring discussions.

I am also grateful to Prof. Dr. Christian Bauckhage and Prof. Dr. Oliver Gruss who kindly agreed to serve as members of my thesis committee. Furthermore I would like to acknowledge that my research was supported in between 2009 and 2012 by the federal state of NRW with a scholarship in the scope of the B-IT Research School and from 2012 to 2015 by the Deutsche Forschungsgesellschaft (DFG) within the priority program SPP 1335 Scalable Visual Analytics.

Finally, I want to thank my family. My parents and my brother made me always believe in myself. And my wonderful partner Sabine supported me through all turbulences and especially the stressful times before paper deadlines.

Part I

Introduction and background

Chapter 1

Introduction

A measure of an effective visualization can also be its ability to generate unpredicted new insights, beyond predefined data analysis tasks. After all, visualization should not only enable biologists to find answers but also to find questions that identify new hypotheses.

— Purvi Saraiya et al. 2005 [119]

This introductory chapter aims to motivate the usefulness of visualizations for morphometric studies and introduce our visual analytics approach for exploration of shape variability in a biomedical image ensemble. The needed concepts of shape analysis for the presented approach are sketched and some background information is given on the exemplary application to rodent skull morphology. The chapter closes with the thesis outline and a tabular overview of the introduced methods.

1.1 Motivation

Evolution has spawned a fascinating variety of species, each with its characteristic outwards appearance as well as unique structure and form of internal parts – all of which are the subject of study in *morphology*. The form of organs, i.e. their anatomy or *shape*, has always intrigued biologists because of the sheer variation on all scales and throughout all orders – even within a species or between closely related ones, no two individuals are alike. Identifying the determining factors that affect organ shape and its variability is a major objective of *morphometrics*. Its

tools are used in medicine and biology to correlate shape variation to phylogenetical, geographical, epidemiological or disease factors. Example applications in biology include identification of traits that characterize evolutionary relationship for taxonomic classification and deriving quantitative shape parameters that can be used as phenotypical markers for genetic studies. Application in medicine is targeted at image based diagnostics, for instance to differentiate pathological from healthy organ shape or to predict the progress of a degenerative disease. Currently both topics are actively being researched.

Unfortunately a fully automated morphometrical analysis is often not feasible due to the complexity of the problem, i.e. the many degrees of freedom in shape versus the multitude of influencing factors. In order to enable conclusive statistical tests a careful study design is thus necessary where the expert a-priori selects relevant shape features and formulates specific hypotheses on their affecting factors. Apparently, for this the expert has to be well informed on the shape variation in question and its potential relations to extrinsic factors. The required detail knowledge is usually gained via an initial inspection of the shape ensemble, i.e. an upstream *exploratory* study is conducted on the same dataset that is used subsequently for detailed statistical analysis.

Traditionally, visualizations play a major role for this kind of exploration. On the one hand there are abstract statistical displays that are good to illustrate higher order relationship in-between the individuals of the ensemble, e.g. in form of a scatter plot, while on the other hand there are more concrete 2D/3D displays of shape and shape differences, e.g. by depicting a deformed shape or multiple shapes superimposed. The latter method takes advantage of human shape perception capabilities and communicate easily particular shape features. Although visualizations play a critical role during morphometric studies, so far their designs are primarily driven by the intent of communicating final results of a statistical analysis.

The main motivation behind this thesis is to provide tools for an interactive exploration that assists experts in getting *insight* into the variability of shape with respect to its potential sources. More specifically, the envisaged exploration should stimulate hypothesis generation and support subsequent qualitative assessment. To accomplish this, a *visual analytics* approach is employed.

1.2 A visual analytics approach to shape analysis

Visual analytics, as Daniel Keim et al. [76] put it, *combines automated analysis techniques with interactive visualizations for an effective understanding, reasoning and decision making on the basis of very large and complex datasets*. Applying these concepts to shape analysis is what drives the methods developed in this thesis – methods that therefore might be subsumed under the term *visual shape analytics* [1].

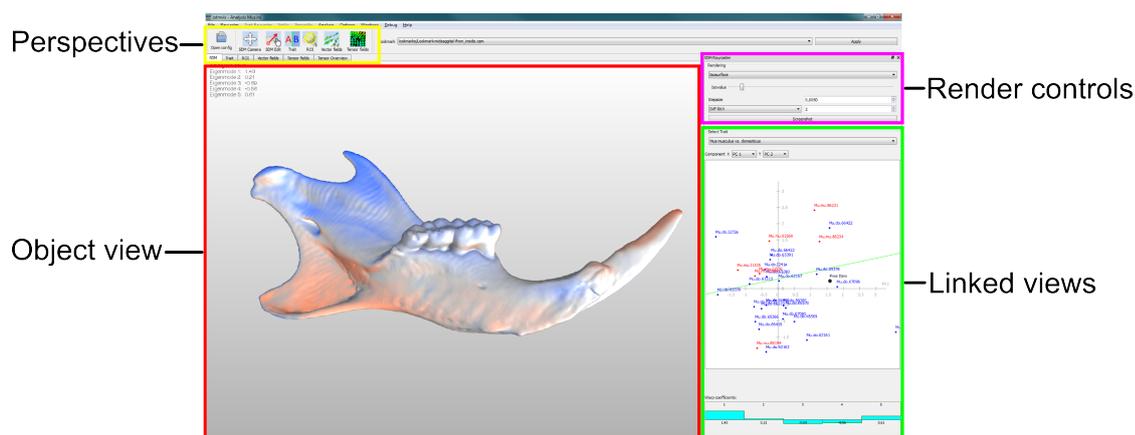


Figure 1.1: Screenshot of the proposed visual analytics system. While the object space view shows different visualizations of shape variability, the linked views provide abstract representations of shape space like the shown interactive scatter plot and coefficient histogram.

A central task of these methods is to establish a link between the two main data spaces encountered in shape analysis: 3D *object space* that serves the geometric representation of shape and *shape space*, a high-dimensional space where statistical modeling and analysis of shape ensembles is carried out. The essential duality between shape space and object space is also reflected in the user interface of our software prototype as shown in Fig. 1.1. An effective link between these spaces is realized by combining automated analysis techniques and interactive visualizations to facilitate *navigation* of shape space. Fig. 1.2 shows an example where the shape difference between two groups of a shape ensemble is characterized by a direction in shape space that is computed automatically and visualized subsequently in an animation of an accordingly deforming shape.

Although there are several works that deal with navigating shape spaces for 3D modeling and content generation [25, 129, 40, 158], only few exploration systems that support morphometric studies are described yet. Probably the first such system was proposed by Busking et al. [37] focusing on unconstrained navigation of a statistical shape model, without considering additional attributes. In a recent

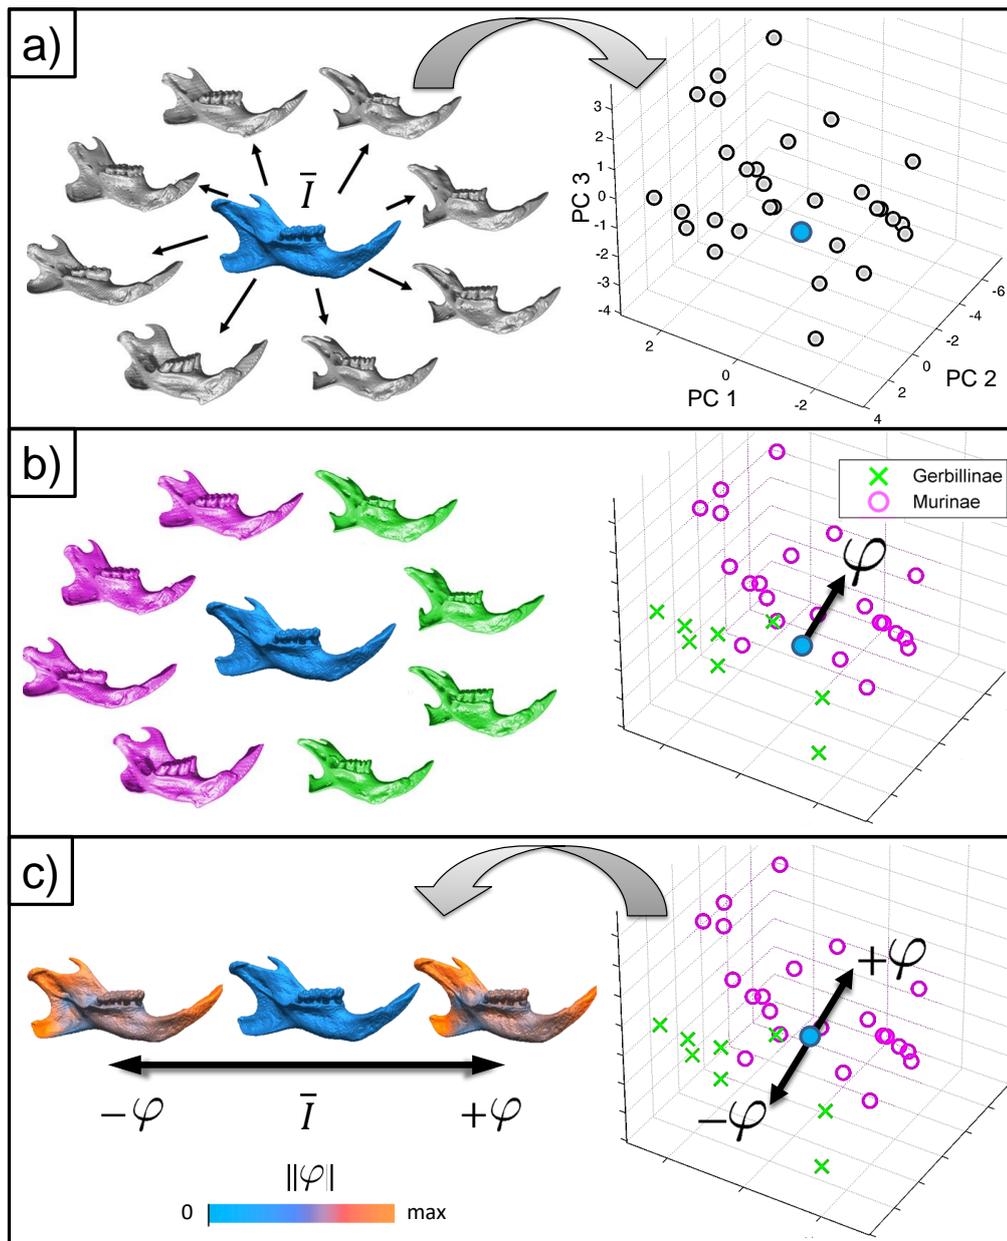


Figure 1.2: Exemplary visual shape analytics of a rodent mandible dataset, illustrating the pipeline from (a) automatic pre-processing over (b) user-interaction to (c) final visualization, shown both in 3D object space (left) and a 3D projection of shape space (right). In the end, a characteristic difference in size and orientation of the rear processes between two groups of rodents is apparent from the visualization and can also be quantified locally on the mandible. **(a)** The image ensemble, that is registered against a template \bar{I} , is mapped into a linear representation of shape space. **(b)** The user interactively classifies the ensemble into two phylogenetic groups. Based on that, a separating hyperplane with normal φ is derived. **(c)** An interpolation of images, achieved by synthesis along direction $\pm\varphi$, results in an animation that reveals the characteristic shape difference between the two rodent families. Color coded is the deformation magnitude $\|\varphi\|$ where warmer color indicates larger absolute difference from the template that is shown in blue.

work Klemm et al. [85] demonstrate how to effectively relate large amounts of heterogeneous attribute data with a statistical shape model using multiple linked views to assist epidemiological studies.

In this thesis we extend these original works, that nicely demonstrate the feasibility and effectiveness of an interactive navigation for shapes represented as 3D surface meshes, to the computationally more demanding volumetric setting that provides additional detail and makes non-surface structures accessible. Further, we expand the manual navigation technique of Busking et al. to a portfolio of techniques that support automated analysis in several ways. A noteworthy novelty in this regard is navigation via model-based editing that allows steering in shape space via direct user manipulation in 3D object view.

1.2.1 Workflow

Let us compare the visual analytics workflow against the established one in shape analysis based on the graphical overview of the according pipelines provided in Fig. 1.3. Both approaches have in common that they digitize a physical collection of specimens, e.g. via a computed tomography (CT) scans, yielding virtual representations of the physical individuals. The following preprocessing subsumes estimation of a template shape and determining transformations that encode the shape variation of the ensemble. For landmark methods, corresponding feature points have to be selected manually for this while for image based methods, required correspondences are automatically established via image registration algorithms. Up to this point there is no difference between the two pipelines despite that our approach is solely image based.

The main difference lies in subsequent analysis. Traditionally, a particular statistical analysis is performed and its results are visualized afterwards. In that way the analysis has to be rerun on any adjustment, e.g. change of hypothesis, prior information, parameters of a particular analysis or on choosing a different statistical method. This means that for an exploratory investigation the latter part of the pipeline has to be repeated multiple times, often in a cumbersome, manual way, involving separate software packages.

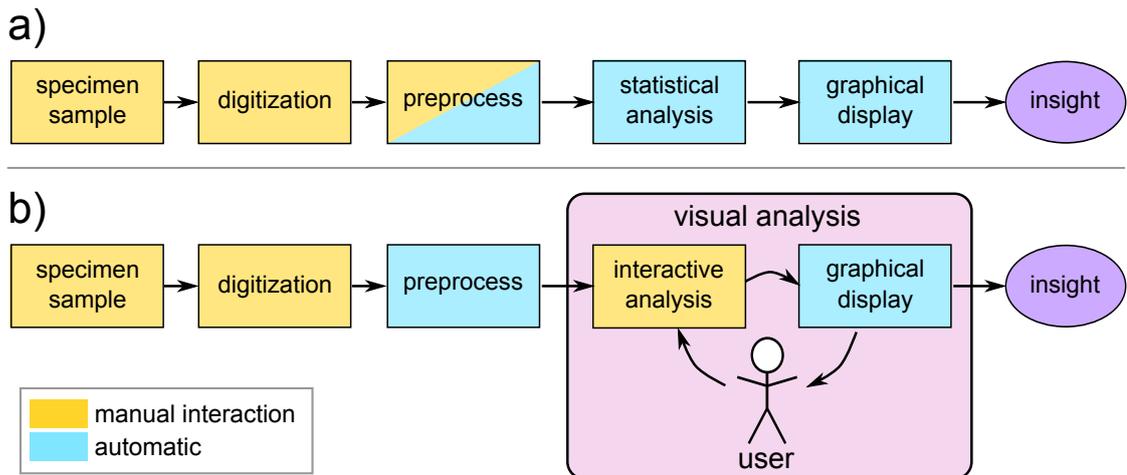


Figure 1.3: Comparison of work flows between (a) traditional statistical shape analysis and (b) the proposed visual analytics based approach.

1.2.2 Navigating shape space

As pointed out, key to our exploration system is an efficient link between the abstract shape space and 3D visualization of shape variation. The link is established by mapping an arbitrary point from shape space to its corresponding 3D shape, or more precisely speaking, an accordingly deformed template image. In the end this facilitates the already mentioned navigation in shape space, that is, sampling shapes at particular points, along any direction or even arbitrary trajectories in shape space. This *synthesis* of deformation provides the basic exploration facility of our visual analytics approach.

All exploration metaphors introduced in the work at hand are based on navigating shape space one way or another, based on different kinds of user input. Integral to our approach is the development of visualization techniques designed to effectively convey shape variation, because they facilitate the main feedback channel to the user input. To this end several visualization techniques are employed, ranging from dynamic volume and indirect isosurface rendering over vector field visualization to the use of tensor glyphs in order to encode higher order attributes of shape variation. Since these kinds of 3D representations in object space communicate the main information to be realized by the user, they are always supplied in a central view of the visual analytics system. Additional information is provided in linked views that are also available for user input.

A completely manual navigation in the spirit of Busking et al. [37] is implemented for instance via a linked scatter plot view where the user can specify trajectories in 2D projections of shape space. By that, an initial unconstrained

exploration of the shape variability contained in an ensemble can be performed, where synthesized shapes are dynamically displayed in the object view in an animation.

For a more targeted analysis, navigation is assisted by methods designed to operate on higher level semantic input from the user. This is accomplished by means of interactive statistical methods that derive interesting trajectories in shape space automatically. An example is the interactive classification illustrated in Fig. 1.2.

1.2.3 Efficiency considerations

A key requirement for every interactive approach is efficiency in order to provide instantaneous feedback to user actions. This constrains the choice of available shape analysis methods to those that can be realized interactively.

For the shape space employed in this thesis, namely that emerging from diffeomorphic transformations, two efficient strategies for linearization can be identified in this regard. In case of relatively small deformations one can resort to a well known ad-hoc approach that operates on linear displacement vector fields. This is the approach taken in Chap. 5 and 6. However, for ensembles exhibiting larger deformations a more elaborate strategy is required. To this end a recently developed parameterization of diffeomorphisms based on stationary velocity fields is adopted in Chap. 4, where an efficient implementation is derived, custom tailored for interactive visualization.

For statistical modeling the normal distribution model is considered throughout the work at hand as it is probably the most widespread model, with many success stories in applications in computer graphics, computer vision and medical image analysis as outlined in Chap. 2. In the second part of this thesis it is shown that this model permits several efficient navigation techniques.

1.3 Contributions

The main contribution of this thesis lies in the provided portfolio of automated analysis and interactive visualization methods for shape analysis in a morphometric context. An illustrative overview of the introduced methods is given in tables 1.1, 1.2, 1.3 at the end of this chapter. Put together and seen as a whole these methods constitute a framework for visual shape analytics. In particular this is accomplished to a great extent by the following technical contributions:

- **Fast and accurate rendering of image deformations based on a non-linear deformation model.** Based on recent advances in image analysis a reliable interactive visualization of pronounced shape differences is enabled as they typically arise in inter-species comparison. (Related methods: VIS1–4.)
- **Novel methods for visual analytics of shape co-variates.** Given a user specified labeling of the ensemble according to potential influence factors, corresponding representatives are computed on the fly. Visualizing a smooth interpolation in-between the representatives reveals the impact of the different factors on shape and, for the special case of two labels, displays their characteristic shape difference. (Related methods: VIS2, NAV2.)
- **A novel tensorial description of covariance between points on the shape.** Based on the model-based editing framework of Blanz et al. [24], a first set of visual analytics methods is derived for the research on modularity and integration, see Sec. 1.5.2 below. (Related methods: VSA1–4.)

As is common practice in computer graphics, most of the work presented in this thesis has been published previously [1, 2, 3, 4, 5].

1.4 Background on concepts of shape analysis

Exploration of shape variability builds upon a few central concepts of shape analysis, i.e. *transformations* that describe shape differences and allow to establish a template or mean shape, *shape spaces* that model the entirety of shapes representable for a specific class of transformations w.r.t. to such a template, and *statistical shape models* that provide concise representations of shape variability in shape space. Each of these concepts is briefly introduced now.

1.4.1 Transformations of shape

A century ago D’Arcy W. Thompson [134] recognized that the

essential task [of morphology] lies in the comparison of related forms rather than in the precise definition of each; and the deformation of a complicated figure may be a phenomenon easy of comprehension [...].

This statement holds several insights that are key to the exploration of shape variability, namely that one is actually interested in the *differences* between shapes, that these can be encoded via transformations that deform shape, and that the

inherent complexity of the latter encoding can be effectively communicated via visualization.

Thompson’s fundamental idea to encode shape difference via transformations still pervades modern day computerized morphometrics. A first thorough formalization of this approach is found in geometric morphometrics (GM) pioneered in the late 1970’s by Kendall [77, 78], Bookstein [27] and many others, that is later even celebrated as a “revolution” in morphometrics [117, 7]. In contrast to traditional quantitative methods based on ruler measurements, GM describes the shape’s geometry explicitly by sampling the organ outline or contour at so called landmark positions. These are defined as homologous feature points that are shared across all shapes of an ensemble thereby establishing point-wise correspondences between the shapes. A major achievement of GM is the development of shape statistics for this kind of landmark point configurations, with the *mean shape* as one of its central concepts. The mean shape, or *template* as we will refer to it, is computed by superimposing all shapes in a common reference frame and averaging the coordinates at each landmark position over all shapes in the ensemble. The crucial step in this construction, i.e. the superimposition, is achieved by applying transformations that remove differences in position and orientation, because those are arbitrary and therefore of no further interest.¹

Shortly after GM, the discipline of computational anatomy (CA) emerged based on the general pattern theoretical framework of Grenander [64, 65]. CA contributes shape statistics founded on *diffeomorphisms*, i.e. the set of smooth deformations that possess a smooth inverse. Based on diffeomorphisms, plausible descriptions and models of local shape differences for anatomies as different as bone and brain structures are successfully derived. A key technique of CA is the modeling of shape variability from 3D images at image resolution without the need for manually selected feature points. Hence this technique, that is also referred to as deformation based morphometry [15], is of special interest for medical image analysis.

¹Depending on the particular study sometimes also differences in isotropic scale and/or shear are factored out as well.

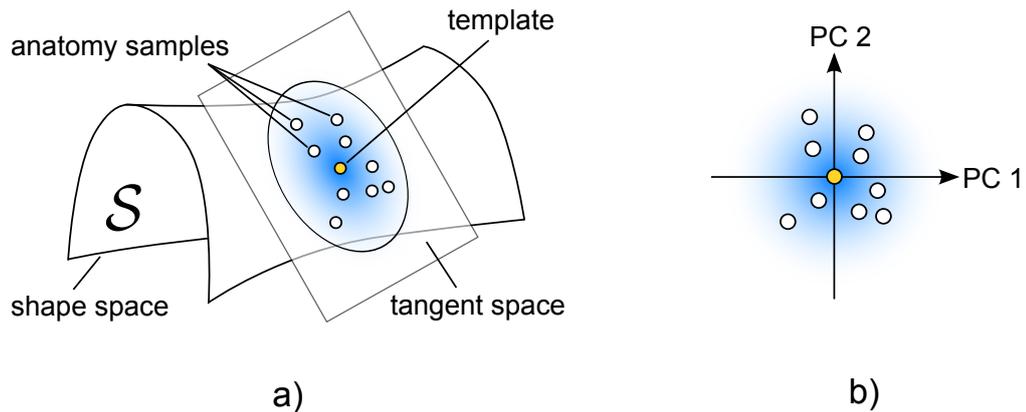


Figure 1.4: Illustration of shape space and its linearization. **(a)** Shape space is a high-dimensional non-linear, i.e. curved, manifold while a tangent space provides a template-centered linearization that permits multivariate analysis. (Figure adapted from [20].) **(b)** A statistical deformation model is derived via principal component analysis in tangent space.

1.4.2 Shape spaces

A *shape space* \mathcal{S} comprises the set of all possible shapes or images representable with respect to a particular class of transformations. Each point in \mathcal{S} corresponds to a specific transformation which applied to the template results in an actual shape or image. For measuring (dis)similarity of two shapes, shape space is endowed with a distance measure or metric that allows to define a shortest path or *geodesic* connecting any two points. An important property is that a deformation along a path in \mathcal{S} is required to be smooth and invertible, such that a shape is altered in a continuous fashion. The metric also facilitates definition of a generalized average, a Fréchet mean, that corresponds to a mean or template shape.

Unfortunately, practical computations on \mathcal{S} are computational expensive in general because of its global non-linear structure. In order to enable standard linear techniques of multivariate statistical analysis, \mathcal{S} has to be linearized. To this end a tangent space is spanned leading to a local flattening of shape space around a particular reference point, corresponding to a template shape, see Fig. 1.4(a).

1.4.3 Statistical shape models

Probably the most popular statistical model to explain the spread of samples in (linearized) shape space is that of a multivariate Gaussian or normal distribution. It is defined by first and second order moments, i.e. mean and variance. Since shape spaces are high dimensional, the second order moment is estimated as a large sample variance-covariance matrix. A much more compactly parameterized model is achieved via a principal component analysis (PCA) that provides a canonical coordinate system in terms of principal modes of shape variation, see Fig. 1.4(b).

Sampling from the PCA model allows to synthesize novel, virtual shapes that are useful to illustrate shape variability. Efficient synthesis is one of the key techniques to convey particular aspects of shape variation in the visual analytics approach proposed herein.

1.5 Applications to rodent skull morphology

The application to a real world dataset is always a great opportunity to demonstrate the potential of novel analysis methods. We had the great luck to develop our methods in a project together with Dr. Anja C. Schunke from the MPI Plön, a morphometrics expert who provided us not only with a very unique collection of high quality CT datasets of rodent skulls but also with her invaluable feedback and supporting discussions during development. The provided data offers in particular challenges like shape differences at interior, volumetric structures and large scale variations that are hard to assess with previous methods which are either limited to surface features or small scale image deformations.

Rodents in general and the mouse in particular is a popular model system in biology because it is a mammal with a very high diversity. Rodent skull and especially its mandible is a standard model for morphometric analysis which is also emphasized by its use as worked example throughout a popular introductory textbook on morphometrics [159].

For this thesis five datasets have been compiled as described in App. A, four of them consisting of rodent mandibles and one dataset of the upper skull.

1.5.1 Morphological structure of the rodent mandible

The rodent mandible is a particular interesting model structure, partly due to its relatively simple structure, e.g. without articulations. Nonetheless, it comes with several functional units, like incisor teeth for gnawing, cheek teeth for chewing, one articular and several muscle bearing processes. Irregardless of its relative simplicity, finding modules and modeling their interaction is considered a difficult task [90, 88] as described in the next section. Additionally, the mandible provides also some interior structures like tooth roots and incisor arch that are only accessible with 3D imaging.

Fig. 1.5 gives an overview of rodent mandible anatomy. Each separate half of the mandible consists of a single bone with three processes in the back, a row of molar teeth and the single incisor, whose posterior end lies far back in the mandible, below or behind the coronoid process. A common functional segmentation is the separation into two parts, the frontal region, bearing incisor and molar teeth, and the rear processes with muscle attachments.

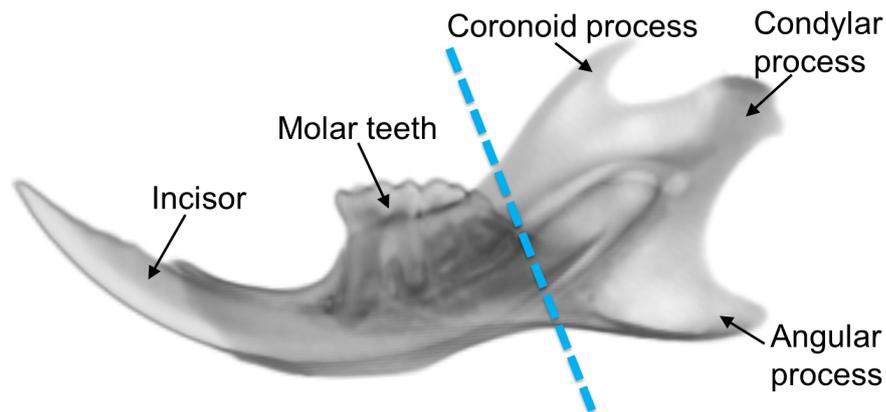


Figure 1.5: Anatomical parts of the rodent mandible referred to in this work. A common subdivision into two functional subunits is indicated [90].

1.5.2 Applications

The high diversity of rodents results in many interesting traits potentially related to phylogenetic and ecological factors. However, many sources for shape variation on the mandible are still not fully understood and there is a need for exploratory approaches to come to some useful hypotheses. With this in mind, datasets have been purposefully compiled to reflect different factors of phylogeny at subfamily and genus level as well as diet, distinguishing omni-, herbi- and carnivory.

Corresponding traits are investigated separately in Chap. 5, including allometric effects that are related to size. A simultaneous investigation of several factors is demonstrated in Chap. 4 for phylogeny and diet, comparing and disentangling their particular influences on shape variation. Another exploration in the same chapter also hints at a geographic gradient in the skull dataset.

A topic that steered some attention in recent years is that of modularity and integration. Integration refers to the degree that particular shape characteristics interdepend, and modules are parts of the shape that are tightly integrated, but are relatively independent from other modules [86]. Naively speaking, traits inside an integrated module probably share some of the processes or actions that shaped them during their evolution while separate modules may have developed more or less independent of each other. Based on the assumption that the amount of covariation between different shape features is a cue to their integration, a visual shape analytics pipeline is devised in Chap. 6 to investigate such covariation at multiple levels of detail.

1.6 Outline

The remainder of the thesis is structured as follows. Related work on shape space exploration and visualization is reviewed in Chap. 2. The necessary basics about modeling of form and its variability from images are introduced in Chap. 3. Key aspects covered in this chapter are a group-wise image registration algorithm and the calculation of a statistical deformation model. Initially, these aspects are introduced for the classical linear deformation model.

After introducing the general visual shape analytics approach in the current chapter and settling the basic modeling techniques in the subsequent two chapters, the second part, consisting of Chap. 4, 5 and 6, comprises our portfolio of visual shape analytics methods. The chapters of the second part should be sufficiently self-contained such that they can be read in any order.

In Chap. 4 a non-linear extension of the standard deformation model is presented, together with four methods taking specific advantage of this novel model, see table 1.1. Particularly the introduced group browser allows for a rapid investigation of the influence of different external factors on shape. The chapter further describes the implementation of a raycaster to render deformed images that is also used in the methods described in subsequent chapters.

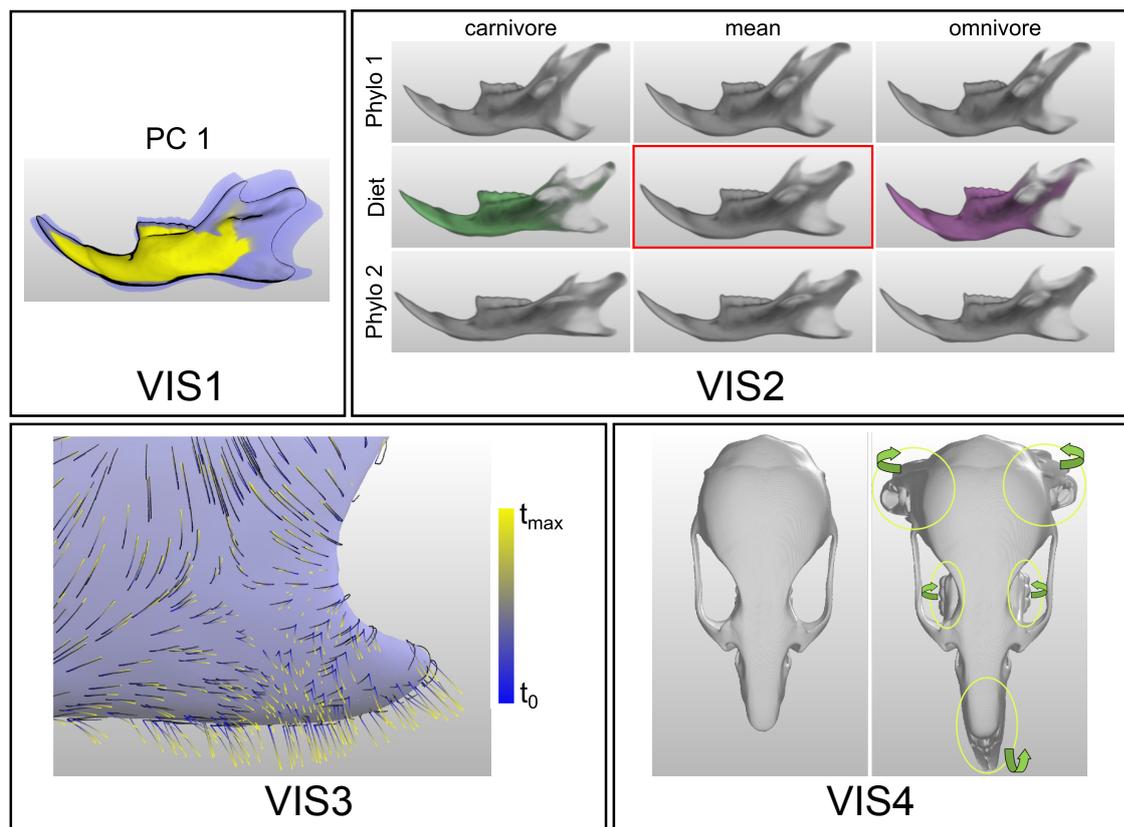
Chap. 5 shows two methods that, when used in combination, allow for a hierarchical navigation of shape space, see table 1.2. This is accomplished with respect

to prior knowledge on classification and regions of interest (ROI). Selecting a ROI allows to switch to a sub-space of shape space by recomputing PCA statistics with respect to the chosen ROI. Classification on the other hand allows to navigate to subspaces with particular shape differences filtered out, i.e. orthogonal to the characteristic difference between the two selected groups.

Unlike the two previous chapters, the focus of Chap. 6 is not on correlating shape variation against external attributes, but rather on the covariation between different parts *on* the shape. As described above, this plays a crucial role in research on modularity and integration. In order to investigate this kind of covariation and identify hypotheses on module delimits, three tightly interlinked visualizations are developed, see table 1.3. Together, they facilitate an effective interactive analysis of covariation at different levels of detail. All of these methods are based on a novel definition of a linear interaction operator that leads to two tensors, summarizing different aspects of covariation. Further, an additional automatic segmentation method is introduced based on one of this tensors. The resulting segmentation is anatomically meaningful and provides an additional way to inspire hypotheses on module delimits.

Each chapter in the second part devotes at least one substantial section to applications of the developed methods. To this end real-world datasets of rodent mandible and skull were investigated in collaboration with the domain expert Dr. Schunke.

Finally, Chap. 7 concludes the thesis with a brief summary and discussion of future prospects.



VIS1: Likelihood volume

Brief: Integrated visualization of a whole trajectory in PCA space, e.g. to get an overview of variation along principal modes.

Sample: Likelihood volume for the first principal mode (PC1) of a rodent mandible dataset shows that most variability is found in the rear three processes.

VIS2: Group browser

Brief: Comparative visualization of impact of multiple factors by interpolating between group mean shapes that are selected interactively, allowing to sift quickly through multiple comparisons.

Sample: Interpolation between several group mean shapes arising from phylogeny and diet shows that the impact of diet on shape is more emphasized in the second phylogenetic group (Phylo 2) compared to the first one (Phylo 1).

VIS3: Projected streamlines

Brief: Visualization of tangential part of shape variation.

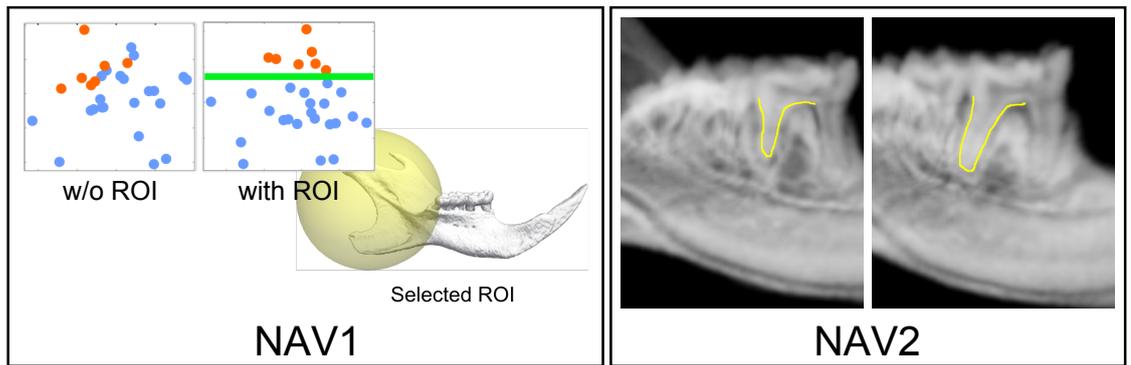
Sample: Streamlines of a shape variation (blue to yellow) and its tangential part (desaturated) show differentiated perpendicular and tangential trends in the process to the lower right versus the central structure.

VIS4: Reformation

Brief: Render otherwise occluded structures in a single view.

Sample: Visualization of a mouse skull viewed from above, left without and right with an applied reformation that makes incisor, auditory bulla and molar teeth become visible in a single view for joint investigation.

Table 1.1: Overview of visualization and rendering methods introduced in Chap. 4.



NAV1: Region of interest

Brief: Define a PCA space w.r.t. to an interactively selected region of interest in order to focus investigation on particular local structures.

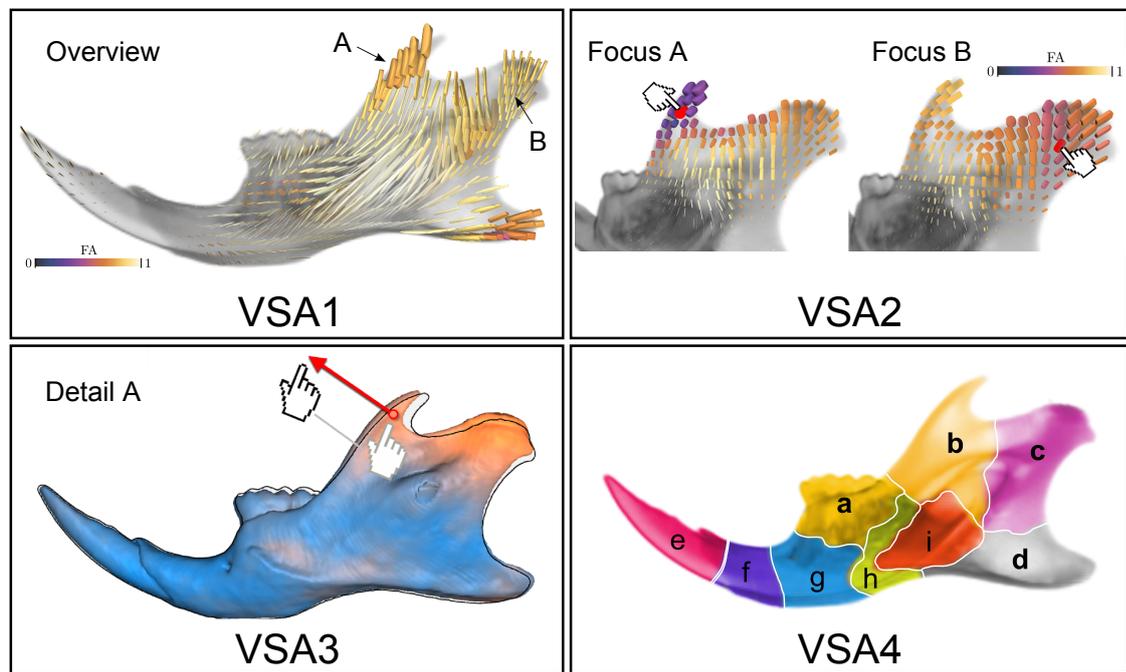
Sample: After selection of a ROI the correspondingly weighted PCA model shows a clear separation between two phylogenetic groups (orange and blue) on a single PC axis and makes it thus more easy to discover in an exploration.

NAV2: Classification

Brief: Define a trait vector as a direction in PCA space corresponding to the characteristic shape difference between two interactively specified groups.

Sample: A volumetric shape difference between two phylogenetic groups, i.e. in size and orientation of a tooth root, becomes apparent from visualizing the extremes of the corresponding phylogenetic trait vector in shape space.

Table 1.2: Overview of navigation methods introduced in Chap. 5.



VSA1: Overview tensor

Brief: Provide guidance to regions of interesting covariation *on* the shape.

Sample: The rear processes at A and B show strong covariation with other parts on the shape (related to size of glyph), although their preferred direction of variation is slightly different (related to orientation and shape of glyph).

VSA2: Interaction tensor

Brief: Visualize the covariation pattern w.r.t. a particular point on the shape.

Sample: The pattern of the interaction tensor fields at the Coronoid (A) and the Condylar process (B) reveals their mutual relationship.

VSA3: Model based editing

Brief: Visualize effects of a specific directional variation at a point interactively.

Sample: Model based editing shows particulars of covariation, e.g. that the Condylar follows the position of the Coronoid process.

VSA4: Anatomic segmentation

Brief: Automatically partition the shape w.r.t. its intrinsic covariance.

Sample: The computed anatomic segmentation nicely correlates with functional units.

Table 1.3: Overview of the methods introduced in Chap. 6 to facilitate visual shape analytics of anatomic covariation, e.g. to identify hypotheses on module boundaries.

Chapter 2

State-of-the-art

A brief review of related work on navigating shape spaces and visualization of shape variability is compiled in this chapter to put this thesis into context.

2.1 Navigation in shape space

A particular focus of application of shape space representations in computer graphics is effective authoring of 3D content by means of interpolation in available 3D model databases, see [25, 9, 10] to cite just a few works in this field. In computer vision [45] and medical image segmentation [69] shape space representations are used to introduce model knowledge. A special advantage in all these works is their combination of statistical analysis and efficient synthesis to generate novel 3D shapes that are plausible w.r.t. a statistical model. This is exactly what is necessary for interactive visual exploration of shape variability in the context of visual shape analytics. However, in order to enable *targeted exploration* of a shape ensemble, additional methods for navigation in shape space are required. A key challenge in this context is to make the high dimensionality of shape spaces accessible.

2.1.1 Navigation along traits

A first idea on this was already given by Blanz and Vetter [25] who parametrized the shape of human faces via regression on semantically motivated traits like age, sex, weight, etc. in PCA space. They demonstrated that exaggerating these traits can be used to create easily understandable caricatures of certain type. Matusik et al. [102] showed that navigation along traits is an effective means of identifying specific appearance characteristics of surface reflectance functions. In the context of visual shape analytics this idea is applied in Chap. 5 to relate shape variation to external attributes.

2.1.2 Navigation via scatter plots

Manual exploration using two dimensional views as interfaces for navigation have been suggested by several authors. Kilian et al. [79] present a shape exploration based on barycentric interpolation between example shapes. To this end a 2D embedding view of the shape ensemble is derived via multidimensional scaling (MDS) followed by a triangulation. By drawing curves in this view, arbitrary interpolations can be explored. Instead of a triangulation, Smith et al. [129] rely on generalized barycentric interpolation inside a convex control polygon that, by clicking a point inside the polygon, allows the user to dial up a particular affine combination of a set of registered car shape models. Additional regression values on specific attributes like *sportiness* are overlaid on the polygon for guidance. For the specific case of mesh animations, Cashman et al. [40] use a combination of MDS and radial basis functions to come up with a 2D map visualization of the animation as a spline curve. On this map, a repetitive motion will for instance show up as a curve with several loops. By manipulating the curve, the animation can be edited in a high level way. Busking et al. [37] use a scatter plot that shows a 2D projection of PCA space. The projection can be adjusted interactively by manipulating 2D representations of a set of axes or vectors in shape space [23]. For synthesis of shapes in-between sample points in the 2D projection natural neighbors interpolation is used, based on a Voronoi tessellation that is computed efficiently on the GPU. Klemm et al. [85] use multiple linked views to explore medical population data for epidemiology, e.g. to identify disease-specific risk factors. Aim of their interactive visual analysis is parameter and group selection for subsequent statistical analysis. The data also includes MR images from which 3D surface models of the lumbar spine are semi-automatically extracted. During exploration, mean shapes of selected groups are displayed, colored according to

their difference to the global mean shape.

2.1.3 Direct manipulation approaches

An interesting alternative to interaction with abstract 2D views and scatter plots are direct manipulation approaches to explore and generate shape variations. Probably one of the first approaches in this regard is model based editing introduced by Blanz et al. [24]. Based on the user modifying the position of just a few feature points their approach optimizes the most likely shape that matches the user input as closely as possible. Thanks to the linearity of the PCA model this optimization turns out to be a simple least squares problem that can be solved efficiently. Lewis and Anjyo [97] pick up the same idea for editing facial blendshape models while Tena et al. [133] and Berner et al. [22] present generalizations of this approach to part based shape models. Coffey et al. [44] present an interactive manipulation interface to navigate the space of simulation outputs in order to refine the design of a mechanical biopsy device, taking into account its functionality. Interestingly, the metaphor of direct spatial manipulation has been recently applied also to time-varying scatter plots [92], where dragging around a point facilitates navigation in time by matching the input to an existing point and its temporal trajectory. In Chap. 6 model based deformation is used to analyze covariance on shapes.

2.1.4 Navigation of subensembles

For industrial CT images comparative visualizations were made for the analysis of defects for material sciences [112]. In order to visualize the shape distribution of a set of feature objects, pores or other material defects in form of an uncertainty cloud the concept of mean objects was introduced. Clustering of mean objects provides a hierarchical representation well suited for exploration.

A common task for exploratory morphometric analysis is to disentangle the factors that determine shape variation. A visual analytics method to accomplish this is introduced in Chap. 4. Categorical factors decompose the shape ensemble into subsets, for instance into several phylogenetic or dietary groups. In order to unveil the impact of each factor on shape variation, mean shapes of the corresponding subsets are derived on the fly, enabling interpolation in-between group means and the ensemble template.

2.2 Visualization of shape variability

Although visualization plays such a central role in shape analysis, there seem to be only two articles published yet that give sort of a survey [89, 38]. Klingenberg [88] critically discusses common visualization methods for landmark analyses in geometric morphometrics and provides helpful guidelines for practitioners. Different visualization options for statistical deformation models used in computational anatomy are compared by Caban et al. [38] and evaluated in a small user study. Both works contribute valuable insights about effectiveness and limitations of many important visualization techniques. However, some often encountered visualizations such as color coded isosurfaces or vector fields are missing in the mentioned surveys, and animation is not discussed either.

In the following a brief review of visualization techniques is given, organized by their primary underlying visual paradigm: *Superimposition* and *side-by-side comparison* relate to spatial layout, *direct visualization* focuses on ways to display deformation by warping methods, *encoded visualization* is about the use of color-coding and glyphs to communicate higher order information and finally, *animation* deals with the temporal dimension.

2.2.1 Superimposition

The original shape samples are shown superimposed in a reference coordinate system, e.g. given by Procrustes alignment. This kind of display is quite common and effective for 2D landmark and contour data [89] and is used in many publications and textbooks in geometric morphometrics. An advantage is, that it does not require a deformation or statistical model per se. Nevertheless, plotting for instance superimposed landmarks yields point clouds whose distributions reveal the local covariance structure at each landmark. Superimposing contour data gives a good overview of global variability but quickly becomes cluttered for many contours. In our experience, this cluttering becomes even worse when superimposing 3D surfaces [18], because of the additional occlusion interfering with the superimposition. In practice we observe that at most three surfaces are shown superimposed using alpha blending and contrasting colors, see e.g. Abbasloo et al. [6].

Superimposition is also used to assess results of pairwise registration of surfaces or images. The interactive 3D volume registration system of Smit et al. [128] makes use of multi-volume rendering to superimpose fixed and moving volume, color-coded and opacity blended to reveal areas of mis-registration. The checkerboard method is an alternative way of superimposing two 2D images (or slices of

a volume) that does not require blending. The white squares of the checkerboard offer a view onto one of the images, the black squares onto the other. A generalization of this technique to more than two images was presented by Malik et al. [100] and a generalization to tensor field visualization was recently given by Zhang et al. [161].

Likelihood volumes [38, 71] can be understood as a generalized superimposition of 3D images by means of blending more than two images. An efficient implementation of a likelihood volume for a non-linear deformation model is presented in Chap. 4 where it is used as an overview visualization. When sampling a deformation densely, likelihood volumes produce a visualization resembling motion blur. A similar approach was taken to visualize the uncertainty of estimated isosurfaces [109, 108].

2.2.2 Side-by-side comparison

Instead of superimposing one or more shapes in a single view, multiple views can be employed as well. This provides an alternative in cases where superimposition is not applicable or would lead to a cluttered display. Unfortunately, small scale shape variations are hard to recognize in a side-by-side display. Following Tufte's small multiples [139], a small-scale shape rendering can serve as an iconic representation that allows comparative displays showing many shapes at once. This technique is used for instance to overlay small shape renderings on a scatter plot showing a 2D projection of shape space [37] or to visualize mean shapes of different sub groups of the dataset [85].

2.2.3 Direct visualization

This paradigm subsumes approaches that depict deformations explicitly by deforming a graphical representation of the shape or the embedding 3D space. A specific appeal of direct space warping techniques is their generality, that they can be applied to landmark and surface data in 2D or 3D in the exact same manner.

Showing a distorted Cartesian grid is amongst the classic methods to illustrate anatomic differences, as it was made popular by D'Arcy Thompson [134] and used even earlier by Artists like Dürer and Da Vinci in their anatomical studies. While these early examples were hand-crafted, the first automatic graphics procedure was introduced by Bookstein [26] based on thin-plate spline (TPS) interpolation of space in between landmarks. TPS remains one of the dominant visualizations in morphometrics to this day [89], not least because its efficiency. Wiley et al. [151] use

TPS for instance to interpolate between known sample shapes from an evolutionary tree to generate hypothetical ancestral shapes.

Somewhat a hybrid between direct and encoded visualization (see below) are the deformable grids [41, 38]. Initially developed for 2D uncertainty data [41] they were generalized to show anatomic variation from statistical deformation models in 3D by Caban et al. [38]. A very coarse grid is overlaid onto the image and deformation is visualized by modulating the depiction of grid edges, e.g. by drawing an edge as a sinusoid curve with the local deformation magnitude mapped to its amplitude.

2.2.4 Encoded visualization

In contrast to direct visualization, methods that fall under this paradigm visualize particular aspects of deformation implicitly by means of color coding or glyph rendering. Scalar attributes are easily visualized via color coding by applying a transfer function that maps the scalar value range to some color gradient. In computational anatomy one often encounters variability and probability maps that color code magnitude of local variability and outcome of statistical tests respectively [135]. Hamarneh [67] use color coding to highlight “hot spots” of localized shape variation. Lüthi et al. [99] use color coding to visualize the remaining flexibility of a statistical shape model after parts of it have been fixed, for instance by a semi-automatic model based registration procedure. Zollikofer and Ponce de Léon [163] show a successful combination of color coding and vector field visualization on 3D surfaces to communicate deformation decomposed into directions parallel (vector field) and perpendicular (color) to the surface. Kirschner and Weisarg [83] present an implementation of this kind of visualization in an interactive system for active shape models.

Kindlmann et al. [81] visualize anatomic covariance tensor fields using superquadric tensor glyphs that summarize the local covariance structure at each sample point on the surface of a mean shape. For each point a 3×3 sample covariance matrix on the set of displacement vectors from the mean to each individual is computed. Additional scalar measures derived from the covariance tensor data like fractional anisotropy and Frobenius norm are used for color coding glyphs and shape surface respectively. The same glyph visualization is used for the covariance tensors described in Chap. 6. Van Golen [141] uses custom glyphs to show the influence of each landmark on an active shape model, i.e. how strongly the overall shape variation described by the model depends on a particular landmark.

When dealing with image based shape models, deformations are often rep-

resented as dense vector fields. This enables vector field visualization methods like color coding of Jacobians [115], detection of critical points and display via glyphs [136] or color coding custom tailored scalar flow measures [36]. Streamline rendering is another vector field visualization method [124] that is used in Chap. 4 to uncover the tangential component of non-linear shape variations.

2.2.5 Animation

Showing a particular variation as image deformation in an animated way is an ideal presentation to the human eye [135, 96]. It allows to utilize the excellent motion perception capabilities of humans that renders small deformations much better perceivable than from a set of static images. Therefore, animation is one of the preferred visualizations in many approaches. It is an obvious choice when illustrating dynamic processes like respiratory motion of lungs and inner organs in humans. Handels and Hacker use animation to present an interactive anatomical atlas [68], exemplary modeling the kidney via a medial representation [62].

Real time animation, while easy to achieve in principle for 3D surface models, poses a challenge for 3D image models. This results from the fact that 3D image warping involves the *inverse* mapping, that is computationally expensive to approximate. An in-depth discussion of that fact is given in Chap. 4 where advantage is taken of the log-domain framework to efficiently realize the inverse deformation.

Chapter 3

Modeling of form and its variability from images

Even as our understanding of the encoded parts of biology has exploded since Thompson's time, our appraisals of form lagged far behind, awaiting the technologies of medical imaging and interactive computer graphics. Once those technologies arrived, the driver of studies of biological form became not biophysics but the cognitive sciences: the same evolved perceptions of form and form-comparison now filtered through the formalisms of quantitative pattern engines and the representations of uncertainty that, for lack of a better word, we call statistics.

— Fred L. Bookstein 2011

The goal of the visual analytics methods developed in this thesis is to enable the user to gain insight into the variability of form contained in a biomedical image ensemble. Pre-requisite to this goal is mathematical modeling of form and its variability from images. This chapter reviews the used modeling techniques and algorithms for analysis and synthesis of variation in form and settles basic notation and definitions. The statistical deformation model employed throughout this thesis is introduced together with a custom algorithm for group wise image registration that provides an average template against which variation is measured.

3.1 Introduction

3.1.1 What is form? What is shape?

In order to model variation in form, one has to specify first what is meant by the term “form”. In an anatomic context it is intuitive to define form as the spatial layout of an organ, or introducing yet another term, its *shape*. This conception is in agreement with Websters college dictionary [48]:

***form** (fôrm), n., v., **formed, form-ing.** –n. **1.** external appearance of a clearly defined area, as distinguished from color or material [...]. **2.** the shape of a thing or person. [...] **5.** something that gives or determines shape; a mold.*

Commonly the term “shape” is related to the outer surface of an object, as perceived by the human eye. In addition to this outer surface view, in medical imaging also internal structures are made visible that constitute the organs spatial layout. In our holistic approach this additional information is taken explicitly into account and accordingly, the terms form and shape are meant here to encompass also interior structure. Both terms, form and shape, will be used interchangeably in the following. However, one should be aware of potentially different connotations in the cited literature. For instance, in geometric morphometrics the term shape is used to designate specifically the geometry of organs after factoring out differences due to location, orientation and size.

3.1.2 Transformations to define shape and its variation

A mathematical definition of shape can be given in terms of its invariance to specific linear transformations; or as David George Kendall [77] succinctly puts it:

The idea is to filter out effects resulting from translations, changes of scale and rotations and declare that shape is ‘what is left’.

This does not only apply to surface and point data that is usually associated with the term shape, but also to the volumetric structure of an anatomy as represented in biomedical images. Before the analysis of variation, at least rigid transformations, i.e. translation and rotation, are factored out because position and orientation are arbitrary, depending solely on the choice of some external coordinate frame. Sometimes the class of rigid motions is extended to similarity transformations, including isotropic scaling, or even to the fully affine case, depending on the

study at hand. After filtering out linear transformations the remaining difference in form is subsequently considered as the shape variation of interest. To characterize the difference between shapes another class of transformations comes into play, namely non-rigid ones that provide additional degrees of freedom allowing for local deformations. Although the separation of shape and its variation according to different classes of transformations may be ambiguous in general, see for instance the nice discussion in Yezzi and Soatto [155], it is especially sound in the case of stiff bone structures such as the rodent skull data considered in this work.

As a consequence of these considerations, whenever comparing two shapes via a transformation φ that maps one shape onto the other, φ is decomposed into two parts

$$\varphi = \varphi_{\text{global}} \circ \varphi_{\text{local}}. \quad (3.1)$$

where \circ denotes concatenation of mappings. The global part φ_{global} accounts for non-shape differences and will be realized by a linear transformation as discussed. When comparing an ensemble of shapes against some template shape, all global parts will be factored out first in a preprocessing step. The particular procedure to do this is referred to as *alignment*. Thereby a common coordinate frame between the anatomies of an ensemble is established, i.e. the one of the template. After the alignment procedure, the remaining local parts φ_{local} are pivotal to further analysis, as they represent the shape variation of the ensemble. In summary, one can say that φ_{global} defines what shape is, while φ_{local} encodes shape difference and variation.

3.1.3 Modeling pipeline

Equipped with a concept of shape and its variation, the modeling pipeline used in this work can be outlined. This will also provide an overview of the relevant definitions and methods introduced in the remainder of the chapter. Figure 3.1 illustrates the pipeline. The starting point (a) is the input dataset \mathcal{I} that consists of n images denoted as $\mathcal{I} = \{I_1, \dots, I_n\}$. Variation of form is modeled on this through a set of deformations $\{\varphi_1, \dots, \varphi_n\}$ with respect to a template image \bar{I} . Each deformation φ_i maps the template image \bar{I} onto the i -th individual. Deformations and template image are both estimated simultaneously in the registration stage (b). This also includes the separation of each deformation into its global and local part, i.e. the registration stage subsumes image alignment. The resulting aligned image set is denoted as $\mathcal{I}' = \{I'_1, \dots, I'_n\}$. In the subsequent analysis stage (c) these deformations are considered as samples from an underlying normal distribution and

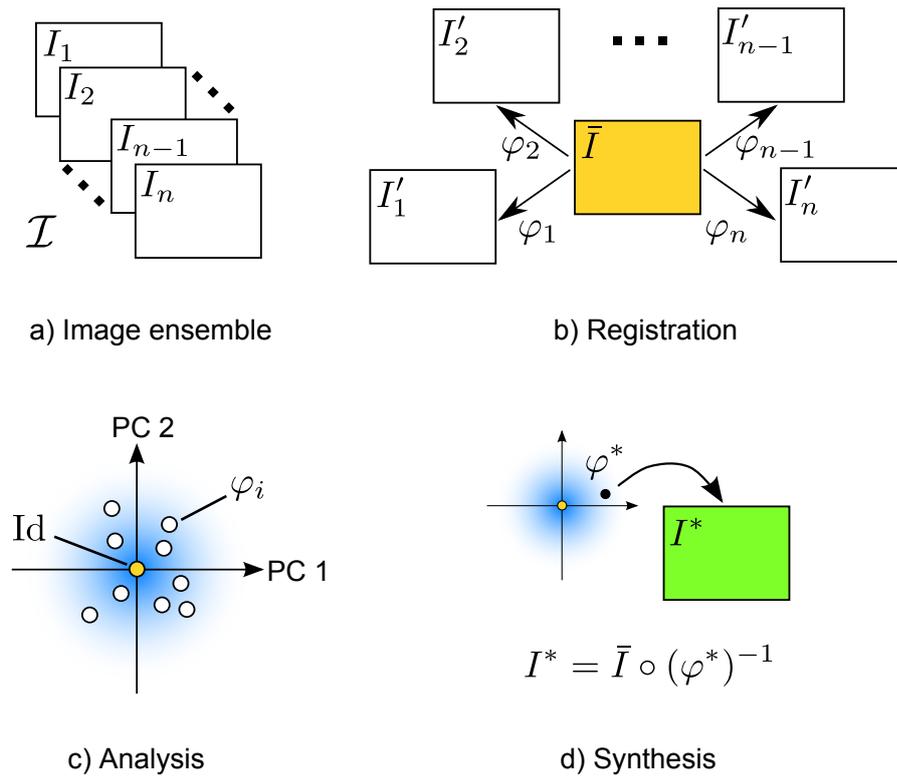


Figure 3.1: The modeling pipeline.

a statistical deformation model is set up accordingly. The model provides principal modes of variation that compactly encode the gross of variability. Later, interactive analysis will be performed in this reduced model space, where each individual is represented as a point. One can think of the distance between two points in this space to reflect the (dis)similarity in form of the corresponding anatomies, while the distance of a point to the origin is interpreted as the probability of the particular deformation represented by that point under the normal distribution model. For interactive analysis and visualization, arbitrary points in model space will be synthesized as deformations, illustrated in the last stage (d). In order to actually produce an image of the correspondingly deformed anatomy, a synthesized deformation is applied to the template via image warping.

Obviously, defining the template image \bar{I} is a crucial step in model building. To this end a custom group wise registration algorithm is introduced in Sec. 3.4. Before that, some general concepts of image registration are required that are given in Sec. 3.3. Finally, the statistical deformation model and its efficient computation is described in detail in Sec. 3.5, setting the stage for the next chapters. But first, some basic methods and notation regarding 3D images have to be settled.

3.2 Biomedical images and their deformation

Common image modalities that capture the 3D geometry of an organismic anatomy at high quality are x-ray computed tomography (CT) and magnetic resonance imaging (MRI). Both modalities are represented by 3D images with scalar values that relate to tissue type specific characteristics, namely x-ray absorption in CT and nuclear resonance behavior in MRI. While CT excels in assessment of bone structures, brain imaging is a prominent domain of MRI. In case of CT data, Hounsfield units provide a calibrated measurement scale that eases the classification of air, fat tissue, muscle tissue and bone. In the following, images and their deformations are defined independent of a particular modality. The methods of this thesis are demonstrated on ensembles of CT images but apply also to MRI data, although this requires the choice of different metrics for image registration as discussed below.

3.2.1 Images

Throughout this work a 3D image I is treated as a scalar function

$$I: \Omega \rightarrow \mathbb{R} \quad \text{with} \quad \Omega \subset \mathbb{R}^3 \quad (3.2)$$

representing some sort of *intensity* over a compact domain Ω . It is important to note that in practice the image domain Ω is discretized, commonly over a regular grid, holding intensity information only at discrete positions, i.e. the image *voxels*. We will use the symbol Ω to denote continuous as well as discretized domain, depending on context. Addressing intensity in the discrete domain at positions in between voxels requires interpolation. Although fairly easy and efficient to realize on a regular grid, special treatment is required in case of an arbitrary sampled or deformed domain as described below.

3.2.2 Transformations and deformations

In order to relate the geometric difference between the same anatomy in different images or to express a particular geometric variation we will make use of image *deformations*. Mathematically a deformation φ of an image can be understood as a mapping between two image domains $\Omega, \Omega' \subset \mathbb{R}^3$, i.e.

$$\varphi: \Omega \rightarrow \Omega'. \quad (3.3)$$

Usually the target domain Ω' is identified with $\Omega' = \Omega$ for notational simplicity. In practical implementations one has of course to deal with a different extent of the target domain when applying the deformation to an image. However, for our image ensembles we will devise a canonical parameterization domain Ω below. Particular transformation classes will be discussed below in Sec. 3.3.

A general representation of a deformation φ is to prescribe a displacement $u(x)$ at each position $x \in \Omega$ to encode the mapping

$$\varphi(x) = x + u(x) \quad \text{with} \quad u: \Omega \rightarrow \mathbb{R}^3, \quad (3.4)$$

where u is accordingly called *displacement field*. Note that u shares the same parameterization as the image domain, i.e. it is discretized at voxel positions. Algebraic operations on sets of deformations represented by (3.4) are the subject of Sec. 3.5. To this end, a discretized displacement field u will be written as column vector $\mathbf{u} \in \mathbb{R}^{3N}$, where N denotes the number of voxels. An alternative representation to (3.4) based on a differential equation is presented in Chap. 4.

3.2.3 Image warping

For visualization, and also during image registration, deformed images have to be computed. Applying a deformation φ to an image I is also referred to as *image warping* and will result in a deformed image I^* , sometimes denoted as

$$I \xrightarrow{\varphi} I^*. \quad (3.5)$$

Naively mapping each voxel position x_{voxel} of I will result in general in positions $x^* = \varphi(x_{\text{voxel}})$ that are not lying on voxels in the target image I^* . That is because I^* is discretized over the same regular grid as I . In order to avoid scattered data interpolation when computing the deformed image, an established procedure in image warping is to iterate over the voxel positions of the target image I^* instead of the voxel positions in I . Thereby, for each voxel in I^* a corresponding position in the undeformed image I is found, where standard interpolation techniques can be applied. This is realized by applying the *backward* or inverse mapping φ^{-1} , i.e.

$$I^*(x) = (I \circ \varphi^{-1})(x) = I(\varphi^{-1}(x)). \quad (3.6)$$

Obviously, (3.6) requires a bijective mapping such that an inverse exists and is well defined. Fortunately that is by definition the case for the diffeomorphisms considered throughout the thesis.

Computing the inverse of a mapping

For linear transformations analytical inverses are easy to define. Unfortunately, for non-linear transformations, computation of an inverse mapping is very expensive in general. Therefore many efficient approximations and heuristics have been developed, especially for 2D image warping [152].

A heuristic that is commonly used in interactive raycasting systems [114, 33] is to simply negate the displacement field

$$\varphi^{-1}(x) \approx x - u(x). \quad (3.7)$$

Note that the approximation error of this heuristic grows quadratically with the (maximum) displacement magnitude [114], see App. C for a proof. This behavior makes the heuristic suitable only for very smooth and overall small deformations. In Chap. 4 we will see a much better approximation that can also deal with large scale deformations.

3.3 Image registration

Image registration is the problem of spatially superimposing two or more images. For a pair of images the task at hand is to find a *reasonable* transformation that deforms one image into the other, such that corresponding structures are spatially superimposed afterwards as best as possible, rendering the deformed image *similar* to the other one. Choosing different definitions for what a reasonable transformation is and how to measure image similarity, a multitude of registration approaches for particular applications arise. Together with segmentation, image registration is one of the core techniques in medical image analysis [156]. It is applied for instance to register medical images of the same patient organ from multiple modalities, e.g. CT and MRI, or between consecutive points in time, e.g. pre- and post-operative, in order to combine the different information contained by removing difference due to patient movement between scans. The population studies targeted in the present work face yet another scenario in that images usually stem from different subjects and, even more important, the transformation is not considered a nuisance parameter that has to be removed. Instead, one is not so much interested in the superimposed image per se, but much more in the transformation itself as it encodes difference in form [15].

Formally, pairwise image registration superimposes a source (moving) image I_1 onto a target (fixed) image I_0 by applying a transformation φ to I_1 . The task

of finding an optimal transformation φ can be stated as an optimization problem with a cost or energy function $E(\varphi, I_0, I_1) \mapsto \mathbb{R}$ of the general form

$$E(\varphi, I_0, I_1) = \text{Sim}(I_0, I_1 \circ \varphi) + \alpha \text{Reg}(\varphi) \quad (3.8)$$

that is minimized with respect to φ . The energy (3.8) comprises two terms: An image similarity term $\text{Sim}(\cdot, \cdot)$ that measures the quality of registration and a regularization term $\text{Reg}(\cdot)$ that is used to control properties of the transformation and render the problem well-posed. Both terms are trade-off using a weighting parameter $\alpha \in \mathbb{R}_+$. In general, finding an optimal transformation minimizing the similarity measure alone is an ill-posed problem [59] because of many local minima arising from ambiguities in superimposition. Therefore different kind of regularization terms are employed to restrict the set of solutions.

Image similarity measure

The similarity measure is a data term that assures for the superimposed images to match locally, for instance favoring similar intensity values. A common image metric is the sum of squared intensity differences (SSD)

$$\text{Sim}_{SSD}(I, I' \circ \varphi) = \frac{1}{2} \int_{\Omega} (I(x) - I' \circ \varphi(x))^2 dx. \quad (3.9)$$

The assumption underlying this metric is that homologous points exhibit the same intensity in both images, modulo Gaussian noise. Since this is a valid assumption for the CT images considered throughout this work, Sim_{SSD} will be used exclusively herein.

For other imaging modalities or when considering mixed modalities, alternative similarity measures are advised. In order to overcome the assumption that homologous points exhibit the exact same intensity, alternative measures are modeled after different models on the relationship between intensities. For instance, the correlation coefficient assumes a linear relationship, while mutual information assumes a more general, probabilistic one. A comprehensive overview of image similarity measures can be found in the survey by Crum et al. [50].

Transformation classes

In order to distinguish non-form and form related variation as in (3.1), two different classes of transformations are considered. Affine mappings are used for initial global alignment of the image ensemble while subsequent non-linear registration is based on diffeomorphisms. The latter step is also referred to as deformable registration since it allows for local deformations to superimpose corresponding structures. For an overview on deformable registration we refer to the recent survey by Sotiras et al. [130].

In deformable registration one distinguishes two different settings, that of small and that of large deformations. In the small deformation setting the concatenation of two deformations is approximated by addition of the corresponding displacement fields. The advantage of this assumption is that it enables to perform linear multivariate analysis directly based on displacement fields, as shown in Sec. 3.5. Obviously this first order approximation is only valid for small variations in shape, like for instance the ones observed in Chap. 5 and 6. The more general large deformation setting requires more involved non-linear methods. A greater flexibility is for instance achieved by modeling deformation as temporal flow, whose properties are governed by a differential equation [43, 55]. Thereby a large scale deformation is assembled by concatenating many small scale deformations over time. An efficient parameterization of diffeomorphisms will be discussed in Chap. 4.

3.3.1 Global alignment

The transformation for global alignment is parameterized via a matrix $A \in \mathbb{R}^{3 \times 3}$ and translation vector $t \in \mathbb{R}^3$ leading to

$$\varphi_{global}(x) = Ax + t. \quad (3.10)$$

Following the above discussion on non-form related transformations in Sec. 3.1.2, the mapping A is restricted to the class of similarity transformations, i.e. rotation and isotropic scaling. This means that A can be decomposed as $A = R\Lambda$ into an orthogonal matrix R , i.e. $R^T R = I$, that represents a pure rotation that does not include a reflection, i.e. $\det(R) = 1$, and a diagonal matrix $\Lambda = \begin{pmatrix} \lambda & & \\ & \lambda & \\ & & \lambda \end{pmatrix}$ with a uniform scale factor $\lambda \in \mathbb{R}_+$.

Rotation and translation have each three DoF while uniform scale contributes a single one. The parameterization of a similarity transform thus ends up with seven DoF. This constrains the transformation space reasonably and no explicit

regularization term in (3.8) is enforced.

The optimal parameters for (3.10) are estimated by minimizing the cost function (3.8) with respect to the seven DoF here. To this end, iterative optimization procedures are employed that operate on a multiresolution scheme. The most basic method would be a gradient descent. An overview of optimization methods for global alignment can for instance be found in the paper by Klein et al. [84].

3.3.2 Deformable registration

The transformation for deformable registration is parameterized using the already introduced general representation (3.4) via a displacement vector field

$$\varphi_{local}(x) = x + u(x). \quad (3.11)$$

This kind of deformation has as many DoF as the image voxel count and thus requires a regularization to render a solution tractable. For this task a diffusion regularization is assumed throughout this work.

$$\text{Reg}_{diff}(u) = \frac{1}{2} \int_{\Omega} \|\nabla u\|^2 dx \quad (3.12)$$

In particular, the symmetric log-domain diffeomorphic demons algorithm by Vercauteren et al. [142] is used for deformable registration here. Although the described optimization algorithm does not directly minimize (3.8) with above regularizer (3.12), it can conceptually be understood within this general framework [58, 143].

Diffusion registration belongs to the class of physics-based deformation models [105, 130]. Alternatives within this class include elastic body and viscous fluid flow models [58]. Also from this class, and very related to the log-domain method, are flows of diffeomorphisms that are implemented in the framework of large displacement diffeomorphic metric mappings (LDDMM) [21]. These are prominently introduced in computational anatomy and are especially suited to study anatomical variability [103]. Because performing synthesis in LDDMM requires computationally expensive algorithms like geodesic shooting [104], these methods are out of reach for interactive applications in the foreseeable future. Very promising alternative representation based on stationary velocity fields (SVF) recently emerged [13, 14]. As shown in Chap. 4, these allow for efficient visualizations.

Another class of deformation models very popular for studies of anatomic variability are derived from interpolation theory. In contrast to the above displace-

ment representation (3.11), these approaches are parameterized over interpolating functions that provide a more compact representation amenable to efficient optimization schemes. Important examples are free form deformations (FFD) and thin plat spline (TPS) interpolation. For FFD the deformation is represented as low-degree b-splines on a coarse control grid [19, 122]. Rueckert et al. [118] introduced statistical deformation models based on FFD by applying PCA to the b-spline coefficients. TPS take a special role in geometric morphometrics where they serve two purposes. Since TPS interpolate smoothly between given control points by minimizing bending energy [146], it is a suitable way to augment the result of a landmark analysis to the space in between landmarks for visualization purposes. In addition, TPS has led to an alternative shape analysis framework where PCA is not performed on the landmark covariance matrix but on the TPS bending energy matrix, a theory introduced by Bookstein et al. [26]. Drawbacks of the parametric TPS and FFD approaches are, that they are not inherently diffeomorphic. FFD easily produces self-overlaps while TPS interpolation often yields implausible deformations away from its control points. Further, both methods provide only a limited resolution determined by the grid size in FFD and control point placement in TPS. In order to pertain the full information at image resolution we avoided these parametric approaches and chose to directly operate on dense representations like (3.11).

3.4 Estimating an average template

This section describes a method that, based on pair-wise image registration, will simultaneously register an ensemble of images and estimate a template image that serves as mean shape.

3.4.1 Template estimation in shape space

A central notion in studying shape variability is the average or mean shape against which differences and eventually variation is measured. The definition of such a mean shape can be understood in analogy to the computation of a centroid of a set of points, i.e. the ordinary Euclidean mean¹. Recall that in the context of this work, shape is represented by images and consequently a *template image* takes the role of a mean shape here.

¹Equally, one can consider as physical analogy the barycenter of a set of point sized particles bearing unit mass.

The Euclidean mean has a minimum squared error property that gives rise to a construction via optimization. Consider for this a set of n points in d dimensions in a Euclidean space, i.e. $p_i \in \mathbb{R}^d$ for $i = 1, \dots, n$. In this situation the mean \bar{p} is defined as

$$\bar{p} := \arg \min_{p^*} \frac{1}{n} \sum_{i=1}^n \|p_i - p^*\|^2. \quad (3.13)$$

The point \bar{p} is the unique point that has minimal squared distance on average to all other points.

In computational anatomy metrics for the space of diffeomorphism are employed to define distance or dissimilarity between anatomies as geodesics, i.e. shortest connections with respect to the particular metric. In order to define a statistical mean on such smooth manifolds a generalization of the above formula is invoked. For a metric space (X, d) , consisting of a smooth manifold X endowed with a metric d , the Fréchet mean of a set $\{x_1, \dots, x_n\}, x_i \in X$, if it exists and is unique, is defined as

$$\bar{x} := \arg \min_{x^*} \frac{1}{n} \sum_{i=1}^n d(x_i, x^*)^2. \quad (3.14)$$

In order to compute average templates for computational anatomy, several group-wise registration approaches have been developed [66, 17, 73] based on (3.14). Starting from a candidate image, the general concept is to iteratively update this image such that the average of squared geodesic distances between template and each individual approaches a minimum.

3.4.2 Linearized situation

The optimization algorithms involved in above described template estimation are quite complex and computationally very expensive. Fortunately, if an adequate linearization of shape space is available, the optimization can be performed in an Euclidean setting [98]. Ideally, this would lead to a closed form solution. However, since imperfections in image registration remain, still an iterative procedure is required.

In this thesis two different linearizations will be treated:

- For small deformations, the trivial linearization in form of linear displacements provides an adequate approximation. Therefore one can simply operate linearly on the displacement vector fields, without any additional mapping between a tangent space and the manifold of diffeomorphisms. This approach is pursued in the remainder of this chapter as well as Chap. 5 and 6.

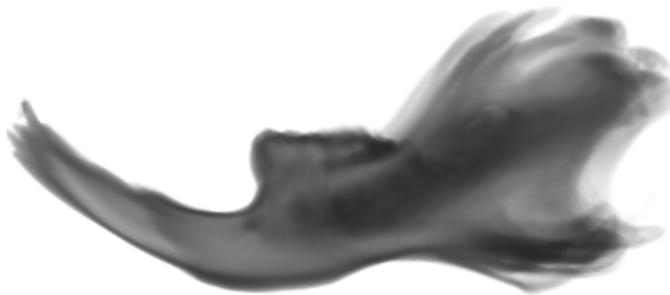


Figure 3.2: Superimposition of a rodent mandible ensemble consisting of 16 representatives, aligned via similarity transforms. The display illustrates the result of a global alignment computed by the group-wise image registration algorithm.

- For large deformations a better approximation is required. To this end we choose the efficient log-domain representation of Arsigny et al. [13]. The log-domain provides a particular tangent space that parameterizes (part of) the manifold of diffeomorphisms. In log-domain one can operate linearly while results have to be explicitly mapped to diffeomorphic deformations via an exponential map. This approach is discussed and utilized in Chap. 4.

An alternative characterization of the Euclidean centroid that will be used in our group-wise registration algorithm is introduced now. It is based on the property that the average over the vectors $(p_i - \bar{x})$ from the average to all other points $i = 1, \dots, n$ turns out as zero vector,

$$\frac{1}{n} \sum_{i=1}^n (p_i - \bar{p}) = 0. \quad (3.15)$$

One can think of $(p_i - \bar{p})$ as a set of force vectors from \bar{p} to p_i that is in balance. It is obvious that the arithmetic mean \bar{p} fulfills this property² and clearly it is the unique point doing so. Although (3.15) is a rather trivial observation in \mathbb{R}^d , it will turn into a useful stopping criterion in our algorithm for estimating an average template in the space of admissible deformations. Similarly to the centroid, the mean shape can be characterized as that shape exhibiting *on average* a null deformation to all other individuals.

²Plugging in the definition yields $\frac{1}{n} \sum_{i=1}^n (p_i - \bar{p}) = \frac{1}{n} \sum_{i=1}^n p_i - \frac{1}{n} \sum_{i=1}^n \bar{p} = \frac{1}{n} \sum_{i=1}^n p_i - \bar{p} = 0$.

3.4.3 Group-wise registration algorithm

We will now describe a common iterative bootstrapping procedure for the estimation of a template \bar{I} . The introduced algorithm is based on pair-wise image registration and described here for the small deformation setting. Starting with an initial estimate for the template, all individual images are registered pairwise against this. The estimate is then updated to better approximate a suitable ensemble average. These two steps, registration and update, are iterated until a sufficiently good approximation is found.

The general procedure is borrowed from generalized Procrustes analysis (GPA) that serves the same purpose of estimating a mean shape in landmark analysis [63]. Both algorithms work iteratively and refine an initial estimate, usually one of the individuals from the dataset, until convergence against some characterization of ensemble average. A key difference is that in the landmark setting, point wise correspondences are already known, and thus a deformable registration is not required. Alignment in GPA is implemented via Procrustes superimposition to factor out similarity transformations, while this is accomplished using global image registration here.

Algorithm 3.1 conceptualizes our implementation. In each iteration, all images are registered against the current template estimate \bar{I} in two steps. First, a global image registration is performed yielding aligned images I'_i . Second, dense correspondences encoded as deformations φ_i are established via non-linear image registration. Subsequently the average template \bar{I} is updated and an average deformation $\bar{\varphi}$ is computed. Following the reasoning of the previous section and particularly (3.15), an average template is found as soon as $\bar{\varphi}$ becomes the identity mapping Id. Due to non-perfect dense registration, the average deformation $\bar{\varphi}$ will in practice not arrive at Id and a soft convergence criterion $\|\bar{\varphi}\| < \epsilon$ with a threshold $\epsilon \ll 1$ is used. The remaining bias in $\bar{\varphi}$ is finally removed from the template and the deformations are accordingly centered to this unbiased template.

The central step of the algorithm is to average deformations (line 6), and this is the only step depending on the particular deformation model in use, at least explicitly. In the small displacement setting a set of deformations $\{\varphi_1, \dots, \varphi_n\}$ is averaged as

$$\bar{\varphi} = \text{Id} + \bar{u} \quad \text{with} \quad \bar{u} = \frac{1}{n} \sum_{i=1}^n u_i \quad \text{where} \quad \varphi_i = \text{Id} + u_i. \quad (3.16)$$

Implicitly also the breaking condition $\|\bar{\varphi}\| < \epsilon$ depends on the deformation model

and is computed here simply via the Frobenius norm on the average displacement field $\|\bar{\varphi}\| := \|\bar{u}\|$. All other steps of the algorithm are expressed in terms of warping images, i.e. applying a transformation to an image, denoted by \circ . Additionally, expensive computations of inverse mappings φ_i^{-1} are required in each iteration when computing the intensity average (line 5). In Chap. 4 the algorithm is generalized to log-domain via an average over stationary velocity fields (4.6). Additionally, the log-domain approach will simplify the computation of inverse mappings.

Algorithm 3.1 Template construction

Input: Set of images $\{I_1, \dots, I_n\}$.

Output: Average template \bar{I} and a set of deformations $\{\varphi_1, \dots, \varphi_n\}$ with $\bar{I} \circ \varphi_i^{-1} \approx I'_i$.

1: Initialize $\bar{I} \leftarrow I_1$.

2: **repeat**

3: Find global linear transformations A_i for $i = 1, \dots, n$ such that

$$\bar{I} \approx I'_i = I_i \circ A_i.$$

4: Find local non-linear deformations φ_i for $i = 1, \dots, n$ such that

$$I'_i \approx \bar{I} \circ \varphi_i.$$

5: Compute average intensity image

$$\bar{I} \leftarrow \frac{1}{n} \sum_{i=1}^n I'_i \circ \varphi_i^{-1}.$$

6: Compute average deformation

$$\bar{\varphi} \leftarrow \begin{cases} \text{use Eq. (3.16)} & \text{for displacement models in Chap. 5, 6} \\ \text{use Eq. (4.6)} & \text{for SVF model in Chap. 4.} \end{cases}$$

7: **until** $\|\bar{\varphi}\| < \epsilon$

8: Center deformations

$$\varphi_i \leftarrow \varphi_i \circ \bar{\varphi}^{-1} \text{ for } i = 1, \dots, n.$$

9: Remove remaining bias from template

$$\bar{I} \leftarrow \bar{I} \circ \bar{\varphi}^{-1}.$$

3.5 Statistical deformation model

The variability contained in a shape ensemble represented by a set of (scalar) images $\mathcal{I} = \{I_1, \dots, I_n\}$ is often described using first and second order moments, i.e. mean and covariance. The mean is implicitly provided by the template image that results from the group-wise registration presented in the previous section. The same algorithm computes the (shape) differences between template and each individual, that are encoded as deformations given by the set $\{\varphi_i\}$. From these deformations a statistical deformation model (SDM) can be set up. Please note that the model is only about deformations and as such is only concerned with spatial and *not* intensity variation.

3.5.1 Mean and covariance

In a statistical deformation model [118] (shape) difference is encoded via displacement vector fields $u: \Omega \rightarrow \mathbb{R}^3$. Analysis of the vector field data u can be reduced to multivariate statistics by treating each voxel and each dimension separately. For this purpose it is convenient to consider a vector field as a long column vector $\mathbf{u} \in \mathbb{R}^{3N}$ where N denotes the number of voxels in the discretized image domain Ω .

Using this notation, the first moment is simply the arithmetic average

$$\bar{\mathbf{u}} = \frac{1}{n} \sum_{i=1}^n \mathbf{u}_i. \quad (3.17)$$

According to the group-wise registration the displacement fields are already centered, i.e. they encode deformations in the coordinate system of the template image that represents the average anatomy. Consequently the mean deformation should be the identity transformation implying a zero displacement vector everywhere $\bar{\mathbf{u}} = \mathbf{0}$. The set of centered displacement fields constitute the data matrix

$$\mathbf{X} = [\mathbf{u}_1, \dots, \mathbf{u}_n] \in \mathbb{R}^{3N \times n} \quad (3.18)$$

for further analysis.

The second moment is estimated from the $3N \times 3N$ sample covariance matrix

$$\Sigma = \frac{1}{n-1} \sum_{i=1}^n (\mathbf{u}_i - \bar{\mathbf{u}})(\mathbf{u}_i - \bar{\mathbf{u}})^T. \quad (3.19)$$

In the case of a zero first moment (3.19) simplifies to $\Sigma = \frac{1}{n-1} \sum_{i=1}^n \mathbf{u}_i \mathbf{u}_i^T$, or

written in matrix form $\Sigma = \frac{1}{n-1} \mathbf{X}\mathbf{X}^T$.

3.5.2 Linear generative model

In order to ease interpretation of covariance, a principal component analysis (PCA) is applied. It provides an uncorrelated basis \mathbf{B} for the vector field data, i.e. in this basis the covariance matrix becomes diagonal, and each displacement field gains a new representation

$$\mathbf{u}_i = \mathbf{B}\mathbf{c}_i \quad \text{for } i = 1, \dots, n \quad (3.20)$$

in terms of a coefficient vector \mathbf{c}_i . The coefficient vectors give the coordinates of each sample \mathbf{u}_i in PCA space. How many dimensions does this space have? Since in our setting $n \ll 3N$ the rank of the covariance matrix is at most $n' = n - 1$. This means that there exist only $n' = \dim(\Sigma) < n \ll 3N$ many basis vectors and therefore $\mathbf{c}_i \in \mathbb{R}^{n'}$. For a unique representation, the n' pairwise orthogonal columns of \mathbf{B} are ordered descending according to their variances σ_i^2 , given by their length. Note that each basis vector represents itself a deformation encoded as displacement vector field and is also termed *mode of (shape) variation*.

Taking linear combinations of these modes constitutes a generative model

$$\mathbf{u} = \mathbf{B}\mathbf{c} \quad (3.21)$$

where $\mathbf{c} = (c_1, \dots, c_{n'})^T$ should be chosen with $c_i \in [-3, +3]$ conforming to a range of three standard deviations σ_i of the underlying normal distribution model, as introduced in Sec. 3.5.4 below. In the navigation process, novel deformations φ are synthesized from (3.21) via

$$\varphi = \text{Id} + u = \text{Id} + \mathbf{B}\mathbf{c} \quad (3.22)$$

where the vectorial representation (3.21) is implicitly interpreted again as displacement vector field u .

3.5.3 Principal component analysis

The linear model (3.21) is computed by diagonalizing the covariance matrix **Sigma**. By that, a change of basis is found such that in the new basis the covariance matrix has only diagonal entries. Recall that the diagonal entries σ_{ii} of a covariance matrix give the variance in direction of the i -th basis vector, while off-diagonal entries σ_{ij} with $i \neq j$ give the covariance between dimension i and j . Because covariance is proportional to linear correlation this procedure in effect decorrelates the data.

The change of basis is realized by an eigenvalue decomposition. Omitting the normalization $\frac{1}{n-1}$ for now, the decomposition has the form

$$\mathbf{X}\mathbf{X}^T = \mathbf{U}\mathbf{S}^2\mathbf{U}^T \quad (3.23)$$

where $\mathbf{U} \in \mathbb{R}^{3N \times n'}$ is a set of orthonormal eigenvectors, i.e. $\mathbf{U}^T\mathbf{U} = \mathbf{I}$, and \mathbf{S}^2 is a diagonal matrix containing the corresponding eigenvalues. The notation \mathbf{S}^2 is used to indicate that the eigenvalues correspond to the squared singular values of \mathbf{X} . Note that for symmetric matrices such a decomposition always exists, see for instance [132], and that the covariance matrix is symmetric by construction.

The diagonalization (3.23) can be computed via singular value decomposition (SVD) [137] of the data matrix

$$\mathbf{X} = \mathbf{U}\mathbf{S}\mathbf{V}^T \quad (3.24)$$

where \mathbf{U} is as before, \mathbf{S} a diagonal matrix of singular values and \mathbf{V} is an additional orthogonal matrix, i.e. $\mathbf{V}\mathbf{V}^T = \mathbf{I}$, that will be identified with the coefficients in the new basis below. Expanding (3.23) with (3.24) shows that valid eigenvectors and eigenvalues are found

$$\mathbf{X}\mathbf{X}^T = \mathbf{U}\mathbf{S}\mathbf{V}^T (\mathbf{U}\mathbf{S}\mathbf{V}^T)^T = \mathbf{U}\mathbf{S}\mathbf{V}^T\mathbf{V}\mathbf{S}\mathbf{U}^T = \mathbf{U}\mathbf{S}^2\mathbf{U}^T \quad (3.25)$$

and clarifies that the eigenvalues in \mathbf{S}^2 are the squares of the singular values in \mathbf{S} .

In order to recover the basis of the linear model (3.21) the matrix of eigenvectors \mathbf{U} has to be scaled by their eigenvalues and the previously omitted normalization $\frac{1}{n-1}$ has to be taken into account again:

$$\mathbf{B} = \frac{1}{\sqrt{n-1}}\mathbf{U}\mathbf{S} \quad (3.26)$$

Dimensionality reduction can be conducted by forming a smaller basis \mathbf{B}_k , which contains only the first $k < n'$ scaled eigenvectors corresponding to the largest

eigenvalues. Thus, \mathbf{B}_k captures most of the variance of the dataset.

The representation \mathbf{C} of the input data in terms of the new basis are that coefficients of the unique linear combinations of columns of \mathbf{B} such that $\mathbf{X} = \mathbf{BC}$. Inverting the last equations yields

$$\mathbf{C} = \mathbf{B}^{-1}\mathbf{X} = \frac{1}{\sqrt{n-1}}\mathbf{S}^+\mathbf{U}^T\mathbf{X}. \quad (3.27)$$

where \mathbf{S}^+ denotes a pseudo-inverse that only inverts the non-zero entries along the diagonal matrix \mathbf{S} . Expanding further via Eq. (3.24) shows that the coefficients are actually given by the matrix \mathbf{V}^T with a scale factor

$$\mathbf{C} = \frac{1}{\sqrt{n-1}}\mathbf{V}^T. \quad (3.28)$$

In Chap. 5 a weighted PCA variant will be presented that allows to decompose variance with respect to particular parts on the shape, i.e. regions of interest. Partial least squares (PLS) is an alternative linear analysis based on SVD that will be described in Chap. 6.

3.5.4 Normal distribution

PCA estimates a normal distribution on the displacement fields modeled as a random variable [72]. By the above construction, the corresponding coefficients \mathbf{c} follow a centered normal distribution of unit variance and the probability density takes the form

$$p(\mathbf{c}) = (2\pi)^{-n'/2} e^{-\frac{1}{2}\|\mathbf{c}\|^2}.$$

In Chap. 6 we will make use of the negative log-likelihood of the above

$$-\log p(\mathbf{c}) = \frac{1}{2}\|\mathbf{c}\|^2 + \text{const}. \quad (3.29)$$

as a measure for penalizing unlikeliness of a displacement field.

3.5.5 Efficient computation

As explained above, performing PCA is essentially a diagonalization of the sample covariance matrix. Computing a full diagonalization has a computational time complexity in $O(m^3)$ where m is the number of rows/columns of the square matrix [110]. Unfortunately the size of the covariance matrix is quite large for typical image sizes, i.e. $3N \times 3N$ in terms of number of voxels N , leading to a computationally prohibitive runtime. However, looking closely at the singular value decomposition (3.24), a much more efficient implementation is possible.

We exploit the well-known fact that the eigen decomposition of the two matrices $\mathbf{X}\mathbf{X}^T$ and $\mathbf{X}^T\mathbf{X}$ share the same eigenvalues. This can be seen by comparing the expansions of (3.24) for each matrix. The former matrix $\mathbf{X}\mathbf{X}^T$ is proportional to the covariance estimate with the expansion (3.25) given above. Similarly, the latter matrix $\mathbf{X}^T\mathbf{X}$, that is simply referred to as scatter matrix, expands to a valid eigen decomposition

$$\mathbf{X}^T\mathbf{X} = \mathbf{V}\mathbf{S}^2\mathbf{V}^T. \quad (3.30)$$

with identical eigenvalues \mathbf{S}^2 .

The key observation is that above scatter matrix is of size $n \times n$ and that is in our setting always much smaller than $3N \times 3N$ because there are much fewer images in the ensemble than voxels in each image, i.e. $n \ll N$. Performing a direct eigen decomposition of the smaller scatter matrix thereby leads to a more efficient way to compute coefficient matrix \mathbf{V}^T and eigenvalues \mathbf{S}^2 , required to describe PCA space. In order to recover the original eigenvectors of the (non-normalized) covariance matrix from this, a single matrix multiplication suffices, solving (3.24) for

$$\mathbf{U} = \mathbf{X}\mathbf{V}\mathbf{S}^{-\frac{1}{2}}. \quad (3.31)$$

In order to achieve the efficiency necessary for an interactive approach, navigation methods and automated analysis will operate on this compact representation \mathbf{V} . Where applicable, computations will consequently be mapped from high-dimensional shape space, i.e. where the eigenmodes \mathbf{U} live, to the smaller n -dimensional space of coefficients \mathbf{V} .

3.6 Summary

This chapter introduced necessary concepts and techniques to set up our visual analytics methods presented in the following part of this thesis. A specific formalization of an ensemble's shape variability is obtained by group-wise image registration. Its result is a template image and an associated set of deformations, mapping the template to each individual in the common reference frame. From this data a statistical deformation model is computed via PCA.

Concluding this chapter, the whole pipeline is shown again in Fig. 3.3. Further details on the preprocessing of the datasets are given in App. A.

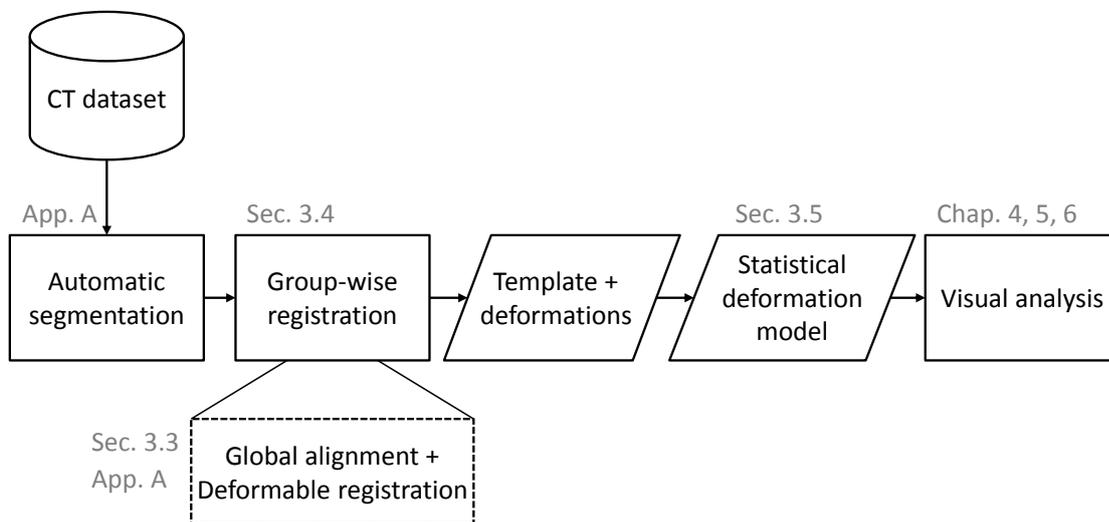


Figure 3.3: The preprocessing pipeline.

Part II

Visual shape analytics

Chapter 4

Accurate interactive visualization of large deformations and variability in biomedical image ensembles

As we have seen [...], modeling images leads to objects and these objects have shape – so we need stochastic models of shape, the ultimate non-linear sort of thing.

— David Mumford 2002

Large image deformations pose a challenging problem for the visualization and statistical analysis of 3D image ensembles. Simple linear interpolation in the tangent space of the ensemble introduces artifactual anatomical structures that hamper the application of targeted visual shape analytics techniques. In this chapter we make use of the theory of stationary velocity fields to facilitate interactive non-linear image interpolation and plausible extrapolation for high quality rendering of large deformations and devise an efficient image warping method on the GPU. This does not only improve quality of existing visualization techniques, but opens up a field of novel interactive methods for shape ensemble analysis. Taking advantage of the efficient non-linear 3D image warping, we showcase four visualizations: 1) browsing on-the-fly computed group mean shapes to learn about shape differences between specific classes, 2) interactive reformation to investigate complex morphologies in a single view, 3) likelihood volumes to gain a concise overview of variability and 4) streamline visualization to show variation in detail, specifically uncovering its component tangential to a reference surface. Evaluation on

a real world dataset shows that the presented method outperforms the state-of-the-art in terms of visual quality while retaining interactive frame rates. A case study with a domain expert was performed in which the novel analysis and visualization methods are applied on standard model structures, namely skull and mandible of different rodents, to investigate and compare influence of phylogeny, diet and geography on shape. The visualizations enable for instance to distinguish (population-)normal and pathological morphology, assist in uncovering correlation to extrinsic factors and potentially support assessment of model quality.

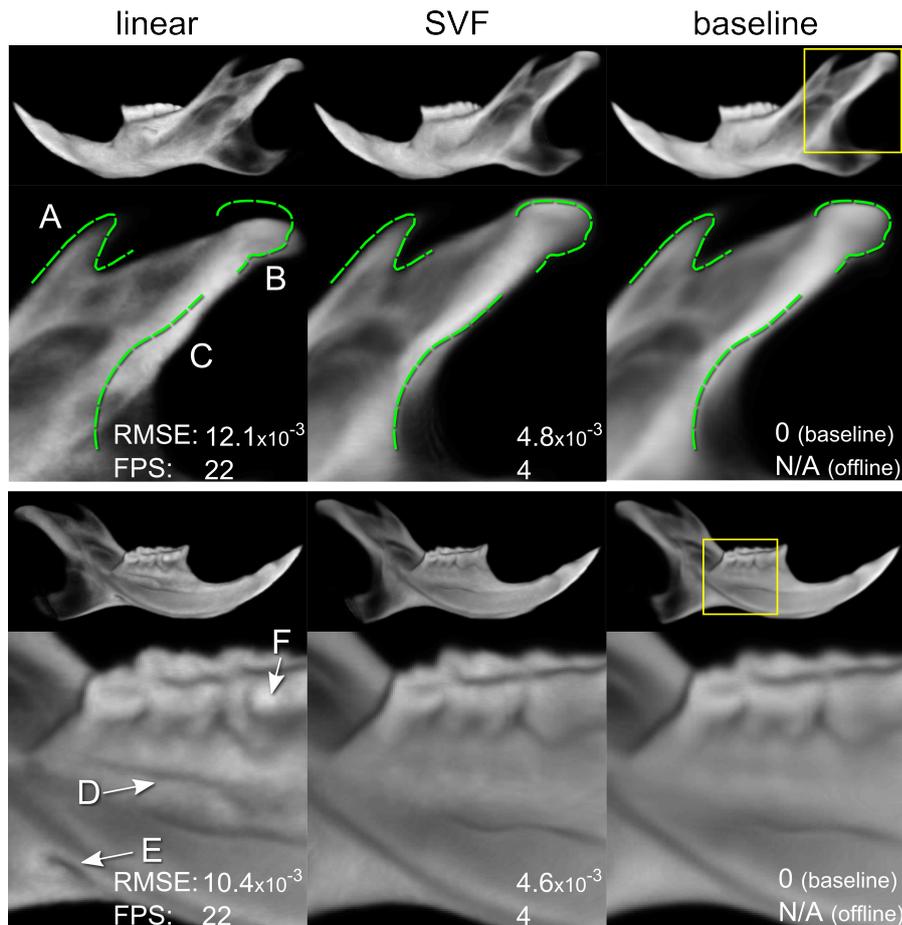


Figure 4.1: Interactive visual reconstruction of two rodent mandible anatomies using previous (linear) and proposed SVF approach. Linear reconstruction fails at A-C w.r.t. global structures while, even more severe for interpretation, at D-F artifacts that resemble anatomical features are introduced. Both kinds of error are remedied by our approach, decreasing the RMSE by about a factor of two on average.

4.1 Introduction

As described in the first part of this thesis, only recently the potential of interactive visual analysis for understanding anatomic variability by interpolating shapes or images has been recognized [35, 37, 85]. In these works, the visualization itself becomes an analytical tool enabling targeted inspection of different aspects of variability. Navigating the space spanned by an ensemble of anatomies via interpolation serves as an initial exploration vehicle before applying quantitative analysis, e.g. for feature selection or hypothesis generation.

To this end non-linear deformations are quintessential that describe the morphological differences in a biomedical image ensemble. An intuitive and widely used visualization of these differences (and the induced variability) is achieved by interpolation. Unfortunately, interpolation between 3D images at a high quality requires computationally expensive algorithms. Therefore, to be sufficiently fast, previous interactive applications resort to a heuristic [114, 33] whose error grows quadratically with deformation magnitude. A severe drawback of this heuristic is that it produces visual artifacts that unfortunately are similar to plausible anatomical structures. This impedes the interpretation and the usefulness of the visual analysis to a great extent as is confirmed by a morphometric expert we collaborated with. Fig. 4.1 shows a typical example.

The key to accurately represent large deformations as well as avoid mentioned artifacts is a non-linear deformation model. Inspired by the great success of diffeomorphic transformations in computational anatomy, in this chapter we describe a fast and accurate image warping algorithm based on the theory of stationary velocity fields (SVF) [13], tailored for interactive visualizations. The SVF based approach has shown to successfully model variability of anatomies as different as human brain [14, 30] and femur bone [124], but so far it has not been applied in an interactive application. A huge benefit is that it provides a sound tangent space, the log-domain, for the group of diffeomorphisms whose elements are lifted onto the non-linear group structure by simple integration. This does not only allow us to trivially carry over classical, linear shape statistics, but also extends the design space of visual analysis methods. In particular, the below presented techniques take advantage of the following properties, not available in any previous interactive system: 1) linear combinations of deformations are robust to evaluate and artifact free, 2) stable extrapolation is possible as well and 3) for each deformation an inverse mapping exists and its computation comes at no extra cost. The last property is pivotal for fast and accurate image warping, while the conceptual

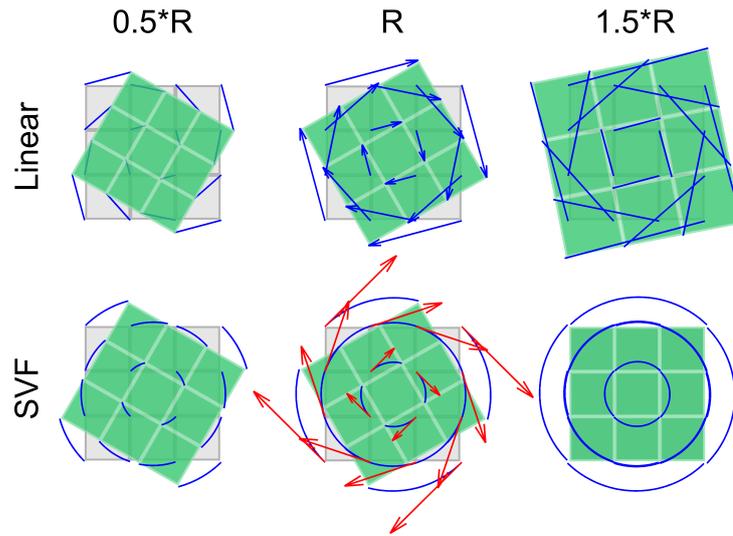


Figure 4.2: Transforming a rectangular grid by a displacement field (top) and a velocity field (bottom) representing each a clockwise rotation R around $\theta = 60^\circ$ (center). Arrows indicate vector fields and deformation trajectories are shown in blue. Evaluating scaled versions of the vector fields demonstrate inter- and extrapolation. While displacements follow linear paths, resulting in shrinkage and expansion of the transformed grid, integrating in log-domain leads to curved trajectories, perfectly reproducing expected rotations of $0.5\theta = 30^\circ$ and $1.5\theta = 90^\circ$.

importance of the first two properties is illustrated in Fig. 4.2.

A great motivation for our research is rooted in the emerging application area of population imaging [31], where large deformations are inherent. Recent advances in image registration now permit the creation of detailed models for inter-subject and population studies and we expect to see many more of these in the future. We designed our methods with the complexity of such studies in mind, leading to tools for rapid investigation of complex shape variation in dependence of various influencing factors. To that effect the high degree of image detail in our visualizations nicely complements the elaborate integrated visual analysis pipeline of existing approaches [85, 131].

4.1.1 Overview and contributions

To the best of the authors' knowledge, the presented methods comprise the first interactive system for visual analysis of a *non-linear* volumetric deformation model. This chapter is split into two parts. The first part introduces our novel image warping method and showcases four visualizations taking explicit advantage of the SVF model (Sec. 4.3). One novel and three established techniques are described, adapted and extended in several ways to perform interactive visual analysis of shape variability:

- **Group browsing** is a novel technique to render mean shapes of subgroups of the ensemble, that are interactively selected to facilitate comparison of shape difference between different classes related to extrinsic factors.
- **Volumetric reformation** [93, 144] enables investigation of variation on complex morphologies in a single view.
- **Likelihood volumes** [38, 71] display the probability of shape variation along principal modes for an overview.
- **Projected streamlines** visualize details of shape variation, paying special attention to its component tangential to a surface.

The picked visualizations are not exhaustive, but each highlights one of the new design possibilities gained through using the SVF model. Using the methods in combination allows us to perform visual analysis of shape ensembles containing large deformations. This is demonstrated on two datasets of rodent skull anatomy, investigating influence of phylogeny, diet and geography on shape (Sec. 4.5).

The second part of the chapter is devoted to an in depth evaluation of the introduced methods (Sec. 4.4). A thorough benchmark on a real world dataset confirms that in fact a higher visual quality is achieved. Furthermore, a detailed comparison substantiates that artifacts of the linear method are actually remedied.

4.2 Related work

4.2.1 Image registration with diffeomorphisms

Recall from Chap. 3 that variation between two shapes is assessed via image registration algorithms. In particular, deformable image registration is an ill-posed problem and a well defined solution is achieved by regularization, restricting the set of allowed transformations to the problem at hand [59]. When no additional a-priori information is available, as in our case, choosing a set of diffeomorphisms are considered a “good working framework” [143]. Because of their definition as smooth mappings with a smooth inverse they exhibit no singularities, i.e. no holes are created nor do folds occur and thus anatomy is kept intact. In fact, many successful image registration algorithms rely on the theory of large deformation diffeomorphic metric mapping (LDDMM) that is well established in computational anatomy [21, 103], not only for registration, but also for statistics on deformation features; see [140] and references therein.

In the LDDMM framework a diffeomorphic mapping φ is constructed as a time-varying flow $\phi(x, t)$, $x \in \Omega$, $t \in [0, 1]$ with $\varphi(x) = \phi(x, 1)$. The flow is given as a solution to the transport equation

$$\frac{d}{dt}\phi(x, t) = w(\phi(x, t), t) \quad (4.1)$$

with initial condition $\phi(x, 0) = x$, parameterized over vector fields w in the tangent space of the Riemannian manifold of diffeomorphisms. Integrating (4.1) over time yields the final mapping. This theory has turned out to be very powerful and versatile in many applications. However, computations based on time-varying flows are quite expensive and memory intensive and thus this setting is not suitable for interactive visual analysis. Fortunately, interesting alternatives using a stationary, i.e. time independent, parameterization are available [13, 14] that admit fast synthesis with low memory footprint, as is demonstrated here.

4.2.2 Stationary velocity fields

Arsigny et al. [13] propose a Log-Euclidean framework for diffeomorphisms in which Eq. (4.1) is treated using a stationary vector field $v(x) \equiv v(x, t)$ instead of a time-varying one,

$$\frac{d}{dt}\phi(x, t) = v(\phi(x, t)) \quad (4.2)$$

with $\phi(x, 0) = x$ and the final mapping $\varphi(x) = \phi(x, 1)$ achieved by integration to time $t = 1$, as before. Interestingly, (4.2) covers (almost) the entire space of diffeomorphisms and is thereby deemed as expressive as (4.1), improving performance of registration algorithms [14, 142] and also proved successful in modeling the variability of a range of different anatomies [14, 30, 124]. The term Log-Euclidean alludes to the fact that v is considered the logarithm of φ , and integration now takes the role of an exponential map of v [13]. Considering an analogy to the group of rotations $SO(3)$ should give some intuition. There, the matrix logarithm provides a so called infinitesimal generator. Using generators, interpolation of rotations reduces to simple algebraic manipulation (component wise addition and scaling) of their logarithms, retrieving the final rotation via matrix exponential. In this spirit, the exponential of the vector field v defined as

$$\varphi(x) = \exp(v)(x) := \phi(x, 1) = \phi(x, 0) + \int_0^1 v(\phi(x, t)) dt \quad (4.3)$$

will always produce a diffeomorphism. One can imagine that integrating (4.2) describes the trajectory of a particle traveling according to the velocity v prescribed at each location.

So far, SVFs have not been used for interactive visualization. Amongst other things we show how existing systems, that work with displacements (3.4), can be extended to handle SVF's by simply integrating (4.3) in the rendering stage. This can for instance be applied to some of the methods presented in chapters 5 and 6.

4.2.3 Visualization of image deformations

Several GPU raycasting methods have been proposed to render deformed images [95, 114, 47, 46], all based on the paradigm of direct space warping [19]. The general idea is to not trace straight rays through a deformed image, but rather deform the traced rays via the inverse deformation, reading the intensity information from the original, undeformed image. In that way, the rendering will display the correctly deformed image, while the image data itself stays unchanged in memory and interpolation capabilities of graphics hardware can be exploited. Based on this technique, Kurzion and Yagel [95] propose a set of analytical, invertible local mappings which can be combined to achieve volumetric deformations. A fast implementation for texture based volume rendering was introduced by Rezk-Salama et al. [114], deforming a coarser proxy geometry gained by subdivision of the volume into piecewise linear patches with subsequent interpolation. In Brunet et al. [33] a modern, single pass raycasting solution is devised to enable interactive image manipulation via freely movable control points. The costly computation of the inverse mapping is deferred in their approach to the CPU and thus does not scale to complex deformation models. More similar to our approach, Correa et al. [47] derive a deformation by linearly combining (inverse) displacement maps, well designed to enable illustrative volume rendering by introducing cuts [46]. However, the algorithm they describe can only combine two such maps on the GPU.

Caban et al. [38] presented a comprehensive review of visualization techniques for statistical deformation models and also introduced likelihood volumes, that are integrated into our interactive system in Sec. 4.3.6. Kindlmann et al. [81] applied superquadric glyphs to visualize anatomic covariance in a very concise way. This technique is ideal for an overview but does not give access to individual modes of variation and navigation techniques as presented in this and later chapters.

Another class of approaches is based on visualizing deformation properties derived from a displacement vector field. These include color-coding the determinant of the Jacobian [115, 36], revealing growth and atrophy, or displaying critical

points [136]. On top of this, our ODE approach admits further flow visualization techniques like streamline integration.

4.3 Visual analysis methods

4.3.1 A log-domain approach

The basis of our approach is, that we do not consider displacement vector fields to capture shape variation, but instead choose a SVF representation (4.3), i.e. $\varphi = \exp(v)$. A key advantage of this approach is that the log-domain, where velocity fields live, offers a natural linearization

$$\varphi_a \circ \varphi_b = \exp(v_a) \circ \exp(v_b) = \exp(v_a + v_b) \quad (4.4)$$

that is lifted to the diffeomorphic group structure by integration. As a consequence, the exact inverse mapping is found by simply integrating $-v$, i.e.

$$\varphi^{-1}(x) = \exp(-v)(x). \quad (4.5)$$

We also take advantage of the robust extrapolation possible in the SVF model. Considering an arbitrary deformation given as velocity field v , an extrapolation is achieved by evaluating $\exp(\lambda v)$ for a scale factor $\lambda > 1$, resulting in an exaggerated deformation that can be used for caricaturistic visualization [111].

4.3.2 Non-linear SVF model and its navigation

As presented in Sec. 3.4.3, we use a group wise image registration to estimate a template image \bar{I} and a set of deformations. But this time, the latter are given in terms of velocity fields $\{v_i\}$ such that $\varphi_i(x) = \exp(v_i)$ describe non-linear trajectories from the template to each individual anatomy. During registration, the template \bar{I} is estimated so that the set of velocity fields is centered in log-domain, i.e. $\bar{v} = \frac{1}{n} \sum_{i=1}^n v_i = 0$. For this purpose Algorithm 3.1 now makes use of the following average $\bar{\varphi}$ of a set of deformations $\{\varphi_1, \dots, \varphi_n\}$ that is computed as

$$\bar{\varphi} = \exp(\bar{v}) \quad \text{with} \quad \bar{v} = \frac{1}{n} \sum_{i=1}^n v_i \quad \text{where} \quad \varphi_i = \exp(v_i). \quad (4.6)$$

When computing the template image, all images of the ensemble are deformed into the template domain and their intensities are averaged

$$\bar{I} = \frac{1}{n} \sum_{i=1}^n I'_i \circ \exp(v_i). \quad (4.7)$$

Individual images I'_i can thus be reconstructed from the template as

$$I'_i \approx \bar{I} \circ \exp(-v_i). \quad (4.8)$$

Exploiting the linearity in log-domain, the same linear PCA statistics as described in Sec. 3.5 can be applied on the set $\{v_i\}$, where velocity fields are considered for this purpose as column vectors $\mathbf{v}_i \in \mathbb{R}^{3N}$ providing a data matrix $\hat{\mathbf{X}} = [\mathbf{v}_1, \dots, \mathbf{v}_n]$. The result of diagonalizing the sample covariance matrix $\hat{\mathbf{\Sigma}} = \frac{1}{n-1} \hat{\mathbf{X}} \hat{\mathbf{X}}^T$ is a generative model

$$\mathbf{v}_{\hat{c}} = \hat{\mathbf{B}} \hat{\mathbf{c}} \quad (4.9)$$

with a basis matrix $\hat{\mathbf{B}} \in \mathbb{R}^{3N \times n'}$ of principal velocities, where n' denotes the dimensionality/rank of $\hat{\mathbf{\Sigma}}$. Shape variability is explored by synthesis of novel deformations from this model by choosing a coefficient vector $\hat{\mathbf{c}}$ and integrating $\hat{\varphi} = \exp(\hat{\mathbf{B}} \hat{\mathbf{c}})$ where the column vector $\hat{\mathbf{B}} \hat{\mathbf{c}}$ is implicitly identified with its corresponding velocity field \hat{v} for notational simplicity. An image is synthesized to visualize the deformation by warping the template according to $\hat{\varphi}^{-1}$. Interpolation along a specific direction $\hat{\mathbf{c}}$ in PCA space can be shown in an animation, for instance to inspect a particular mode as illustrated in Fig. 4.12. Alternative visualizations to this include the below presented likelihood volumes and projected streamlines.

Other navigation strategies based on choosing a direction $\hat{\mathbf{c}}$ in PCA space can directly be applied to the SVF model as well, including those presented in Chap. 5 and 6. A very common way to assess principal modes of variation is to interpolate a single entry in $\hat{\mathbf{c}}$ over the range $[-3, +3]$ [96]. In Chap. 5 a technique is presented that applies a support vector machine to identify a direction $\hat{\mathbf{c}}$ in PCA space, thereby describing shape difference between two classes of the dataset, and interpolate along this direction. In Chap. 6 a technique for analysis of co-variation between structures on the same shape is introduced. There, the user prescribes a local shape deformation manually and the most likely coefficient set $\hat{\mathbf{c}}$ according to the model is computed, revealing correlated shape changes.

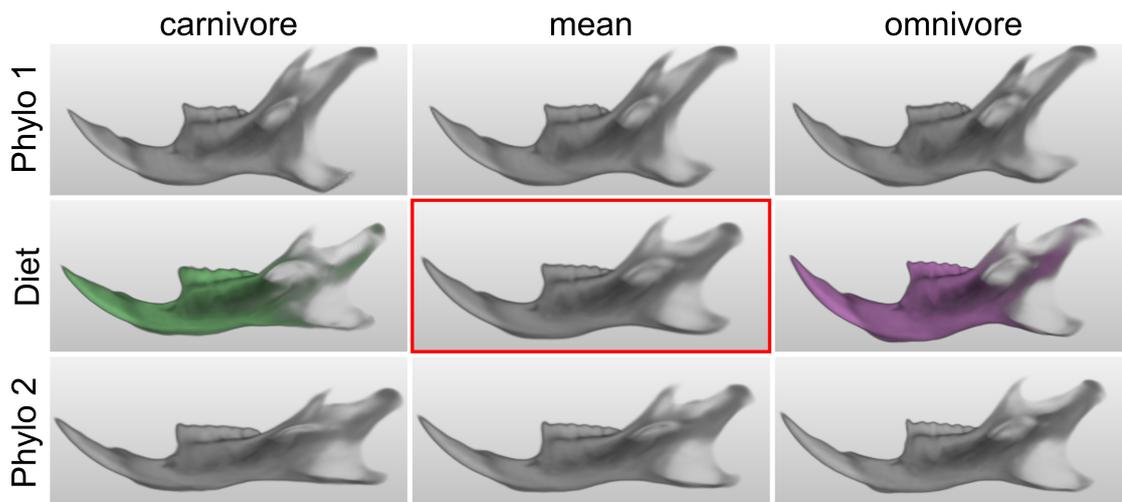


Figure 4.3: Browsing mean shapes of different phylogenetic and dietary groups allows comparative analysis of the impact of these factors onto anatomy. The template is highlighted. Note that dietary mean shapes, shown in color, are extrapolated to emphasize difference.

4.3.3 Browsing group mean shapes

A common task for exploratory analysis is to disentangle the factors that determine shape variation. Categorical factors decompose the shape ensemble into subsets, for instance into several phylogenetic or dietary groups as in our mouse mandible ensemble. In order to unveil the impact of each factor on shape variation, we derive mean shapes of the corresponding subsets, enabling interpolation in-between group means and the ensemble template. This way it is possible to investigate the difference between particular groups defined by their associated factor while considering their distinctiveness within the shape ensemble by comparison to the template. Precomputing mean shapes of all possible (in the worst case exponentially many) subsets is not practical and thus we opt for an online computation, allowing the user to flexibly define groups on the fly. An example of browsing group means, in combination with extrapolation, is given in Fig. 4.3.

We define a group mean \bar{I}_G for a particular subset $G \subseteq \{1, \dots, n\}$ of the ensemble in the same way as the template \bar{I} for all individuals. The characteristic property of a group mean is that it has on average a zero deformation to all other images in its group, like a barycenter of mass points. We thus have to find a set of velocity fields $\{v'_i\}_{i \in G}$ that is centered and maps \bar{I}_G to each group individual $I_{i \in G}$. This is achieved by subtracting the (non-zero) average

$$\bar{v}_G = \frac{1}{|G|} \sum_{i \in G} v_i \quad (4.10)$$

from the available v_i , leading to $v'_i = v_i - \bar{v}_G$. It is easily confirmed that the average of $v'_{i \in G}$ is zero. From this centered set of velocity fields we can now define the group mean similar to (4.7) as

$$\bar{I}_G = \frac{1}{|G|} \sum_{i \in G} I'_i \circ \exp(v_i - \bar{v}_G). \quad (4.11)$$

Unfortunately, a naïve evaluation of \bar{I}_G would require $|G|$ image warps, which is not feasible for interactive exploration. However, by sacrificing the correct intensity average of the group (but not its shape!) a more efficient computation can be devised. To this end we replace each individual image I'_i occurring in (4.11) by its reconstruction (4.8), leading to the following approximation:

$$\bar{I}_G \approx \frac{1}{|G|} \sum_{i \in G} \bar{I} \circ \exp(-v_i) \circ \exp(v_i - \bar{v}_G) \quad (4.12)$$

$$= \frac{1}{|G|} \sum_{i \in G} \bar{I} \circ \exp(-\bar{v}_G) \quad (4.13)$$

$$= \bar{I} \circ \exp(-\bar{v}_G) \quad (4.14)$$

With (4.14) we arrive at an estimate for the group mean that requires only a single image warp based on the velocity field \bar{v}_G . Note that here the expensive multiple image warps of (4.11) are exchanged with a simple arithmetic average over velocity fields included in the group G , providing a deformation that is subsequently removed from the ensemble template \bar{I} via a single integration of the sum given in (4.10).

4.3.4 Volumetric reformation

Instead of showing anatomical structures in their original spatial configuration, in a reformation the shape is purposefully deformed to aid interpretation. A classical example is curved planar reformation [74] where tubular structures, like blood vessels and colon, are “unrolled” by resampling the input image onto curved planes following the centerline of the particular structure. The thus reformed image provides a more concise visualization and minimizes camera interaction. Our approach does not use planar but volumetric reformation targeting the same goal. Inspired by the work of Kretschmer et al. [93] our reformation approach is based on local, volumetric deformations of structures of interest. But instead of using an as-rigid-as-possible deformation model like they do, we derive velocity fields for local transformations, that are blended. In spirit, this is very similar to the orig-

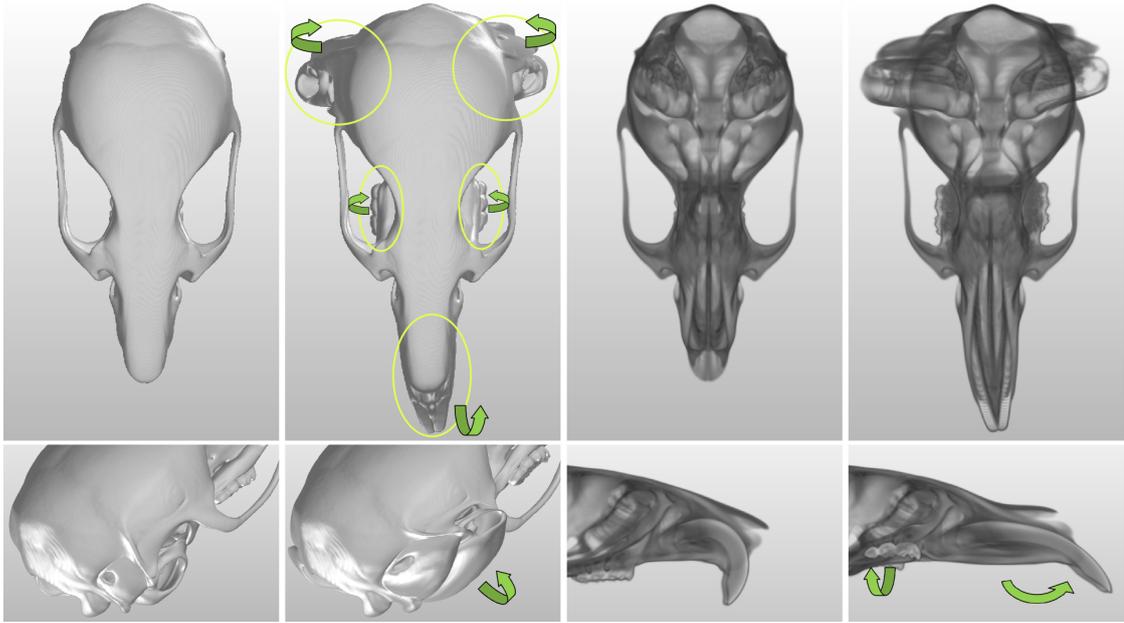


Figure 4.4: Volumetric reformation of five regions on a mouse skull to assess otherwise occluded structures in a single view. Efficient image warping enables realtime manipulation and simultaneous browsing of PCA space (see Fig. 4.12). **Top:** Isosurface and DVR rendering from top view without and with reformation. **Bottom:** Alternate views.

inal approach by von Funck et al. [144] for mesh editing based on integration of divergence-free vector fields. The application we demonstrate is to reform a complex anatomy such that otherwise occluded structures become visible and can be investigated in a single view. This particularly reduces cognitive load in dynamic displays of shape variability, where one can stay focused on the information encoded in the animation without having to adjust camera or synchronize (mentally) between multiple views. See Figs. 4.4 and 4.12 for an example.

Since the reformation should be applicable in combination with our other visualizations, we will represent it via a SVF that can be added to any further deformation. To this end an arbitrary diffeomorphic transformation can be deduced via the scaling and squaring method [13] or, more efficiently, we could make use of a polyaffine deformation [12]. In order to preserve spatial context, we restrict ourself in this showcase to stiff transformations and choose to setup our reformation from local rotations in a region of interest (ROI). Consider that a rotation of angle $\theta < \pi$ around an arbitrary axis is given as a 3×3 matrix R together with its center of rotation $x_{pivot} \in \mathbb{R}^3$. In order to incorporate this transformation into the integration (4.2) we take the matrix logarithm $L = \log(R)$, yielding an anti-symmetric matrix [70], commensurate with the log-domain representation. With

that, the velocity field of the local rotation is defined as

$$v_{rot}(x) = w(x)L(x - x_{pivot}) \quad (4.15)$$

where $w: \Omega \rightarrow [0, 1]$ is a weighting function that smoothly interpolates between $w(x) = 1$ inside the ROI, and $w(x) = 0$ outside of it. In the resulting deformation the ROI is rotated, while w ascertains a continuous transition between the rotated ROI and the non-rotated remainder of the volume. Note that eventually the rotation v_{rot} is *added* to a deformation v for combined visualization.

Simple and efficient choices of ROI weights w include spheres and bounding boxes. For these geometries the weights can be trivially parameterized over 1D functions, either based on the distance to the rotation center or to the planes of the box. Let t_{in} and t_{out} denote the distance range for the transition between inside and outside of the ROI, for instance $0 < t_{in} < t_{out} \leq r$ for a spherical ROI with radius r . For smooth interpolation of weights in the transition area $t_{in} \leq t \leq t_{out}$ the following standard cubic function is used

$$w_{1D}(t) = 1 - \left(\frac{t - t_{in}}{t_{out} - t_{in}} \right)^2 \left(3 - 2 \left(\frac{t - t_{in}}{t_{out} - t_{in}} \right) \right) \quad (4.16)$$

while weights interior $t < t_{in}$ are set to 1 and exterior $t > t_{out}$ to 0. The choice of t_{in} and t_{out} controls the extent of the transition area. Its particular choice is a trade off between global and local spatial context. Choosing a large area increases smoothness of the transition, providing global context, but disturbs local structures, since angle and volume preservation are abandoned in the transition area. In our experience, the choice depends on the complexity of the structure. For a structure exhibiting local variations, like for instance the ear bulla of a mouse skull, a sharper transition is advised than for reforming structures with large scale variation, like the incisor of the mouse skull.

A complex volumetric reformation is realized by adding several user specified rotations and their ROIs as in Eq. (4.15). For m such rotations $\{v_{rot}^{(1)}, \dots, v_{rot}^{(m)}\}$ the final reformation velocity field reads

$$v_{reform} = v_{rot}^{(1)} + \dots + v_{rot}^{(m)}. \quad (4.17)$$

After choosing axis and center for a single rotation, the rotation angle can be tuned by just scaling the logarithm $L \rightarrow \lambda L$ where $\lambda \in \mathbb{R}$ controls the interpolation from no rotation ($\lambda = 0$) to the specified rotation ($\lambda = 1$) and can even be extrapolated ($\lambda > 1$) or inverted ($\lambda < 0$) robustly. Together with the presented raycasting

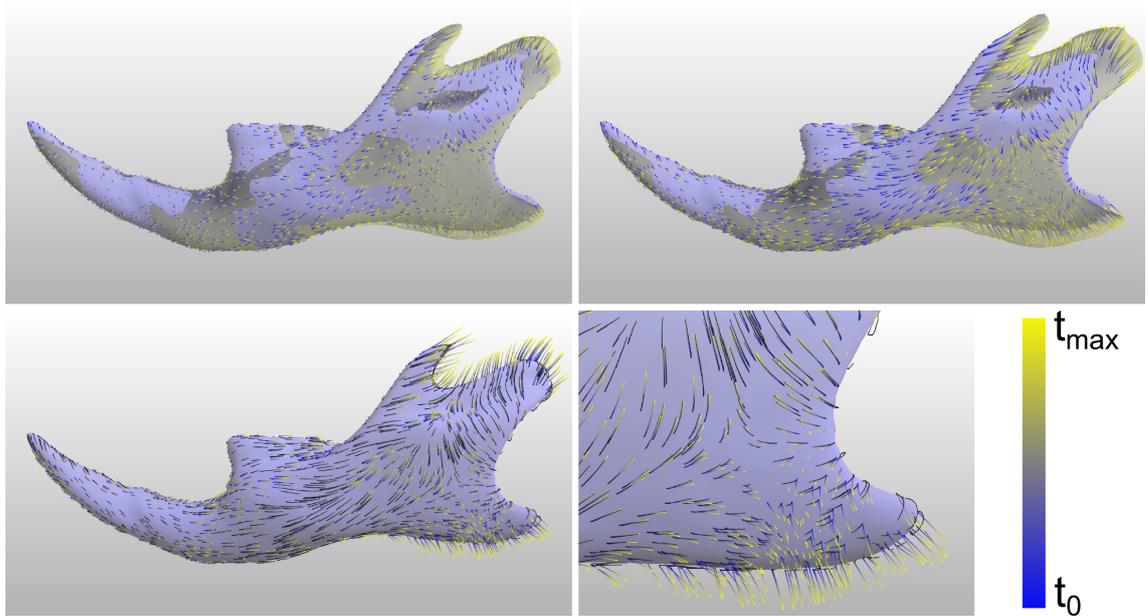


Figure 4.5: Streamline visualization of PC2. Shown are the mean shape in blue, a displaced version in yellow and streamlines color-coded according to integration time $[t_0, t_{max}]$. **Top:** $t_{max} = 1.5$ and $t_{max} = 3.0$ respectively. **Bottom:** Projection onto mean shape (grayscale) provides additional information about the tangential component of shape variation.

procedure, this facilitates a user interface for interactive volume editing, where every aspect of the reformation is adjustable in realtime. Interestingly, being able to interpolate from normal to reformed anatomy turned out to be quite useful in establishing (“learning”) the mental mapping between the two, which is essential to interpret this kind of visualization.

4.3.5 Projected streamlines

A natural new opportunity for analysis of variation arises from the chosen representation (4.2) of deformations as ODE, for which a rich set of vector field and flow visualizations are available. A particular method, nicely displaying the now curved trajectories, is streamline integration, that can readily be applied to our data and has been already used for static illustration of SVFs [124]. In order to convey specifically the surface parallel component of shape variation, we describe an extension of streamline integration that computes a projection onto an isosurface of the template. Displaying tangential variation in this way provides an advantage over direct visualization of a deforming shape that is affected by the aperture problem, i.e. that motion can only be perceived in direction perpendicular to contours. An example of our final visualization is shown in Fig. 4.5.

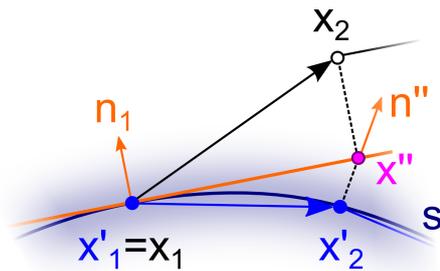


Figure 4.6: Streamline projection onto an isosurface (see text for details).

Streamlines are generated by evaluating the integration of a velocity field at a monotone sequence of time points $t_1 < \dots < t_k$ resulting in a trajectory of positions x_1, \dots, x_k , that are connected with lines for final display. The projection of the streamline onto an isosurface S is illustrated in Fig. 4.6. We assume that the seed point x_1 already lies on the surface. For each subsequent point x_i , $i > 1$ on the original trajectory a projected point x'_i lying on the surface S has to be found. We found a projection via gradient descent to be unstable on our CT images, because intensity variation outside the anatomy is not smooth and contains singularities. Therefore we propose a more robust two step procedure. In the first step, x_i is projected onto the tangent plane of its (projected) predecessor, defined by the normal at x'_{i-1} . This yields an intermediate point x'' , already closer to S , but still an unsatisfying approximation in areas with local curvature. Therefore, in the second step, a normal direction n'' is estimated and a Newton iteration on the intensity function along $-n''$ is executed to find x'_i as the closest point on S . Normal estimation $n'' = \nabla I(x'') / \|\nabla I(x'')\|$ is stabilized in the vicinity of S by operating on a prefiltered (smoothed) image. Let S be defined by the isovalue γ , then a Newton iteration on $f(\alpha) = I(x'' + \alpha n_i)$ can be formalized as

$$\alpha_{n+1} = \alpha_n - \frac{f(\alpha_n) - \gamma}{f'(\alpha_n)}. \quad (4.18)$$

It will eventually end up in a zero-crossing converging after say l steps, defining the projected point $x'_i = x''_i + \alpha_l n_i$. In our implementation we choose a fixed number of $l = 10$ steps to avoid conditional branching in the shader program.

In the example in Fig. 4.5 seed points are chosen uniformly sampled over a mesh, corresponding to an outer isosurface of the template. Projection is done onto a surface with slightly lower isovalue, resulting in a small offset of the trajectories from the rendered mesh surface, enhancing visibility. Streamlines are color coded according to integration time, while the projection is shown in a desaturated gradient, evoking the impression of a shadow. Note that the projection remains

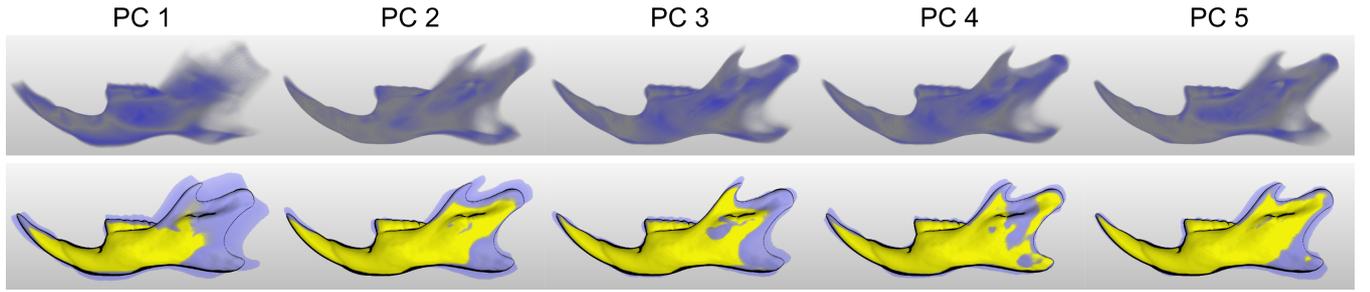


Figure 4.7: Likelihood visualization. The motion of a particular variation mode over $[-3\sigma, +3\sigma]$ is accumulated in a single likelihood volume, color coded from blue to yellow. **Top:** Direct volume rendering of likelihood function. **Bottom:** Alternative visualization that provides context via contour lines and isosurface. The transfer function is adjusted to bi-partition areas of low (yellow) and high (blue) variability more clearly.

visible also in areas where streamlines penetrate the surface and get occluded. This is a desirable feature when investigating the tangential component of shape variation.

4.3.6 Likelihood volumes

Likelihood volumes were previously used by Jiao et al. [71] to visualize uncertainty of orientation distribution functions in HARDI. Caban et al. [38] introduced a similar technique to display the probabilistic properties of a shape ensemble in a single, static visualization. To this end they sample images according to the probability density of the PCA model and integrate them in a single volume. More precisely, their likelihood volume counts for each voxel how many times it was covered by the deformed anatomy, classified by thresholding the sample images. The result of a direct volume rendering of this is reminiscent of a motion blurred picture, where areas of higher shape variability appear smoothed out. While the original likelihood volume is precomputed once offline, our SVF model allows an online evaluation. This enables to interactively derive displays of variability in user chosen subspaces of the PCA model. Fig. 4.7 displays for instance the principal modes of variation for an overview.

In order to derive a likelihood volume, we synthesize a discrete sample of k velocity fields from our model along a particular direction $\hat{\mathbf{c}}_{dir}$, $\|\hat{\mathbf{c}}_{dir}\| = 1$ in PCA space and average the correspondingly warped images. For a uniform sampling $\gamma_1, \dots, \gamma_k$ of the domain $[-3, 3]$, covering three times standard deviation, this yields a likelihood image

$$I_{likelihood} = \frac{1}{k} \sum_{i=1}^k \bar{I} \circ \exp(\gamma_i \hat{\mathbf{B}} \hat{\mathbf{c}}_{dir}). \quad (4.19)$$

Considering that the integration $\exp()$ dominates GPU runtime, k should be small, especially for non sparse coefficient sets where computing $\hat{\mathbf{B}}\hat{\mathbf{c}}_{dir}$ requires additional texture lookups. In our experiments already a minimal sampling using $k = 6$ produced informative summaries at interactive framerates, without requiring any preprocessing.

4.3.7 Implementation and workflow

Each described visual analysis method prescribes a different velocity field v , based on user input. But instead of being restricted to deploy one method at a time, the log-domain representation enables seamless integration of several methods at once by simply adding their corresponding velocities. In this way, algorithm 4.1 comprises either PCA navigation or group browsing, both combined with a reformation that is specified in advance; likelihood volumes are treated separately via (4.19). Navigation in shape space is facilitated by choosing a direction $\hat{\mathbf{c}}$ using one of the methods from Sec. 4.3.2 or interactive selection of a group G . The corresponding animation is controlled via a scalar parameter λ , mapped for instance to a GUI slider. On camera or parameter change, a deformation is computed via numerical integration of velocities v_{fused} , as detailed next. The impact of different integration schemes on quality and performance are compared in Sec. 4.4.3.

Algorithm 4.1 Linear fusion of velocity fields

Input: Position $x \in \Omega$.

Output: Velocity $v_{fused} \in \mathbb{R}^3$.

1: Evaluate reformation v_{reform} at x via (4.17).

2: Evaluate v_{model} at x by interpolating...

...using (4.9) along direction $\hat{\mathbf{c}}$ in PCA space with $\lambda \in [-3, 3]$:

$$v_{model} = \lambda \hat{\mathbf{B}}\hat{\mathbf{c}}$$

...using (4.10) between \bar{I} and mean of group G with $\lambda \in [0, 1]$:

$$v_{model} = \lambda \bar{v}_G$$

3: Return $v_{fused} = v_{reform} + v_{model}$.

Direct volume rendering

For direct volume rendering *backward* integration is performed during raycasting at each sample position x along a ray, resulting in a displaced position $x' = \varphi^{-1}(x) = \exp(-v_{fused})(x)$. The remaining implementation is that of a standard raycaster for direct volume rendering, with the only difference that instead of reading intensity information from $\bar{I}(x)$ it is now read from $\bar{I}(x')$. We use a single pass fragment shader implementation, similar to Brunet et al. [33].

Vertex displacement

An alternative to the described direct space warping in volume rendering is to apply the *forward* deformation $\varphi(x) = \exp(v_{fused})(x)$ on 3D point samples x , either representing vertices of a precomputed isosurface mesh or seed points for streamline integration. This can be efficiently realized by implementing algorithm 4.1 and numerical integration in a geometry shader, either displacing mesh vertices resulting in a deformed surface, or outputting a line strip connecting intermediate positions of the integration, resulting in a streamline.

4.4 Evaluation

In order to assess visual quality of SVF and linear approach we consider PCA reconstruction results of the mouse mandible dataset, because these represent real shapes with known anatomical structure, aiding the discussion. Since direct space warping is employed, final renderings have to be compared. We benchmark on four different views, that are typical choices in exploration and expose variation in several directions. A baseline is created for each specimen using a high quality image warping algorithm in advance, that is subsequently rendered with exactly the same raycaster and transfer function as the online reconstruction. The error of a particular rendering is measured as intensity difference to the baseline, accumulated over all views. The root mean square error (RMSE) is calculated over renderings of size 512×512 , yielding a single number measuring the overall visual quality, with higher values indicating stronger deviations from the baseline.

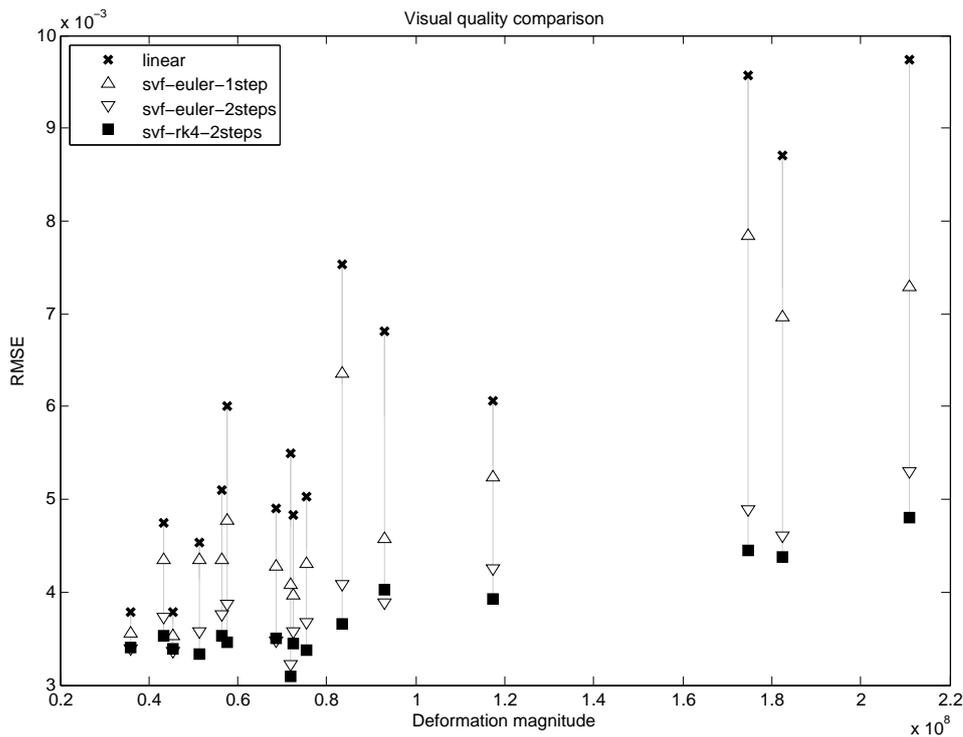


Figure 4.8: Evaluation of visual quality between linear and SVF model using different integration schemes. (See text for details.)

4.4.1 Visual quality

A summary of our benchmarks is given in Fig. 4.8, comparing different SVF integration methods to solve (4.3) against the heuristic inverse (3.7) of the linear model. The selected integration methods all provide interactive performance; see Sec. 4.4.3. Apparently the Runge-Kutta 4 scheme with two steps (RK4-2) produces the best results while the linear method performs worst. Choosing more steps and/or a higher order scheme than RK4-2 does not lead to a noticeable increase in quality, probably because of the used quantization; see Sec. 4.6. Overall one can observe that increase in error in dependence of the deformation magnitude has a substantially lower slope for RK4-2 than for the linear method, leading to the assumption that also datasets with even larger deformations than the chosen one can be treated with the proposed method. Notably a single Euler step outperforms the linear method, even though both have identical computational complexity.

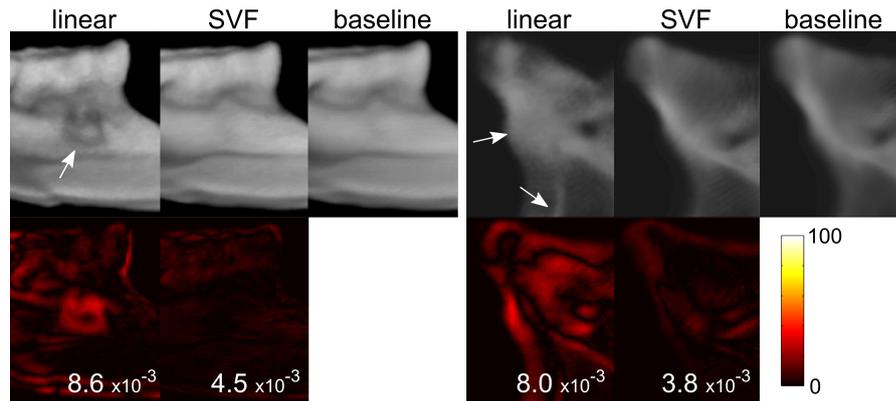


Figure 4.9: Additional examples of apparent anatomical artifacts produced by the linear method that are successfully resolved in the proposed SVF approach. Heat maps tell the absolute 8-bit intensity difference to the baseline, inset numbers denote RMSE.

Discussion of artifacts

Detailed examples of the benchmark are shown for *Mer-Ger* and *Pachyuromys* in Fig. 4.1 and *Oxy-Sca* and *Deo-Ura* in Fig. 4.9. On a first glance the figures illustrate that both, linear and SVF method, produce renderings matching the baseline at coarse scale, although the RMSE shows about a factor of 2 increased quality of the SVF over the linear method. Closer examination in the areas of higher local error in the linear reconstruction reveals apparent anatomical differences in the depiction. We can distinguish between two typical kinds of errors in the linear approach. In Fig. 4.1 at locations B and C approximation errors are visible prominently on contours of shape structures, that do not match the ones of the baseline. This can partly be attributed to the errors accumulated in the linear combination of the displacement fields in the PCA reconstruction, similar to the motivating example in Fig. 4.2, and partly to the approximation error of the heuristic inverse. Another, even more severe class of error occurs in areas of locally varying deformation, as shown in the closeups at locations D-F in Fig. 4.1 and in additional examples in Fig. 4.9. In such less smooth areas of the deformation the heuristic inverse produces severe artifacts that, although locally confined, unfortunately resemble plausible (but non-existent) anatomical structure, as confirmed by the morphometric expert of our team. One could say that the linear method “hallucinates” virtual anatomical structures that render any analysis based on this visualization at least uncertain. Even datasets not exhibiting strong, global deformation may still feature locally varying changes, that lead to the second class of errors when rendered using the linear approach.

4.4.2 Comparison of PCA spaces

By replacing the usually employed displacement vector fields with velocity fields in statistical analysis, the impact on the particular PCA space is of interest. In general, different PCA spaces are not compatible, but nevertheless we can observe that in our case the PC coordinates of both models are very close, as illustrated in Fig. 4.10. This of course results from the fact, that both inputs, displacement and velocity fields, encode the same deformation. Actually the displacements are gained by integrating the corresponding velocity field. Consequently it is not surprising that the principal modes of variation match each other closely. In a scatter plot matrix, the greatest difference is visible in mode pair PC4/5, that is explainable by very close eigenvalues leading to a rotation, as shown in the inset in Fig. 4.10(b). Overall, the two models differ only marginally in their PCA spaces, but much more in their visualization of synthesis results.

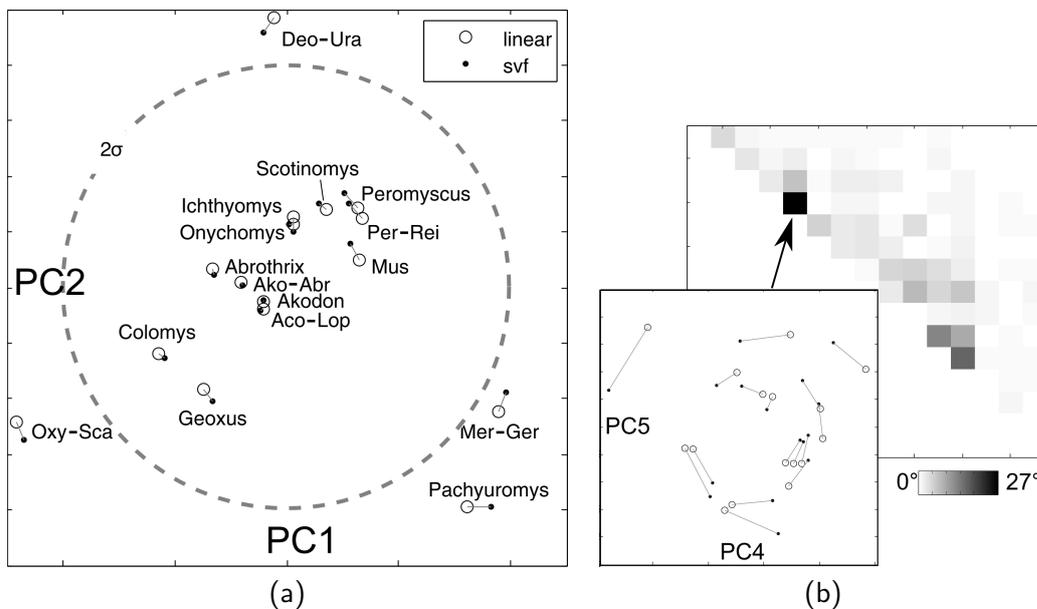


Figure 4.10: Comparison between PCA spaces for SVFs and linear displacement fields shows only a marginal difference for rodent mandible dataset. **(a)** Scatter plot of first two modes, where same shapes in both models are connected by a line and the region covering two standard deviations is marked. **(b)** Difference measured as rotation angle between PC basis pairs. The average rotation angle of 3.04° reflects high agreement. Inset scatter plot shows PC pair with greatest discrepancy.

4.4.3 Performance

Since we aim at an interactive system, performance is obviously a critical point. Fig. 4.11 shows an evaluation in this regard for navigating the PCA shape model. Because integrating a linear combination of principal modes requires many texture lookups, the navigation scenario is a computationally representative test case. The diagram shows that even in the worst case, all integration methods (taking always two steps in this evaluation) provide interactive performance at a minimum of 4 frames per second (FPS). In order to achieve higher framerates during user interaction and animation, for instance when manipulating coefficients or adjusting camera, we switch in this situations to a preview rendering at a lower resolution of 256×256 . Please note that the just described worst case behaviors only occur when considering all modes of our dataset at once. For many relevant application scenarios much fewer modes are sufficient, because most of the variability is concentrated on few modes, e.g. in the rodent mandible dataset 90% of variability is explained by the first six modes. In some situations, when dealing with moderate variations, the Midpoint scheme is an acceptable trade off between efficiency and quality.

The dominating factor on performance is the number of integrations (with according texture lookups) executed per fragment, illustrated in Fig. 4.11 on the right hand side. A potential optimization would be to avoid tracing most of the empty space by replacing the currently used very coarse bounding box by a much tighter bounding geometry, say a precomputed isosurface mesh of the template, that is warped very efficiently in the geometry shader as described in Sec. 4.3.7.

All tests were undertaken on a GeForce GTX 780 graphics card, rendering to a framebuffer texture of size 512×512 , averaging FPS over 10 seconds for each integrator / number of modes combination to reduce CPU and operating system side effects, limited at 50Hz marking realtime performance.

4.5 Visual analysis of rodent anatomy

In order to evaluate the usefulness of the presented visual analysis methods, a small case study was performed with a morphometric expert of our team. In advance, two ensembles of rodent anatomy were chosen, both with extrinsic factors with partly known and partly unknown impact on shape. One is a set of mandibles, a morphologically relatively simple structure, from diverse rodents, e.g. old and new world mice, gerbils etc., displaying a great variability based on different degrees

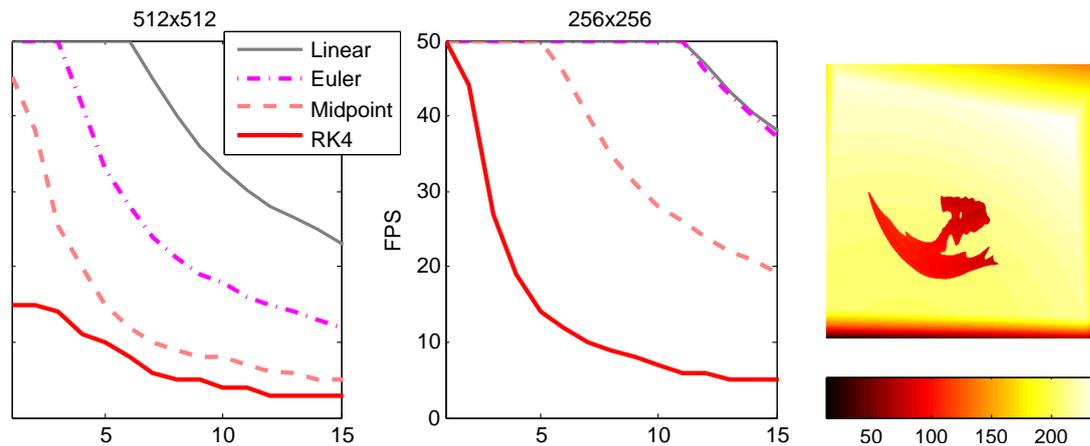


Figure 4.11: Rendering performance for quality and preview framebuffer size, plotted against number of terms in linear combination (measurement limited at 50Hz). Heat map shows number of integrations per pixel.

of relationships and diet types. The other comprises a set of skulls from a single species, with the skull having a considerably larger complexity than the mandible.

The two ensembles were analyzed in two separate sessions of each about 30 minutes, not including time for explaining the visualizations and user interface. Overall, the visualization used most of the time was to navigate single modes of variation. This is not surprising since it comes closest to the usually employed static illustrations of mode extrema. In comparison, the dynamic visualization was described to provide a “better understanding” in the sense that the particular variation becomes more obvious, especially for smaller structures. For the mandible ensemble, most of the time was spent browsing group means, which was directly accessible to the expert, who appreciated that groups could be defined interactively and used this feature intensively. A small hindrance was that group selection required to manually look up additional information on phylogeny. A linked view showing a taxonomy tree for interactive selection would provide a more intuitive browsing interface, but is outside the scope of our present work.

4.5.1 Inspection of skull ensemble

For exploration of the skull ensemble, no hypotheses were assumed a-priori, and only PCA navigation was employed, illustrated in Fig. 4.12. The first mode was quickly related to shape variation due to size (allometry), observing the typical difference in proportions between extent of brain case and length of snout, i.e. larger brain case and shorter snout for smaller specimen and vice versa. Further exploration revealed eventually a geographic signal in third and fourth mode, after

switching back-and-forth between different modes. In advance of the session, the anatomic reformation was shown to the expert who found especially the animated interpolation between normal and reformed anatomy helpful to understand how to read the latter. The reformed view was mostly used for overview, for instance to decide upon which structures to consider for further inspection. It served well to detect presence of variation at relevant structures like the two ear bullae, that otherwise would have required at least two different camera views. For final inspection an unreformed view was preferred, providing the correct morphological context required for interpretation.

4.5.2 Analysis of mandible ensemble

The aim of the mandible case study was to verify if known hypotheses about impact of diet and phylogeny could also be observed in our volumetric model, which could actually be confirmed by browsing group means as illustrated in Fig. 4.3. Differences in the respective signals are clearly visible for the two phylogenetic groups. While the Gerbillinae (Phylo 1) are known to have a unique mandible morphology (see also Fig. 4.10(a), *Mer-Ger* and *Pachyuromys*) and thus display a stronger phylogenetic signal (greater similarity along the row), the second group (Phylo 2) shows a stronger dietary signal (stronger similarity along the columns). Between carni- and omnivore group mean, one can see the characteristic difference in height of the central structure, potentially attributed to different mechanical demands in chewing different kinds of food.

At the beginning of the session likelihood volumes as in Fig. 4.7 were shown as overview. Even though not being accustomed to this kind of visualization, the expert could directly relate to it, and discovered an interesting correlation at the three rear processes in PC2, closely related to the dietary signal. While the overview does not inform on the direction of the observed correlation, this becomes apparent when displaying streamlines as in Fig. 4.5. Noteworthy, the expert recognized candidate regions for morphological modules [88], delineated by areas of different flow characteristics in tangential variation (see also Chap. 6). This is a promising insight that awaits further investigation.

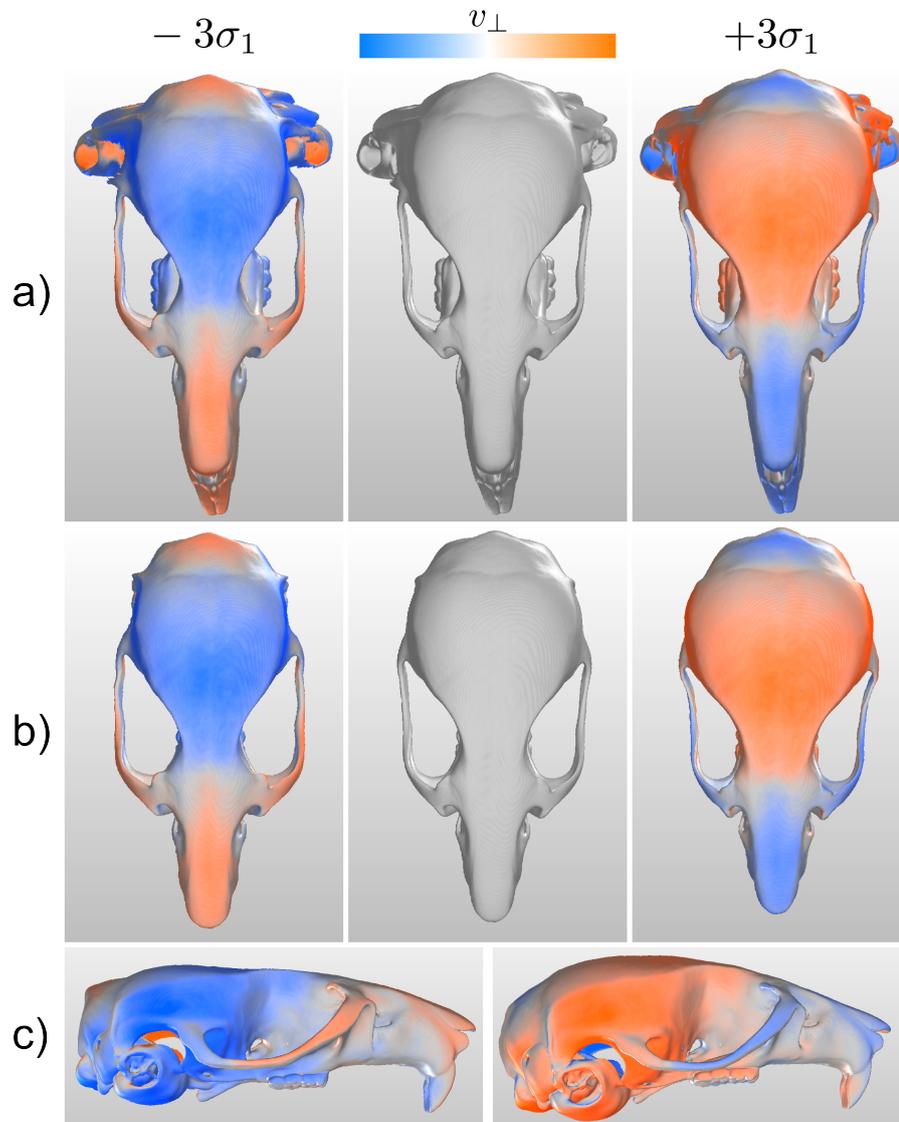


Figure 4.12: Browsing PCA shape space. Columns show the mean shape (center) deformed according to the principal mode of variation (PC1) synthesized for different standard deviations σ_1 . Color coded is deformation strength perpendicular to the surface where blue/red indicates shrinkage/expansion. Note that reformation and PCA synthesis are applied simultaneously, while only the latter is color coded. **(a)** Reformed anatomy, **(b)** original anatomy, **(c)** alternate view on original anatomy.

4.6 Discussion

Limitations

Two main limitations of the presented approach arise with respect to scalability and representation of discontinuous mappings. Sec. 4.4.3 shows that our current implementation stays interactive when combining up to 15 velocity fields during integration. The factor limiting GPU performance here are the number of texture lookups involved. To circumvent this bottleneck, additional strategies are required; for instance based on a multi-pass approach or sets of precomputed sums. GPU memory can become another limiting factor in the view of large population studies that cover several hundred individuals. However, thanks to smoothness of velocity fields and lower dynamic range compared to displacements, downsampling in resolution and quantization can be effectively applied to reduce memory consumption. In our experiments velocity fields were reduced from $200^2 \times 400$ at 32-bit (183MB/field) to $100^2 \times 200$ at 8-bit (6MB/field) without significant loss in quality. Still, further investigation of optimal interpolation schemes for downsampling and additional compression methods is appropriate. For compression it should be possible to exploit the fact that most of the ensemble variance is concentrated on a few principal modes, irregardless of sample size. A limitation inherent to the SVF approach is that it can not model discontinuous mappings nor changes in topology. For that purpose a combination of our approach with the discontinuity maps of Correa et al. [47] seems a promising endeavor.

Future directions

The presented methods provide a starting point for interactive analysis of shape variation in presence of large deformations. So far the focus was on efficient synthesis of deformed images and the shown visualizations are prototypical and can be improved in many regards. Subject to future work are custom-tailored adaptations that guide the user to deformations of interest and convey specific local aspects of variability. A related issue in this regard is to assess model quality. To this end it is planned to enhance the presented visualizations to explicitly take registration uncertainty into account. For instance, the streamline visualization could help to uncover tangential drift that indicates a suboptimal group wise registration result [52]. Another aspect not considered here is that variation occurs at different scales. In this regard it would be interesting to see our methods extended to the hierarchical polyaffine model of Seiler et. al [123]. Due to the complexity inherent

in such a hierarchy, interactive displays of variability at multiple scales provide an ideal exploration vehicle.

For group browsing an efficient computation of group averages has been devised. It is yet unclear how this procedure can be generalized to also estimate group covariances on the fly; that would allow to interactively browse complete PCA spaces of each and every group.

We also see potential applications of our interactive SVF approach beyond shape analysis. For instance, image registration algorithms could be monitored during runtime using our fast image warping method, enabling debug visualizations for parameter tuning in the spirit of RegistrationShop [128], adding diffeomorphisms to its class of transformations. Concerning 3D modeling it would be interesting to see if the easy and efficient modeling of weighted local rotations, used in our reformation technique, does provide benefits for interactive editing of 3D volumes [42, 32] and meshes [144].

4.7 Conclusion

In this chapter we presented an efficient algorithm for 3D image warping in a GPU raycaster, applied to visualize statistical deformation models, based on the theory of stationary velocity fields. Albeit simple in implementation, it improves visual quality significantly compared to previous work, that relies on a linear model of displacements and a heuristic algorithm for image warping. In our approach, image warping is naturally combined with a non-linear deformation model of diffeomorphisms, linearized in log-domain via SVFs. We showed that this combination extends the design space for visual analysis of anatomic variability and introduced a novel technique for targeted browsing of subgroups of an ensemble in relation to extrinsic factors. A comparative study on rodent mandibles and a population study on mouse skulls, performed by a morphometrics expert of our team, attested the applicability of the presented methods.

Chapter 5

Visual analysis based on prior classification knowledge and regions of interest

For a semantically steered exploration of shape space we identify in this chapter two prominent objectives for visual investigation. The first objective is automatic detection of distinguished shape variations between anatomically different, a-priori known groups of individuals. The second is to take into account and exploit expert knowledge about relevant regions on the shapes. In order to meet the first objective, we advocate the use of dimensionality reduction methods combined with a parameterization defined on user specified classifications. This idea was already successfully applied in data-driven reflectance models and also turns out to be valuable in the context of morphometry, as it allows for efficient assessment of characteristic shape differences between groups. The second objective can be achieved by an appropriate weighted linear analysis that facilitates navigation to sub-spaces in which local shape variations appear more pronounced with respect to a user defined region of interest.

5.1 Introduction

In this chapter we utilize semantic concepts to facilitate visual analytics of shape spaces built from biomedical image ensembles. We present two interaction methods to guide the exploration of shape space *semantically*. These methods extend related existing approaches as summarized schematically in Fig. 5.1. The first method extracts an interesting direction (trait vector) in shape space describing the characteristic shape change between two groups, see Fig. 5.2 for an example. In the second method relevant sub-spaces are considered corresponding to regions of interest on the shape. These methods prove especially valuable when confronted with conceptual difficulties in visual analysis of high dimensional shape spaces as they emerge especially for 3D image data. In order to enable semantic-driven interactive exploration in this high dimensional setting, efficient techniques to compute the underlying statistical deformation model as outlined in Chap. 3 are employed.

The semantically steered visual analysis is demonstrated on a datasets of rodent mandibles, investigating shape differences with respect to subfamily, genus, diet and size. First promising results in the context of morphometrics are described.

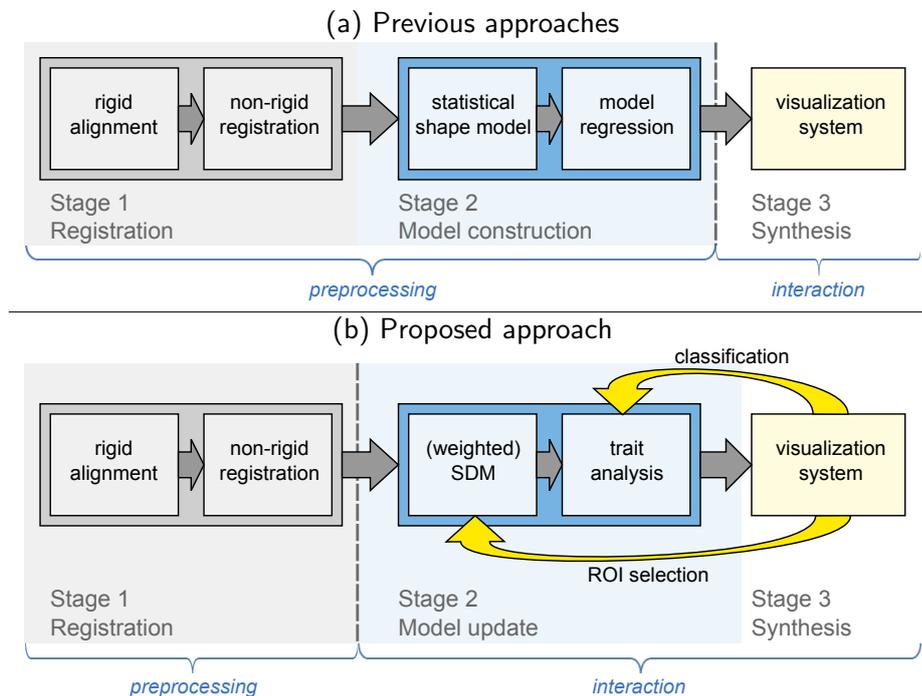


Figure 5.1: Schematic workflow. **(a)** In previous shape space exploration systems such as [25, 9, 10, 129, 37] interaction is restricted to the visualization stage. **(b)** The approach presented in this chapter features two novel methods (yellow arrows) that facilitates an interactive feedback to the model stage.

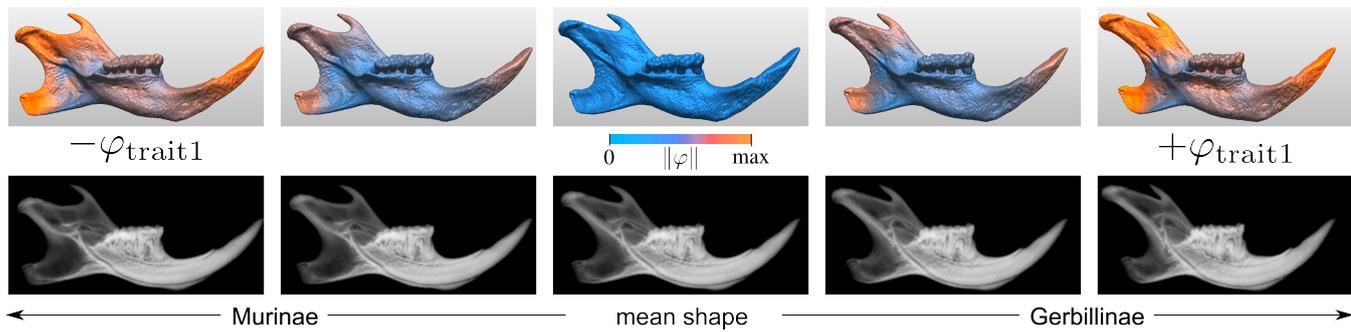


Figure 5.2: Smooth interpolation in shape space along a semantic trait vector. The visualization shows the characteristic shape difference between two subfamilies in the Gerbillinae-Murinae dataset. **Top row:** Isosurface rendering, warmer colors (orange) indicate difference to the mean shape (blue). **Bottom row:** Direct volume rendering showing internal structure changes.

5.1.1 Semantic concepts for shape space exploration

We identified two kinds of basic semantic information in morphometric studies. The following two observations are not only basic requirements of standard visual analysis techniques, but will also be of central importance in the exploration of complex shape spaces arising from 3D image data:

- S1** Knowledge about classifications of individuals into different groups, e.g. according to biological taxonomy such as the phylogenetic tree or other factors such as diet, size or ecological factors. A known classification into two classes or a scalar labeling (e.g. size in mm) will be used to compute a corresponding direction in shape space.
- S2** Prior hypotheses about the importance of specific areas on the shape for a specific study. Restricting exploration to a sub-space of shape space corresponding to an important area enables the investigation of local shape covariances and will make smaller shape details accessible.

The idea of how to include classification knowledge (S1) is inspired by the work of Matusik et al. [102], who adopt user supplied classifications for a meaningful parameterization of the space of reflectance functions. We transfer this idea to the domain of statistical shape analysis. Relevant sub-parts of shapes (S2) have been considered in the context of shape analysis to the best of the authors knowledge only for non-interactive applications so far, see Sec. 5.2.

5.1.2 PCA based exploration of high dimensional datasets

There are conceptual differences in the way the different results of principal component analysis applied to low dimensional landmark and high dimensional volume datasets need to be interpreted. The first problem is due to the higher dimensionality of the input data. For landmark data it is often sufficient to examine the principal axes of the PCA to find correlations between shape and investigated trait. However, for higher dimensional data a PCA is often ambiguous and the correlation under investigation is not captured by a single axis but spread over several axes (see Sec. 2.3.1 in [162] for an illustrative example). This renders examination of warps along the principal axes alone inadequate. A solution would be to consider arbitrary directions and not only principal axes. Since exhaustive manual search for meaningful directions is not effective (and rather prohibitive) one could use some additional information about the investigated trait to compute a corresponding direction automatically. The second problem arises from the fact that in the high dimensional setting we have much fewer input shapes than dimensions in shape space (whose dimensionality is the number of voxels in each volume). However, the dimensionality of the PCA model is limited by the number of samples, i.e. input shapes. Accordingly smaller shape variations which are not captured in this lower dimensional PCA model are lost and not accessible. A trivial observation here is that if one performs a PCA on a part of the initial volume, shape variation with respect to this part would be captured in more detail (at the cost of omitting some global characteristics). Again, manual inspection of all possible volume partitionings is prohibitive.

5.2 Related work

Related works for navigating shape spaces have been pointed out in Chap. 2.

The idea to build a statistical shape model based on a weighted analysis was already presented by Blanz et al. [24] for 3D meshes. Although the authors note that the weights may reflect an importance assigned to specific feature point, the weights are chosen proportional to measurement noise to achieve a robust surface reconstruction.

Busking et al. [37] presented a first interactive system for shape space exploration of a set of registered 3D meshes. Similar to our approach, smooth interpolation of shape changes along linear paths in shape space can be performed. In contrast to our approach, there is no possibility of including semantic knowledge

in the course of an exploration, since in [37] the shape model stays fixed.

Capturing small scale shape variations for detailed shape models built from small sample sets has initially been considered for 2D contour data in the literature. Davatzikos et al. [51] develop a hierarchical shape model using the wavelet transform to partition the contour simultaneously in space as well as in frequency. Alcantara et al. [8] propose a special localized PCA which not only optimizes the explained variation per component but also its sparsity and local connectedness. A similar idea but slightly different technical realization using sparse PCA is followed in [127]. All these approaches have in common that they automatically partition the shape to optimize the model and do not take into account semantic information. In contrast to this we rely on a user selected region of interest on which basis we adapt our shape model using weighted PCA.

5.2.1 Canonical variate analysis

A widespread method in morphometrics to separate known groups is canonical variate analysis (CVA). CVA can be considered as a rotation of the PCA principal axes [39] such that the first rotated axis maximizes the ratio of between- to within-group variation. Therefore the efficient PCA techniques presented in this work also apply for CVA. Nevertheless for reasons of robustness and improved separation quality we rely here on support vector machine classification to separate groups and derive trait vectors.

5.3 Weighted analysis of a region of interest

The goal of this section is to include information about the importance of specific regions on the shape into the computation of the SDM introduced in Sec. 3.5. The resulting weighted SDM will provide a basis of eigenmodes which explains the variability of the specific, important regions.

We assign to each voxel $i = 1, \dots, N$ an importance weight $\psi_i \in [0, 1]$. This weighting is derived from a region of interest (ROI), which is selected interactively in a visualization of the template image. A hard threshold is applied by setting

$$\psi_i = \begin{cases} 1 & \text{if voxel } i \text{ is inside ROI} \\ 0 & \text{if voxel } i \text{ is outside ROI.} \end{cases}$$

Choosing a hard threshold is equivalent to performing a local PCA restricted to the selected region. More general weighting functions can easily be defined in this setting, e.g. applying a smooth transition near the border of the ROI. Moreover, different importance weights can be assigned to separate regions of the shape, reflecting semantic prior assumptions about relative importance of these specific parts on the undertaken morphometric investigation.

In order to carry out weighted analysis, the SDM is re-evaluated using a weighted PCA. For this purpose we define a diagonal $3N \times 3N$ weighting matrix $\Psi = \text{diag}(\psi_1, \psi_1, \psi_1, \dots, \psi_N, \psi_N, \psi_N)$, where each weight is repeated three times to account for the 3D displacement vectors encoded in the columns of \mathbf{X} . Accordingly, the spectral decomposition of the smaller scatter matrix for the weighted PCA is now given by

$$\mathbf{X}^T \Psi \mathbf{X} = \mathbf{V}_\Psi \mathbf{S}_\Psi^2 \mathbf{V}_\Psi^T. \quad (5.1)$$

and is computed efficiently in the same way as described before in Sec. 3.5.5.

Note that for synthesis we still reconstruct the principal modes \mathbf{U}_Ψ using the original, unweighted deformation fields

$$\mathbf{U}_\Psi = \mathbf{X} \mathbf{V}_\Psi \mathbf{S}_\Psi^{-\frac{1}{2}}. \quad (5.2)$$

This ensures that complete individuals are synthesized and not only partial anatomy inside the ROI.

5.4 Robust classification using machine learning

To integrate classification knowledge in the exploration process, we will employ machine learning methods to automatically compute a trait vector in shape space, pointing in the direction of the greatest difference between two groups. By adding or subtracting multiples of this trait vector to the mean shape, the corresponding shape difference can be visualized. See Fig. 5.2 for an example. The trait vector is computed in the same way as Matusik et al. [102] computes trait vectors in the space of reflectance distribution functions.

We consider here binary classification of the dataset into two groups given in the form of a label vector $l \in \{-1, +1\}^n$. A multitude of algorithms exists for learning the classification, and we choose to use a linear Support Vector Machine (SVM) here (see [49] for an introduction). The advantages of using a SVM in our setting are that it produces reliable results also for small sample sets, exhibits a maximum margin criterion which stabilizes the trait vector direction in our experiments, and can also separate groups which are not strictly linearly separable by using a soft margin criterion.

For efficiency, the SVM classifier is trained on the rows of the smaller matrix \mathbf{V} . Recall that the j -th row contains the coefficients corresponding to the j -th displacement field in reconstruction of the eigenmodes (3.31). Note that we might also use the reduced matrix \mathbf{V}_k for training, which contains only the first k eigenvectors. This can be more robust to misclassification due to uninformative or noisy parts of the eigenmode spectrum. In our application such a problematic behaviour wasn't observed, but this could be of relevance for other datasets.

As a result the linear SVM classification returns a separating hyperplane in the n -dimensional \mathbf{V} -space, given by its normal $w \in \mathbb{R}^n$ and its distance to the origin $b \in \mathbb{R}$. The classification function for a point $x \in \mathbb{R}^n$ in the row space of \mathbf{V} is given by the sign of its distance to this hyperplane, $\text{sgn}(x \cdot w + b)$. In order to recover a deformation mode $\varphi_w = \text{Id} + \mathbf{u}_w$ corresponding to the learnt direction w , we have to project w into the column space of \mathbf{V} and reconstruct a displacement vector field \mathbf{u}_w from this, i.e.

$$v_w = \mathbf{V}w \quad \text{and} \quad \mathbf{u}_w = \mathbf{X}v_w. \quad (5.3)$$

Based on the trait vector, we can also compute a new linear model where the modes are chosen orthogonal to the trait vector while further fulfilling the PCA criterion to maximize the explained deformation variance per mode. This new

model can be derived by projecting \mathbf{X} into the subspace orthogonal to \mathbf{u}_w by

$$\mathbf{X}_\perp = \mathbf{X} - \mathbf{u}_w \mathbf{u}_w^T \mathbf{X}$$

and performing PCA on \mathbf{X}_\perp . Note that by this procedure the influence of the shape difference between the two groups is greatly reduced in the resulting model.

5.5 Semantic exploration

Actual exploration of shape space is facilitated in a dynamic visualization. To this end, deformations are synthesized from the SDM as described in Sec. 3.5.2 and applied to the mean shape, which in turn is displayed using conventional volume rendering techniques [33]. To perform visual analysis of shape space the proposed methods are integrated in the following ways.

For assigning importance weights to specific parts of the shape, a simple interactive selection mechanism can be used to specify the ROI. Subsequently a new SDM is computed using weighted PCA, and the exploration can be continued, constrained to the corresponding sub-space.

Given a classification into two groups, a corresponding trait vector φ_w is computed. To visualize the dominant shape difference between the two groups, we synthesize images $I = \bar{I} \circ \alpha \varphi_w$, where $\alpha \in [-1, 1]$ is interactively adjusted to produce a smooth animation like the one illustrated in Fig. 5.2. Additionally, a new SDM can be derived by projecting into the sub-space orthogonal to φ_w .

In the course of an exploration session, ROI and trait vector can be used in combination. Applying the respective sub-space projections iteratively, a new way to navigate in shape space emerges. This navigation is not restricted to a fixed shape space representation, but can be used to narrow down the exploration to sub-spaces, which provide a more compact representation of the specific shape covariation under investigation.

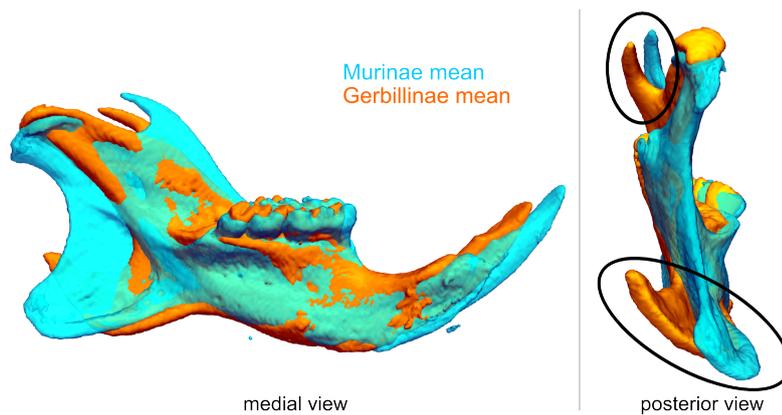


Figure 5.3: Prominent structural differences in Gerbillinae-Murinae dataset.

5.6 Visual analysis of traits in Gerbillinae and Murinae

In this section we demonstrate the proposed trait-based navigation by exploring the shape space of the Gerbillinae-Murinae dataset described in more detail in App. A. The dataset is chosen from two closely related subfamilies, namely mice and rats (Murinae) and gerbils and jirds (Gerbillinae). Both groups contain species that vary in diet as well as in size, thus enabling analyses of shape changes based on different factors. The analysed traits are in particular:

Trait 1: Subfamily Murinae vs. Gerbillinae

Trait 2: Genus *Apodemus* vs. *Mus*

Trait 3: Diet herbivore vs. carnivore

Trait 4: Size

Note that shape differences between Gerbillinae and Murinae subfamilies are quite characteristic as can be seen for instance by overlaying the respective mean shapes of each subfamily in Fig. 5.3. Particular strong differences are visible at the back of the mandible in location and shape of the upper coronoid and the lower angular process. These processes seem to tilt outwards from Murinae to Gerbillinae, as can be seen from the posterior view.

In the following the visual exploration of each of the aforementioned traits are described.

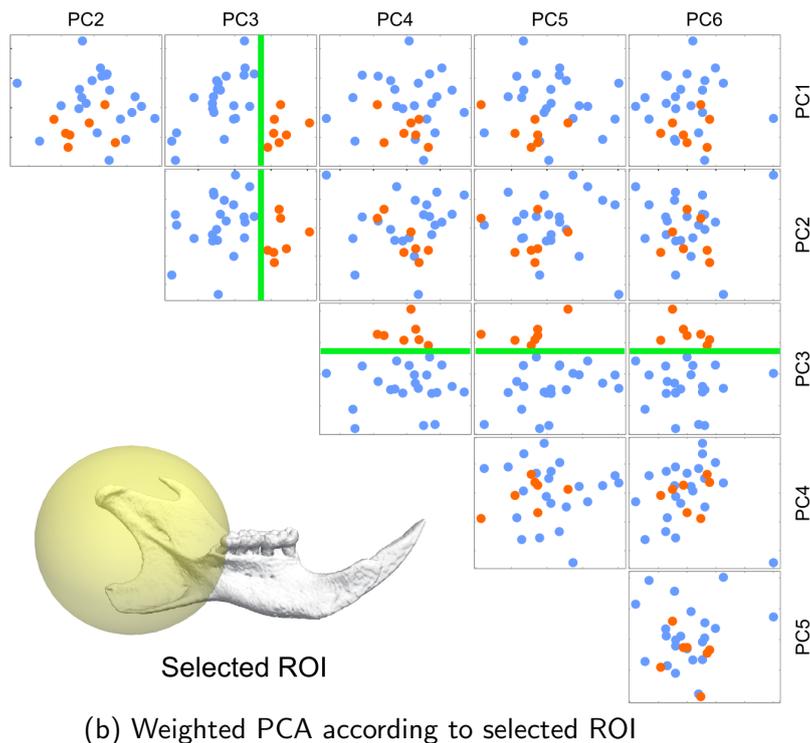
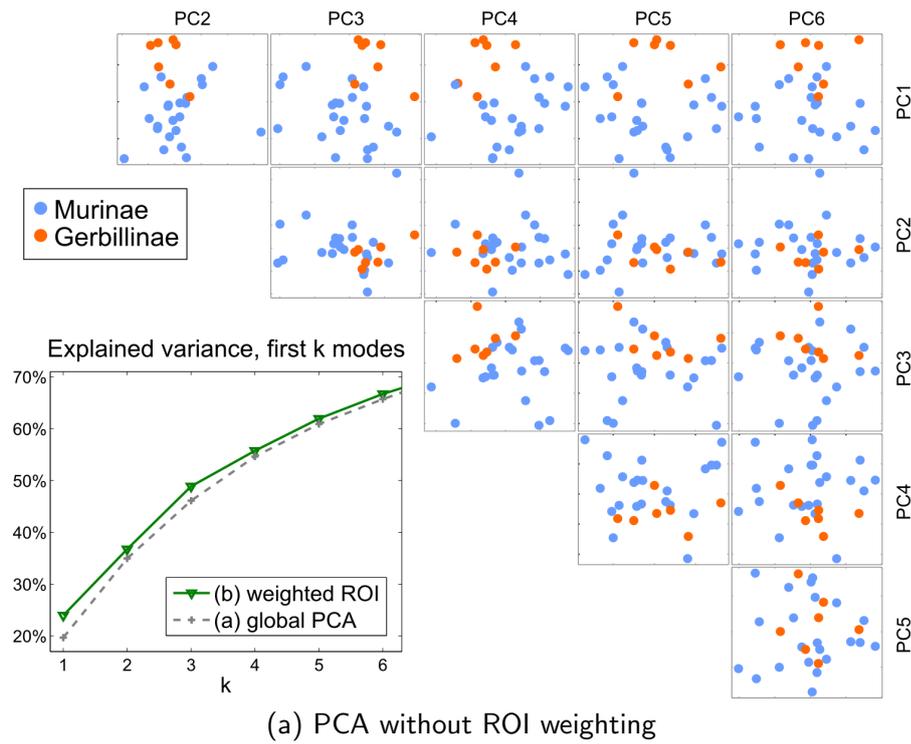


Figure 5.4: Comparison of global and weighted analyses of subfamily classification based on scatter plot matrices of the first 6 PC's which capture over 65% of the shape variance, see inset in (a). This is an example where the proposed ROI selection method provides a more concise PCA model better suited for visual analysis. **(a)** In global analysis no combination of PCs separate the subfamily classes. **(b)** A weighted analysis provides clear subfamily separation along the third PC (green line). The spherical ROI (yellow) is chosen to include the three rear mandible processes which are supposed to be important shape features in discerning the subfamilies.

Trait 1: Subfamily Murinae vs. Gerbillinae

An initial visual exploration of PCA plots based on all specimens shows that none resulted in complete separation of the two subfamilies, see Fig. 5.4a. After applying a ROI selection on the posterior part of the mandible (inset in Fig. 5.4b), the third PC calculated according to this selection provides a good separation. Additionally, a relatively good separation of the subfamilies is apparent in a combination of first and third PC for the complete dataset (Fig. 5.6), which is also reflected by the computed trait vector. The hyperplane normal to the learnt trait vector φ_{trait1} separates both groups without error, i.e. a cross validation shows no misclassification. Fig. 5.5 illustrates the shape changes as captured by the learnt trait vector via a smooth interpolation from the mean shape in directions of Murinae and Gerbillinae. The trait vector features most of the expected differences, mainly in the posterior part of the mandible, where the three processes display characteristic deformations, and a less well known shape difference in the incisor. Direct volume rendering provides additional information on changes in the roots of the molar teeth, as illustrated in Fig. 5.7. The roots tend to be longer and slightly more slender and point further downwards in the Gerbillinae while they are shorter and more bulky in the Murinae, where the molar roots also point downwards and backwards (highlighted in the close-ups), except for the anteriormost root of the first molar. Differences in the cheek teeth of the two subfamilies are actually more pronounced in the roots than in the crowns.

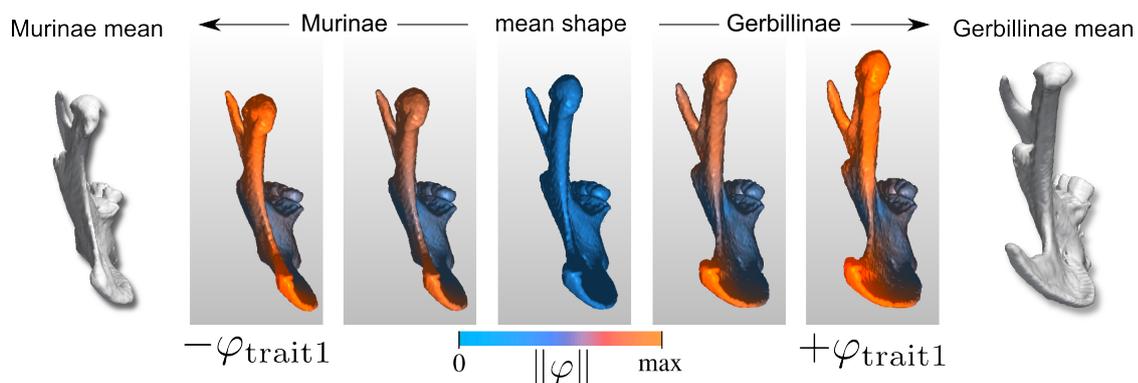


Figure 5.5: Smooth interpolation along trait vector φ_{trait1} learnt from classification into Gerbillinae and Murinae (posterior view). An isosurface is shown color-coded according to local deformation strength. For comparison the respective mean shapes per group are depicted to the left and right.

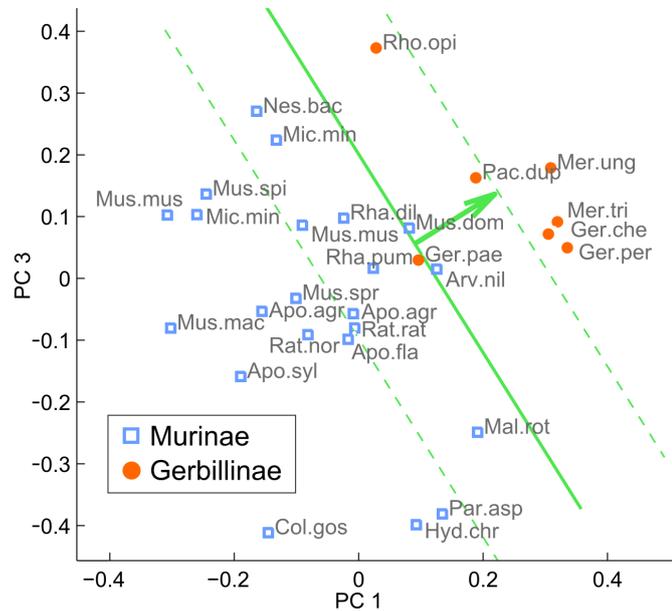


Figure 5.6: PCA plot of best separating axes for analysis Gerbillinae vs. Murinae. The learnt trait vector φ_{trait1} is shown (green arrow) together with its separating hyperplane and margins (green lines).

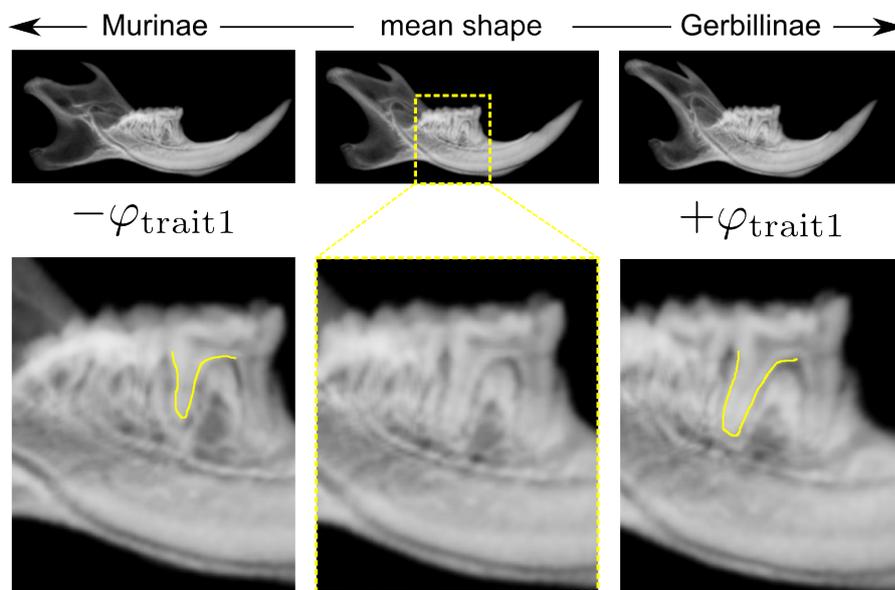


Figure 5.7: Volumetric shape difference in roots of molar teeth as captured by the subfamily trait vector φ_{trait1} . Part of the root of the anteriormost molar is highlighted (yellow contour).

Trait 2: Genus *Apodemus* vs. *Mus*

Given that the dataset used here is extremely heterogeneous we also ran an analysis using only a small, much more homogeneous subset. It consists of the members of the genera *Apodemus* and *Mus*, which differ little within and among groups. Nonetheless it is possible to separate the genera as visible in Fig. 5.9. The learnt trait vector $\varphi_{\text{trait}2}$ is visualized in Fig. 5.8. It shows the typical slight elongation in *Apodemus* vs. the shortening in *Mus* and an additional slight shift in the position of the molar roots from rather vertical in *Mus* to slightly backwards inclined in *Apodemus*.

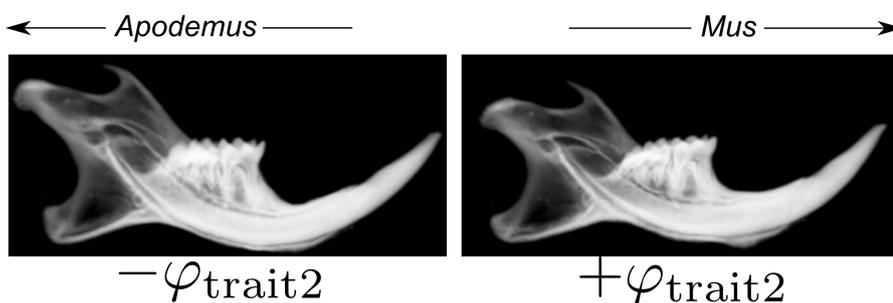


Figure 5.8: Trait vector $\varphi_{\text{trait}2}$ for genus *Apodemus* vs. *Mus*.

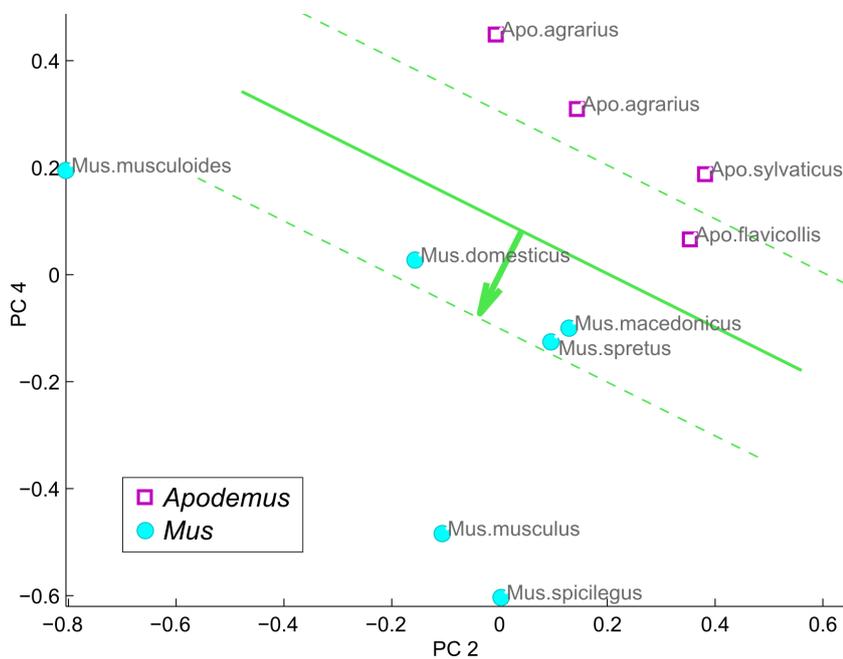


Figure 5.9: PCA plot of best separating axes for analysis *Apodemus* vs. *Mus*. The learnt trait vector $\varphi_{\text{trait}2}$ is shown (green arrow) together with its separating hyperplane and margins (green lines).

Trait 3: Diet

The following analyses are restricted to the murine species, as in this subgroup at least three samples for each character state are available. We start with an analysis of dietary effects, using only animal-eating (carnivorous) and plant-eating (herbivorous) species, because omnivory is a diet mixed of carnivory and herbivory and would thus only add noise to the results. Again the classification results in a trait vector (shown in Fig. 5.10) that corresponds to shape changes matching the general observations. For dietary effects these are an elongated mandible in carnivorous murines and a short, higher one in herbivores. Teeth lie closer to the articular process in the latter, and the incisor describes a tighter arc of a circle.

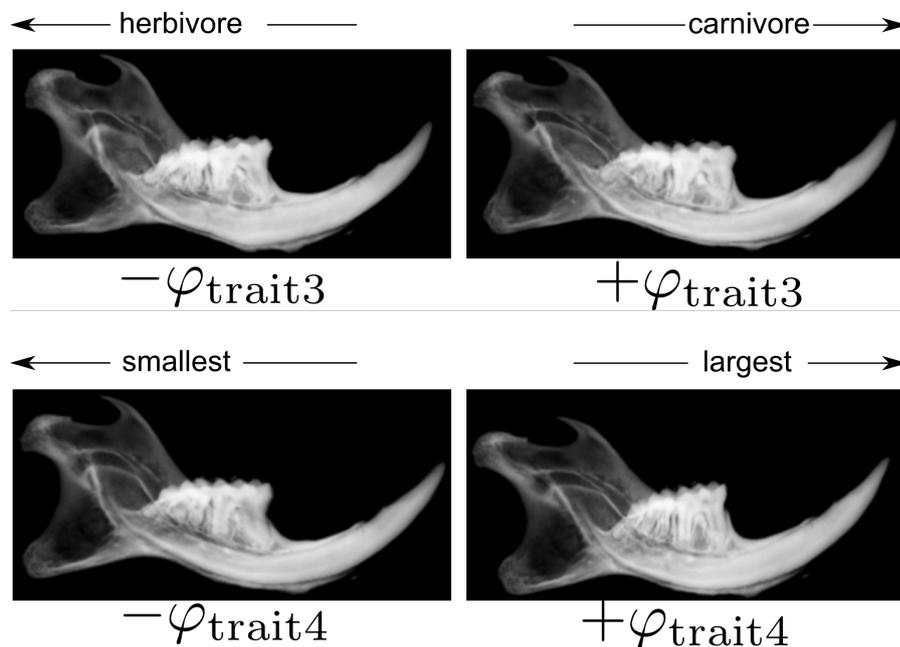


Figure 5.10: Traits 3 and 4. **(Top:)** Trait φ_{trait3} for diet herbivore vs. carnivore. **(Bottom:)** Trait φ_{trait4} for smallest vs. largest individuals.

Trait 4: Size

Finally we investigate allometric effects by using the smallest and largest species. Fig. 5.10 shows the according visualization for trait vector φ_{trait4} . In this case, which also achieves a good separation of the two groups, relatively complex shape changes can be observed. From small to large species the molar teeth row becomes relatively smaller, the tooth roots more narrow, the arc of circle of the incisor smaller, and the angular process shorter and it moves forward, to name just a few. It is already known that the size of the molar teeth row changes less with

size changes of the whole mandible; the other characters await investigations with larger sample size to validate their generalization.

Summary of visual analysis

Altogether the results match those observed for usual 3D landmark analysis (Schunke & Tautz, in prep.), which also showed a good separation of Gerbillinae and Murinae, carnivorous and herbivorous as well as large and small species. Although the differences observed here tended to be less pronounced than those for the set of twenty landmarks, the traits gave valuable additional information of the complete shape changes also in between extreme points. Additionally it was possible to analyse data inside the mandible, namely the tooth roots, which already provided some promising observations that need further investigation.

5.7 Conclusion

We showed in this chapter that by using efficient linear models to describe shape space, interactive visual analysis based on additional semantic information can be conducted on dense deformation models retaining the full information as acquired in medical imaging. However, due to the additional information of this deformation model and its higher dimensionality compared to landmark based models, characteristic traits may not coincide with principal component axes, as is often the case in landmark analysis, and the searched trait is spread across several PCs. Therefore we considered an optimal linear combination of the PCs to extract the trait in question, which is found in a robust way via a support vector machine.

Since shape spaces are inherently non-linear, linear models tend to capture only the prominent shape variations reliably. To face this concern we proposed and implemented a method to derive linear models for sub-spaces of a shape space in order to obtain a more reliable representation of shape variability.

The usefulness of each of these methods was demonstrated on a dataset of rodent mandibles, including one combined application for trait 1 of classification and ROI. We are convinced that, for larger datasets exhibiting many more complex volumetric structures, the advantages of a combined application of both methods will become most apparent. Such a complex example displaying many volumetric structures is the rodent skull (see Fig. 5.11), which is considered for further studies.

In summary we presented in this chapter two methods for exploiting semantic knowledge to facilitate a meaningful visual analysis of morphological shape spaces.

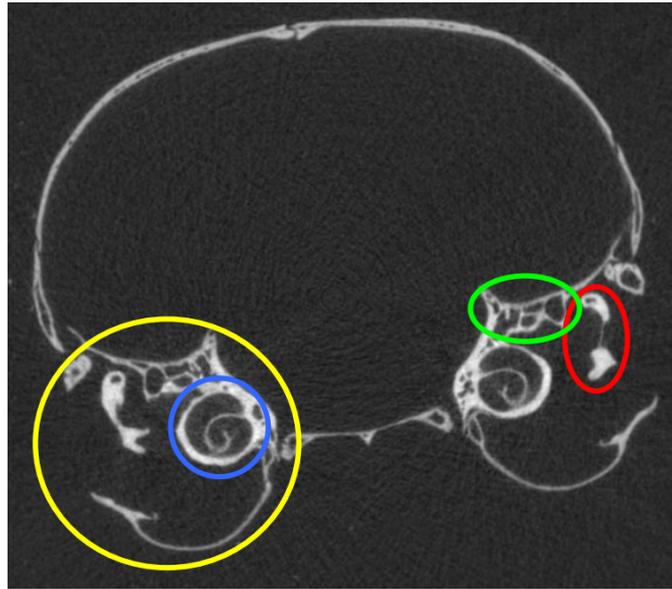


Figure 5.11: Examples of volumetric structures in rodent skull shown on a CT slice of a *Rattus norvegicus* individual: ear ossicles (red), cochlea (blue), semicircular canals (green), bulla (yellow).

These allow us to integrate relevant expert knowledge in form of classification, e.g. according to biological taxonomy, and to highlight regions of interest on the shape. We demonstrated the usefulness of these methods on a dataset of rodent mandibles. Further evaluation studies are subject to future work.

Chapter 6

Visual analysis of anatomic covariation

Visualization is an essential component of understanding anatomic variability. Effective visualizations provide feedback for verifying the various algorithmic steps in computing variability, and offer flexible means of displaying and interacting with the data so as to form and refine hypotheses about the structure and origins of anatomic variability.

— Gordon L. Kindlmann et al.[81]

Gaining insight into anatomic covariation helps the understanding of organismic shape variability in general and is of particular interest for delimiting morphological modules. Generation of hypotheses on structural covariation is undoubtedly a highly creative process, and as such, requires an exploratory approach. In this chapter we propose a generalization of the known anatomic covariance tensor. While the latter is a direct representation of the local covariance structure at a particular point, the methods introduced in this chapter encode the covariation *in-between* different points on the shape, a concept referred to as *interaction* for better disambiguation. A novel interaction operator, interaction tensor and overview tensor facilitate exploration of interaction at different levels of detail, stimulating rapid formation and (qualitative) evaluation of hypotheses. Considering interaction as a generalized decomposition of covariation, the visualizations of this chapter show an unprecedented detail in representation of covariation.

An exemplary case study about modularity of the rodent mandible is conducted, comparing the proposed approach against state-of-the-art methods. Results substantiate the effectiveness of our interactive visualization in producing

module hypotheses that are in agreement with recent morphological theories, but provide a much more detailed model of covariation and interaction.

6.1 Introduction

In this chapter we are particularly concerned with shape covariation and novel visual analytics methods to reveal the often complex interdependency between specific structural parts and the shape as a whole. In biology and anthropology the concept of morphological integration and modularity, particularly of skull and mandible, has gained increasing interest over the last decade, see for instance the work of Klingenberg [88] and references therein. Differences in covariation within and between certain parts of a morphological structure are described, aiming at a better understanding of morphology in general as well as of underlying developmental, functional, or genetic constraints.

At the core of this chapter stands a new covariance tensor, the interaction tensor, which is derived from the model-based deformation framework of Blanz et al. [24]. The introduced tensor summarizes covariation between a specific point and the remainder of the shape and thus generalizes the anatomic covariance tensor used by Kindlmann et al. [81]. A key observation made here is that for linear shape models, that are probably the most common type in morphometric studies, the covariation between a point p and all other points q on the shape is a linear relationship itself, expressed as interaction operator below. While the latter comprises all details of *interaction*, i.e. covariance between two points, a second not so obvious observation will lead us to another abstraction. It will allow us to summarize relevant aspects of all interaction tensor fields in a single so called overview tensor field. Based on that, an automatic segmentation into anatomically meaningful parts is possible, providing an additional kind of overview.

Although linear analysis yields efficient visualization algorithms for a single interaction tensor field, there are as many such fields as voxels in the dataset. Even for image data as small as 64^3 voxels, manual examination of all possible tensor fields becomes prohibitive. On the other hand, a single such tensor field is already a simplifying summary of the underlying linear relationship. Therefore additional (visual) navigation strategies are required.

Following Shneiderman’s information visualization mantra “overview, zoom & filter, then details-on-demand” [126] we provide three different visualizations with decreasing level of abstraction. By going forth-and-back between the different visualizations the massive amount of detailed information on covariation can be

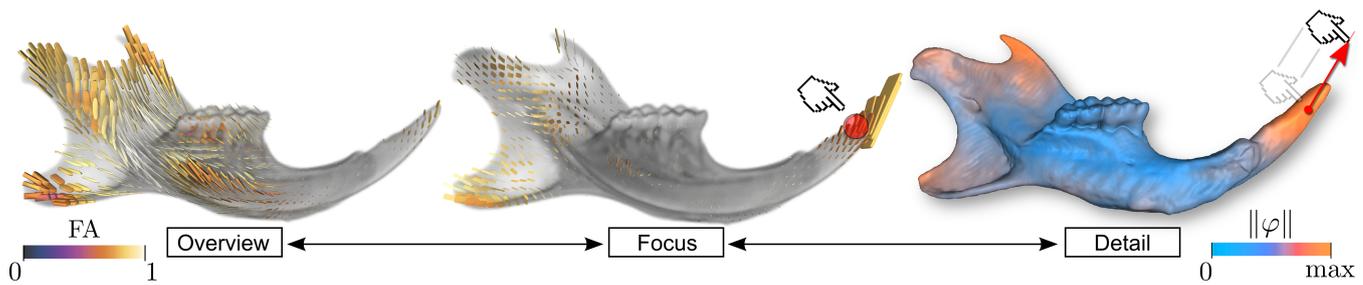


Figure 6.1: The proposed visual analysis of covariation is performed on three levels of abstraction, demonstrated on a mouse mandible: A static overview (left) provides guidance to candidate points exhibiting non-trivial covariation patterns with the remaining shape. For a particular point p a focus visualization (center) reveals the underlying covariation pattern to p . By interactively dragging around p , details of the correlated shape variation can be investigated in a dynamic animation (right), uncovering also specific directional dependencies of covariation.

effectively investigated. The complete pipeline is illustrated in Fig. 6.1. An initial overview visualization highlights strength and directional dependency of covariation associated with certain areas on the structure. Thereby it provides guidance to potentially interesting substructures for further exploration. Each interaction tensor field of a candidate region can in turn be looked at in greater detail in a second glyph visualization. This gives insight into the overall strength and proportions of covariation between a candidate region and global shape, but omits directional dependencies. To examine the latter at the level of individual shape variations, a third visualization provides the possibility to perform model-based deformation interactively in a *click-and-drag* style, inspired by recent works in facial animation.

In combination, the presented three visualizations provide an intuitive and informative interface to visual analysis of shape covariance. Results on our model organism, the rodent mandible, document its utility. Note that the rodent mandible is a particularly suitable choice for demonstration, because it received considerable attention in research on biological modularity in recent years. This allows us to perform a comparison of our findings against results from literature that were achieved with state-of-the-art methods.

6.2 Related work

6.2.1 Interactive approaches to investigate anatomic covariance

In the last years a rich set of interactive visual analysis tools for landmark data and linear shape models has been released [87, 37, 107, 94, 147]. All of these systems feature a PCA model for exploration, while some additionally provide support for canonical variates analysis (CVA) and regression tools like partial-least squares (PLS). CVA can be used to analyze covariation with respect to specific groups of individuals in the dataset, providing an orthogonal parameter space similar to PCA for exploration. A PLS regression can be calculated between two a priori selected landmark groups, i.e. shape parts, to visually analyze covariation between them in a 2D correlation plot [87, 147]. Once a hypotheses on co-varying groups of individuals or landmarks is formed, CVA and PLS methods provide a valuable tool for further exploration. The generation of hypotheses on shape covariation in the first place is so far only supported by browsing the PCA model.

Inspiration for the presented interaction method are the direct manipulation approaches described in Chap. 2, stemming primarily from works in facial animation [97, 133]. Although these works are a great source of inspiration and we will internally use the same least-squares minimization [24], the ultimate goals of artistic animation and visual analysis of shape variation are quite contrary.

6.2.2 Anatomic covariance field

The method introduced by Kindlmann et al. [81] provides a very concise way to convey the overall variability contained in a dataset. To this end the local covariance at a sample point on the shape is evaluated, resulting in an anatomic covariance tensor that is visually encoded in an easy to read glyph.

Transferred to our volumetric setting, the tensor is defined for a particular point $p \in \Omega$ in image domain as the sample covariance matrix of the set of associated displacement vectors $\{\mathbf{u}_{p,1}, \dots, \mathbf{u}_{p,n}\}$ where $\mathbf{u}_{p,i} \in \mathbb{R}^3$ is the difference between the correspondence of p in the i -th ensemble image to the template image. Using this notation, the 3×3 anatomic covariance tensor is defined as

$$\mathbf{T}_{\text{global}}(p) = \frac{1}{n-1} \sum_{i=1}^n \mathbf{u}_{p,i} \mathbf{u}_{p,i}^T \quad (6.1)$$

The glyph-based visualization is based on a spectral analysis that yields prin-

principal directions and variances along these for each tensor. A graphical primitive, typically a sphere, is oriented and scaled accordingly to illustrate the covariance structure. Since a covariance tensor \mathbf{T} is symmetric it can be diagonalized as

$$\mathbf{T} = R\Lambda R^T \quad (6.2)$$

into a rotation R from the unit to an eigenvector basis and a diagonal matrix $\Lambda = \begin{pmatrix} \lambda_1 & & \\ & \lambda_2 & \\ & & \lambda_3 \end{pmatrix}$ of corresponding eigenvalues, sorted in descending order $\lambda_1 \geq \lambda_2 \geq \lambda_3$. The ratios between the eigenvalues delineate the anisotropy of local shape variation [148].

The same spectral analysis and glyph-based visualization is adapted for our interaction and overview tensors.

6.2.3 Partial least-squares analysis

A state-of-the-art approach to assess shape covariation is to perform a partial least-squares analysis (PLS), pioneered by Tucker [138] and Wold [153] and introduced to morphometrics by Rohlf and Corti [116]. It requires a precise hypothesis about a two-block separation of the shape. By reordering the rows of our original data matrix we can represent the two blocks as $\mathbf{X}^T = [\mathbf{X}_1^T \ \mathbf{X}_2^T]$. PLS is technically similar to PCA, but instead of the covariance matrix $\Sigma = \mathbf{X}\mathbf{X}^T$, the cross-covariance matrix $\Sigma_{12} = \mathbf{X}_1\mathbf{X}_2^T$ is diagonalized. Since Σ_{12} is no longer symmetric, this yields not an eigenvalue but a singular value decomposition $\Sigma_{12} = \mathbf{L}\mathbf{D}_{12}\mathbf{R}^T$ with different left and right singular vectors. Where the first PCA eigenvector \mathbf{u} maximizes the covariance $|\text{cov}(\mathbf{X}^T\mathbf{u}, \mathbf{X}^T\mathbf{u})|$, the first pair of singular vectors (\mathbf{l}, \mathbf{r}) maximizes the cross-covariance between the two blocks $|\text{cov}(\mathbf{X}_1^T\mathbf{l}, \mathbf{X}_2^T\mathbf{r})|$. The same holds for further pairs of singular vectors in the respective orthogonal subspaces. Extensions to three and more blocks exist [29], although lacking a closed form solution. By stacking singular vector pairs, full deformation fields for the complete shape can be constructed for visualization purposes. The different size of the blocks has to be taken into account though [91].

In our literature research we found PLS analysis to be virtually the only visual tool used to study shape covariation, see for instance [90] and references in Sec. 4.5. However, it does not allow inspection of covariation at a finer level than the blocks, which have to be selected a-priori, and individual covariation patterns remain superimposed in the singular vectors.

6.2.4 Shape segmentation

Decomposing an anatomy into modules can also be understood as a shape segmentation problem. Unfortunately, standard approaches to this problem do not take deformation statistics into account. For an overview of state-of-the-art approaches for mesh segmentation one can refer to the survey of Shamir [125]. The segmentation of multiple meshes with dense correspondence is considered in analysis and compression of mesh animations. There, similar temporal vertex trajectories are clustered [120]. De Smet and van Gool [53] present an approach to define optimal regions for model-based 3D face reconstruction. In their work, an ad-hoc statistic on normalized displacement magnitudes is taken into account.

In Sec. 6.3.5 we will introduce an approach for segmentation of the template image into anatomically meaningful parts based on a statistical deformation model. This is accomplished by considering the information encoded in our overview tensor that provides a much richer statistic compared to the one utilized by De Smet et al. [53].

6.3 Inter-point covariance

Our approach builds upon the least-squares framework introduced by Blanz et al. [24] for model-based deformation, which is briefly summarized first before introducing interaction operator, interaction tensor and overview tensor in the following subsections.

6.3.1 Model-based deformation

For linear shape models Blanz et al. [24] propose a model-based deformation framework which allows the user to produce plausible shapes via manual displacement of single vertices. The “plausibility” of a shape is defined via its probability according to the Gaussian distribution underlying all PCA shape models, see Sec. 3.5.

In a least-squares optimization, the coefficients \mathbf{c} in PCA space of a plausibly deformed shape are estimated where the displaced vertices \mathbf{u}_p (the user edit) are considered as soft-constraints and the optimization is additionally regularized by the shape’s probability, penalizing improbable outcomes. Using a squared data term and the negative log-likelihood (3.29) the optimization can be expressed via the following energy functional

$$E(\mathbf{c}) = \frac{1}{2} \|\mathbf{u}_p - \mathbf{B}_p \mathbf{c}\|_2^2 + \frac{\gamma}{2} \|\mathbf{c}\|_2^2 \quad (6.3)$$

where \mathbf{u}_p is the stacked vector of displacements and \mathbf{B}_p is the shape space basis matrix reduced to the rows corresponding to the displaced vertex coordinates in \mathbf{u}_p .

At the constrained positions p the minimizer of $E(\mathbf{c})$ will try to match the edited shape \mathbf{u}_p while the remaining vertices are deformed to yield a highly probable shape, depending on the choice of the regularization parameter $\gamma \in \mathbb{R}$. In this work we take on the view that the result of the above optimization gives us the shape variation correlated to the edit \mathbf{u}_p , at least for a fixed γ . The regularization parameter γ is chosen via cross-validation in a leave-one-out fashion benchmarked on the average reconstruction error, see App. B for details. In our applications the resulting error function is found to be smooth and in particular stable around the optimum allowing for a robust choice of γ .

An example application of model-based editing in order to uncover directional dependency in covariation on the shape is shown in Fig. 6.2.

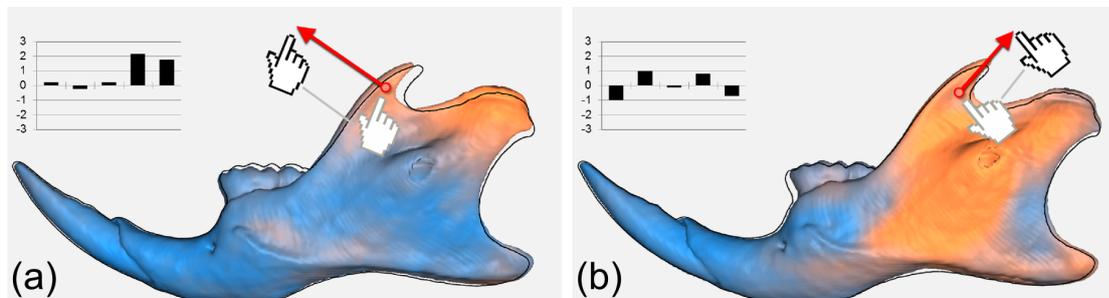


Figure 6.2: Model-based editing on *Mus* dataset at coronoid process. The visualization shows an isosurface of the template displaced by \mathbf{u}_q , color coded from cool to warm (0 max) with the magnitude of the displacement. The templates silhouette is overlaid in black for comparison. Inset bar plots describe the first 5 PC coefficients in units of standard deviations. Observe that varying the coronoid position is strongly associated with the condylar process (a), while changing its length impacts the rear of the mandible at a larger scale (b).

6.3.2 A linear operator for inter-point covariation

The same model-based deformation framework (6.3) can directly be applied to statistical deformation models, where instead of vertices now displacement vectors are constrained. Restricting to the special case of a single constrained position p , we now derive a linear operator for covariation between the point p and any other point $q \in \Omega$. The displacement predicted from (6.3) at q to a variation at p is what we interpret in this paper as the correlated shape change, or *interaction*, between p and q .

Once, optimal coefficients \mathbf{c}_{opt} that minimize Eq. (6.3) are found, the corresponding displacement \mathbf{u}_q at q co-varying with the change \mathbf{u}_p at p is given by

$$\mathbf{u}_q = \mathbf{B}_q \mathbf{c}_{\text{opt}} \quad \text{where} \quad \mathbf{c}_{\text{opt}} := \arg_{\mathbf{c}} \min E(\mathbf{c}). \quad (6.4)$$

We follow the default ansatz to solve this least squares optimization problem and start from the necessary optimality criterion of a vanishing gradient $\frac{\partial}{\partial \mathbf{c}} E \stackrel{!}{=} 0$:

$$\begin{aligned} \frac{\partial}{\partial \mathbf{c}} E(\mathbf{c}) &= 0 \\ \frac{\partial}{\partial \mathbf{c}} \left[\frac{1}{2} \|\mathbf{u}_p - \mathbf{B}_p \mathbf{c}\|_2^2 + \frac{\gamma}{2} \|\mathbf{c}\|_2^2 \right] &= 0 \\ \frac{1}{2} \frac{\partial}{\partial \mathbf{c}} (\mathbf{u}_p - \mathbf{B}_p \mathbf{c})^T (\mathbf{u}_p - \mathbf{B}_p \mathbf{c}) + \frac{\gamma}{2} \frac{\partial}{\partial \mathbf{c}} \mathbf{c}^T \mathbf{c} &= 0 \\ \frac{1}{2} \frac{\partial}{\partial \mathbf{c}} [\mathbf{u}_p^T \mathbf{u}_p - 2\mathbf{c}^T \mathbf{B}_p^T \mathbf{u}_p + \mathbf{c}^T \mathbf{B}_p^T \mathbf{B}_p \mathbf{c}] + \gamma \mathbf{c} &= 0 \\ \mathbf{B}_p^T \mathbf{B}_p \mathbf{c} - \mathbf{B}_p^T \mathbf{u}_p + \gamma \mathbf{c} &= 0 \\ \underbrace{(\mathbf{B}_p^T \mathbf{B}_p + \gamma \mathbf{I})}_{\mathbf{A}} \mathbf{c} &= \underbrace{\mathbf{B}_p^T \mathbf{u}_p}_{\text{const.}} \end{aligned}$$

Considering that the last line has the canonical form $\mathbf{A} \mathbf{c} = \text{const.}$ with a symmetric $n \times n$ system matrix \mathbf{A} that is guaranteed to be of full rank for $\gamma > 0$, a unique inverse exists and the solution is given by

$$\mathbf{c}_{\text{opt}} = \mathbf{A}^{-1} \mathbf{B}_p^T \mathbf{u}_p. \quad (6.5)$$

Expanding (6.5) into (6.4) we arrive at a displacement vector

$$\mathbf{u}_q = \mathbf{B}_q (\mathbf{B}_p^T \mathbf{B}_p + \gamma \mathbf{I})^{-1} \mathbf{B}_p^T \mathbf{u}_p. \quad (6.6)$$

With that we define now our **interaction operator** as exactly that linear relationship between displacements at p and q given by the 3×3 matrix

$$\mathbf{Z}_{pq} := \mathbf{B}_q (\mathbf{B}_p^T \mathbf{B}_p + \gamma \mathbf{I})^{-1} \mathbf{B}_p^T. \quad (6.7)$$

Note that the reflexive operator is not necessarily the identity in general, i.e. $\mathbf{Z}_{pp} \neq \mathbf{I}$, because of the regularization term that penalizes unlikely edits. It can also happen that the edit is not contained in the span of \mathbf{B} at all, meaning it has zero probability. One should further keep in mind that the relationship is not symmetric and $\mathbf{Z}_{pq} \neq \mathbf{Z}_{qp}^{-1}$ in general.

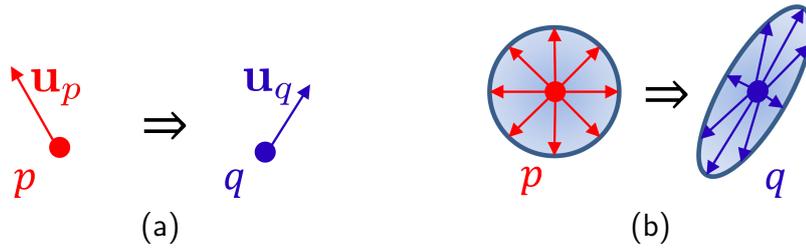


Figure 6.3: Illustration of interaction operator and tensor. **(a)** The interaction operator predicts a displacement \mathbf{u}_q at a point q given a displacement \mathbf{u}_p at a point p . **(b)** The interaction tensor describes the covariance response at q with respect to an unbiased unit covariance perturbation at p based on the interaction operator.

6.3.3 A new shape covariance tensor

We will now summarize the covariance structure of the interaction operator in the so-called interaction tensor. The connection between operator and tensor is illustrated in Fig. 6.3. For a given displacement at point p in the shape, Eq. (6.7) predicts the most probable corresponding displacement at q . To visualize this relationship encoded in \mathbf{Z}_{pq} we resort to statistical covariance analysis. There, covariance structure of a random vector \mathbf{x} is defined as $\Sigma(\mathbf{x}) = \mathbb{E}\{(\mathbf{x} - \mathbb{E}(\mathbf{x}))(\mathbf{x} - \mathbb{E}(\mathbf{x}))^T\}$ with expectation \mathbb{E} . Assuming that input displacements \mathbf{u}_p are random vectors drawn from a distribution with covariance $\Sigma(\mathbf{u}_p)$, the covariance structure at q follows directly from the linearity of expectation:

$$\Sigma(\mathbf{u}_q) = \Sigma(\mathbf{Z}_{pq}\mathbf{u}_p) = \mathbf{Z}_{pq}\Sigma(\mathbf{u}_p)\mathbf{Z}_{pq}^T \quad (6.8)$$

For our directed interaction we want an unbiased estimate of $\Sigma(\mathbf{u}_q)$ and assume thus a prior of isotropic covariance at p . By setting therefore $\Sigma(\mathbf{u}_p) = \mathbf{I}$ we arrive at the following simple formula for a novel covariance tensor. We define the **interaction tensor** at q for a fixed p as

$$\mathbf{T}_p(q) := \mathbf{Z}_{pq}\mathbf{Z}_{pq}^T. \quad (6.9)$$

From the properties of the interaction operator, i.e. $\mathbf{Z}_{pp} \neq \mathbf{I}$ and $\mathbf{Z}_{pq} \neq \mathbf{Z}_{qp}^{-1}$, it follows that $\mathbf{T}_p(q)$ describes a one-sided, directed interaction from p towards q and in general a posterior covariance $\mathbf{T}_p(p) \neq \mathbf{I}$ at p .

An example application of the interaction tensor unveiling patterns of covariation is shown in Fig. 6.4.

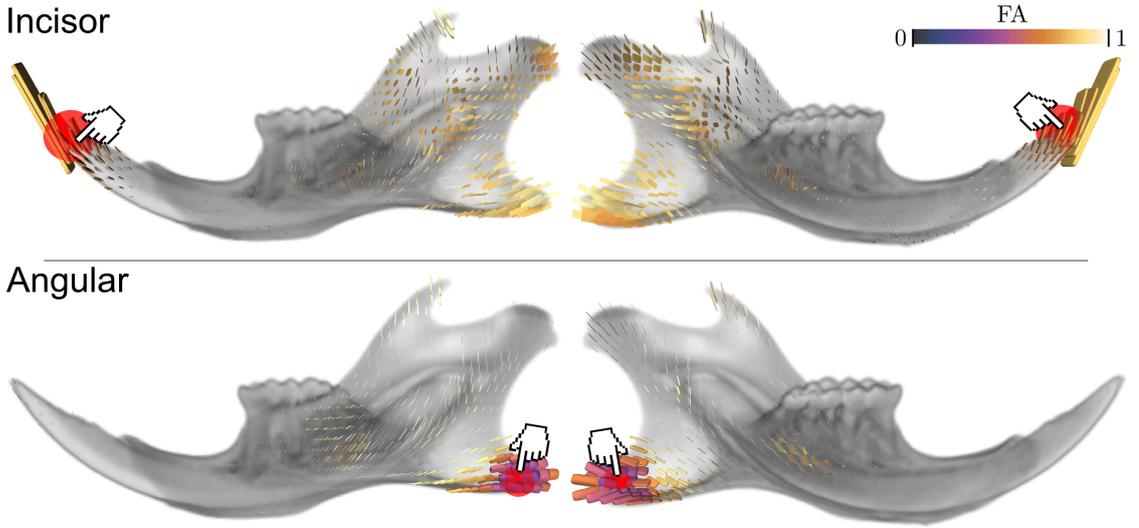


Figure 6.4: Interaction tensor fields \mathbf{T}_p of *Mus* dataset for probes at incisor and angular process, visualized using super quadric glyphs where the probing location p is highlighted. Note that the incisor tensor field is scaled by a factor of 2.5, attributed to the smaller overall interaction strength (compare Fig. 6.5). While the incisor probe produces a global covariation pattern, the response at the angular process is concentrated locally.

Efficient tensor sampling

To speed up the computation of local tensor fields we split \mathbf{Z}_{pq} into two factors, depending each solely on p and q , respectively:

$$\mathbf{Z}_{pq} = \mathbf{B}_q \mathbf{Z}_p \quad \text{where} \quad \mathbf{Z}_p = (\mathbf{B}_p^T \mathbf{B}_p + \gamma \mathbf{I})^{-1} \mathbf{B}_p^T$$

Thereby the matrix inversion in \mathbf{Z}_p has only to be done once on positioning the probe, while the associated tensor field (6.9) can be sampled at the cost of a matrix multiplication per sample.

Sample-based computation

An alternative derivation of (6.9) can be given using a sample estimated covariance matrix. Thereby it could also be applied for other, even non-linear, deformation models. Given a set of sample displacement vectors $\hat{\mathbf{u}}_{p,i}$ with $i = 1, \dots, k$ and associated displacement predictions $\hat{\mathbf{u}}_{q,i}$ from (6.7) or another model, the local covariance is estimated as

$$\hat{\mathbf{T}}_p(q) = \frac{1}{k-1} \sum_{i=1}^k \hat{\mathbf{u}}_{q,i} \hat{\mathbf{u}}_{q,i}^T$$

Choosing unit length displacements $\hat{\mathbf{u}}_{p,i}$ with orientations normally distributed on the sphere, $\hat{\mathbf{T}}_p(q)$ approaches $\mathbf{T}_p(q)$ for $k \rightarrow \infty$.

6.3.4 Overview tensor

For an overview visualization we want to assess which directions of displacement \mathbf{u}_p at a point p will result in large deformations on the remaining shape, describing potentially interesting covariation. The strength of interaction inflicted by \mathbf{u}_p can be measured by averaging the squared magnitude of the displacement responses $\mathbf{u}_q = \mathbf{Z}_{pq} \mathbf{u}_p$. Such a scalar measure is provided by

$$\eta_p(\mathbf{u}_p) = \frac{1}{|\Omega|} \sum_{q \in \Omega} \|\mathbf{Z}_{pq} \mathbf{u}_p\|^2 = \frac{1}{|\Omega|} \sum_{q \in \Omega} \mathbf{u}_p^T \mathbf{Z}_{pq}^T \mathbf{Z}_{pq} \mathbf{u}_p. \quad (6.10)$$

Notice the occurrence of the quadratic forms $\mathbf{Z}_{pq}^T \mathbf{Z}_{pq}$ in the definition of this scalar measure. These terms carry the relevant global information about interaction that we are interested in. Based on this observation, a first *tensorial* description of global interaction can be defined by averaging the quadratic forms, resulting in

$$\mathbf{\Gamma}'(p) = \frac{1}{|\Omega|} \sum_{q \in \Omega} \mathbf{Z}_{pq}^T \mathbf{Z}_{pq}. \quad (6.11)$$

We would like to stress that this is not an average of interaction tensors $\mathbf{Z}_{pq} \mathbf{Z}_{pq}^T$ but one of the form $\mathbf{Z}_{pq}^T \mathbf{Z}_{pq}$. The advantage of the latter form is that the spectrum of its average $\mathbf{\Gamma}'$ informs on which editing directions \mathbf{u}_p impact the shape stronger (large λ_i) and weaker (small λ_i), while the magnitude $\|\mathbf{\Gamma}'\|_{Frob}$ relates to the overall interaction strength between a particular point and the whole shape.

However, this relation is not scaled correctly yet because the scalar measure η_p from which (6.11) is derived treats all local perturbations \mathbf{u}_p as equally important, irregardless of their probability due to the shape model. In order to take into account that the effects of unlikely perturbations are of lesser interest for an overview, we weight $\mathbf{\Gamma}'$ with the interaction tensor $\mathbf{T}_p(p)$ of a point towards itself because that describes exactly the posterior local covariance and thereby the likelihood of local perturbations at p with respect to the shape model. Including this weighting the final definition of the **overview tensor** reads

$$\mathbf{\Gamma}(p) = \mathbf{T}_p^T(p) \mathbf{\Gamma}'(p) \mathbf{T}_p(p). \quad (6.12)$$

An example overview visualization based on (6.12) is shown in Fig. 6.5. Note

that in order to exaggerate tensors with smaller trace glyphs are scaled by the square-root of tensor eigenvalues here.

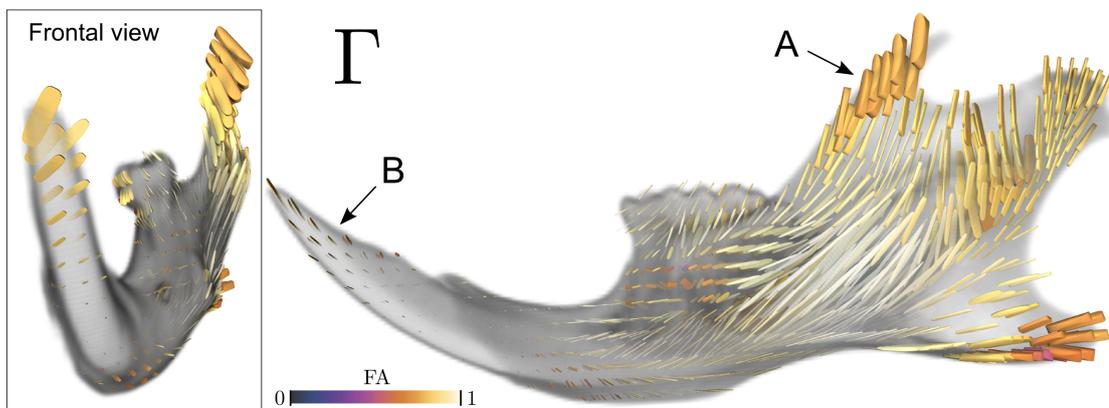


Figure 6.5: Overview tensor visualization of *Mus* dataset. Here glyph size encodes overall strength of the underlying interaction pattern \mathbf{T}_p (compare Fig. 6.4 to point B) while glyph shape informs on directional dependency of variation (compare Fig. 6.2 to point A).

6.3.5 Anatomic segmentation

Interestingly, based on the overview tensor 6.12 a novel *automatic* segmentation into anatomically meaningful parts can be defined, see Fig. 6.6. To this end the tensor field is clustered based on a distance function that is a weighted average between a metric on covariance matrices and spatial distance

$$d(p_i, p_j) = \|\mathbf{\Gamma}(p_i) - \mathbf{\Gamma}(p_j)\|_{Frob} + \alpha \|p_i - p_j\|_2 \quad (6.13)$$

where a weight $\alpha > 0$ favors spatially localized clusters. Spatial closeness is simply measured via Euclidean distance because we are dealing with 3D volumes and not surfaces. As tensor metric Frobenius norm $\|\cdot\|_{Frob}$ is used. Alternative choices include the Log-Euclidean metric [60] that is popular for analyzing diffusion tensors. A comparison of different distance measures between covariance matrices and their invariance properties is given by Dryden et al. [54].

Once, a dissimilarity like the one above is defined, any distance based clustering algorithm can be employed to compute a segmentation. In our experiments, results were chosen as the best of 100 runs of a k-medoids algorithm [75]. By selecting a relational clustering approach, that re-uses original data points as cluster centers, the computation of tensor averages is avoided. In case of Frobenius norm the average can be far off the data, that really represents only positive semi-definite matrices, while for other metrics the computation of average requires gradient

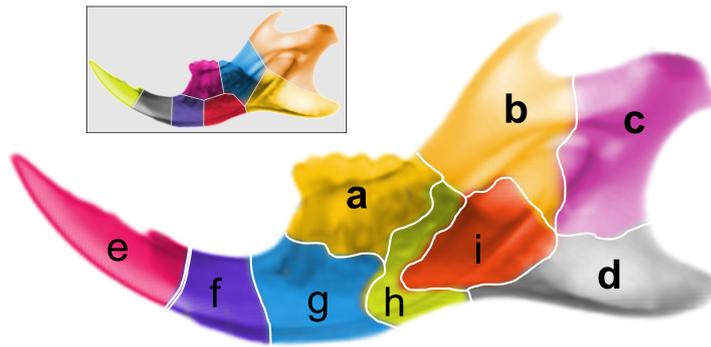


Figure 6.6: Anatomic segmentation of *Mus* dataset ($\alpha = 0, k = 9$) compared against plain k -means result, shown in the inset. Note how clusters a to e nicely match the functional anatomy of molar teeth, rear processes and tip of the incisor.

descent or similar iterative procedures, increasing time complexity of the clustering procedure.

6.4 Visualization system

6.4.1 Interactive model based editing using a rubber band metaphor

We use a direct visualization approach that renders deformed shapes in a raycaster as described in Chap. 4. Deformations are synthesized as displacement fields from the PCA model based on coefficients \mathbf{c}_{opt} that minimize Eq. (6.5) based on a user edit \mathbf{u}_{edit} .

The user interface for editing utilizes the metaphor of a *rubber band* that is being pulled at one end, after the other end was attached to a fixed position on the shape. In this picture the amount of stretching of the rubber band that is required to achieve a noticeable effect in deformation gives information about the “rigidity” of the structure the band is attached to, i.e. the probability of a local shape variation into pull direction. Note that also for nearly rigid points, strong covariations in other parts of the shape may show up, see for instance Fig. 6.17(g)-(h). Rubber band deflection and attachment position specify the edit vector \mathbf{u}_{edit} and location p for the underlying model-based deformation. During interaction with the rubber band, the energy minimization (6.3) is solved perpetually and the raycaster will show the resulting computed deformation of the template in real time.

The selection of a starting point for interactive editing is performed on a user defined isosurface via ray-picking, implemented directly in the raycaster. Selected is the intersection point of a ray cast from under the mouse cursor with the isosur-

face, which is closest to the viewer. Alternatively, a visibility-driven ray-picking technique like that of Wiebel et al. [149] can easily be integrated into our system, thus removing the requirement to specify an isovalue explicitly. After the starting point p_0 is selected, the user inputs the edit displacement vector by dragging around the mouse. The 2D mouse position is then mapped to another point q to define the edit vector $\mathbf{u}_{edit} = q - p_0$. In a naive approach, q is computed by projecting the mouse position onto a plane through p_0 , parallel to the screen. This has the drawback that the specified edit vector depends on the exact camera position and viewing angle. Since shape variation is eventually split up into tangential and normal part in many analyses, we project the mouse position into these directions depending on the surface normal at the starting point. As long as the projected point is not near a contour of the shape, we project onto the tangent plane that in this situation is roughly parallel to the screen. Near contour points the normal is nearly perpendicular to the viewing direction and a projection onto a screen parallel plane is preferred, allowing the user to edit the contour in normal direction.

The raycaster provides direct volume rendering (DVR) and indirect isosurface rendering. While DVR makes interior structures accessible, it requires additional techniques to be employed for color-coding due to color-mixing along a cast ray. For simplicity we encode additional information like displacement magnitude only in the isosurface rendering. Our implementation achieves on average 11 fps during editing on a Intel Core2 Q6600 CPU at 2.4GHz equipped with a Nvidia GTX 460 graphics card.

6.4.2 Tensor field visualization

For the visualization of covariance tensor fields we use the same glyph-based technique as Kindlmann et al. [81]. In a glyph based approach, an effective visualization is achieved by scaling and rotating a geometrical primitive, typically a sphere, according to R and Λ from the tensors diagonalization (6.2). Instead of a sphere, we follow Kindlmann et al. and use a superquadric glyph [80] because of its qualities for visual disambiguation of linear, planar and spherical shaped tensors, that is relevant to our visual analysis.

The geometric encoding of tensor properties is accompanied by color-coding the glyph according to local anisotropy. In deformation models the anisotropy informs on the amount of directional dependency of shape variation and is particularly helpful for our local analysis in understanding the exact interaction between two

points. A suitable measure in this context is fractional anisotropy, defined as

$$FA := \left(\frac{3(\lambda_1 - \mu)^2 + (\lambda_2 - \mu)^2 + (\lambda_3 - \mu)^2}{2(\lambda_1 + \lambda_2 + \lambda_3)^2} \right)^{1/2}. \quad (6.14)$$

FA interpolates between the spherical case ($\lambda_1 = \lambda_2 = \lambda_3$ and $FA = 0$) and the linear one ($\lambda_1 \gg \lambda_2 \approx \lambda_3 = 0$ and $FA = 1$), irregardless of tensor norm $\|\mathbf{T}\|_F$. The final visualization is produced by sampling the tensor field on a regular grid and placing a corresponding superquadric glyph with color mapped FA at each sample position.

For the local covariance visualization the tensor field \mathbf{T}_p serves as input, for which a probe is positioned in 3D at a point p of interest in advance. The overview visualization is generated from the precomputed tensor field $\mathbf{\Gamma}$. For all visualizations and computations the image domain Ω is thresholded to the Hounsfield range of bone structures. Additional methods like glyph packing [82] and halos [121] could be applied to enhance visual disambiguation of the often observed highly anisotropic patterns.

6.5 Experiment on a toy dataset

In order to illustrate the utility of our model-based editing method for exploring shape covariation, we designed a synthetic dataset of smiley figures with known covariation between parts. The dataset contains 49 faces of size 128×128 exhibiting two independent patterns of covariation in facial expression and ear size, see Fig. 6.7. Registration and template estimation is performed with the same methods as used later for the scientific datasets. We will now shortly describe each of the properties we regard as relevant for the intended navigation.

Continuity

For the rubber band metaphor it is important that the reaction of the edited shape is proportional to the deflection of the band, as it is the case in model-based deformation. Fig. 6.8 (top row) illustrates that continuously changing the edit vector results in a continuous deformation of the shape. This enables the user to judge the rigidity of a manipulated structure by means of deflection. Further, small scale variations can be emphasized by pulling the rubber band farther away.

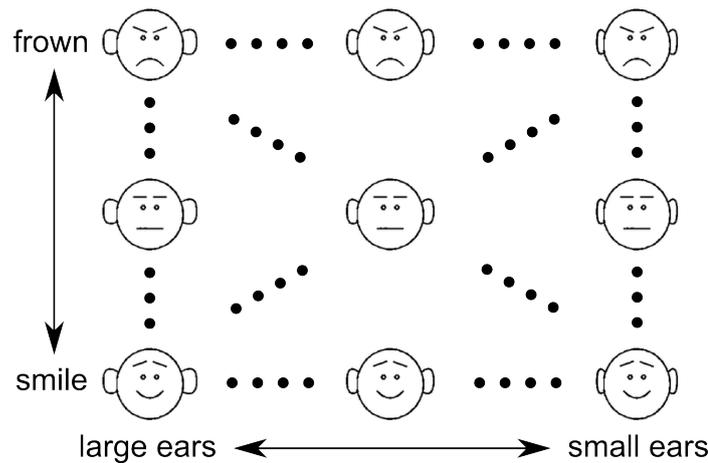


Figure 6.7: The synthetic dataset is sampled from two independent variations: Horizontally, a smile turns into a frown. Note that the curvature of the mouth is correlated to orientation of the brows. Vertically, size of the ears changes bilateral.

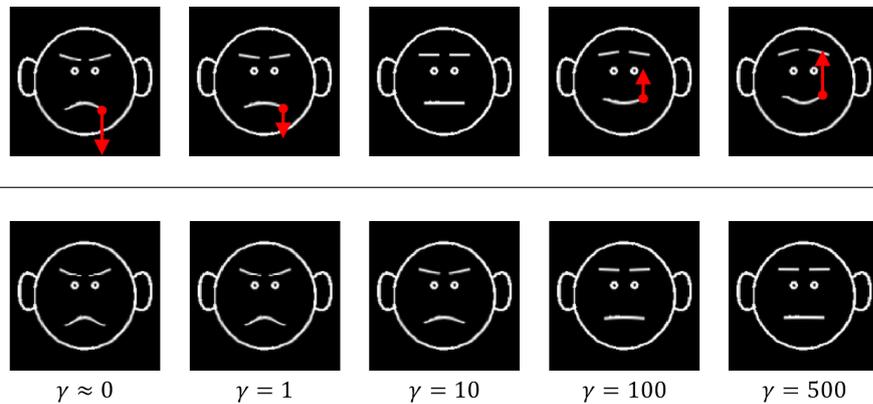


Figure 6.8: Illustration of continuity (top) and influence of γ (bottom) on synthetic dataset.

Influence of γ

The scale of γ can be understood from the Bayesian derivation of the energy minimization (6.3), given in Blanz et al. [24]. There, the coefficients \mathbf{c} to an displacement field \mathbf{u} are found by maximizing the posterior probability $P(\mathbf{c}|\mathbf{u})$. Modeling \mathbf{u} as a measurement subject to independent Gaussian noise in shape space of variance σ_N in all dimensions, an equivalent minimization is found where $\gamma = \sigma_N^2$. In our applications we measure displacements in image units, i.e. number of pixels/voxels. Hence, a choice of $\gamma = 1$ amounts in the Bayesian setting to a measurement error in the input displacement of about ± 3 pixels. In this case the solution can not be expected to deviate more than 3 pixels from the edit displacement \mathbf{u}_{edit} , irregardless how improbable the edit is. Hence, choosing γ too small will lead to over-fitting artifacts. With increasing γ the solution coefficients

will be shifted more and more towards the mean shape until ultimately $\mathbf{c} = 0$, as illustrated in Fig. 6.8 (bottom row). This shows the importance of cross-validation to automatically select an optimal γ value for real world datasets.

Correlated shape change

Fig. 6.9 illustrates, that the response of the editing system to different local edits shows reasonable covariation on a global scale. For instance, raising a single corner of the mouth results in a transition to a smiling face with raised eyebrows (a). As expected, the PCA explains the two independent variations on different principal components. Thus, the sparse PC coefficients \mathbf{c}_{opt} in (a)-(c) verify that the solutions also separate between the two variations nicely, depending on the edited structure. Edit (b) shows the robustness of the system against pulling into arbitrary directions, where the result only reflects the probable component of the movement, leading to enlarged but not distorted ears. Edit (d) shows a sanity check that the system does not answer with spurious solutions to totally improbable inputs.

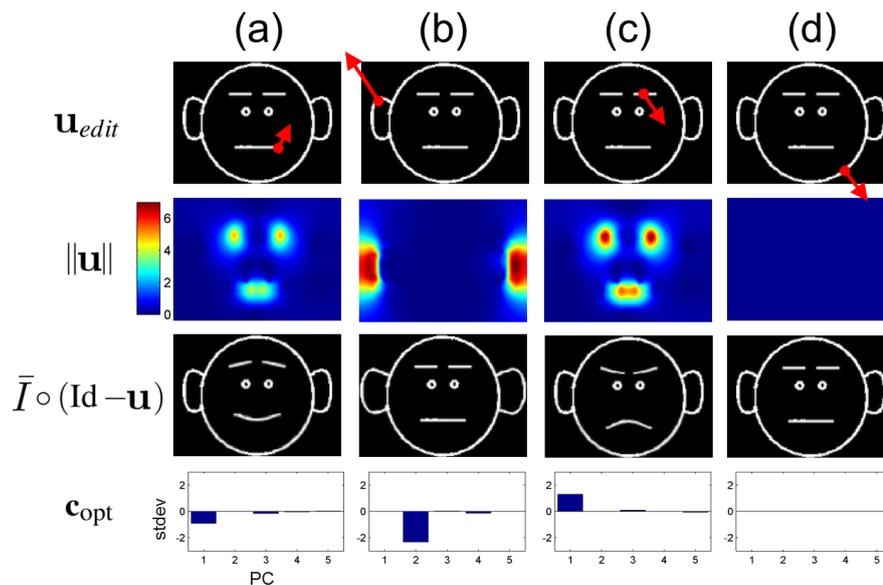


Figure 6.9: Editing results on synthetic dataset, rows from top to bottom show: The performed edit (red arrow) on the template \bar{I} , the magnitude of the resulting displacement field \mathbf{u} , the warped template and the PCA coefficients \mathbf{c}_{opt} . Edits (a)-(c) illustrate that the different correlated structures can be edited separately while (d) shows the rigidity of the system in case of deformations not seen in the example dataset. Note that edit vectors are slightly scaled for illustration purposes.

6.6 Visual analysis of modularity in *Mus*

A comparative modularity analysis performed with our approach is described now. This case study on a real world dataset is undertaken on the backdrop of morphological integration and modularity, for which the mouse mandible is a standard model [90, 160, 34, 113]. As mentioned in Sec. 1.5, finding modules and modeling their interaction is considered a difficult task. In the majority of studies so far, hypotheses of the existence and position of modules have been formulated a priori and tested subsequently. So far, PLS analysis is the primary method to visualize covariation on the shape in such studies.

Modular structure of the mouse mandible

The structure of the mouse mandible anatomy is described in Sec. 1.5 together with the common functional segmentation into frontal region and rear processes. However it is speculated that there could be an additional set of modules at a smaller scale [90] and in recent work several finer scale subdivisions [160, 34, 113] are addressed. For convenience, the illustration from Sec. 1.5 is repeated in Fig. 6.10.

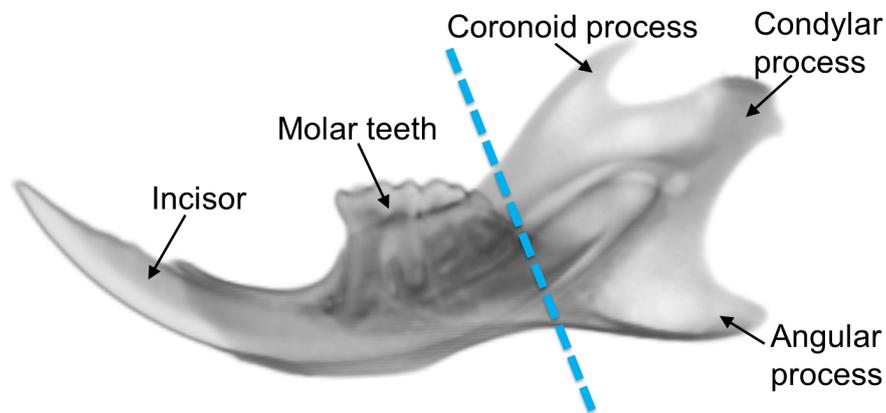


Figure 6.10: Anatomical parts of the rodent mandible referred to in this work. A common subdivision into two functional subunits is indicated [90].

PCA model

The analysis is conducted on the *Mus* dataset consisting of left mandibles of 30 house mice (*Mus musculus*), described in more detail in Appendix A. In the PCA model 93% of shape variability is captured by the first 5 components, see Fig. 6.11. For synthesis and model-based editing the PCA model was reduced to these first five modes, while all tensor computations were performed on the full model.

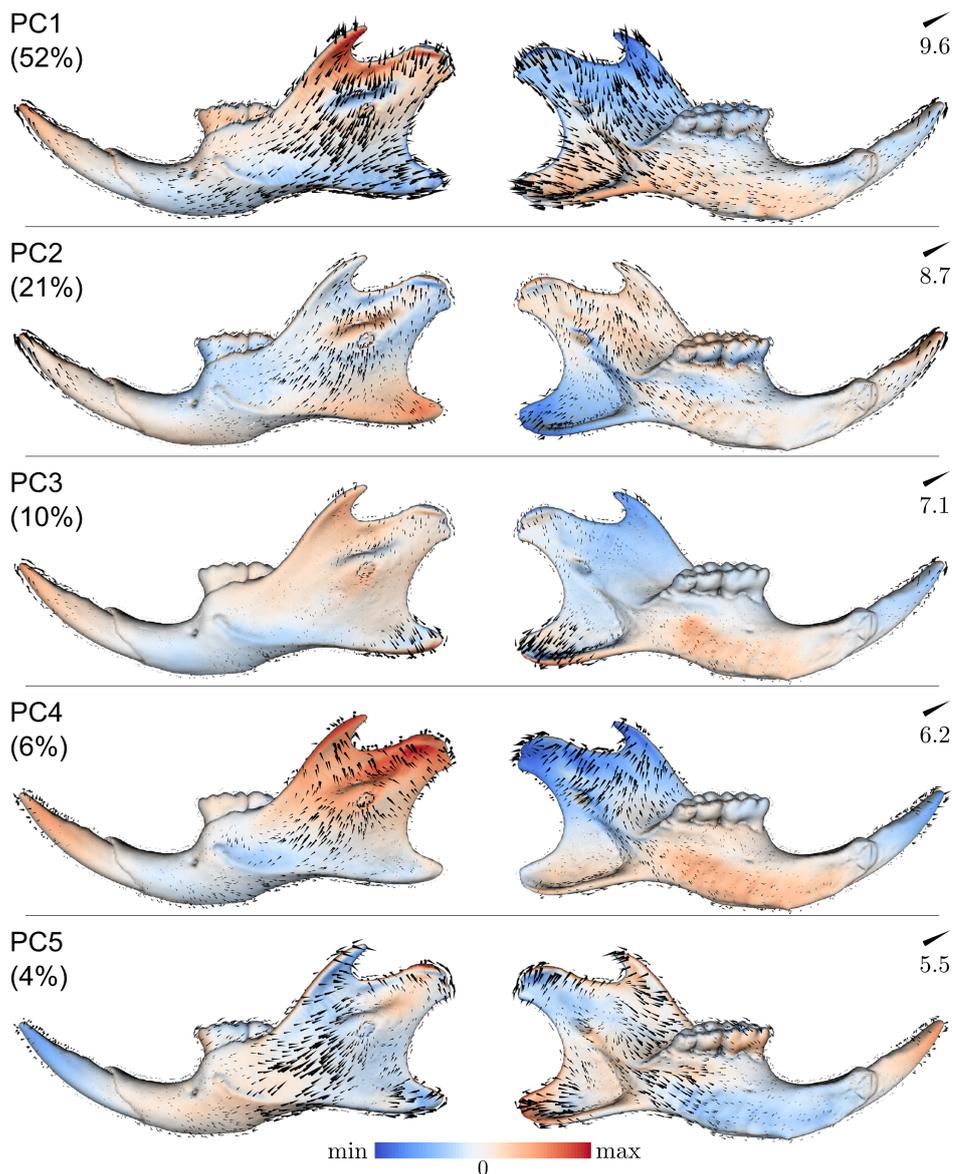


Figure 6.11: PCA eigenmodes of *Mus* dataset capturing 93% of the total variance. Shown are the vectorfields on a representative iso-surface decomposed into a surface orthogonal and tangential part, visualized color-coded and as vector glyphs respectively with glyphs of max. magnitude scaled to same length. (Visualization based on [163].)

6.6.1 Global analysis and PLS

For later comparison, we will first describe results achieved with previous methods of related work in morphometrics, applied to our dataset. Particularly PCA [160, 106], PLS [34] and a combination of both [113, 90] was used. Additionally we will apply the anatomic covariance visualization of Kindlmann et al. [81].

Global methods like PCA give us an impression of the overall variability contained in the dataset. Fig. 6.11 shows the relevant PCA eigenmodes. We observe that processes and incisor are each influenced by several modes, so the PCA does not reveal a clear separation into modules. A more concise overview of shape variation is provided by the anatomic covariance field in Fig. 6.12. It can be read from the glyph pattern that the three processes have different principal directions of variation. At the posterior part of the incisor, inside the mandible, a strongly anisotropic region is visible following nicely the incisor curvature. Note that the anatomic covariance field shows point-wise variability of the dataset but does *not* convey information on covariance *in-between* different points on the shape.

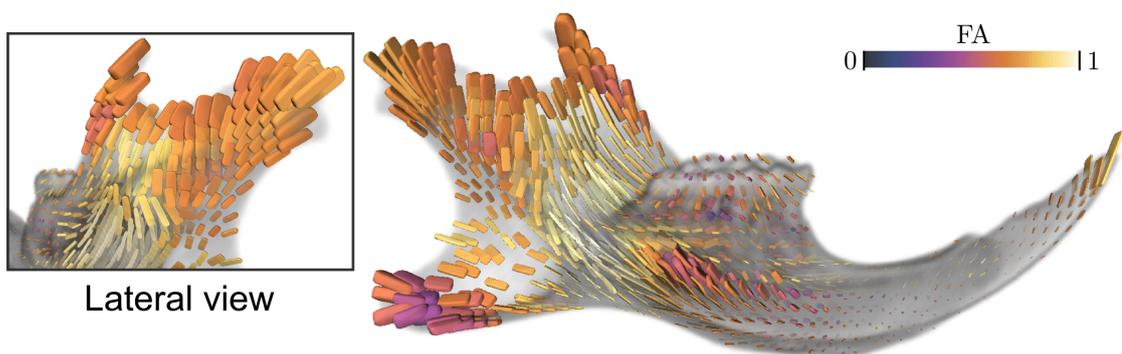


Figure 6.12: Anatomic covariance field [81] for *Mus* dataset. A lateral view is shown in the inset for comparison with Fig. 6.15.

Results of a PLS analysis are visualized in Fig. 6.13. The visible similarity in PLS1 of the rear processes to (minus) PC1 is also encountered in landmark analyses [90] and can probably be attributed to the fact that the processes constitute a great proportion of global variation. The PLS2 modes specifically indicate an interaction between coronoid and condylar process and between incisor and angular process. Modes beyond PLS1 are harder to interpret because they only represent covariation orthogonal to previous mode pairs.

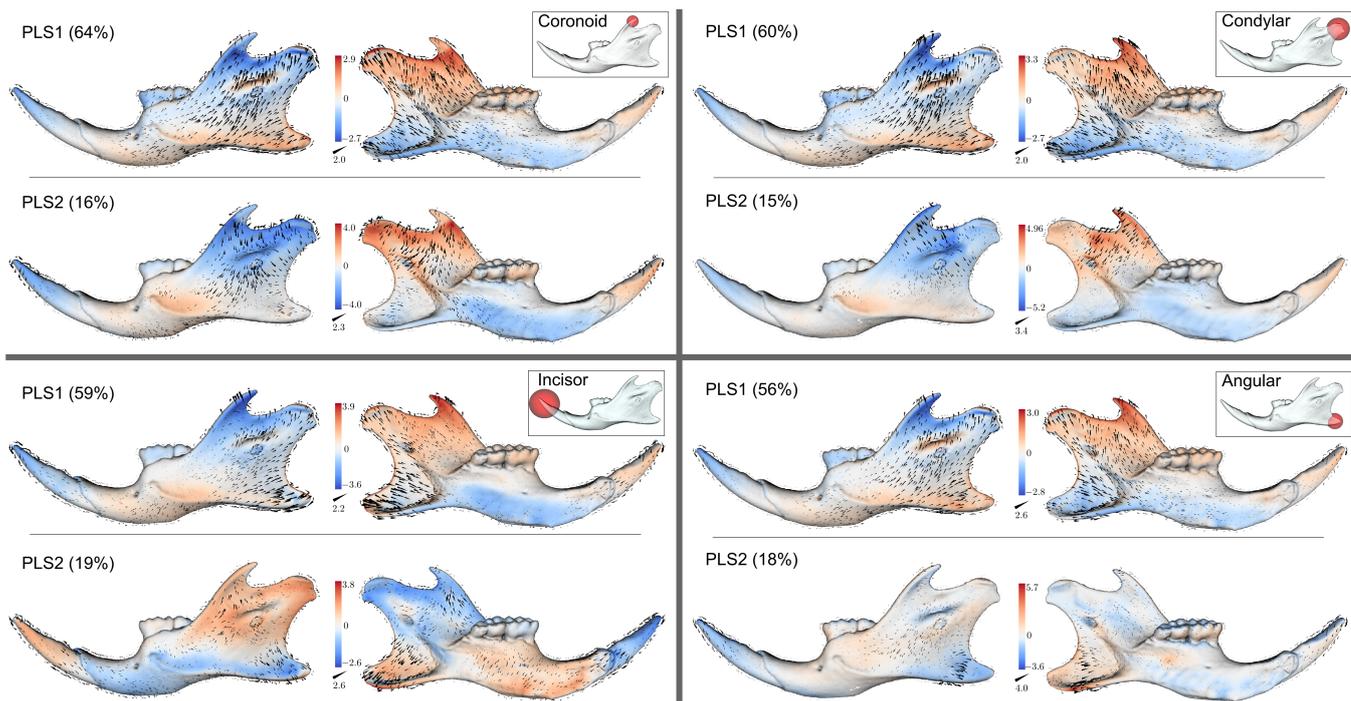


Figure 6.13: PLS analysis on *Mus* dataset with four different segmentations into white and red blocks as indicated in the insets, visualization design as in Fig. 6.11. Note the similarity of PLS1 for the rear processes which resembles the first mode PC1 (up to sign) of the global analysis. The variation patterns of PLS2 shows strong covariation between coronoid and condylar processes (top row) while the incisor affects more the angular process (bottom row).

6.6.2 Results on *Mus* dataset

In this section we validate our approach by reproducing hypotheses on module delimitations and interactions from literature [90, 160, 34, 113]. Additionally new observations of our analysis are described, pointing to a finer-scaled hypothesis on module segmentation. Our analysis follows the pipeline illustrated in Fig. 6.1, going back-and-forth between the different views.

Starting from the overview shown in Fig. 6.5 one can identify several candidate regions. Among them, in agreement with literature [90], the rear processes and tip of the incisor (visible in frontal view) exhibit strong impact on covariation, as indicated by glyph sizes. Additionally the posterior end of the incisor inside the mandible also shows up prominently. The latter could not be observed in previous approaches based on landmark data and focusing on the outline of mandible shape; therefore it is excluded in the following as it can not serve for validation.

A main advantage of our approach is that one can assess covariation at the level of individual displacements. This enables to distinguish directional dependencies in covariation as for instance found at the coronoid process. The overview suggests at least two different covariation patterns: Principal direction of the corresponding glyph at (A) in Fig. 6.5 predicts a stronger response on variation of length (vertical) versus position (horizontal) of the process. The specific shape variations are quite different as an investigation in detail view shows, illustrated in Fig. 6.2. Editing the coronoid process also reveals a coupled interaction to the adjacent condylar process. The existence of an interaction between these processes was known already [160] and is confirmed by PLS analysis, but the specific pattern of covariation visualized in focus view in Figs. 6.14 and 6.15 provides a new level of detail not seen previously.

Looking further at the condylar process in the overview suggests a separation between a tip part and a proximal part, closer to the mandible center. While the pattern near the tip is quite homogeneous, the proximal one is more diverse. In fact, editing the condylar process shows different reactions depending on the exact position of the probe on the process. This suggests a more fine-scaled analysis, proposed also in recent work [34, 113], with the here presented methods.

Examining incisor and angular process in focus view provides good examples of two extremes of global and local patterns, visible in Fig. 6.4. The angular probe is locally concentrated but still exhibits highly directed ($\lambda_1 \gg \lambda_2 \approx \lambda_3 \approx 0$) interaction with the upper two processes. The strength of interactions for incisor and angular probes again reproduce PLS findings. Since the complete incisor structure spans throughout the mandible, an effect on the rear processes is expected

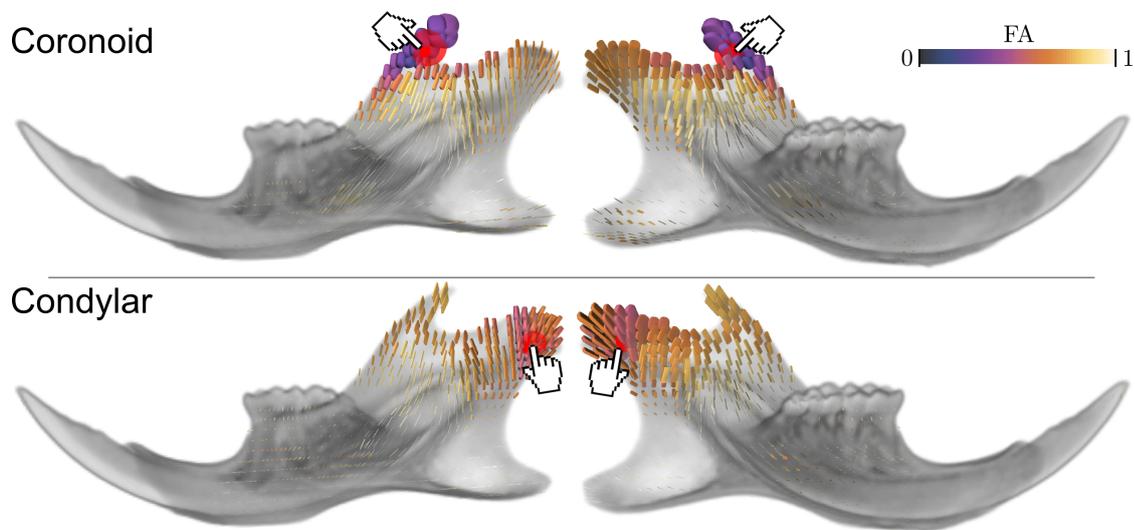


Figure 6.14: Local covariance fields for probes at coronoid and condylar processes, probing points highlighted in red. A strong interaction in-between the two processes becomes apparent and shows the diminishing impact on angular process and incisor. This finding is in agreement with PLS analysis and results of Zelditch et al. [160].

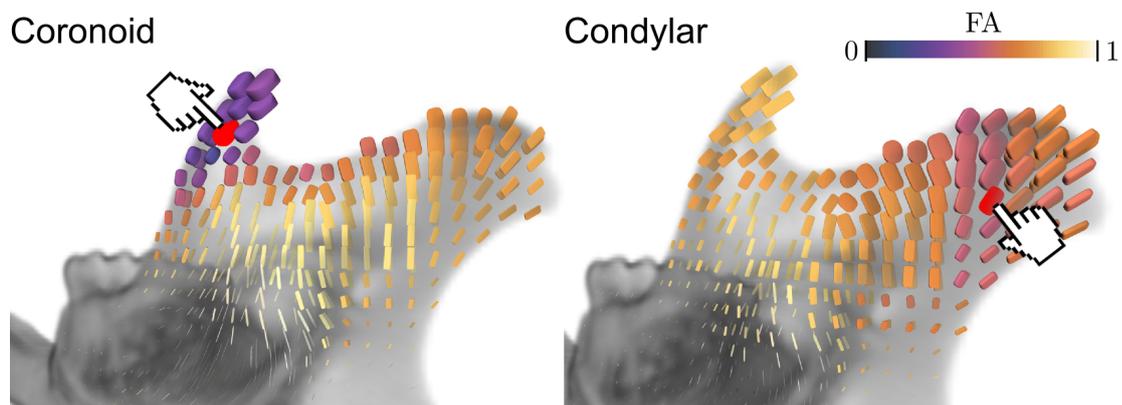


Figure 6.15: Close-ups of local covariance fields for coronoid and condylar processes show their consistent interaction. While the local region close to each probe (highlighted in red) exhibits more isotropic covariation, the reaction pattern at the other process is more directed. This reveals which part of the anatomic covariance field can be attributed to this particular interaction between the two processes.

on variation in the incisor tip. Editing the tip of the incisor as shown in Fig. 6.16 reveals further that the associated variation can be decomposed into two parts, depending roughly on the principal axes of the local tensor near the incisor tip. Pulling in directions of the principal axis (a),(b) inflicts a deformation at the upper coronoid and the lower angular process, while pulling orthogonal (c) to the axis has associated variation at the central condylar process and the rear part of the mandible. This is another example of directional dependencies, which can not be explored with previous methods.

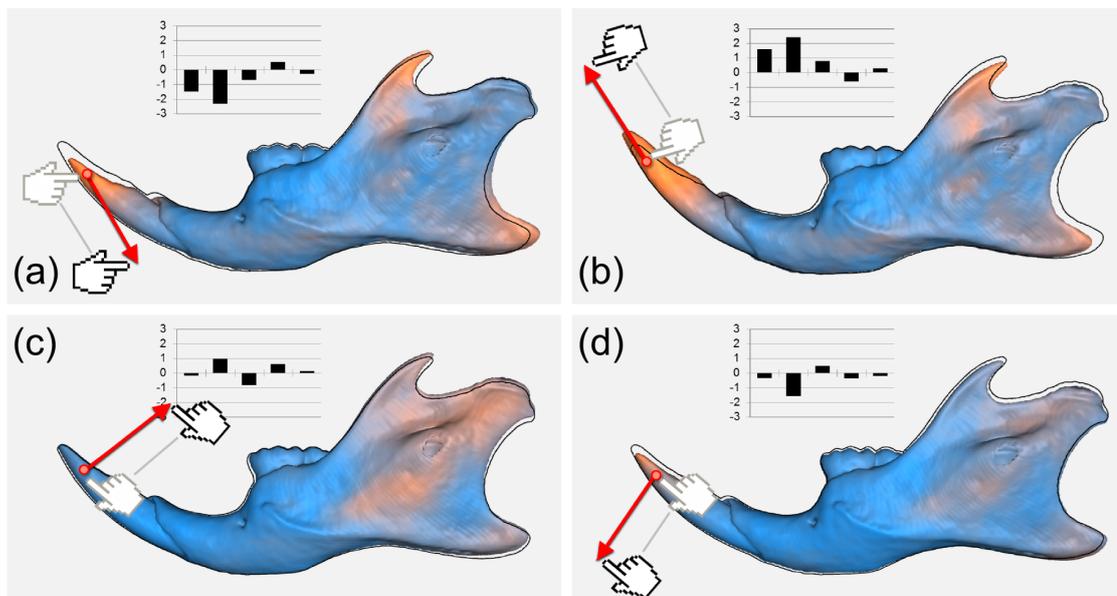


Figure 6.16: Model-based editing on *Mus* dataset at incisor, same legend as in Fig. 6.2. Pulling the incisor in directions (a) and (b) leads to the expected elongation / foreshortening, seemingly correlated with the shape of the tips of the rear processes. Dragging the incisor orthogonally (c)-(d) keeps the front nearly rigid showing minor interaction in the posterior part.

The anisotropy pattern in the focus view in Fig. 6.15 also nicely illustrates some properties of our local tensor, setting it apart from the anatomic covariance field of Kindlmann et al. [81]. Anisotropy near the probe is expected to be similar to that of the probe, indicating strong local covariation (typical for stiff bone structures). That anisotropy is higher in the process opposite of the probe indicates a structured covariation. Note that these patterns are distinct from the anatomic covariation field shown in the inset in Fig. 6.12. From this we conclude that our system allows one to focus on particular components of the global pattern attributed to specific inter-part interactions.

6.7 Visual analysis of group differences in *Cricetinae*

This section describes a trait and covariance analysis conducted on the *Cricetinae* dataset. This is done to illustrate a combined use of the interactive model-based editing presented in this chapter with the classification method from Chap. 5. The investigated dataset consists of nine mandibles of two hamster species, see App. A.

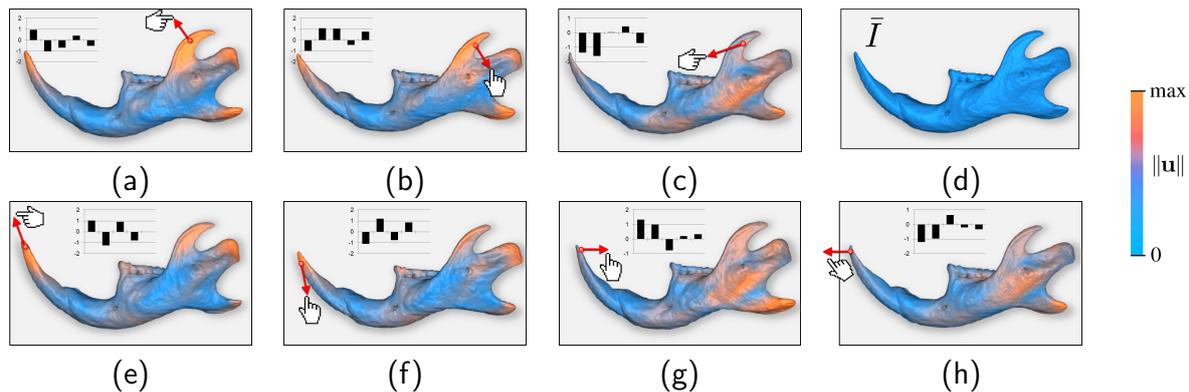


Figure 6.17: Rubber band edits on *Cricetinae* dataset at coronoid process (a)-(c) and incisor (e)-(h). Color coded is the magnitude of the displacement from the template (d). Inset bar plots describe the first 5 PC coefficients in units of standard deviations. Edits (a)-(b) show clearly a covariation of the rear processes, while (c) produces a more diffuse pattern. Pulling the incisor in directions (e) and (f) leads to the expected elongation / foreshortening, seemingly correlated with the shape of the tips of the rear processes. Dragging the incisor vertically (g)-(h) keeps the front of the incisor nearly rigid but inflicts changes in the complete posterior part of the mandible.

Fig. 6.17 depicts an interactive editing session investigating two prominent structures on the mandible. Edits (a) to (c) focus on the coronoid process, while (e) to (f) manipulate the incisor. The investigation reveals interesting covariation between the two groups in the dataset and gives an impression on how to interpret results of the rubber band method.

Coronoid process

Pulling the rubber band orthogonal to the contour of the process in (a) and (b) leads the contour to follow the edit direction. This induces a shape change on the process itself from *Cricetus*-like to *Cricetulus*-like. Globally, a covariation in the tips of the other two processes and the incisor can be observed. The behavior drastically changes when pulling for instance in direction (c). Here, the process responds more rigidly and provides a totally different pattern of covariation. Note that the more pronounced rigidity can be recognized in the small deviation from

the mean shape (d), despite the longer stretching of the rubber band than in (a) and (b). A presumption following from these findings is, that (a) and (b) describe shape differences between the two sub-groups of the dataset, while (c) does not. This is supported by the PCA plot shown in Fig. 6.18, where (a) and (b) are close to the barycenter of each group. The green line in the PCA plot shows a separating hyperplane between the groups, computed with a support vector machine as described in Chap. 5. The angle between the hyperplane and linear trajectories from the mean shape at (0, 0) to each of the edits shows that the trajectories to (a) and (b) will run nearly orthogonal to the hyperplane, but almost parallel towards (c). This suggests that (c) does not carry shape covariation relevant to discriminate between the two groups, at least not in the first two components.

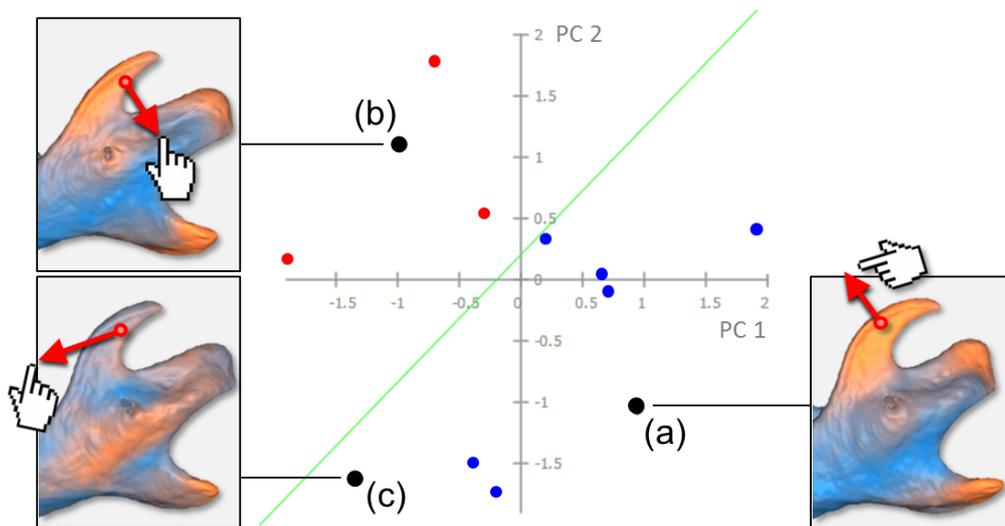


Figure 6.18: PCA plot of Cricetinae dataset and edits (a)-(c) from Fig. 6.17 (black dots) onto first two components, where individuals are color coded by their corresponding sub-groups *Cricetus* (red) and *Cricetulus* (blue), being clearly separated as illustrated (green line). Edits (a) and (b) navigate between group-wise shape characteristics at the rear processes, while results of an edit (c) orthogonal to (a) and (b) stays indifferent.

Incisor

Editing the tip of the incisor in directions (e) and (f) produces covariations similar to (a) and (b). Looking at the PC coefficients in the insets, we can see that the similarity is mainly due to a congruent behavior of the first two principal components, while the third and fourth mode vary oppositely. Note that the tip of the incisor follows nicely along with the edits, making it longer and shorter. This tells us, that this is a typical, in the sense of likely, shape variation in the hamster dataset. Trying to change the position vertically in (g) and (h) induces

much less deformation on the front of the incisor, although nearly the same amount of “force”, i.e. magnitude of the edit, was used. Interestingly, a strong covariation can be observed in the central region, between the processes. This effect can be understood from the anatomy of the incisor, whose bone structure describes an arc throughout the mandible, starting from the tip of the incisor and ending exactly in the just observed covarying region between the rear processes. This shows that also long-range interactions on far-away parts on the mandible are captured in our model.

6.8 Conclusion

Starting from an existing method for shape animation and reconstruction, we developed in this chapter a complete visual analysis pipeline for effective exploration of shape covariation in statistical deformation models. To this end a new local covariance tensor was derived and integrated, yielding two novel visualizations accompanying the model-based deformation of Blanz et al. [24]. Together, all three methods complement each other in a visual analysis of interactions between different parts of the shape at varying degrees of detail. In addition, a novel automatic segmentation procedure could be derived based on the overview tensor, that decomposes the shape into anatomic meaningful parts.

Exploratory analysis on a scientific dataset was carried out on the backdrop of morphological integration, illustrating the utility of the visual analysis in the search for module boundaries. Comparison to state-of-the-art techniques, namely PLS, PCA and anatomic covariance tensor, showed that the presented approach enables an exploration at a finer level of detail including for the first time directional dependencies. Reproducing several hypotheses from recent works in morphometry can be taken as sanity check here.

Further applications in morphometric studies are expected a) on datasets with higher diversity consisting of different species and b) on more complex structures like the rodent skull. It would also be interesting to see how effective the approach is in uncovering other, non-anatomic sources of shape variation like (structured) registration or reconstruction errors. Realtime sampling and interactive filtering of local covariance fields is subject to future work.

Part III
Closing

Chapter 7

Conclusion

The main challenge to be overcome for multivariate morphometric studies concerns visualizing and communicating the findings, which is probably why so many geometric morphometric studies of the brain have focused on simple two dimensional structures such as the corpus callosum.

— John Ashburner and Stefan Klöppel [16]

7.1 Summary

In this thesis we applied a visual analytics approach to study shape variability in biomedical image ensembles in relation to extrinsic as well as intrinsic attributes. To this end a portfolio of interactive visualization and analysis methods along the lines of Busking et al. [37] and Klemm et al. [85] was introduced that expand the state-of-the-art in at least three respects:

- So far, interactive shape analysis was restricted to landmark and surface models. The methods developed in this thesis show efficient ways to operate on much denser deformation models that describe shape variation at image resolution.
- The previously available set of exploration methods is extended by interactive classification, weighted local analysis, a first direct manipulation approach and novel tensor visualizations.

- An efficient and accurate visualization of non-linear deformations is facilitated by a GPU raycasting algorithm.

The very unique collection of high quality CT datasets of rodent skulls contributed by Dr. Schunke provided an ideal testbed for our methods. The continuous feedback from a morphometrics expert was invaluable during development and evaluation of the novel visual analytics methods. Together with Dr. Schunke, exemplary investigations on the influence of phylogeny, diet and geography on the form of rodent skull and mandible were conducted that demonstrate the potential of our visual shape analytics methods.

7.2 Future prospects

Visual analytics perspective

Visual analytics methods have found application in nearly every domain that requires the analysis of large and high dimensional data sets, ranging from financial market to climate research. This thesis is among the first works that apply the visual analytics paradigm to problems in shape analysis and morphometrics. Therefore we concentrated on central problems that are of general interest to many practitioners in the field, like unconstrained navigation and targeted analysis with respect to specific factors. Of course, there is additional potential in providing custom tailored linked views based on a geographic map, a phylogenetic tree or a Manhattan plot.

We believe that the previous works as well as what is presented in this thesis represents a “critical mass” of visual shape analytics methods that calls for a comparative study. Such a study can serve practitioners as a guide to the best method for a particular problem and help the visual analytics researchers to uncover potential shortcomings that require additional attention in future research. Of course, the design of a such a study is a challenge on its own. A possible approach could be to somehow measure the performance of experts in discovering shape variations and correlations in a data set that was purposefully designed after representative morphometric tasks.

Another noteworthy point is that visual shape analytics is often implemented in tight cooperation between a computer scientist and a domain expert. This makes it an interesting instance for *pair analytics* [11] in order to understand the scientific reasoning process in morphometry.

In the future we hope to see applications of the presented methods in studies in morphometrics and computational anatomy. Especially population studies could provide an ideal application domain because of their exploratory nature [31, 85]. However, the ensemble sizes treated in these studies is at the order of several hundreds of individuals while the largest dataset treated with our methods had less than 50 specimen. To this end scalability of the presented methods has to be considered, either improving the current implementation, e.g. by further parallelization or distributed computing, or by investigating other ingenious means to represent the variability of larger ensembles by fewer representatives.

Shape analysis perspective

From a technical point of view we see several directions for future work. Although our methods were developed to work with image data and statistical deformation models, transferring them to landmark and surface data and corresponding PCA based statistical models is rather straight forward. Nevertheless, adaption to other successful shape representations and statistical models beyond that is presumably challenging and would broaden the class of anatomy ensembles our methods can be applied to. Medial representations [62] for instance provide another very powerful non-linear statistical model, namely principal geodesic analysis [61]. A further challenging example is the recent model of Durrleman et al. [56] that describes dense deformations with sparse parameters and can thereby also handle varying topologies, including cases that do not allow a perfect registration.

A recently very active topic are hierarchical shape models that allow investigation of shape variation at multiple scales. There exist several promising approaches utilizing different decompositions of shape variation, either based on wavelet theory [51] (particularly popular in medical image analysis [154, 157, 57]), sparse PCA [127], polyaffine transformation tree [123] or deflation of principal warps [28]. However, effective means of navigating such complex multiscale representations have only sparsely been addressed so far, partly because the complexity of some of the methods rules out an interactive approach. At least two of the methods developed in this thesis could provide promising starting points for navigation and efficient implementation of hierarchical models: The local analysis of a region of interest, described in Chap. 5, allows already for a manual navigation at different scales, while the tensor based segmentation, introduced in Chap. 6, identifies anatomic meaningful regions automatically.

The methods of this thesis were designed to investigate shape variability based on statistical shape models. A crucial step in establishing such models is to define

correspondences across all shapes by means of a registration algorithm. We believe that some of our visualizations could be adapted to investigate the quality of registration and that of the resulting statistical model. For instance do the projected streamlines reveal tangential parts of deformation that is otherwise also used as an indicator for mis-registration, as excessive tangential drift increases the entropy of the statistical model unnecessarily [52].

Appendix A

Data sets and preprocessing

A.1 Acknowledgements

All rodent skull data sets used throughout this thesis are courtesy of Dr. Anja C. Schunke, who assembled and digitized this truly unique collection at the Max Planck Institute for evolutionary biology, Plön. The original specimens were kindly loaned by Dr. Rainer Hutterer, Zoological Research Museum Alexander Koenig, Bonn. Preprocessing of the CT images was supported by Vitalis Wiens and Krishna P. Soundararajan that were employed as student researchers under a grant by Deutsche Forschungsgesellschaft (DFG) within the priority program SPP1335 Scalable Visual Analytics.

A.2 Acquisition

Each individual was scanned by Dr. Schunke using a Scanco VivaCT-40 μ CT scanner located at the Max Planck Institute for evolutionary biology, Plön.

A.3 Data sets used in this thesis

A list of the data sets used in this thesis is given in Table A.1. Some of the data sets are illustrated in Figs. A.1 and A.2.

Dataset		n	Description
<i>Apodemus Flavicollis</i>	(S)	22	Samples of <i>Apodemus Flavicollis</i> throughout Europe. (Chap. 4)
<i>Cricetinae</i>	(M)	9	Samples from two hamster species, 6 <i>Cricetus cricetus</i> (common hamster), 3 <i>Cricetulus migratorius</i> (dwarf hamster). (Chap. 6)
Diet and phylogeny	(M)	16	Representatives of two diets (8 omnivorous, 8 carnivorous) and phylogeny (8 diet/genus pairs), averaged over a sample of 48 specimens. (Chap. 4)
Gerbillinae-Murinae	(M)	29	Samples from two subfamilies, 22 Murinae, 7 Gerbillinae. (Chap. 5)
<i>Mus</i>	(M)	30	Samples from single species <i>Mus musculus</i> (house mice). (Chap. 6)

Table A.1: Data sets used in this work. Legend: (M)=Mandible, (S)=Upper skull, n =number of input images to PCA analysis (not necessarily the number of individuals, see text).

A.4 Preprocessing

All data sets are semi-automatically segmented into three parts, the upper skull and two mandibles [150]. For all analyses except the dataset on phylogeny and diet (see below), only a single mandible was considered for analysis, out of symmetry reasons. We chose canonically the left one if it was available. If the left mandible specimen was missing or damaged, it was replaced with its mirrored right counterpart where possible. Differences due to translation, rotation and scale are factored out via an image based similarity alignment. The alignment is optimized based on L2 intensity error on histogram equalized images using the elastix toolbox [84]. Elastic registration is performed with the symmetric log-domain diffeomorphic demons algorithm [142]. Group-wise registration was initialized with a typical representative individual chosen manually by Dr. Schunke. For the data sets at hand, three to five iterations of the registration algorithm presented in Chap. 3 were sufficient to reach convergence.

Prior to registration, input images were resampled to a resolution of $200 \times 200 \times 400$ isotropic voxels for all data sets except the dataset Gerbillinae-Murinae that was resampled to $135 \times 173 \times 280$ voxels. The resulting deformation fields, irregardless if they represent displacements or velocities, were subsequently resampled by a factor of four to $50 \times 50 \times 100$ and a factor of two to $68 \times 87 \times 140$ in case of the Gerbillinae-Murinae dataset. The downsampling factor for deformation fields was estimated empirically on several examples in such a way, that PCA on the resampled fields resulted in the same coefficient as for the original resolution and no visual difference between PCA reconstructions from PCA models, set up from original and resampled images, could be perceived. See also the discussion in Chap. 4 on downsampling deformation fields. Note that for visualization, resampled deformations are interpolated and applied to a full resolution template image.

Special treatment of dataset on phylogeny and diet

The dataset on phylogeny and diet that is analyzed in Chap. 4 received some additional preprocessing steps. This particular dataset is compiled from 48 specimen in order to investigate and compare influence of diet and phylogeny (on a genus level) on mandible shape. To alleviate a bias due to different number of samples in the phylogenetic groups, one representative mandible is considered per genus, averaged over all of its specimens. Left and right mandible of each specimen were averaged in advance to remove asymmetric effects. Eventually, 16 representatives are analyzed, 8 omnivorous and 8 carnivorous. The closest omni- and carnivore relatives supply 8 diet/genus pairs.

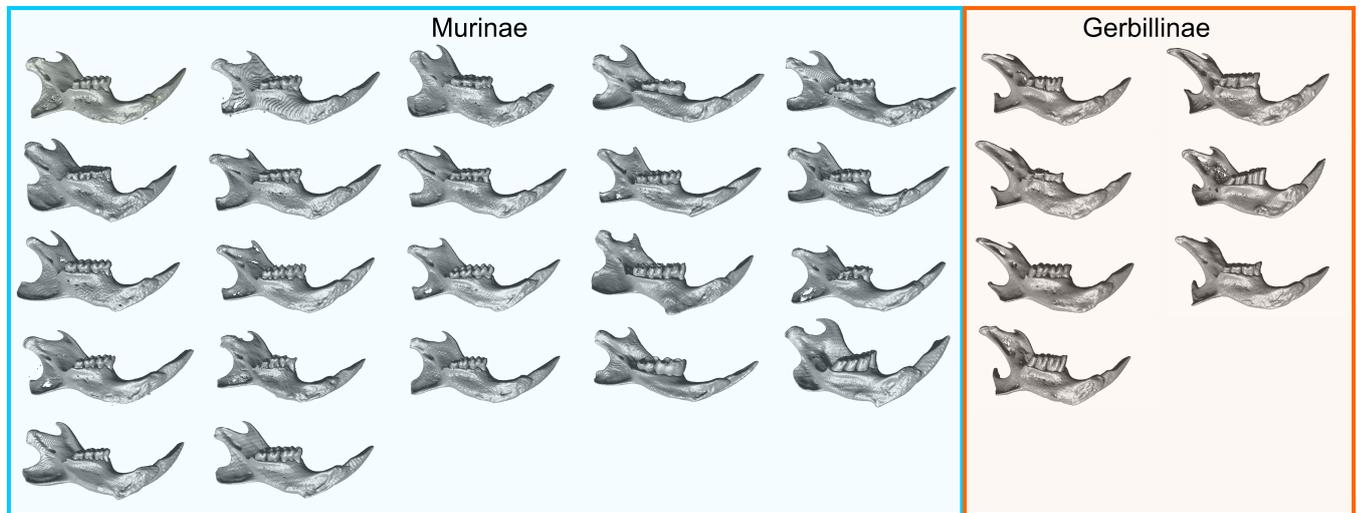


Figure A.1: Dataset **Gerbillinae-Murinae** consists of 22 Murinae and 7 Gerbillinae mandibles. The shown isosurfaces illustrate the shape variability present in this dataset.

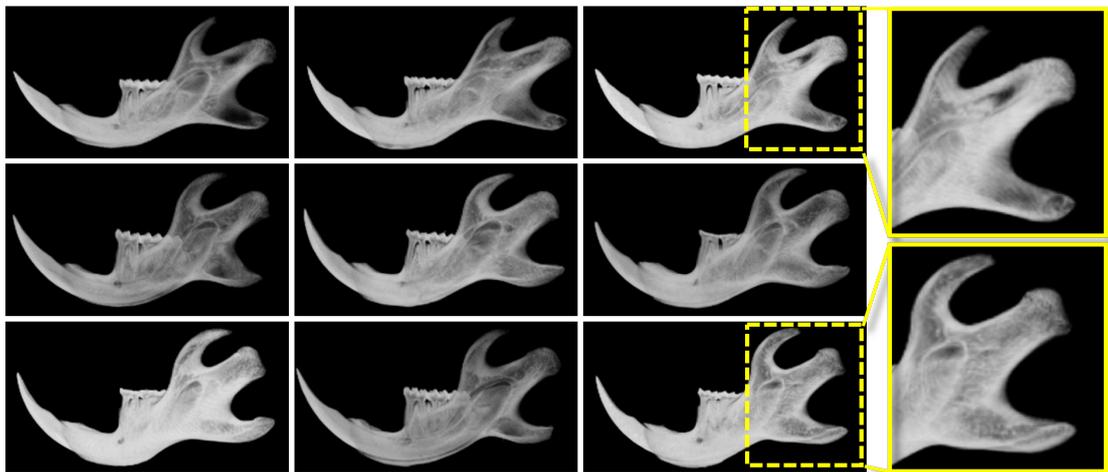


Figure A.2: Dataset **Cricetinae** consists of 9 individuals from the two sub-groups *Cricetulus* (first row) and *Cricetus* (second and third row). Shape characteristics at the rear processes are highlighted.

Appendix B

Cross-validation of regularization parameter for model-based deformation

B.1 Cross-validation algorithm

The model-based editing in Chap. 6 requires the minimization of the energy (6.3), that is repeated here for convenience:

$$E(\mathbf{c}) = \frac{1}{2} \|\mathbf{u}_p - \mathbf{B}_p \mathbf{c}\|_2^2 + \frac{\gamma}{2} \|\mathbf{c}\|_2^2 \quad (\text{B.1})$$

Only due to Tikhonov regularization with a parameter γ has this energy a well defined unique minimum. Therefore, for the methods described in Chap. 6 to work, it is important to reason about the choice of γ . Fortunately, a suitable value for γ can be found automatically via a leave-one-out cross-validation procedure that is given in algorithm B.1. For each value γ it computes the average reconstruction error for a left-out displacement field when expressed in the PCA model according to Eq. (6.3), where the PCA model is set up over the remaining displacement fields.

Using the same notation as in the algorithm, the ordinary cross-validation function that is minimized can be stated as

$$V_0(\gamma) = \frac{1}{n} \sum_{r=1}^n \left\{ \frac{1}{|\Omega|} \sum_{p \in \Omega} \|\mathbf{u}_r - \mathbf{B}_p^{(r)} \mathbf{c}_{\text{opt}}^{(r)}(\gamma) + \bar{\mathbf{u}}^{(r)}\| \right\}.$$

Note that here the reconstruction error of the *complete* vectorfield \mathbf{u}_r is measured

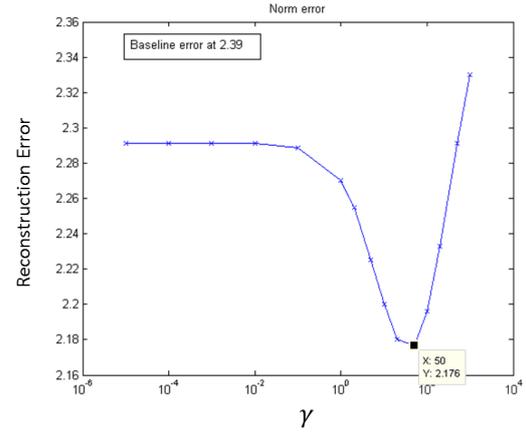


Figure B.1: Cross-validation result for the *Mus* dataset with an optimum at $\gamma = 50$.

and not only at the displacement vector $\mathbf{u}_{r,p}$ at position p , else wise V_0 would trivially achieve its minimum for $\gamma = 0$.

Algorithm B.1 Cross-validate γ parameter

Input: Set of displacement fields $[\mathbf{u}_1, \dots, \mathbf{u}_n] \in \mathbb{R}^{3|\Omega| \times n}$

Output: Sampling of ordinary cross-validation function $V_0(\gamma)$

- 1: **for** each left out datum \mathbf{u}_r **do**
 - 2: $\bar{\mathbf{u}}^{(r)} \leftarrow \frac{1}{n-1} \sum_{i \neq r} \mathbf{u}_i$
 - 3: $\mathbf{X}^{(r)} \leftarrow [\mathbf{u}_1 - \bar{\mathbf{u}}^{(r)}, \dots, \mathbf{u}_{r-1} - \bar{\mathbf{u}}^{(r)}, \mathbf{u}_{r+1} - \bar{\mathbf{u}}^{(r)}, \dots, \mathbf{u}_n - \bar{\mathbf{u}}^{(r)}]$
 - 4: Compute eigenvectors $\mathbf{B}^{(r)}$ of $(\mathbf{X}^{(r)})^T \mathbf{X}^{(r)}$
 - 5: **for** each $\gamma \in \{\gamma_1, \dots, \gamma_k\}$ **do** \triangleright *Sampling of γ , e.g. perform a line search*
 - 6: **for** each voxel position $p \in \Omega$ **do**
 - 7: $\mathbf{Z}_p^{(r)}(\gamma) \leftarrow \left((\mathbf{B}_p^{(r)})^T \mathbf{B}_p^{(r)} + \gamma \mathbf{I} \right)^{-1} (\mathbf{B}_p^{(r)})^T$ \triangleright *Depends on γ*
 - 8: $\mathbf{c}_{\text{opt}}^{(r)}(\gamma) \leftarrow \mathbf{Z}_p^{(r)}(\gamma) (\mathbf{u}_{r,p} - \bar{\mathbf{u}}_p^{(r)})$
 - 9: $e_{r,\gamma}(p) \leftarrow \|\mathbf{u}_r - \mathbf{B}^{(r)} \mathbf{c}_{\text{opt}}^{(r)}(\gamma) + \bar{\mathbf{u}}^{(r)}\|$ \triangleright *Reconstruction error*
 - 10: **end for**
 - 11: $\bar{e}_\gamma(r) \leftarrow \frac{1}{|\Omega|} \sum_{p \in \Omega} e_{r,\gamma}(p)$
 - 12: **end for**
 - 13: $V_0(\gamma) \leftarrow \frac{1}{n} \sum_{r=1}^n \bar{e}_\gamma(r)$
 - 14: **end for**
-

B.2 Smoothness of cross-validation function

Noteworthy, the cross-validation function turned out to be smooth with a unique minimum for all considered data sets. Therefore, a robust choice for the parameter γ could be made automatically in all cases. See Fig. B.1 for an example.

Appendix C

Proof of error estimate for heuristic inverse

C.1 Linear approximation of inverse

We consider a continuous mapping between open sets $\Omega \subset \mathbb{R}^d$ where usually $d = 2, 3$ for the application of (interactive) image warping. The mapping $\varphi: \Omega \rightarrow \Omega$ is realized as

$$\varphi(x) = x + u(x) \tag{C.1}$$

using a displacement vector field $u: \mathbb{R}^d \rightarrow \mathbb{R}^d$.

A computationally very efficient approximation of the inverse mapping φ^{-1} is to consider a first order linearization $\tilde{\varphi}^{-1} \approx \varphi^{-1}$ that simply negates the displacement field, i.e.

$$\tilde{\varphi}^{-1}(x) = x - u(x). \tag{C.2}$$

Theorem

For a maximum displacement magnitude η in d the above approximation has an error in the order of $O(\gamma^2)$, i.e. (C.2) is in fact a linearization and it holds

$$\tilde{\varphi}^{-1}(\varphi(x)) = x + O(\gamma^2). \tag{C.3}$$

Proof

Applying the definition results in

$$\begin{aligned}\tilde{\varphi}^{-1}(\varphi(x)) &= \varphi(x) - u(\varphi(x)) \\ &= x + u(x) - u(x + u(x)).\end{aligned}$$

A Taylor expansion of the last term yields

$$\begin{aligned}\tilde{\varphi}^{-1}(\varphi(x)) &= x + u(x) - (u(x) + \nabla u(x)u(x) + O(|u(x)|^2)) \\ &= x - \nabla u(x)u(x) + O(|u(x)|^2).\end{aligned}$$

where $\nabla u(x)$ denotes the Jacobi matrix of $u(x)$ at x .

By definition we know that the maximum displacement has a length of γ , i.e. $|u(x)| \leq \gamma$. This allows us to rewrite our displacement field u using an auxiliary field g as

$$u(x) = \gamma g(x).$$

It follows that g has the property $|\nabla g(x)g(x)| \leq 1$. Expanding this into the Taylor expansion gives an upper bound on the approximation error and concludes the proof:

$$\begin{aligned}\tilde{\varphi}^{-1}(\varphi(x)) &\leq x - \gamma^2 \nabla g(x)g(x) + o(\gamma^2) \\ &= x + o(\gamma^2)\end{aligned}$$

□

Appendix D

Mathematical notation & acronyms

D.1 Mathematical notation

Basics

- $\mathbb{R}, \mathbb{R}^d,$
 $\mathbb{R}^{d_1 \times d_2}$ Set of real numbers and its associated d and $d_1 \times d_2$ dimensional tuple sets, from which scalars, vectors and matrices are drawn respectively.
- $\|\cdot\|$ Standard L^2 norm on \mathbb{R}^d , i.e. $\|(a_1, \dots, a_d)^T\| := \sqrt{a_1^2 + \dots + a_d^2}$ if not noted otherwise.

3D vector and matrix algebra

- \cdot Scalar product, e.g. $\begin{pmatrix} a_x \\ a_y \\ a_z \end{pmatrix} \cdot \begin{pmatrix} b_x \\ b_y \\ b_z \end{pmatrix} = a_x b_x + a_y b_y + a_z b_z$.
- ∇ First order differential operator, e.g. $\nabla = (\frac{\partial}{\partial x}, \frac{\partial}{\partial y}, \frac{\partial}{\partial z})$.
- A Linear map $A: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ in matrix representation $A \in \mathbb{R}^{3 \times 3}$.
- Ax Matrix vector product that applies linear mapping represented by A to a vector x .
- I Identity matrix $I = \begin{pmatrix} 1 & & \\ & 1 & \\ & & 1 \end{pmatrix}$.
- Λ Diagonal matrix $\Lambda = \begin{pmatrix} \lambda_1 & & \\ & \lambda_2 & \\ & & \lambda_3 \end{pmatrix}$ and for uniform scaling $\lambda_1 = \lambda_2 = \lambda_3$.
- R Rotation matrix $R \in \text{SO}(3) \subset \mathbb{R}^{3 \times 3}$, i.e. $RR^T = I$ and $\det(R) = +1$.
- t Translation vector $t \in \mathbb{R}^3$ used in combination with a linear map $A \in \mathbb{R}^{3 \times 3}$ to represent a rigid, similarity or fully affine transformation.

Dimensions

n	Number of images in an ensemble.
N	Number of voxels in (template) image.
n'	Rank of sample covariance matrix and dimensionality of shape space as represented by a statistical deformation model.

Images

Ω	Compact rectangular image domain $\Omega \subset \mathbb{R}^3$.
x	3D position $x \in \mathbb{R}^3$ as column vector, usually a position in an image domain Ω that is treated either as continuous or voxel position, depending on context.
p, q	Two 3D positions in image domain Ω that relate to each other, used specifically in analysis of covariance between two different points in Chap. 6.
\mathcal{I}	Input image ensemble as a set of n input images $\mathcal{I} = \{I_1, \dots, I_n\}$.
I, I_i	3D intensity image $I: \Omega \rightarrow \mathbb{R}$ parameterized over a compact domain $\Omega \subset \mathbb{R}^3$, discretized over a rectilinear grid into voxels. Indexed version I_i represent individual images from the image ensemble \mathcal{I} .
I'_i	Aligned image, i.e. input image I_i after global registration of image ensemble.
I^*	A synthesized or otherwise deformed image.
\bar{I}	Template image that represents an ensemble average.

Deformations

φ	Transformation/deformation that maps the image domain $\Omega \subset \mathbb{R}^3$ onto itself, i.e. $\varphi: \Omega \rightarrow \Omega$, usually parameterized over 3D positions $x \in \Omega$ that are either treated continuously or discretely at voxel positions, depending on context.
$\varphi_1 \circ \varphi_2$	Concatenation of mappings, i.e. $\varphi_1 \circ \varphi_2(x) = \varphi_1(\varphi_2(x))$.
Id	Identity mapping $\text{Id}: \Omega \rightarrow \Omega$ with $\text{Id}(x) := x$.
φ^{-1}	Inverse mapping to φ defined by $\varphi \circ \varphi^{-1} = \text{Id}$.
u, \mathbf{u}	Displacement vector field $u: \Omega \rightarrow \mathbb{R}^3$ that is used in the general representation of a deformation $\varphi(x) = x + u(x)$. Bold face identifies an encoding of the field as single long column vector $\mathbf{u} \in \mathbb{R}^{3N}$.
v, \mathbf{v}	Velocity field $v: \Omega \rightarrow \mathbb{R}^3$ that is assumed to be sufficiently smooth to generate a diffeomorphic deformation when integrated via $\exp(v)$. Bold face identifies an encoding of the field as single long column vector $\mathbf{v} \in \mathbb{R}^{3N}$.
$\exp(\cdot)$	Either denotes a matrix exponential, see $\exp(A)$, or an exponential map of a stationary velocity field, see $\exp(v)$.
$\log(\cdot)$	Either denotes a matrix logarithm, see $\log(A)$, or a principal logarithm of a diffeomorphic deformation, see $\log(u)$.

$\exp(v)$	Integration of a stationary velocity field v whose result is a diffeomorphic deformation. For interactive rendering integration is computed via standard numerical schemes like Runge-Kutta, otherwise a scaling and squaring approach [13] is used.
$\log(u)$	Stationary velocity field as an infinitesimal generator of a diffeomorphic deformation represented by a displacement field u . It is computed via an inverse scaling and squaring algorithm [13] and can be interpreted as a kind of principal logarithm of the diffeomorphism.
$\exp(A)$	Matrix exponential [70].
$\log(A)$	Matrix logarithm [70].

Statistical deformation model

\mathbf{X}	Data matrix of displacement vector fields $\mathbf{X} = [\mathbf{u}_1, \dots, \mathbf{u}_n] \in \mathbb{R}^{3N \times n}$. Usually \mathbf{X} is considered to be centered, i.e. of zero column mean $\bar{\mathbf{u}} = \frac{1}{n} \sum_{i=1}^n \mathbf{u}_i = \mathbf{0}$.
Σ	Sample covariance matrix, for a centered data matrix $\Sigma = \frac{1}{n-1} \mathbf{X}\mathbf{X}^T$.
\mathbf{B}	Basis of linear model $\mathbf{B} \in \mathbb{R}^{3N \times n'}$ with principal modes of variation encoded in its columns.
\mathbf{B}_p	Rows of \mathbf{B} corresponding to displacements at voxel $p \in \Omega$, i.e. $\mathbf{B}_p \in \mathbb{R}^{3 \times n'}$.
\mathbf{c}, \mathbf{C}	Coefficient vector $\mathbf{c} \in \mathbb{R}^{n'}$ used for reconstruction/synthesis of displacements $\mathbf{u} = \mathbf{B}\mathbf{c}$ and coefficient matrix $\mathbf{C} = [\mathbf{c}_1, \dots, \mathbf{c}_n]$ such that $\mathbf{X} = \mathbf{B}\mathbf{C}$ for a centered data matrix.
\mathbf{T}	Tensor field $\mathbf{T}: \Omega \rightarrow \mathbb{R}^{3 \times 3}$ assigning a 3×3 positive semidefinite matrix $\mathbf{T}(p)$ to each (sample) point in image domain Ω . Used for the global covariance tensor \mathbf{T}_{global} and the novel interpoint covariance tensor $\mathbf{T}_p(q)$ in Chap. 6.
\mathbf{Z}_{pq}	Interaction operator $\mathbf{Z}_{pq} \in \mathbb{R}^{3 \times 3}$ introduced in Chap. 6.
Γ	Overview tensor field, parameterized as \mathbf{T} above.

For the statistical model based on velocity fields that is introduced in Chap. 4 the according accented symbols $\hat{\mathbf{B}}$, $\hat{\mathbf{c}}$ and $\hat{\Sigma}$ are used for better disambiguation.

D.2 List of acronyms

Acronyms are introduced on their first occurrence in each chapter.

AAM	Active appearance model.
ASM	Active shape model.
CA	Computational anatomy.
CVA	Canonical variate analysis.
CPU	Central processing unit.
CT	Computed tomography.
DoF	Degrees of freedom.
FA	Fractional anisotropy.
FFD	Free form deformation.
FPS	Frames per second.
GM	Geometric morphometrics.
GPU	Graphics processing unit.
GUI	Graphical user interface.
HARDI	High angular resolution diffusion imaging.
LDDMM	Large displacement diffeomorphic metric mapping.
MDS	Multi-dimensional scaling.
MRI	Magnetic resonance imaging.
ODE	Ordinary differential equation.
PCA	Principal component analysis.
PC	Principal component.
PLS	Partial least squares.
RK4-2	Fourth order numerical Runge-Kutta integrator applied with two steps.
RMSE	Root mean square error.
ROI	Region of interest.
SDM	Statistical deformation model.
SVD	Singular value decomposition.
SVF	Stationary velocity field.
SVM	Support vector machine.
TPS	Thin plate spline.
VA	Visual analytics.

Author's publications related to this thesis

- [1] Max Hermann and Reinhard Klein. A visual analytics perspective on shape analysis: State of the art and future prospects. *Computers & Graphics*, 53, Part A:63–71, September 2015.
- [2] Max Hermann, Anja C. Schunke, Thomas Schultz, and Reinhard Klein. Accurate interactive visualization of large deformations and variability in biomedical image ensembles. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 22(1):708–717, 2016.
- [3] Max Hermann, Anja C. Schunke, and Reinhard Klein. Semantically steered visual analysis of highly detailed morphometric shape spaces. In *Proceedings of the 1st IEEE Symposium on Biological Data Visualization (BioVis 2011)*, pages 151–158, October 2011. (**Honorary mention**).
- [4] Max Hermann, Anja C. Schunke, Thomas Schultz, and Reinhard Klein. A visual analytics approach to study anatomic covariation. In *Proceedings of the IEEE Pacific Visualization Symposium (PacificVis)*, pages 161–168, March 2014.
- [5] Max Hermann, Anja C. Schunke, and Reinhard Klein. Anatomic segmentation of statistical shape models. Poster at Symposium on Statistical Shape Models & Applications (SHAPE2014) in Delémont, Switzerland, June 2014. (**Best poster award**).

Bibliography

- [6] ABBASLOO, A., WIENS, V., HERMANN, M., AND SCHULTZ, T. Visualizing tensor normal distributions at multiple levels of detail. *IEEE Trans. on Visualization and Computer Graphics* 22, 1 (2016), 975–984. (Cited on page 24)
- [7] ADAMS, D. C., ROHLF, F. J., AND SLICE, D. E. Geometric morphometrics: Ten years of progress following the revolution. *Italian Journal of Zoology* 71, 1 (2004), 5–16. (Cited on page 11)
- [8] ALCANTARA, D. A., CARMICHAEL, O., HARCOURT-SMITH, W., STERNER, K., FROST, S. R., DUTTON, R., THOMPSON, P., DELSON, E., AND AMENTA, N. Exploration of shape variation using localized components analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 31 (2009), 1510–1516. (Cited on page 85)
- [9] ALLEN, B., CURLLESS, B., AND POPOVIĆ, Z. The space of human body shapes: reconstruction and parameterization from range scans. *ACM Transactions on Graphics (Proc. SIGGRAPH)* (2003), 587–594. (Cited on pages 21 and 82)
- [10] ANGUELOV, D., SRINIVASAN, P., KOLLER, D., THRUN, S., RODGERS, J., AND DAVIS, J. SCAPE: shape completion and animation of people. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 24, 3 (2005), 408–416. (Cited on pages 21 and 82)
- [11] ARIAS-HERNANDEZ, R., KAASTRA, L. T., GREEN, T. M., AND FISHER, B. Pair analytics: Capturing reasoning processes in collaborative visual analytics. In *Proceedings of the 44th IEEE Hawaii International Conference on System Sciences (HICSS)* (2011), pp. 1–10. (Cited on page 128)
- [12] ARSIGNY, V., COMMOWICK, O., AYACHE, N., AND PENNEC, X. A fast and log-euclidean polyaffine framework for locally linear registration. *Journal of Mathematical Imaging and Vision* 33, 2 (2009), 222–238. (Cited on page 64)
- [13] ARSIGNY, V., COMMOWICK, O., PENNEC, X., AND AYACHE, N. A log-euclidean framework for statistics on diffeomorphisms. In *Proceedings of the 9th international conference on Medical Image Computing and Computer-Assisted Intervention 2006*

- (*MICCAI'06*), *Part I* (2006), Springer, pp. 924–931. (Cited on pages 38, 41, 55, 58, 64 and 141)
- [14] ASHBURNER, J. A fast diffeomorphic image registration algorithm. *NeuroImage* 38, 1 (2007), 95–113. (Cited on pages 38, 55 and 58)
- [15] ASHBURNER, J., HUTTON, C., FRACKOWIAK, R., JOHNSRUDE, I., PRICE, C., AND FRISTON, K. Identifying global anatomical differences: Deformation-based morphometry. *Human Brain Mapping* 6, 5-6 (1998), 348–357. (Cited on pages 11 and 35)
- [16] ASHBURNER, J., AND KLÖPPEL, S. Multivariate models of inter-subject anatomical variability. *NeuroImage* 56, 2 (2011), 422–439. (Cited on page 127)
- [17] AVANTS, B., AND GEE, J. C. Geodesic estimation for large deformation anatomical shape averaging and interpolation. *NeuroImage* 23 (2004), S139–S150. (Cited on page 40)
- [18] BAIR, A., AND HOUSE, D. Grid with a view: Optimal texturing for perception of layered surface shape. *IEEE Transactions on Visualization & Computer Graphics (TVCG)* 13, 6 (nov 2007), 1656–1663. (Cited on page 24)
- [19] BARR, A. H. Global and local deformations of solid primitives. In *Proceedings of the 11th Annual ACM Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'84)* (1984), pp. 21–30. (Cited on pages 39 and 59)
- [20] BAUER, M., BRUVERIS, M., AND MICHOR, P. W. Overview of the geometries of shape spaces and diffeomorphism groups. *Journal of Mathematical Imaging and Vision* 50, 1-2 (2014), 60–97. (Cited on page 12)
- [21] BEG, M. F., MILLER, M. I., TROUVÉ, A., AND YOUNES, L. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International Journal of Computer Vision* 61, 2 (2005), 139–157. (Cited on pages 38 and 57)
- [22] BERNER, A., BURGHARD, O., WAND, M., MITRA, N. J., KLEIN, R., AND SEIDEL, H.-P. A morphable part model for shape manipulation. Tech. Rep. MPI-I-2011-4-005, MPI Informatik, Dec. 2011. (Cited on page 23)
- [23] BLAAS, J., BOTHA, C. P., AND POST, F. H. Interactive visualization of multi-field medical data using linked physical and feature-space views. In *Proceedings of the 9th Joint Eurographics/IEEE VGTC conference on Visualization* (2007), pp. 123–130. (Cited on page 22)
- [24] BLANZ, V., MEHL, A., VETTER, T., AND SEIDEL, H.-P. A statistical method for robust 3D surface reconstruction from sparse data. In *Proceedings of the 2nd International Symposium on 3D Data Processing Visualization and Transmission (3DPVT'04)* (2004), IEEE, pp. 293–300. (Cited on pages 10, 23, 84, 98, 100, 102, 112 and 123)
- [25] BLANZ, V., AND VETTER, T. A morphable model for the synthesis of 3D faces. In *Proceedings of the 26th annual ACM conference on computer graphics and interactive techniques (SIGGRAPH'99)* (1999), pp. 187–194. (Cited on pages 5, 21, 22 and 82)

- [26] BOOKSTEIN, F. L. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 11, 6 (1989), 567–585. (Cited on pages 25 and 39)
- [27] BOOKSTEIN, F. L. *Morphometric tools for landmark data: Geometry and Biology*. Cambridge University Press, 1991. (Cited on page 11)
- [28] BOOKSTEIN, F. L. Integration, disintegration, and self-similarity: Characterizing the scales of shape variation in landmark data. *Evolutionary Biology* (2015), 1–32. (Cited on page 129)
- [29] BOOKSTEIN, F. L., GUNZ, P., MITTERÖCKER, P., PROSSINGER, H., SCHÄFER, K., AND SEIDLER, H. Cranial integration in *Homo*: singular warps analysis of the midsagittal plane in ontogeny and evolution. *Journal of Human Evolution* 44, 2 (2003), 167–187. (Cited on page 101)
- [30] BOSSA, M., HERNANDEZ, M., AND OLMOS, S. Contributions to 3D diffeomorphic atlas estimation: application to brain images. In *Proceedings of the 10th international conference on Medical Image Computing and Computer-Assisted Intervention 2007 (MICCAI'07), Part I*. Springer, 2007, pp. 667–674. (Cited on pages 55 and 58)
- [31] BOTHA, C. P., PREIM, B., KAUFMAN, A., TAKAHASHI, S., AND YNNERMAN, A. From individual to population: Challenges in medical visualization. In *Scientific Visualization*. Springer, 2014, pp. 265–282. (arXiv:1206.1148v1). (Cited on pages 56 and 129)
- [32] BRUCKNER, S., AND GRÖLLER, M. E. Exploded views for volume data. *IEEE Transactions on Visualization & Computer Graphics (TVCG)* 12, 5 (2006), 1077–1084. (Cited on page 79)
- [33] BRUNET, T., NOWAK, K. E., AND GLEICHER, M. Integrating dynamic deformations into interactive volume visualization. In *Proceedings of the Eurographics/IEEE VGTC conference on Visualization* (2006), pp. 219–226. (Cited on pages 35, 55, 59, 70 and 88)
- [34] BURGIO, G., BAYLAC, M., HEYER, E., AND MONTAGUTELLI, X. Exploration of the genetic organization of morphological modularity on the mouse mandible using a set of interspecific recombinant congenic strains between c57bl/6 and mice of the *Mus spretus* species. *G3: Genes—Genomes—Genetics* 2, 10 (2012), 1257–1268. (Cited on pages 114, 116 and 118)
- [35] BUSKING, S. *Visualization of Variation and Variability*. PhD thesis, Delft University of Technology, December 2014. (Cited on page 55)
- [36] BUSKING, S., BOTHA, C. P., AND POST, F. H. Direct visualization of deformation in volumes. *Computer Graphics Forum* 28, 3 (2009), 799–806. (Cited on pages 27 and 59)
- [37] BUSKING, S., BOTHA, C. P., AND POST, F. H. Dynamic multi-view exploration of shape spaces. *Computer Graphics Forum* 29, 3 (2010), 973–982. (Cited on pages 5, 8, 22, 25, 55, 82, 84, 85, 100 and 127)

- [38] CABAN, J. J., RHEINGANS, P., AND YOO, T. An evaluation of visualization techniques to illustrate statistical deformation models. *Computer Graphics Forum* 30, 3 (2011), 821–830. (Cited on pages 24, 25, 26, 57, 59 and 68)
- [39] CAMPBELL, N. A., AND ATCHLEY, W. R. The geometry of canonical variate analysis. *Systematic Zoology* 30, 3 (1981), 268–280. (Cited on page 85)
- [40] CASHMAN, T. J., AND HORMANN, K. A continuous, editable representation for deforming mesh sequences with separate signals for time, pose and shape. *Computer Graphics Forum* 31 (2012), 735–744. (Cited on pages 5 and 22)
- [41] CEDILNIK, A., AND RHEINGANS, P. Procedural annotation of uncertain information. In *Proceedings of IEEE Visualization'00* (2000), pp. 77–83. (Cited on page 26)
- [42] CHEN, M., SILVER, D., WINTER, A. S., SINGH, V., AND CORNEA, N. Spatial transfer functions: a unified approach to specifying deformation in volume modeling and animation. In *Proceedings of the 2003 Eurographics/IEEE TVCG Workshop on Volume graphics* (2003), ACM, pp. 35–44. (Cited on page 79)
- [43] CHRISTENSEN, G., RABBITT, R., AND MILLER, M. Deformable templates using large deformation kinematics. *IEEE Transactions on Image Processing* 5, 10 (Oct. 1996), 1435–1447. (Cited on page 37)
- [44] COFFEY, D., LIN, C.-L., ERDMAN, A. G., AND KEEFE, D. F. Design by dragging: An interface for creative forward and inverse design with simulation ensembles. *IEEE Transactions on Visualization & Computer Graphics (TVCG)* 19, 12 (2013), 2783–2791. (Cited on page 23)
- [45] COOTES, T. F., TAYLOR, C. J., COOPER, D. H., AND GRAHAM, J. Active shape models: Their training and application. *Computer Vision and Image Understanding* 61, 1 (Jan. 1995), 38–59. (Cited on page 21)
- [46] CORREA, C. D., AND SILVER, D. Programmable shaders for deformation rendering. In *Proceedings of the ACM Siggraph/Eurographics symposium on Graphics hardware* (2007), pp. 89–96. (Cited on page 59)
- [47] CORREA, C. D., SILVER, D., AND CHEN, M. Discontinuous displacement mapping for volume graphics. In *Eurographics/IEEE VGTC Workshop on Volume Graphics* (2006), pp. 9–16. (Cited on pages 59 and 78)
- [48] COSTELLO, R. B., Ed. *Webster's College Dictionary*. Random House, 1991. (Cited on page 30)
- [49] CRISTIANINI, N., AND TAYLOR, J. S. *An Introduction to Support Vector Machines*. Cambridge University Press, 2000. (Cited on page 87)
- [50] CRUM, W. R., HARTKENS, T., AND HILL, D. Non-rigid image registration: theory and practice. *The British Journal of Radiology* (2004), S140–S153. (Cited on page 36)

- [51] DAVATZIKOS, C., TAO, X., AND SHEN, D. Hierarchical active shape models, using the wavelet transform. *IEEE Transactions on Medical Imaging* 22, 3 (2003), 414–423. (Cited on pages 85 and 129)
- [52] DAVIES, R. H., TWINING, C. J., COOTES, T. F., WATERTON, J. C., AND TAYLOR, C. J. A minimum description length approach to statistical shape modeling. *IEEE Transactions on Medical Imaging* 21, 5 (2002), 525–537. (Cited on pages 78 and 130)
- [53] DE SMET, M., AND VAN GOOL, L. Optimal regions for linear model-based 3D face reconstruction. In *Proceedings of the 10th Asian Conference on Computer Vision (ACCV)– Volume Part III* (2010), Springer, pp. 276–289. (Cited on page 102)
- [54] DRYDEN, I. L., KOLOYDENKO, A., AND ZHOU, D. Non-euclidean statistics for covariance matrices, with applications to diffusion tensor imaging. *The Annals of Applied Statistics* (2009), 1102–1123. (Cited on page 108)
- [55] DUPUIS, P., GRENDER, U., AND MILLER, M. I. Variational problems on flows of diffeomorphisms for image matching. *Quarterly of applied mathematics* 56, 3 (1998), 587. (Cited on page 37)
- [56] DURRLEMAN, S., PRASTAWA, M., CHARON, N., KORENBERG, J. R., JOSHI, S., GERIG, G., AND TROUVÉ, A. Morphometry of anatomical shape complexes with dense deformations and sparse parameters. *NeuroImage* 101 (2014), 35–49. (Cited on page 129)
- [57] ESSAFI, S., LANGS, G., AND PARAGIOS, N. Hierarchical 3d diffusion wavelet shape priors. In *Proceedings of the 12th IEEE International Conference on Computer Vision* (2009), pp. 1717–1724. (Cited on page 129)
- [58] FISCHER, B., AND MODERSITZKI, J. A unified approach to fast image registration and a new curvature based registration technique. *Linear Algebra and its applications* 380 (2004), 107–124. (Cited on page 38)
- [59] FISCHER, B., AND MODERSITZKI, J. Ill-posed medicine – an introduction to image registration. *Inverse Problems* 24, 3 (2008), 034008. (Cited on pages 36 and 57)
- [60] FLETCHER, P. T., AND JOSHI, S. Riemannian geometry for the statistical analysis of diffusion tensor data. *Signal Processing* 87, 2 (2007), 250–262. (Cited on page 108)
- [61] FLETCHER, P. T., LU, C., PIZER, S. M., AND JOSHI, S. Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Transactions on Medical Imaging* 23, 8 (2004), 995–1005. (Cited on page 129)
- [62] FLETCHER, P. T., PIZER, S. M., THALL, A., AND GASH, A. G. Shape modeling and image visualization in 3D with m-rep object models. Tech. rep., University of North Carolina at Chapel Hill, Dept. of Computer Science, 2000. (Cited on pages 27 and 129)

- [63] GOWER, J. Generalized procrustes analysis. *Psychometrika* 40 (1975), 33–51. (Cited on page 42)
- [64] GRENDER, U., AND MILLER, M. I. Computational anatomy: an emerging discipline. *Quarterly of Applied Mathematics* LVI, 4 (Dec. 1998), 617–694. (Cited on page 11)
- [65] GRENDER, U., AND MILLER, M. I. *Pattern theory: from representation to inference*. Oxford university press, 2007. (Cited on page 11)
- [66] GUIMOND, A., MEUNIER, J., AND THIRION, J.-P. Average brain models: A convergence study. *Computer vision and image understanding* 77, 2 (2000), 192–210. (Cited on page 40)
- [67] HAMARNEH, G., WARD, A. D., AND FRANK, R. Quantification and visualization of localized and intuitive shape variability using a novel medial-based shape representation. In *Proceedings of the 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI 2007)*. (2007), pp. 1232–1235. (Cited on page 26)
- [68] HANDELS, H., AND HACKER, S. A framework for representation and visualization of 3D shape variability of organs in an interactive anatomical atlas. *Methods of information in medicine* 48, 3 (2009), 272–281. (Cited on page 27)
- [69] HEIMANN, T., AND MEINZER, H.-P. Statistical shape models for 3D medical image segmentation: A review. *Medical Image Analysis* 13, 4 (2009), 543 – 563. (Cited on page 21)
- [70] HIGHAM, N. J. *Functions of Matrices: Theory and Computation*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2008. (Cited on pages 64 and 141)
- [71] JIAO, F., PHILLIPS, J. M., GUR, Y., AND JOHNSON, C. R. Uncertainty visualization in HARDI based on ensembles of ODFs. In *Proceedings of the IEEE Pacific Visualization Symposium (PacificVis)* (2012), pp. 193–200. (Cited on pages 25, 57 and 68)
- [72] JOLLIFFE, I. *Principal component analysis*. Wiley, 2005. (Cited on page 47)
- [73] JOSHI, S., DAVIS, B., JOMIER, B. M., AND B, G. G. Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage* 23 (2004), 151–160. (Cited on page 40)
- [74] KANITSAR, A., FLEISCHMANN, D., WEGENKITTL, R., FELKEL, P., AND GRÖLLER, M. E. CPR-curved planar reformation. In *Proceedings of the IEEE Visualization Conference* (2002). (Cited on page 63)
- [75] KAUFMAN, L., AND ROUSSEEUW, P. J. Finding groups in data: an introduction to cluster analysis, 1990. (Cited on page 108)
- [76] KEIM, D., KOHLHAMMER, J., ELLIS, G., AND MANSMANN, F., Eds. *Mastering the information age: solving problems with visual analytics*. Eurographics Association, 2010. (Cited on page 5)

- [77] KENDALL, D. G. The diffusion of shape. *Advances in applied probability* (1977), 428–430. (Cited on pages 11 and 30)
- [78] KENDALL, D. G., ET AL. A survey of the statistical theory of shape. *Statistical Science* 4, 2 (1989), 87–99. (Cited on page 11)
- [79] KILIAN, M., MITRA, N. J., AND POTTMANN, H. Geometric modeling in shape space. *ACM Transactions on Graphics (TOG)* 26, 3 (2007), 64. (Cited on page 22)
- [80] KINDLMANN, G. Superquadric tensor glyphs. In *Proceedings of the Eurographics/IEEE Symposium on Visualization (VisSym'04)* (2004), pp. 147–154. (Cited on page 110)
- [81] KINDLMANN, G., WEINSTEIN, D., LEE, A., TOGA, A., AND THOMPSON, P. Visualization of anatomic covariance tensor fields. *Engineering in Medicine and Biology Society (IEMBS'04)*. 1 (Sept. 2004), 1842–1845. (Cited on pages 26, 59, 97, 98, 100, 110, 116 and 120)
- [82] KINDLMANN, G., AND WESTIN, C.-F. Diffusion tensor visualization with glyph packing. *IEEE Transactions on Visualization & Computer Graphics (TVCG)* 12, 5 (2006), 1329–1336. (Cited on page 111)
- [83] KIRSCHNER, M., AND WESARG, S. Interactive visualization of statistical shape models. Poster presentation at Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM), 2010. (Cited on page 26)
- [84] KLEIN, S., STARING, M., MURPHY, K., VIERGEVER, M., AND PLUIM, J. elastix: A toolbox for intensity-based medical image registration. *IEEE Transactions on Medical Imaging* 29, 1 (Jan 2010), 196–205. (Cited on pages 38 and 132)
- [85] KLEMM, P., OELTZE-JAFRA, S., LAWONN, K., HEGENSCHIED, K., VOLZKE, H., AND PREIM, B. Interactive visual analysis of image-centric cohort study data. *IEEE Transactions on Visualization & Computer Graphics (TVCG)* (2014), 1673–1682. (Cited on pages 7, 22, 25, 55, 56, 127 and 129)
- [86] KLINGENBERG, C. P. Morphological integration and developmental modularity. *Annual review of ecology, evolution, and systematics* (2008), 115–132. (Cited on page 15)
- [87] KLINGENBERG, C. P. MorphoJ: an integrated software package for geometric morphometrics. *Molecular Ecology Resources* 11 (Mar. 2010). (Cited on page 100)
- [88] KLINGENBERG, C. P. Cranial integration and modularity: insights into evolution and development from morphometric data. *Hystrix, the Italian Journal of Mammalogy* 24, 1 (2013), 43–58. (Cited on pages 14, 24, 76 and 98)
- [89] KLINGENBERG, C. P. Visualizations in geometric morphometrics: how to read and how to make graphs showing shape changes. *Hystrix, the Italian Journal of Mammalogy* 24, 1 (2013), 15–24. (Cited on pages 24 and 25)
- [90] KLINGENBERG, C. P., MEBUS, K., AND AUFFRAY, J.-C. Developmental integration in a complex morphological structure: how distinct are the modules in the

- mouse mandible? *Evolution & Development* 5, 5 (2003), 522–531. (Cited on pages 14, 101, 114, 116 and 118)
- [91] KLINGENBERG, C. P., AND ZAKLAN, S. D. Morphological integration between developmental compartments in the drosophila wing. *Evolution* 54, 4 (2000), 1273–1285. (Cited on page 101)
- [92] KONDO, B., AND COLLINS, C. Dimpvis: Exploring time-varying information visualizations by direct manipulation. *IEEE Transactions on Visualization & Computer Graphics (TVCG)* 20, 12 (Dec 2014), 2003–2012. (Cited on page 23)
- [93] KRETSCHMER, J., SOZA, G., TIETJEN, C., SUEHLING, M., PREIM, B., AND STAMMINGER, M. ADR - Anatomy-Driven Reformation. *IEEE Transactions on Visualization & Computer Graphics (TVCG)* (2014), 2496–2505. (Cited on pages 57 and 63)
- [94] KRIEGER, J., AND MACLEOD, N. A repository of morphometric tools and online morphometric workbenches. *Webpage*: <http://www.morpho-tools.net/index.html> (Software packages and Mathematica notebooks, page retrieved on 2013/03/16). (Cited on page 100)
- [95] KURZION, Y., AND YAGEL, R. Interactive space deformation with hardware-assisted rendering. *IEEE Computer Graphics and Applications* 17, 5 (1997), 66–77. (Cited on page 59)
- [96] LAMECKER, H., SEEBASS, M., LANGE, T., HEGE, H.-C., AND DEUFLHARD, P. Visualization of the variability of 3D statistical shape models by animation. *Studies in health technology and informatics* (2004), 190–196. (Cited on pages 27 and 61)
- [97] LEWIS, J., AND ANJYO, K.-I. Direct manipulation blendshapes. *IEEE Computer Graphics and Applications* 30, 4 (2010), 42–50. (Cited on pages 23 and 100)
- [98] LOMBAERT, H., GRADY, L., PENNEC, X., PEYRAT, J.-M., AYACHE, N., AND CHERIET, F. Groupwise spectral log-demons framework for atlas construction. In *Medical Computer Vision. Recognition Techniques and Applications in Medical Imaging*. Springer, 2013, pp. 11–19. (Cited on page 40)
- [99] LÜTHI, M., ALBRECHT, T., AND VETTER, T. Probabilistic modeling and visualization of the flexibility in morphable models. In *Mathematics of Surfaces XIII*. Springer, 2009, pp. 251–264. (Cited on page 26)
- [100] MALIK, M. M., HEINZL, C., AND GROELLER, M. E. Comparative visualization for parameter studies of dataset series. *IEEE Transactions on Visualization & Computer Graphics (TVCG)* 16, 5 (2010), 829–840. (Cited on page 25)
- [101] MARTÍNEZ-ABADÍAS, N., MITTEROECKER, P., PARSONS, T. E., ESPARZA, M., SJØVOLD, T., ROLIAN, C., RICHTSMEIER, J. T., AND HALLGRÍMSSON, B. The developmental basis of quantitative craniofacial variation in humans and mice. *Evolutionary Biology* 39, 4 (2012), 554–567.

- [102] MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. A data-driven reflectance model. *ACM Transactions on Graphics (Proc. SIGGRAPH) 22(3)* (2003), 759–769. (Cited on pages 22, 83 and 87)
- [103] MILLER, M. I., TROUVÉ, A., AND YOUNES, L. On the metrics and euler-lagrange equations of computational anatomy. *Annual review of biomedical engineering 4*, 1 (2002), 375–405. (Cited on pages 38 and 57)
- [104] MILLER, M. I., TROUVÉ, A., AND YOUNES, L. Geodesic shooting for computational anatomy. *Journal of mathematical imaging and vision 24*, 2 (2006), 209–228. (Cited on page 38)
- [105] MODERSITZKI, J. *Numerical methods for image registration*. Oxford university press, 2004. (Cited on page 38)
- [106] MONTEIRO, L. R., BONATO, V., AND DOS REIS, S. F. Evolutionary integration and morphological diversification in complex morphological structures: mandible shape divergence in spiny rats (Rodentia, Echimyidae). *Evolution & development 7*, 5 (2005), 429–439. (Cited on page 116)
- [107] O’HIGGINS, P., AND JONES, N. Morphologika 2.5 software package, 2006. *Webpage:* <http://sites.google.com/site/hymsfme/resources> (page retrieved on 2013/03/16). (Cited on page 100)
- [108] PFAFFELMOSE, T., REITINGER, M., AND WESTERMANN, R. Visualizing the positional and geometrical variability of isosurfaces in uncertain scalar fields. In *Computer Graphics Forum* (2011), vol. 30, pp. 951–960. (Cited on page 25)
- [109] PÖTHKOW, K., AND HEGE, H.-C. Positional uncertainty of isocontours: Condition analysis and probabilistic measures. *IEEE Transactions on Visualization & Computer Graphics (TVCG) 17*, 10 (2011), 1393–1406. (Cited on page 25)
- [110] PRESS, W. H. *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge university press, 2007. (Cited on page 48)
- [111] RAUTEK, P., VIOLA, I., AND GRÖLLER, M. E. Caricaturistic visualization. *IEEE Transactions on Visualization & Computer Graphics (TVCG) 12*, 5 (2006), 1085–1092. (Cited on page 60)
- [112] REH, A., GUSENBAUER, C., KASTNER, J., GRÖLLER, M. E., AND HEINZL, C. Mobjects—a novel method for the visualization and interactive exploration of defects in industrial xct data. *IEEE Transactions on Visualization & Computer Graphics (TVCG) 19*, 12 (2013), 2906–2915. (Cited on page 23)
- [113] RENAUD, S., ALIBERT, P., AND AUFFRAY, J.-C. Modularity as a source of new morphological variation in the mandible of hybrid mice. *BMC Evolutionary Biology 12*, 1 (2012), 141. (Cited on pages 114, 116 and 118)
- [114] REZK-SALAMA, C., SCHEUERING, M., SOZA, G., AND GREINER, G. Fast volumetric deformation on general purpose hardware. In *Proceedings of the ACM Siggraph/Eurographics workshop on graphics hardware* (2001), pp. 17–24. (Cited on pages 35, 55 and 59)

- [115] RIDDLE, W. R., LI, R., FITZPATRICK, J. M., DONLEVY, S. C., DAWANT, B. M., AND PRICE, R. R. Characterizing changes in MR images with color-coded Jacobians. *Magnetic resonance imaging* 22, 6 (2004), 769–777. (Cited on pages 27 and 59)
- [116] ROHLF, F. J., AND CORTI, M. Use of two-block partial least-squares to study covariation in shape. *Systematic Biology* 49, 4 (2000), 740–753. (Cited on page 101)
- [117] ROHLF, F. J., AND MARCUS, L. F. A revolution morphometrics. *Trends in Ecology & Evolution* 8, 4 (1993), 129–132. (Cited on page 11)
- [118] RUECKERT, D., FRANGI, A. F., AND SCHNABEL, J. A. Automatic construction of 3D statistical deformation models using non-rigid registration. In *Proceedings of the 4th international conference on Medical Image Computing and Computer-Assisted Intervention 2001 (MICCAI'01)* (2001), Springer, pp. 77–84. (Cited on pages 39 and 44)
- [119] SARAIYA, P., NORTH, C., AND DUCA, K. An insight-based methodology for evaluating bioinformatics visualizations. *IEEE Transactions on Visualization & Computer Graphics (TVCG)* 4, 2 (2005), 443–456. (Cited on page 3)
- [120] SATTLER, M., SARLETTE, R., AND KLEIN, R. Simple and efficient compression of animation sequences. In *Proceedings of Eurographics/ACM SIGGRAPH Symposium on Computer Animation (SCA)* (2005), pp. 209–217. (Cited on page 102)
- [121] SCHULTZ, T., AND KINDLMANN, G. Superquadric glyphs for symmetric second-order tensors. *IEEE Transactions on Visualization & Computer Graphics (TVCG)* 16, 6 (2010), 1595–1604. (Cited on page 111)
- [122] SEDERBERG, T. W., AND PARRY, S. R. Free-form deformation of solid geometric models. In *Proceedings of the 13th Annual ACM Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'84)* (1986), pp. 151–160. (Cited on page 39)
- [123] SEILER, C., PENNEC, X., AND REYES, M. Capturing the multiscale anatomical shape variability with polyaffine transformation trees. *Medical Image Analysis* 16, 7 (2012), 1371 – 1384. (Cited on pages 78 and 129)
- [124] SEILER, C., PENNEC, X., RITACCO, L., AND REYES, M. Femur specific polyaffine model to regularize the log-domain demons registration. In *Proceedings of SPIE Medical Imaging 2011: Image Processing* (Mar. 2011), vol. 7962. (Cited on pages 27, 55, 58 and 66)
- [125] SHAMIR, A. A survey on mesh segmentation techniques. *Computer Graphics Forum* 27, 6 (2008), 1539–1556. (Cited on page 102)
- [126] SHNEIDERMAN, B. The eyes have it: a task by data type taxonomy for information visualizations. In *Proceedings of the IEEE Symposium on Visual Languages, 1996* (1996), pp. 336–343. (Cited on page 98)

- [127] SJÖSTRAND, K., STEGMANN, M. B., AND LARSEN, R. Sparse principal component analysis in medical shape modeling. In *Proceedings of SPIE, Medical Imaging 2006: Image Processing* (Mar. 2006), J. M. Reinhardt and J. P. W. Pluim, Eds., vol. 6144, SPIE. (Cited on pages 85 and 129)
- [128] SMIT, N. N., KLEIN HANEVELD, B., STARING, M., EISEMANN, E., BOTHA, C. P., AND VILANOVA, A. RegistrationShop: An interactive 3D medical volume registration system. In *Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)* (2014), pp. 145–153. (Cited on pages 24 and 79)
- [129] SMITH, R., PAWLICKI, R., KOKAI, I., FINGER, J., AND VETTER, T. Navigating in a shape space of registered models. *IEEE Transactions on Visualization & Computer Graphics (TVCG)* 13, 6 (2007), 1552–1559. (Cited on pages 5, 22 and 82)
- [130] SOTIRAS, A., DAVATZIKOS, C., AND PARAGIOS, N. Deformable medical image registration: A survey. *IEEE Transactions Medical Imaging* 32, 7 (2013), 1153–1190. (Cited on pages 37 and 38)
- [131] STEENWIJK, M. D., MILLES, J., BUCHEM, M., REIBER, J., AND BOTHA, C. P. Integrated visual analysis for heterogeneous datasets in cohort studies. In *IEEE VisWeek Workshop on Visual Analytics in Health Care* (2010). (Cited on page 56)
- [132] STRANG, G. *Introduction to Linear Algebra*. Wellesley-Cambridge Press, 1998. (Cited on page 46)
- [133] TENA, J. R., DE LA TORRE, F., AND MATTHEWS, I. Interactive region-based linear 3D face models. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 30 (2011), 76:1–76:10. (Cited on pages 23 and 100)
- [134] THOMPSON, D. W. *On Growth and Form: A New Edition*. Cambridge University Press, 1942. (Cited on pages 10 and 25)
- [135] THOMPSON, P. M., AND TOGA, A. W. Detection, visualization and animation of abnormal anatomic structure with a deformable probabilistic brain atlas based on random vector field transformations. *Medical image analysis* 1, 4 (1997), 271–294. (Cited on pages 26 and 27)
- [136] TITTEMEYER, M., WOLLNY, G., AND KRUGGEL, F. Visualising deformation fields computed by non-linear image registration. *Computing and Visualization in Science* 5, 1 (2002), 45–51. (Cited on pages 27 and 60)
- [137] TREFETHEN, L. N., AND BAU III, D. *Numerical linear algebra*, vol. 50. Siam, 1997. (Cited on page 46)
- [138] TUCKER, L. R. An inter-battery method of factor analysis. *Psychometrika* 23, 2 (1958), 111–136. (Cited on page 101)
- [139] TUFTE, E. R. *Envisioning information*. Graphics Press, 1990. (Cited on page 25)
- [140] VAILLANT, M., MILLER, M. I., YOUNES, L., AND TROUVÉ, A. Statistics on diffeomorphisms via tangent space representations. *NeuroImage* 23 (2004), S161–S169. (Cited on page 57)

- [141] VAN GOLEN, K. Landmark influence visualization in active shape models. Master's thesis, TU Delft, Delft University of Technology, 2014. (Cited on page 26)
- [142] VERCAUTEREN, T., PENNEC, X., PERCHANT, A., AND AYACHE, N. Symmetric log-domain diffeomorphic registration: A demons-based approach. In *Proceedings of the 11th international conference on Medical Image Computing and Computer-Assisted Intervention 2008 (MICCAI'08), Part I* (2008), Springer, pp. 754–761. (Cited on pages 38, 58 and 132)
- [143] VERCAUTEREN, T., PENNEC, X., PERCHANT, A., AND AYACHE, N. Diffeomorphic demons: Efficient non-parametric image registration. *NeuroImage* 45, 1 (2009), S61–S72. (Cited on pages 38 and 57)
- [144] VON FUNCK, W., THEISEL, H., AND SEIDEL, H.-P. Vector field based shape deformations. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 25 (2006), 1118–1125. (Cited on pages 57, 64 and 79)
- [145] VON LANDESBERGER, T., ANDRIENKO, G., ANDRIENKO, N., BREMM, S., KIRSCHNER, M., WESARG, S., AND KUIJPER, A. Opening up the “black box” of medical image segmentation with statistical shape models. *The Visual Computer* 29, 9 (2013), 893–905.
- [146] WAHBA, G. *Spline models for observational data*. Siam, 1990. (Cited on page 39)
- [147] WEBER, G. W., AND BOOKSTEIN, F. L. EVAN toolbox. *Webpage*: <http://evan-society.org/node/42> (Open Source package, page retrieved on 2013/03/16). (Cited on page 100)
- [148] WESTIN, C.-F., PELED, S., GUDBJARTSSON, H., KIKINIS, R., AND JOLESZ, F. A. Geometrical diffusion measures for MRI from tensor basis analysis. In *Proceedings of the 5th Annual Meeting of the ISMRM* (1997), p. 1742. (Cited on page 101)
- [149] WIEBEL, A., VOS, F. M., FOERSTER, D., AND HEGE, H.-C. WYSIWYP: What you see is what you pick. *IEEE Transactions on Visualization & Computer Graphics (TVCG)* 18, 12 (2012), 2236 – 2244. (Cited on page 110)
- [150] WIENS, V. Volumetric segmentation of complex bone structures from medical imaging data using reeb graphs. In *Central European Seminar on Computer Graphics for Students (CESCG)* (2013), pp. 113–120. (Cited on page 132)
- [151] WILEY, D. F., AMENTA, N., ALCANTARA, D. A., GHOSH, D., KIL, Y. J., DELSON, E., HARCOURT-SMITH, W., JOHN, K. S., ROHLF, F. J., AND HAMANN, B. Evolutionary morphing. In *Proceedings of IEEE Visualization 2005 (VIS'05)* (2005), pp. 431–438. (Cited on page 25)
- [152] WOLBERG, G. Image morphing: a survey. *The Visual Computer* 14, 8 (1998), 360–372. (Cited on page 35)
- [153] WOLD, H. Path models with latent variables: The NIPALS approach. *Quantitative Sociology: International perspectives on mathematical and statistical model building* (1975), 307–357. (Cited on page 101)

- [154] XUE, Z., SHEN, D., AND DAVATZIKOS, C. Statistical representation of high-dimensional deformation fields with application to statistically constrained 3d warping. *Medical Image Analysis* 10, 5 (2006), 740–751. (Cited on page 129)
- [155] YEZZI, A. J., AND SOATTO, S. Deformation: Deforming motion, shape average and the joint registration and approximation of structures in images. *International Journal of Computer Vision* 53, 2 (2003), 153–167. (Cited on page 31)
- [156] YOO, T. S. *Insight into images: principles and practice for segmentation, registration, and image analysis*. A.K. Peters, 2004. (Cited on page 35)
- [157] YU, P., GRANT, P. E., QI, Y., HAN, X., SÉGONNE, F., PIENAAR, R., BUSA, E., PACHECO, J., MAKRIS, N., BUCKNER, R. L., ET AL. Cortical surface shape analysis based on spherical wavelets. *IEEE Transactions on Medical Imaging* 26, 4 (2007), 582–597. (Cited on page 129)
- [158] YUMER, M. E., CHAUDHURI, S., HODGINS, J. K., AND KARA, L. B. Semantic shape editing using deformation handles. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 86. (Cited on page 5)
- [159] ZELDITCH, M. L., SWIDERSKI, D. L., AND SHEETS, H. D. *Geometric morphometrics for biologists: a primer*. Academic Press, 2012. (Cited on page 13)
- [160] ZELDITCH, M. L., WOOD, A. R., BONETT, R. M., AND SWIDERSKI, D. L. Modularity of the rodent mandible: integrating bones, muscles, and teeth. *Evolution & Development* 10, 6 (2008), 756–768. (Cited on pages 114, 116, 118 and 119)
- [161] ZHANG, C., SCHULTZ, T., LAWONN, K., EISEMANN, E., AND VILANOVA, A. Glyph-based comparative visualization for diffusion tensor fields. *IEEE Transactions on Visualization & Computer Graphics (TVCG)* 22, 1 (2016). (Cited on page 25)
- [162] ZIKIC, D., HANSEN, M. S., GLOCKER, B., KHAMENE, A., LARSEN, R., AND NAVAB, N. Computing minimal deformations: Application to construction of statistical shape models. In *Computer Vision and Pattern Recognition* (2008). (Cited on page 84)
- [163] ZOLLIKOFER, C. P. E., AND PONCE DE LÉON, M. S. Visualizing patterns of craniofacial shape variation in *Homo sapiens*. In *Biological Sciences* 269(1493) (2002), pp. 801–807. (Cited on pages 26 and 115)