

**SHAPE MODELING AND ANALYSIS WITH
ENTROPY-BASED PARTICLE SYSTEMS**

by

Joshua Cates

A dissertation submitted to the faculty of
The University of Utah
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in

Computing

School of Computing

The University of Utah

May 2010

Copyright © Joshua Cates 2010

All Rights Reserved

ABSTRACT

Many important fields of basic research in medicine and biology routinely employ tools for the statistical analysis of collections of similar shapes. Biologists, for example, have long relied on homologous, anatomical landmarks as shape models to characterize the growth and development of species. Increasingly, however, researchers are exploring the use of more detailed models that are derived computationally from three-dimensional images and surface descriptions. While computationally-derived models of shape are promising new tools for biomedical research, they also present some significant engineering challenges, which existing modeling methods have only begun to address.

In this dissertation, I propose a new computational framework for statistical shape modeling that significantly advances the state-of-the-art by overcoming many of the limitations of existing methods. The framework uses a particle-system representation of shape, with a fast correspondence-point optimization based on information content. The optimization balances the simplicity of the model (compactness) with the accuracy of the shape representations by using two commensurate entropy metrics and no free parameters. The idea is to maximize both the geometric accuracy and the statistical simplicity of the shape model, in accordance with the principle of parsimony in model selection. The nonparametric representation allows the method to be applied to a larger class of problems than existing methods, including nonspherical surfaces, open surfaces, and sets of multiple surfaces. The relative simplicity of the surface representation and the low number of free parameters results in a framework that is easy to use and can operate directly on image segmentations. In collaboration with scientists from several important areas of biomedicine, I have demonstrated that the proposed method is indeed an effective tool for scientific research.

The specific research contributions of this dissertation are as follows. First, I describe a mathematical framework and a robust numerical algorithm for computing optimized correspondence-point shape models using an entropy-based optimization and particle-system technology. Second, I develop a series of extensions of the framework to more general classes of shape analysis problems, including the analysis of multiple-object complexes, the generalization to correspondence based on generic functions of position, an extension to handle surfaces with open boundaries, and shape modeling with simple regression. Third, I describe the application of statistical hypothesis testing, regression analysis, and multiple-analysis of covariance to the proposed shape models. I also introduce new techniques for visualization and interpretation of these statistics. Finally, in cooperation with biomedical researchers, I present validation of the above research contributions by their successful application to real-world research problems.

To Sarah Cates

CONTENTS

ABSTRACT	iv
LIST OF FIGURES	x
LIST OF TABLES	xiii
ACKNOWLEDGMENTS	xiv
CHAPTERS	
1. INTRODUCTION	1
1.1 A Brief History of Morphometrics	1
1.2 Computationally-Derived Shape Models	4
1.3 Shape Analysis in Modern Biology and Medicine	6
1.3.1 Neurobiology	7
1.3.2 Genetic Phenotyping	7
1.3.3 Other Applications	8
1.4 Technical Challenges for Biomedical Shape Analysis	8
1.4.1 Shape Representation	8
1.4.2 The Model Selection Problem	9
1.4.3 Statistical Analysis and Interpretability	10
1.4.4 Ease-of-Use and Computational Complexity	11
1.5 Research Contributions	11
1.6 Document Organization	13
2. SHAPE MODELING FRAMEWORK	14
2.1 Shape and Shape Representation	14
2.1.1 Shape and its Analysis	14
2.1.2 Shape Representation	16
2.1.3 Surface Description	18
2.1.4 Particle Systems for Surface Representation	19
2.2 Correspondence Shape Models	21
2.2.1 Correspondences	21
2.2.2 Optimized Correspondences	24
2.3 Particle-Based Modeling for Shape Correspondence	26
2.3.1 Motivation	26

2.3.2	Algorithm Overview	28
2.3.3	Surface Representation	29
2.3.4	Adaptive Distributions on Surface Features	31
2.3.5	Surface Constraint	32
2.3.6	The Kernel Width σ for PDF Estimation	32
2.3.7	Numerical Considerations	33
2.3.8	Correspondence: Entropy Minimization in Shape Space	34
2.3.9	Setting Parameters	36
2.4	A Shape Modeling Pipeline	36
2.4.1	Preprocessing Binary Segmentations	37
2.4.2	Initialization and Optimization	38
2.4.3	Alignment of Shape Surfaces	39
2.4.4	Software Implementation	39
2.5	Preliminary Evaluation	39
2.5.1	Comparison with MDL	39
2.5.2	Nonspherical Topologies	40
2.5.3	Shape Analysis of Neuroanatomical Structures	44
3.	MODELING EXTENSIONS	48
3.1	Modeling Shape with Open Surfaces	49
3.2	Modeling Shape Complexes	50
3.3	Correspondence Based on Functions of Position	53
3.4	Correspondence with Regression Against Explanatory Variables	55
4.	STATISTICS AND INTERPRETATION	59
4.1	Summary Statistics	59
4.2	Principal Component Analysis	61
4.3	Hypothesis Testing	62
4.3.1	Challenges	62
4.3.2	Hypothesis Testing Framework	68
4.3.3	Visualizing Significant Group Difference	71
4.4	Linear Models and Regression	75
4.5	Statistical Significance of Regression	76
4.6	Analysis of Nonshape Components: Pose and Scale	78
5.	APPLICATIONS	80
5.1	Group Comparison of the Caudate in Autism	81
5.1.1	Introduction	81
5.1.2	Methods	82
5.1.3	Results and Discussion	82

5.2	Multiobject Analysis of 10 Subcortical Structures in Autism	88
5.2.1	Introduction	88
5.2.2	Methods	88
5.2.3	Results and Discussion	89
5.3	Cortical Thickness Analysis with Correspondence Based on Sulcal Depth	91
5.3.1	Introduction	91
5.3.2	Methods	92
5.3.3	Results	92
5.4	A Regression Model of Pediatric Head Shape with Age	95
5.4.1	Introduction	95
5.4.2	Methods	95
5.4.3	Results and Discussion	96
5.5	Phenotyping in <i>Hoxd11</i> -deficient Mice	100
5.5.1	Introduction	100
5.5.2	Methodology	100
5.5.3	Results and Discussion	101
5.6	Analysis of the Cranial Base in <i>Pax7</i> Mutant Mice	103
5.6.1	Introduction	103
5.6.2	Methodology	105
5.6.3	Results and Discussion	107
6.	CONCLUSION AND FUTURE WORK	116
6.1	Research Contributions	116
6.2	Scientific Impact	117
6.3	Limitations and Future Work	118
	REFERENCES	121

LIST OF FIGURES

1.1	A landmark-based representation of a hypothetical insect wing is constructed by annotating the positions where the veins of the wing intersect.	3
2.1	A sample population of torus shapes that are randomly drawn from a bivariate Gaussian distribution parameterized by large radius R and small radius r	15
2.2	A comparison of closed bone surface (top panel) with an open bone surface (bottom panel).	17
2.3	An illustration of a single slice from a segmentation (left), the corresponding slice of the distance transform (middle), and an antialiased surface rendering of the zero level set of the distance transform (right).	20
2.4	An illustration of surface representation using particle systems. From left to right, increasing the number of particles results in an increased accuracy of the surface reconstruction.	21
2.5	An illustration of the relationship between configuration space and shape space.	23
2.6	An illustration of the basic concepts of a correspondence configuration optimization.	25
2.7	An illustration of the tradeoff between model compactness and the quality of shape representation.	27
2.8	A typical shape modeling pipeline of operations using the PBM correspondence optimization algorithm.	37
2.9	A system of 100 particles on a sphere, produced by a particle splitting algorithm.	38
2.10	The “box-bump” experiment.	41
2.11	The mean and ± 3 standard deviations of the top 3 modes of the hand models.	42
2.12	Comparison of the percentage of total variance in PCA modes of the hand experiment.	43
2.13	An illustration of particle correspondences on several tori from a random distribution on r and R . Corresponding particle positions across the three shapes are indicated by glyphs with matching colors.	43

2.14	The top two PCA modes of variation for the torus shape model, shown in the first and second rows, respectively. Surface reconstructions are from the correspondence positions moved along the axis of each principal component at positions -3 to $+3$ standard deviations from the mean shape. PCA 1 (top row) corresponds to r and PCA 2 (bottom row) correspondence to R	45
2.15	Right hippocampus model mean and ± 3 standard deviations in the top two PCA modes (top and bottom rows, respectively).	45
2.16	P-value maps for the hippocampus and caudate shape analyses, shown on the mean shape. Red (dark in grayscale) indicates significant group differences ($p \leq 0.05$), and blue indicates no significant group differences ($p > 0.05$).	47
3.1	An example of distributing increasing numbers of particles on an open surface, where the boundary is defined by the intersection of an implicit bone surface with a sphere and a cutting plane.	51
3.2	Shape regression on an explanatory variable t correlated with the small radius of a set of tori with $r \sim \mathcal{U}(5, 20)$	57
4.1	Point-wise hypothesis test results for the putamen in a pediatric autism study.	63
4.2	Hotelling T^2 test results with increasing numbers of PCA modes for the amygdala brain structure from a pediatric autism study.	65
4.3	Hotelling T^2 test results with increasing numbers of PCA modes for the putamen brain structure from a pediatric autism study.	66
4.4	Hotelling T^2 test results with increasing numbers of PCA modes for the hippocampus brain structure from a pediatric autism study.	67
4.5	Parallel analysis for the putamen data	70
4.6	LDA visualization of the right putamen from an autism study	72
4.7	A comparison of the linear discriminant visualization (left) with a visualization of the mean group difference (right) on a PBM correspondence model of the left putamen in a shape study of autism versus normal control.	72
4.8	Strain tensors for the right putamen from an autism study. Tensor scale is exaggerated for visualization purposes.	74
5.1	Hypothesized relationships among shape, age, and group in the right caudate.	84
5.2	Shape change with respect to age within TC+DD group and the AUT+PDD group for the right caudate.	85
5.3	Hypothesized relationships among shape, age, and group in the left caudate.	86

5.4	Shape change with respect to age within the TC+DD group and the AUT+PDD group for the left caudate.	87
5.5	Mean brain structure complexes with average pose. Colormap indicates the magnitude and direction of the linear discriminant.	90
5.6	Detail of the right amygdala from Figure 5.5, with the linear discriminant visualized as vectors.	90
5.7	Comparison of the variance of cortical thickness at correspondences from the generalized PBM method (left) and to FreeSurfer (right). Figure by Ipek Oguz.	94
5.8	An illustration of correspondences on the cranial surface for the neonatal head shape study.	97
5.9	Changes in head scale with <i>log</i> of age.	98
5.10	Overview of head shape regression: Changes in head shape with <i>log</i> age. Age is given in months.	99
5.11	LDA visualization of the P1 bone for the <i>HoxD11</i> phenotype study . .	102
5.12	ANCOVA for the PC1 measurements on the P2 bone. Wildtypes are shown as squares and mutants as triangles.	104
5.13	Volume renderings of micro-CT scans of wildtype Pax7(WT/WT), homozygous Pax7(LacZ/LacZ), and heterozygous Pax7(LacZ/WT) median specimens. Images created by Suresh Prajapati.	106
5.14	A visual comparison of the group mean shapes. Shape surfaces are reconstructed from the Euclidean means of the correspondence positions for each group. Mean correspondence positions are indicated by spheres.	108
5.15	Visualization of the linear discriminant from the Hotelling T^2 test for the homozygous shape model (TOP ROW) and the heterozygous shape model (BOTTOM ROW). Arrows indicate direction from the wildtype distribution to the mutant distribution.	109
5.16	Variation in each of the top 5 principal components of the homozygous basioccipital model. Shape (yellow surface) is reconstructed from correspondence positions (green points) at $\pm 1, 2,$ and 3 standard deviations (σ) from the mean shape.	113
6.1	Preliminary results from a PBM study of ventricle shapes. Images courtesy of Manasi Datar.	118
6.2	Preliminary results from a PBM study of femur shapes. Images courtesy of Manasi Datar.	119

LIST OF TABLES

4.1	Results of permutation tests (1000 permutations)	78
5.1	Mean sample variance for cortical thickness and sulcal depth.	93
5.2	Multivariate hypothesis test results for group differences in mean shape for PacZ(LacZ/LacZ), PacZ(LacZ/WT), and combined mutant groups.	108
5.3	Welch two-sample t-tests for mean shape differences in the first 5 principal component loadings for <i>Pax7</i> cohorts.	111
5.4	The percentage of total shape variation represented by each of the first 5 principal modes of variation.	112
5.5	The mean (μ) and standard deviation (σ) in millimeters of the maximum widths for each group of basioccipital bone segmentations.	114
5.6	Hypothesis test results for the groupwise comparison of the simple morphometrics. The metrics used are the maximum widths of the segmentations in the sagittal, coronal, and axial directions, along with the segmentation volume. The multivariate Hotelling T^2 test was conducted using all four metrics together.	115

ACKNOWLEDGMENTS

Several people have made significant contributions to this work. First, I would like to acknowledge and thank my advisor, Ross Whitaker, who has been both a professional and an academic mentor to me for many years. From Ross, I have learned precision and rigor, not to mention volumes about image processing and scientific computing. I would also like to thank Martin Styner, Guido Gerig, and Hal Daumé III for their contributions and advice. I would especially like to acknowledge the significant contributions of P. Thomas Fletcher, whose help has been invaluable, particularly in the area of statistical analysis. Finally, I would like to thank Lara Brewer for all her support and encouragement during my graduate school career.

This work was conducted at the Scientific Computing and Imaging Institute and funded by the Center for Integrative Biomedical Computing (National Institutes of Health (NIH) NCRN Project 2-P41-RR12553-07). This work is also part of the National Alliance for Medical Image Computing (NAMIC), funded by the National Institutes of Health through the NIH Roadmap for Medical Research, Grant U54 EB005149.

CHAPTER 1

INTRODUCTION

This dissertation addresses technical challenges in biomedical shape analysis. It describes a set of novel modeling and analysis tools and validates those tools by their successful application to real-world research problems. *Shape analysis* is defined as both the automatic computation of models of populations of similar shapes and the statistical metrics and tests applied to those shape models. In the latter case, researchers are typically interested in testing biological hypotheses associated with shape and shape differences. The major focus of this dissertation work is the development of simpler, more robust, and more extensible methods for computing *correspondence-point* shape models than are currently available to researchers and the statistical interpretation of those models. This chapter introduces the specific contributions of the dissertation, placing them within the context of the history and major schools of thought of shape analysis research, and motivates them with a discussion of some of the real-world problems in biomedicine for which they are intended.

1.1 A Brief History of Morphometrics

In the study of biological organisms, theories regarding the correlation between the shapes of anatomical structures and their function often arise. As early as the 1600s, for example, Galileo was discussing relationships between bone shapes and sizes in small animals [40]. Charles Darwin famously described a correlation between beak shape and diet in Galapagos island finches as support for his theory of natural selection, hypothesizing that variations in beak shape lead to speciation where the abundance of different food sources varies geographically [26]. D’Arcy Thompson’s 1917 book, *On Growth and Form* [104], did much to encourage the study of how the shape of biological correlates with function, and F. O. Bower made a similar contribution for the field of botany in his 1930 book *Size and Form in Plants* [14]. Both authors helped to advance the notion that perhaps nature could be studied in a more quantitative way, instead of just by the empirical methods that were the norm at the time. Most significantly, Thompson introduced the idea that the variation in similar anatomical shapes might be described by mathematical transformations, and his work is often credited with inspiring the field of *morphometrics*.

Morphometrics is the statistical analysis of the geometry of biological anatomy and employs multivariate methods to analyze measures of shape and size. What are known as traditional morphometric techniques [68] rely primarily on measures of linear distance, angles, and areas of biological forms. Traditional morphometrics

arose as a new branch of statistics by the middle of the twentieth century and was made possible by some of the important statistical advances of the time, such as analysis of variance [37] and principal component analysis [87, 55]. The field was also influenced by the many empirical comparative biology studies that were done in the years after Darwin published his theories on natural selection and took advantage of the relative abundance of data from those studies in formulating many of its basic methodologies [58]. Traditional morphometric tools have been applied extensively to questions of not only evolutionary theory, but also taxonomy, anthropology, and the growth and patterning in organisms (see, for example, the reviews by Adams [1] and Klingenberg, et al. [68]).

While traditional morphometric techniques are an important advance over earlier, empirical methods, they have limited representational and statistical power. Commonly used shape metrics, such as maximal height and the width, do not represent the geometry of a structure in any detailed, or in an unambiguous way, and do not lend themselves well to interpretive visualizations of shape variation [1, 12]. In addition, many biological hypotheses are concerned with allometry, or correlations between the overall size of an organism and the shape of its anatomy [59, 61, 41], and so effectively separating an object's relative scale from its shape is an important consideration. Since linear measures are often correlated with size, for example, allometric questions are better addressed by shape metrics that are invariant to size.

The *geometric morphometric revolution* of the latter part of the 20th century was a series of advances in shape representation and statistics that began to address many of the limitations of traditional morphometrics. Geometric morphometrics incorporate a more rigorous mathematical definition of shape, distinct from size, that is based on homologous anatomical *landmarks*. Landmarks are two- or three-dimensional coordinates of homologous points that lie on the surface of the anatomy [12, 34] and are explicitly defined and measured by the researcher within the context of a particular study. The variability in the shape of an insect wing, for example, might be measured as the variability in the positions where veins intersect (see Figure 1.1, for example) or the variability in animal skulls by the suture points between bones [68]. Landmarks are an important shift from the use of linear measures, which do not explicitly represent anatomical homologies, and may therefore be less repeatable from one observer to the next. They also begin to more completely represent the geometry of the anatomy and can be more easily interpreted in the context of the statistical analysis.

A general framework for landmark-based shape analysis has been developed by several researchers, which is still in widespread use today in the biological and anthropological research communities. In a series of papers, the mathematician David Kendall formulated a rigorous mathematical definition of point-based representations of shape, invariant to scale and orientation [63, 64, 65]. At the same time, zoologist Fred Bookstein did parallel work in the context of biological problems and landmark-based shape [11, 12]. Bookstein and other researchers, notably the mathematicians Kent and Goodall, pioneered the use of generalized Procrustes analysis for normalizing landmark data with respect to size and orientation [57, 66, 45, 12]. Linear distances between shapes in Procrustes space were shown to be related to a tangent-space approximation of geodesic distances in Kendall's shape space [66, 34].

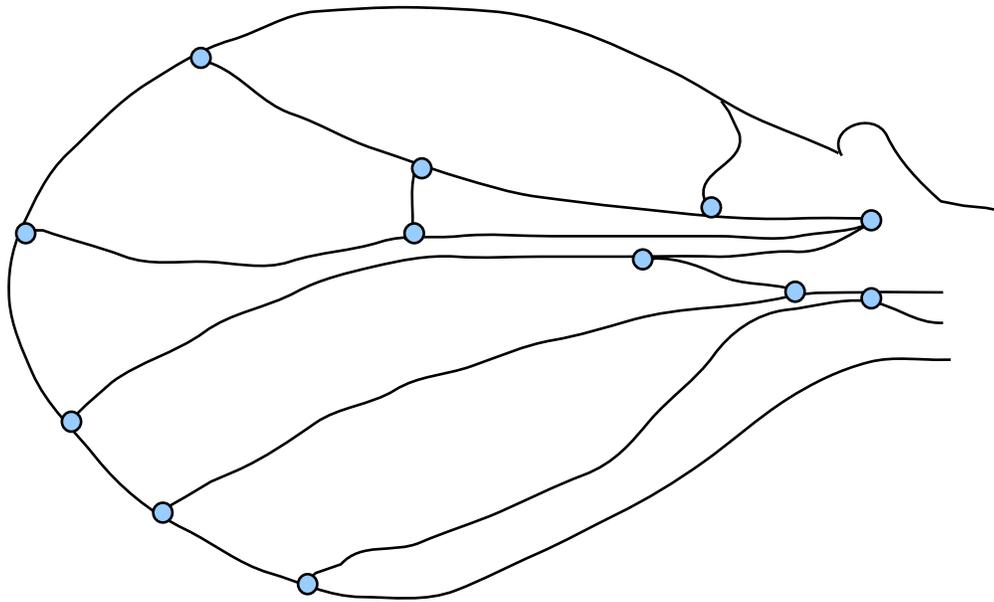


Figure 1.1. A landmark-based representation of a hypothetical insect wing is constructed by annotating the positions where the veins of the wing intersect.

Most geometric morphometric analysis of shape in biomedicine is currently based on these results [1, 34]. The geometric morphometric toolkit also includes methods for directly visualizing the variation in Procrustes points in the space of the geometry from which they were derived, notably by interpolation of a transformation grid (coordinate transformation) from one landmark shape to another using thin-plate splines [11].

1.2 Computationally-Derived Shape Models

Another revolution in shape analysis technology is now underway. This revolution is characterized by the use of computers to automatically construct shape models, usually from three-dimensional digital images of anatomy, such as computed x-ray tomography (CT) imagery and magnetic resonance imagery (MRI). Computationally-derived models of shape offer many advantages over geometric morphometric methods. Much more detailed representations of anatomical geometry are possible than with supervised landmarking, and thus there is the potential to analyze shape with greater statistical power and to capture more subtle shape variations. Shape variation and statistical results can also be visualized more effectively because the geometry of the anatomy is preserved in the analysis more completely. Many clinical and biological applications such as phenotypic screening and patient-specific anatomical modeling require higher throughput and less human intervention than traditional methods allow. Because computationally-derived shape models are unsupervised, they require less work and time for a researcher or clinician to produce. Such models may also be less susceptible to observer bias and error, and the results may be more reproducible. Finally, with modern medical imaging, geometric descriptions of internal anatomy can be acquired noninvasively in living specimens and human subjects. With noninvasive imaging, shape models are now much more viable tools for the study of soft-tissue structures, human anatomy and pathologies, and for longitudinal studies of growth and development.

A variety of computationally-derived shape representations have been proposed, but two major categories of approaches are common. The first strategy is to consider the shape geometry as embedded in the image intensity values at pixels or voxels and then use nonlinear registration to map all sample images to a reference image, or atlas. The variation in shape is considered to be captured by the nonlinear registration parameters. Of the image registration methods, the voxel- and tensor-based morphometry methods and the diffeomorphic methods are in most widespread use. A second approach, and the approach with which this dissertation is concerned, is to extract contours or surface parameterizations from digital images and then sample them in a systematic way to produce an arbitrarily dense set of *correspondence points*, which are used for statistical analysis in a similar fashion as landmark points. This latter category of methods is known variously in the literature as point-distribution models [21], point-based models, or *correspondence models*.

Some early work in shape modeling that is based on image registration was done by Bajcsy and Kovacic, who adapted the thin-plate spline methods of Bookstein in a study of brain structure [7]. A general modeling framework for the study of brain structure shape, called voxel-based and tensor-based morphometry, has

been proposed by several researchers [116, 5]. Ashburner and Friston summarize this work in a 2000 *NeuroImage* paper [4]. These methods deformably register a collection of images to an atlas and then perform statistical analysis in localized regions on either the tensor values describing the transformations or on the actual distribution of voxel values of the transformed images.

Miller, et al. pioneered work in methods that compute diffeomorphic image registrations using fluid-flow models and linear statistics in the space of the initial momentum of the geodesic flows [108, 109, 8]. These methods have been applied to a number of biomedical problems, including the analysis of brain structure [109] and the study of pathologies in the heart [52]. A major feature of the diffeomorphic approach to shape modeling is that diffeomorphic flows are invertible and, thus, a metric may be defined in the space of the deformations. This property has been used by Joshi and Davis to define a mean anatomical image with variation, and they apply these estimators to the problems of unbiased anatomical atlas construction and regression analysis [62, 32].

Correspondence-based shape models are the computational extensions of the geometric morphometric landmark models. They are distinct, however, from landmarks in that they do not explicitly model biological homologies. A landmark model is generally motivated by the hypothesis that landmarks represent positions that are linked by a shared evolutionary pathway or developmental process [34]. A second underlying hypothesis is that the homologous points will exhibit interesting geometric variation relevant to the problem of study. Both of these hypotheses must be formulated in advance by the biologist, based on his or her knowledge of the problem and the anatomy. With correspondence methods, however, there are no such a priori assumptions. The goal is to sample the shape in a sufficiently dense manner so as to capture all of the interesting homologies and geometric variation. Correspondence models attempt to learn the underlying parameters of shape, as opposed to modeling explicitly chosen parameters. A particular set of homologies may then be analyzed after the model has been constructed. The analysis of correspondences, however, can be done in a similar fashion as the analysis of landmark points, using the multivariate techniques and shape space approximations developed by the morphometric community.

While correspondence methods offer many advantages over landmarking methods, the problem of how to automatically choose correspondence positions is a difficult and ill-posed problem. Given a random sample of shapes, the ill-posed nature of the correspondence problem arises from the fact that there are potentially an infinite number of possible configurations that can be chosen for the correspondences on the shapes in the sample set. Some sort of regularization is therefore required. Regularization usually takes the form of a constraint that ensures correspondences faithfully represent the geometries of the samples. Several methods have been proposed in the literature. Many use a purely geometric regularization and construct parameterizations that only consider individual geometries. Other methods combine geometric regularization with a consideration of the statistics of the entire sample to produce a model that is optimized with respect to total variance or information content.

Most methods that use only geometric regularization choose correspondences by constructing a set of shape parameterizations and then sampling surface point

locations in a consistent fashion. Brechbühler, Styner, and Gerig, for example, developed approaches based on spherical harmonic descriptions of shape segmentations. The spherical harmonic representations are first aligned and then sampled into surface meshes to establish correspondence [15, 97]. This idea is an extension of earlier work using Fourier descriptors to parameterize closed contours in two dimensions [88, 71]. Similar approaches have been proposed that use wavelet-based shape representations [81]. Medial-axis parameterizations of shape have also been used to construct nonlinear shape models and applied in a number of neurobiology studies [46, 97]. Dalal, et al. proposed a sliding landmark method which has been shown to compare favorably to the spherical harmonics methods [24, 80]. Their method establishes correspondence by minimizing the bending energy of the thin-plate spline mapping between sample landmarks and a template.

Optimized correspondence models are motivated by the *Occam's razor* principle of parsimony: given a number of possible models for the data, choose the simplest model. The principle of parsimony is widely applied to model selection problems in statistics in order to find models with minimal numbers of parameters and greater predictive power [50]. The idea of an optimized correspondence model of shape was first proposed by Kotcheff and Taylor and more fully developed by Davies, Cootes, and Taylor. Kotcheff and Taylor proposed an algorithm in two dimensions that minimizes the magnitude of the covariance of the correspondences. Davies, et al. [30] proposed an information-theoretic cost function of correspondence positions for two-dimensional shapes based on *minimum description length* (MDL). They later extended the MDL method to three-dimensional shapes [29]. The MDL approach has been shown to produce qualitatively similar results to the method of Kotcheff and Taylor [103, 99]. The idea of an optimized correspondence model represents an important advance in shape modeling technology. The optimized models of Kotcheff, Taylor, and Davies are the inspiration for some of the key ideas presented in this dissertation and are discussed in more detail in Section 2.2.2.

1.3 Shape Analysis in Modern Biology and Medicine

Many important fields of basic research in medicine and biology are beginning to rely on computationally-derived statistical models of shape in their research. Typically, researchers are interested in both characterizing normal shape variation and in hypothesis testing to evaluate or demonstrate shape differences between populations. Two areas of biomedical research are of particular interest as application areas for the work in this dissertation. The first is the study of the human brain, including clinical psychiatry and neuroanatomy. The second is human genetics and, in particular, investigations into relationship between genotype and phenotype using small-animal models of growth and disease. Both of these application areas present unique challenges for computational shape analysis, which are among the topics addressed in this dissertation.

1.3.1 Neurobiology

In vivo imaging of brain anatomy has made shape analysis a feasible tool to study many problems in neuroanatomy and clinical psychiatry, such as longitudinal developmental studies and the progression of neurological disorders [97, 89]. In developmental analyses, such as pediatric neurodevelopment, quantitative magnetic resonance imaging has significantly advanced our understanding of brain development during childhood and adolescence. Many clinical studies, however, still rely on traditional morphometrics, such as intracranial volume and the volumes of brain lobes and subcortical structures [69]. Computational shape analysis in this context therefore promises to give not only new basic insights into the process of development, but also potentially provide new diagnostic measures of individuals against normative models. Regression models of shape versus age, for example, could provide models of growth and age-related variability. The joint analysis of complexes of multiple anatomical structures, or what this dissertation will refer to as *shape complexes*, is also of increasing interest to clinical psychiatrists because certain spectrum disorders, such as autism, for example, are thought to represent a confluence of several underlying abnormalities, impacting the relationships between brain regions [20]. Better statistical models than are currently available are needed to capture those relationships and gain insight into the mechanisms and development of such diseases.

1.3.2 Genetic Phenotyping

Since the invention of gene-targeting technologies, genetically modified mice have become commonplace models for the study of human development and disease [16]. By targeting specific genes for underexpression or overexpression, researchers can create specific alterations in a mouse genome that result in different patterns of anatomical growth and form, or *phenotype*. The phenotype of a mutant mouse population can be contrasted with that of a normal mouse population in order to gain insight into the functionality of the targeted genes.

Small animal phenotyping is a problem for which traditional and geometric morphometric techniques have traditionally been used, but for which computationally-derived shape models hold great promise. To measure differences in phenotype, researchers have typically relied on light microscopy, histology, and traditional morphometric descriptions of anatomy, such as the statistics of length and volume [68, 76]. Using these traditional morphometrics, shape features are typically identified and parameterized in advance of statistical analysis. The analysis is therefore limited by the choice of features and by the feature scales. More comprehensive and detailed representations of shape may allow for observations of genetic expression that have not been possible with morphometric approaches. Additionally, computational models have the advantage that they are volumetric, noninvasive, and more easily adapted to high-throughput phenotypic screening. Several researchers have proposed computational shape models for mouse brain phenotype using deformable registration between groups of mean MRI images [83] and deformable registration to a normal atlas [19]. To date, however, there has been little work done to investigate the use of correspondence models for genetic phenotyping, which is a focus of this dissertation.

1.3.3 Other Applications

Computationally-derived models of shape may potentially impact many other important areas of biomedical study. In orthopedics, for example, researchers have applied shape models to the study of bone shapes and pathologies. In [48], Gregory, et al. use a point-based modeling approach called Active Shape Modeling to study the correlation between proximal femur shapes and the rate of hip fracture in women. As with many other clinical shape studies, the goal of the hip fracture study is to develop an image-based assessment of the risk of the occurrence of the pathology. Active Shape Modeling, which is a technology developed by Cootes and Taylor [21], has also been heavily used in model-guided image segmentation algorithms for a wide variety of biomedical problems (e.g., [92, 110]). Computationally-derived models of shape may also be promising tools for the study of anthropology and paleontology. At the Smithsonian Institution, for example, Mathew Tocheri studies the form and function of primate hands and feet and their evolutionary relationships within the hominid clade [106]. Where previously Tocheri and colleagues have used traditional morphometric approaches, they are now acquiring CT scans of neandertal fossils for computationally-derived models. In Chapter 5, I present an application of the work in this dissertation to another study with implications for anthropology.

1.4 Technical Challenges for Biomedical Shape Analysis

While computational models of shape offer increased power for testing of biological hypotheses of shape and shape differences and are promising new tools for new areas of biomedical research, their development and application present some significant engineering challenges. These challenges are the primary motivation for this dissertation work. First, as introduced in Section 1.2, there are many different possibilities for the representation of the shape, which limits the class of problems to which the model may be applied. Related to shape representation, there is a second issue of model selection and avoidance of statistical pitfalls, such as overfitting of the model to the data. A third challenge is how to statistically analyze the high-dimensional correspondence data and how to effectively visualize statistical results and geometric variation. Finally, a significant challenge is how to reduce the complexity and computational requirements for shape modeling so that these methods may become more accessible to the average biomedical researcher.

1.4.1 Shape Representation

A shape analysis algorithm must adopt a mathematical representation of shape. For the purposes of this dissertation, representations of shape can be roughly classified into two categories: *parametric* models and *nonparametric* models. A parametric representation defines shape as a function of a set of parameters. The SPHARM methodology [15, 98], for example, represents shape boundaries as a linear combination of spherical harmonic basis functions. Other methods have been proposed that rely on wavelet basis functions [81] and Fourier decompositions

[101]. A nonparametric representation, on the other hand, does not assume an underlying mathematical model and the surface is represented implicitly. The interface between the foreground and background pixels in a binary segmentation volume, for example, is a nonparametric surface representation. Landmark models are also a type of undetailed surface representation. Specifically, this dissertation is concerned with nonparametric particle-based surface representations, which are discussed in Section 2.1.4.

The choice of a specific shape representation defines, and may also limit, the class of shapes that can be modeled. To be applicable to the full range of shape analysis problems in biomedicine, a modeling methodology must be capable of representing different topological classes of shapes and even changes in topology. Correspondence-based computational models of shape in the literature rely on spherical parameterizations of shape and can therefore represent only manifold surfaces with spherical topologies. This limitation also precludes the analysis of many important *open-boundary* anatomical structures such as the diaphragm [24] or of localized regions of a structure, such as head of a bone. Another important consideration is that medical or biological shapes are typically derived from the interfaces between organs or tissue types and usually defined implicitly in the form of segmented volumes, rather than explicit parameterizations, triangulations, or surface point samples. Such representations therefore require additional preprocessing steps that may limit the fidelity of the model and introduce error.

Automatic selection of correspondences for nonparametric shape representations has been explored in the context of surface registration [6], but because such methods are typically limited to pairwise correspondences and assume a fixed set of surface point samples, they are not sufficient for the analysis of sets of segmented volumes. While the sliding landmark method of Dalal, et al. does not assume fixed surface correspondences, it still relies on a pair-wise registrations to a fixed template, which must be chosen by another procedure in advance. Image-registration-based methods for shape modeling (Section 1.2) are nonparametric in nature, but they do not deal with the problems of selecting surface landmarks using statistical properties of the set of shape samples and establishing geometrically accurate surface representations.

1.4.2 The Model Selection Problem

As introduced in Section 1.2, the problem of how to automatically choose correspondences that simultaneously encode shape geometry and efficiently represent its variability within a population is fundamentally ill-posed and requires a regularization process. Outside of the shape analysis literature, a standard approach to statistical model selection is to choose a model that explains the sample data in an optimally efficient way, while avoiding overfitting to the data [50, 105]. The spherical harmonic, wavelet, and sampled medial mesh correspondence methods (Section 1.2), however, use a purely geometric solution; they seek only consistently regular parameterizations, not optimal correspondences [15, 97, 81]. The correspondence methods that minimize information content across the shape samples are capable of generating more powerful shape models [70, 27], but a significant

concern is how to regularize the models so that they effectively balance the inherent trade-off between model compactness and good representations of shape geometry.

1.4.3 Statistical Analysis and Interpretability

The geometric morphometrics community has established a standard approach to statistical analysis for landmark data that relies on traditional multivariate statistical methods and generalized Procrustes analysis [1]. The statistical analysis of higher-dimensional, computationally-derived models, however, is more difficult due the larger numbers of degrees of freedom and the challenge of obtaining sufficient numbers of subjects to ensure statistical power. While the correspondence shape modeling literature has proposed some ideas, it has not reached a consensus regarding a systematic approach (e.g. [27, 44, 98, 46]). In addition, many of the proposed methods focus on the special case of hypothesis testing in the absence of explanatory or confounding variables. Thus, there is significant work to be done to investigate the application of standard tools such as regression analysis, multiple-analysis of variance, and multiple-analysis of covariance to correspondence models. Another important consideration is the interpretability of the statistical result in the context of the geometry of the anatomy. Most studies with landmark-based data use Procrustes plots and thin-plate-spline warp visualizations, for example, to understand the variability in shape populations [11]. A similar set of standard visualization techniques is needed for correspondence models in order for these new tools to be more widely accepted by the biomedical community.

Hypothesis testing for group differences using parametric and correspondence models of shape has been most extensively investigated for comparative studies of brain anatomy, where a number of different approaches have been proposed. Among the more popular methods has been that of Styner, et al., who propose point-wise hypothesis testing on shape correspondence points derived from spherical harmonic parameterizations [98, 81]. This approach is appealing because it reveals local shape differences whose significance can be mapped directly onto the shape geometry, but the results require substantial correction for multiple comparison, which can significantly reduce statistical power. Nonlinear statistical metrics have been proposed for medial axis shape models, but these methods are not generally applicable to correspondence models [102, 46]. Golland, et al. propose the use of a support vector machine to construct classifiers based on shape models and then use leave-one-out classification rates as test statistics [44]. A similar method using linear discriminant classifiers has been used by Davies, et al. to investigate hippocampus shape in schizophrenia populations using MDL-based shape models [28]. While the use of classification rates as test statistics is potentially an interesting new approach to high-dimensional statistics, more work is needed to evaluate the effect of the additional free parameters in the analysis and whether these methods offer higher statistical power in general. Another concern is how to interpret and visualize the results for scientists who are not familiar with concepts in machine learning.

1.4.4 Ease-of-Use and Computational Complexity

To be most useful to a biologist or clinician, a tool for shape measurement should be relatively simple to use and robust with respect to the application. Manual landmarking methods, for example, may be time consuming, but the procedure for constructing landmark data is relatively straightforward and typically only requires standard equipment such as a light microscope or camera. In addition, manual landmarking is robust in the sense that it can be applied equally to a wide variety of data types. The challenge for computationally-derived models is therefore to increase the speed and power of the analysis, while maintaining the simplicity and robustness of traditional methods. The repeatability of shape measurements is also of paramount importance. For biologists, proper evaluation of hypotheses often requires that the results be reproducible by other researchers. Computationally-derived models of shape have the potential for greater reproducibility than landmark models, which can be biased by the human factors inherent in the annotation process. Reproducibility, however, may be adversely affected if the modeling process incorporates too many free parameters or too many preprocessing steps. Finally, shape modeling methods that require excessive computational power are not as likely to be used in an average biology laboratory. Ideally, shape modeling should be accessible to researchers on standard desktop computers.

A turn-key correspondence-point shape modeling application for the desktops of the biomedical community will require new, less complicated, and more robust methods. Current approaches are often slow, difficult to implement, and hard to use because of the significant amount of preprocessing involved and the need to tune many free parameters. The three-dimensional MDL correspondence method, for example, relies on spherical parameterizations and subdivisions of an octahedral base shape with smoothed updates that are represented as Cauchy kernels [29]. The initial parameterization must be obtained through another process such as [15], which relaxes a spherical parameterization onto an input mesh. The overall procedure requires significant data preprocessing, including a sequence of optimizations—first to establish the parameterization and then on the correspondences—each of which entails a set of free parameters or inputs in addition to the segmented volumes. Regularization of the solution requires additional free parameterizations and the a-priori choice of anchor shapes [29]. Typical processing times for around 1000 correspondences in small populations ($N \approx 20$) are on the order of hours to days for MDL models [24]. Computation times may be similarly high even for unoptimized correspondence models, such as spherical harmonic models [24].

1.5 Research Contributions

Motivated by the challenges outlined in the previous section, this dissertation proposes a new computational framework for statistical shape modeling and analysis and evaluates that framework in the context of several scientific biomedical studies. The modeling framework uses an optimized, correspondence-point representation of shape. Correspondences are modeled as interacting *particle systems*, a technology from the computer graphics literature, and are positioned on shape surfaces automatically by optimizing the information content of the model. In

this dissertation, the correspondence optimization framework is referred to as the *particle-based modeling* (PBM) framework.

Information content in the PBM optimization is minimized using a variational formulation of the combined entropy of the sample distribution in shape space and the particle distributions on the shape surfaces. The procedure explicitly models the inherent tradeoff between the statistical simplicity of the model (compactness) and the accuracy of the shape representations (good surface samplings), balancing the two commensurate terms with a single free parameter. This entropy-based approach provides a natural equivalence of information content and eliminates the need for ad-hoc regularization strategies based on anchor shapes, templates, and reparameterizations. The method also draws a clear distinction between the objective function and the minimization process and thus can more readily incorporate additional information such as high-order geometric information for adaptive sampling. The proposed particle-system optimization is fully nonparametric; it is therefore directly applicable to nonspherical topologies and can be applied directly to volumetric image data (image segmentations).

The simplicity of the shape representation allows the framework to extend more easily to a larger class of problems than existing methodologies. This dissertation proposes several important extensions to the basic modeling framework, including its application to open surfaces, correspondences based on arbitrary surface features, modeling of shape complexes, and correspondence with regression against explanatory variables. Each of these extensions is evaluated in the context of a research problem in biomedicine.

As discussed in the previous section, much work remains to demonstrate the effectiveness of correspondence models for standard multivariate statistical methods, such as the analysis of variance and shape regression. This dissertation proposes statistical analysis methods for correspondence models that are based on accepted methods currently in use by the morphometrics community. In addition to a straightforward method for hypothesis testing and new methods for visualizing significant group differences, this dissertation also illustrates the application of analysis of variance techniques to correspondence data and the visualization of average forms in these models.

All of the proposed modeling and analysis procedures outlined above are evaluated in the context of real biomedical research problems. Specifically, this dissertation describes results for comparative studies of subcortical brain structures between normal, autistic, and schizophrenic populations, the estimation of cortical surface thickness in human populations, regression modeling of neonatal head shape, phenotyping of the forelimb bones for mice deficient in the *Hox-d11* gene, and phenotyping of the basioccipital bone in mice deficient in the *Pax7* gene.

The following list outlines the specific contributions of the dissertation.

- The particle-based modeling (PBM) framework: A mathematical framework and a robust numerical algorithm implementation for computing optimized correspondence-point shape models using an entropy-based optimization and particle-system technology.
- Extensions to the PBM mathematical framework and implementation for several, more general classes of shape analysis problems, including (a) the analysis of shapes with open surfaces, (b) the analysis of sets of disconnected

surfaces, or *shape complexes*, (c) the generalization to correspondences based on generic functions of position, and (d) modeling shape regression against independent variables (such as size and age).

- Statistical methods for hypothesis testing, multiple analysis of covariance, and regression analysis that are suitable for use with the particle-based correspondence algorithm and several new methods for visualization and interpretation of these statistical results.
- Results from collaborations with biomedical researchers that validate the usefulness of the above methods in several key research areas, including clinical psychiatry and neuroanatomy and phenotyping studies that have applications to human development and anthropology.

1.6 Document Organization

The remainder of this document is organized as follows. Chapter 2 describes the mathematical formulation of the basic PBM shape modeling framework, its implementation, and some preliminary validation experiments on synthetic data. Chapter 3 describes the mathematical extensions to the basic PBM framework along with their implementations. Chapter 4 discusses statistical analysis of correspondence models computed using the PBM algorithm, including some specific approaches to hypothesis testing, regression analysis, and new methods for the visualization and interpretation of significant shape differences. Chapter 5 describes the six different biomedical studies to which I applied the PBM framework and its extensions. Chapter 6 summarizes the major conclusions of this dissertation work and suggests some further related work.

CHAPTER 2

SHAPE MODELING FRAMEWORK

This chapter presents a new methodology, called the *particle-based modeling* (PBM) algorithm, for constructing sets of correspondence points on a collection of shapes. The method uses a particle-system surface representation and an entropy optimization scheme to automatically place the correspondences. The chapter begins with some important mathematical background and related work, including a discussion of shape representation, the correspondence model, and the prevailing ideas for its optimization. Next, the mathematical formulation of the PBM method is developed, followed by a discussion of important considerations regarding data preprocessing and initializing the optimization. The chapter concludes with a discussion of the properties of the PBM optimization, including a comparison with the minimum description length method (MDL) and several validation experiments on synthetic datasets.

2.1 Shape and Shape Representation

2.1.1 Shape and its Analysis

The *shape* of an object is defined mathematically as all the geometric information that is distinct from translation, orientation, and overall size [34, 65]. Stated another way, shape is what remains when you remove location, orientation, and uniform scale. The goal of statistical *shape analysis* is to estimate the variability in the geometric shape of objects in a population using a finite subset of shapes taken from that population. The subset of shapes from the population is called the *sample*, or the *sample set*. In addition to an estimate of variability, a shape model also typically includes an estimate of the mean of the population. Using the estimates of the mean and variability, statistical measurements can be formulated to test hypotheses. Researchers are often interested in testing for significant differences in the mean shapes of two populations, for example.

As an example of a shape analysis problem, consider Figure 2.1, which shows a random sample ($N = 10$) from a population of tori. Tori are convenient examples because their shape can be fully described with only two parameters, the distance from the centerline of the tube to the center of the torus, R , and the smaller radius of the tube, r . The samples shown in the figure are randomly drawn from a Gaussian distribution of mean $R = 30$ and mean $r = 15$, with the standard deviations $\sigma_R = 3$ and $\sigma_r = 5$, respectively. The goal of shape analysis for this population of shapes is to somehow estimate the sample mean shape and the variability in shape from this sample set. In general, the number of shape parameters is also not

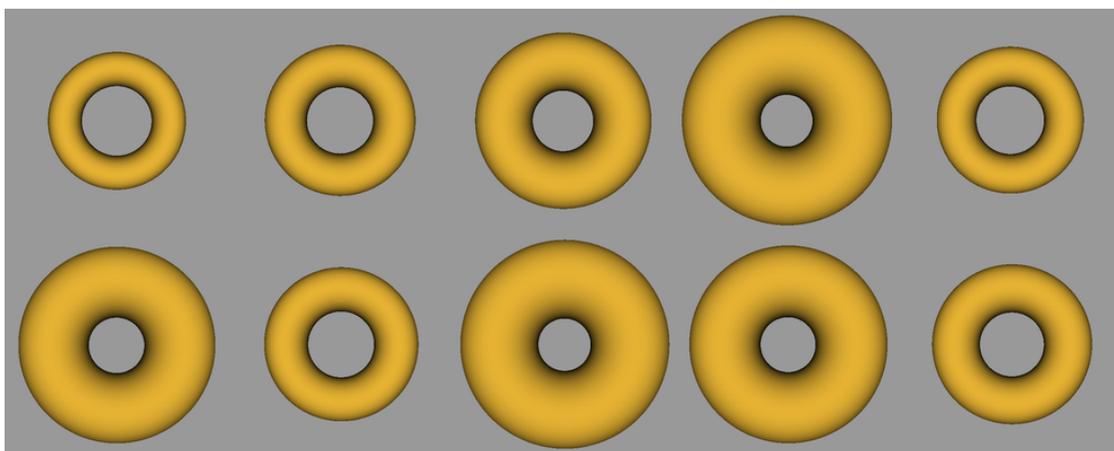


Figure 2.1. A sample population of torus shapes that are randomly drawn from a bivariate Gaussian distribution parameterized by large radius R and small radius r .

known in advance. If we hypothesize that there are a finite number of “true” shape parameters that describe variability, as is the case with the tori, then an additional goal for shape analysis is often to estimate these underlying parameters.

2.1.2 Shape Representation

In biomedicine, shape analysis is applied to discrete anatomical objects whose geometry is defined by either interfaces between tissue types or by functional sub-regions of tissue. Sometimes, both considerations may be factors. Some examples include bones, such as the metacarpal and carpal bones of the hand, muscles, such as the diaphragm, and internal organs like the liver or heart. Of course, anatomical geometry is not often well described by a two-parameter representation, as in the contrived example of Figure 2.1, and therefore a more general mathematical representation for shape must be employed.

Shape representations for computational purposes typically fall into one of two categories: surface representations and region representations. Algorithms that consider shape as embedded in images, such as voxel-based morphometry and the diffeomorphic methods (Section 1.2), are examples of region-based shape representations. A region-based approach is typically used when geometric features of the interior of the shape are of interest, such as with functional brain imaging studies or in other cases where explicit descriptions of the surfaces of anatomical objects are difficult to formulate or are not important for the analysis. By contrast, a surface-based representation of shape only models shape surface geometry. This dissertation is concerned with models of the surfaces of shapes, which are well suited for studying many important classes of biomedical problems, especially those for which the traditional morphometric techniques were originally developed.

A *surface*, for the purposes of this dissertation, is mathematically defined as a smooth manifold of co-dimension one, embedded in a Euclidean vector space. The current discussion assumes a closed surface. Open surfaces are discussed in the next chapter. The dimensionality of the embedding space is typically either two or three. In \mathcal{R}^3 , for example, the surface manifold has a dimension of 2, and in \mathcal{R}^2 it has a dimension of 1, where it is often referred to as a *contour*. This dissertation is primarily concerned with modeling the shape of the surfaces of objects embedded in three-dimensional space, but some examples will be given for models of the contours of objects in two dimensions. It should also be noted, however, that there is a class of biomedical problems for which higher dimensional shape modeling is of interest, such as the analysis of shape variability in four-dimensional time-series data.

A surface may be either *closed* or *open*, depending on whether or not the manifold has a boundary. A closed surface is a manifold with no boundary, and contains only a single connected component. Spheres and tori are good examples of closed surfaces. A biological example of a closed surface is a bone, or a brain structure such as the putamen or caudate. An open surface is a manifold with a boundary, such as a disk, which is a two-dimensional region with a circular boundary. Open surface boundaries in biology are sometimes defined by a functional subregion of a larger structure, such as a section of the cortical surface. Another example of an open surface is the head of a bone. The concepts of open and closed surfaces are illustrated in Figure 2.2, which compares a closed surface representation of a

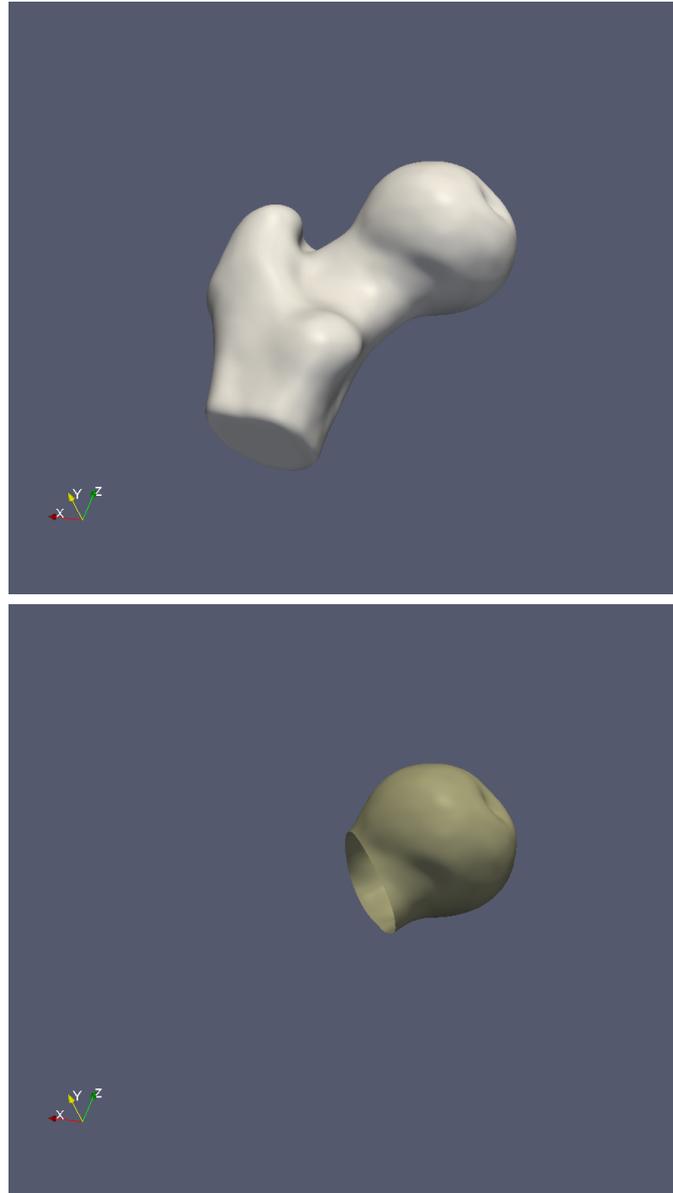


Figure 2.2. A comparison of closed bone surface (top panel) with an open bone surface (bottom panel).

complete femur bone structure (top panel) with an open surface representation of the head of the femur (bottom panel). In the open surface, the boundary is defined by a contour around the bone shaft that is specified as the intersection of a cutting plane with the closed bone surface in the top panel.

2.1.3 Surface Description

A surface must be mathematically described, or encoded, in order to be used computationally. Typically, the geometry of a structure in an image is derived by a process of annotating the structure’s interior or boundary in the space of the image, thus differentiating the object from the background and surrounding structures. The annotation process is known as image *segmentation*, and may be fully automated or incorporate various degrees of human supervision. The output of a segmentation process is a geometric description of the surface of the object of interest, but the format of the description varies, depending on the algorithm. Some segmentation algorithms produce parametric descriptions of the surface geometry, but image-based descriptions, such as *label masks*, are more common. Label masks are binary images whose coordinate grid matches that of the original image and whose pixel values are labeled as either the foreground or the background of the object of interest. In a label mask, the surface geometry is described implicitly as the interface between the foreground and the background.

Often, a surface representation produced by a segmentation process cannot be used as direct input to a shape analysis algorithm. As discussed in Section 1.4, most state-of-the-art methods require preprocessing steps to construct a parametric description, such as spherical harmonics, medial-axis representations, or meshes. In general, label mask segmentations are also not directly useful because they are not suitable representations for finite numerical calculations of surface geometry and features, which assume the image is a sampling of a smooth function. However, a related surface representation, the *signed distance transform* of the label mask, does satisfy this criterion [67]. A distance transform is a graded, membership function, or fuzzy classification of the pixels within an object, with a surface implicitly defined as the zero level set [67]. In particular, distance transforms are useful for constructing particle-based representations of surfaces, which is discussed further in Section 2.1.4.

Like the label mask, the signed distance transform is an implicit, image-based surface representation that retains the grid coordinate system of the original image. Instead of “inside” and “outside” labels for the pixels, however, each position \mathbf{x} in the d -dimensional image is given a scalar value $f(\mathbf{x})$ of its distance from the implicit surface. The implicit surface Γ is given by the locations where $f = 0$. The inside of the surface is defined by $f < 0$ and the outside where $f > 0$. A distance transform can be easily converted back to a label mask format by simply labeling all pixels less than zero as “inside” and all pixels greater than zero as “outside”. Fast algorithms also exist for constructing mesh-based surface visualizations from distance transforms, such as marching cubes [75]. Figure 2.3 is an illustration of the concepts of a label mask segmentation and an associated distance transform image. The left-most panel shows a single slice of a segmentation of a bone structure from a volume. The middle panel depicts the same slice of a distance transform

representation of the structure. The right-most panel shows a marching cubes rendering of the zero level set of the distance transform. For the surface rendering, the binary volume was first processed by an algorithm to remove aliasing artifacts, which is discussed in more detail in Section 2.4.

When properly constructed, a distance transform has nice properties for finite numerical methods for computing surface features and geometry [67]. Several well-known algorithms have been proposed for constructing smooth distance transforms from label masks. The fast-marching method from Sethian, et al. [95], for example, is one of the most common methods. The fast marching method solves the boundary value problem

$$|\nabla f(\mathbf{x})|S = 1, \quad (2.1)$$

where $f = 0$ on the surface Γ . This equation models the propagation of the surface as it moves in its normal direction with speed S . Solving the equation for f gives the arrival time of the surface at a particular location, which is the distance transform function encoded in the image.

2.1.4 Particle Systems for Surface Representation

A related technology to distance transforms, is the work from the computer graphics literature on particle-system surface representations. Particle systems can be used to manipulate or sample implicit surfaces by constructing dynamic, point-set representations. Originally, static point-set surfaces were developed in the context of surface visualization as an alternative to mesh-based representations. In many contexts, point sets allow more flexible control over the accuracy of the representation and the speed of the rendering process [3, 38]. Another advantage of point-set surfaces is that they do not require a specific parameterization and do not impose topological limitations; surfaces can be locally reconstructed or subdivided as needed [10]. Witken and Heckbert introduced the idea of modeling a point set as a system of interacting *particles* that are constrained to lie on an implicit surface. Particles interact with one another with mutually repelling forces, such as electrostatic charge, so that they find distributions that optimally cover, and therefore describe, the surface geometry [115]. Meyer, et al. proposed numerically robust extensions to this approach, including a new class of radial-basis energy functions and methods for curvature-adaptive surface sampling [79].

The method for shape correspondence proposed in this dissertation uses the idea of the particle system surface representation first proposed by Witkin and Heckbert and adapts the numerical approaches of Meyer, et al. The PBM algorithm proposed in this dissertation uses a set of interacting particle systems, one for each shape in the sample, to produce optimal sets of surface correspondences. Adopting a point-based surface representation avoids many of the problems inherent in parametric representations (see Section 1.4.1), such as the limitation to specific topologies and processing steps necessary to construct parameterizations. Another advantage is that, unlike representations that rely on surface meshes, particles do not have fixed neighbors and are free to move past one another to form different neighborhood configurations during the optimization process. This property means that the result

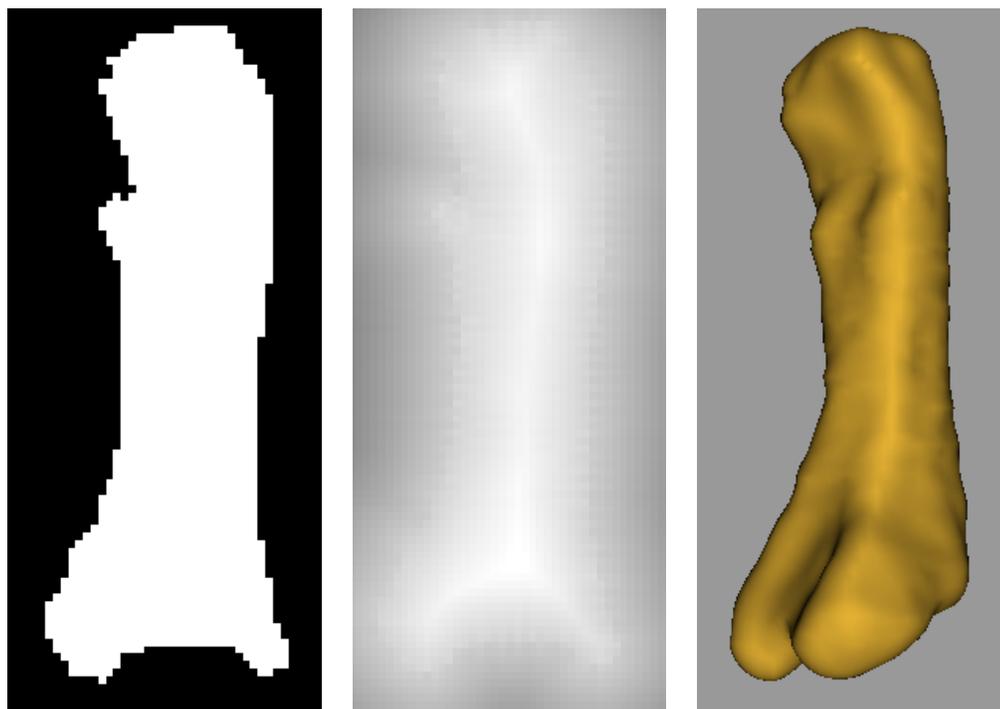


Figure 2.3. An illustration of a single slice from a segmentation (left), the corresponding slice of the distance transform (middle), and an antialiased surface rendering of the zero level set of the distance transform (right).

is less constrained by the initialization and can potentially produce a less biased, more fully optimized model.

Figure 2.4 illustrates the concept of a particle system representation of an implicit surface, in this case, the bone shape from Figure 2.2. The panels from left to right show an increasing number of particles placed on the surface and the resulting surface reconstruction from the particles. The surface reconstruction is done using the method for unorganized sets of points given by Hoppe, et al. [54]. Particles are placed using the method that is described in Section 2.3.3. The number of particles doubles in each panel (256, 512, and 1024). As the particle count increases, so does the detail of the corresponding surface reconstruction.

2.2 Correspondence Shape Models

2.2.1 Correspondences

Correspondence-point models describe shape variation by choosing a discrete set of comparable, or corresponding, sets of points on shape surfaces whose relative positions can be analyzed. The correspondence model is defined as follows. Consider a

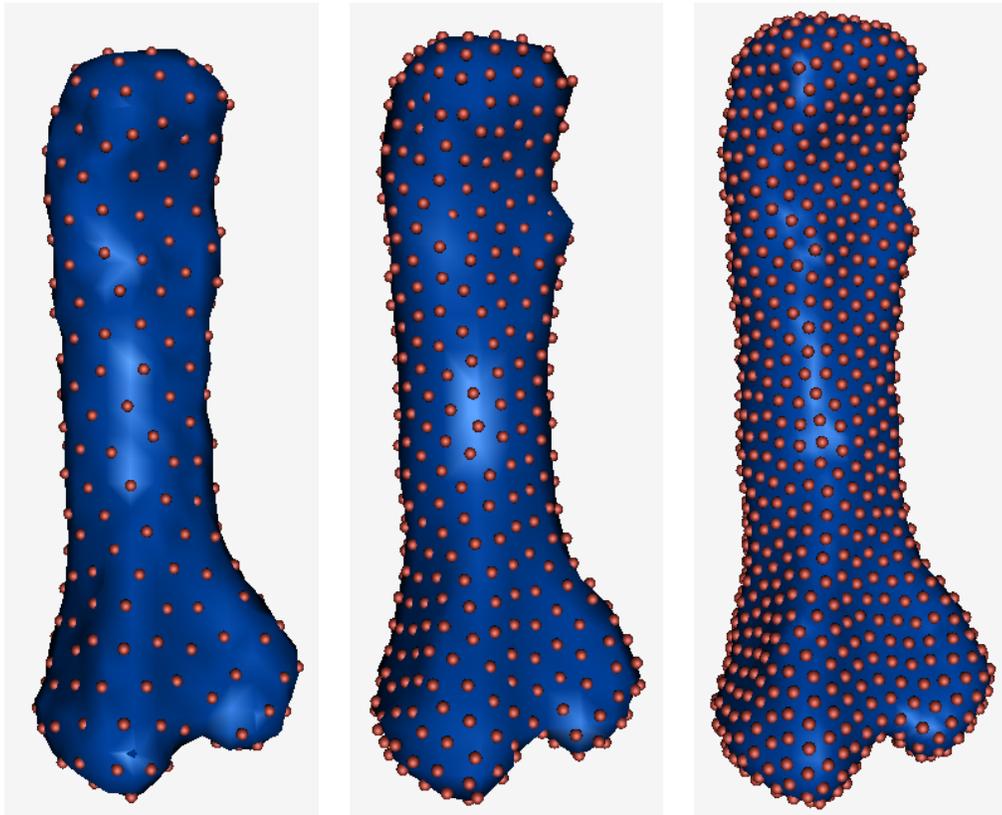


Figure 2.4. An illustration of surface representation using particle systems. From left to right, increasing the number of particles results in an increased accuracy of the surface reconstruction.

statistical sample of N surface representations drawn from a population of surfaces. The surface representations are embedded in a d -dimensional Cartesian space. A model for shape variation is constructed by choosing a set of M , d -dimensional points on each of the N surfaces. Each of the points is called a *correspondence* point. Collectively, the set of M points is known as the *configuration*, after Dryden and Mardia [34], and the space of all possible configurations is the *configuration space*. The *configuration matrix*, \mathbf{C} , is the $M \times d$ matrix of Cartesian coordinates in a configuration. The ordering of the points in the N configurations, and the rows in the configuration matrices, implies correspondence among the surfaces. Row $k \leq M$ in configuration matrix i , for example, corresponds to row k in configuration matrix j . The variation in the rows of the configuration matrices describes geometric variation in shape.

Each configuration can be mapped to a single point X in a $d \times M$ -dimensional *shape space* by concatenating the correspondence coordinate positions into a single vector. The mapping to the dual shape space is invertible. The sample set forms a distribution in shape space, whose statistical properties can be estimated. Figure 2.5 illustrates the mapping from configuration space to shape space and shows a cartoon depiction of the resulting distribution in the high-dimensional shape space. The top of the figure shows a sample set of N hand contours. At the bottom-left of the figure, a configuration for one of the hand contours (sample i) is shown along with its mapping to the shape space vector. The bottom-right corner of the figure shows the (hypothetical) position of sample i relative to the (hypothetical) distribution of the other samples in the shape space. Note that while this cartoon depiction of shape space is two-dimensional, the true shape space has dM dimensions (and is therefore very hard to visualize on a page).

For analysis, sets of configurations are usually aligned within a common d -dimensional coordinate frame by a rotation, translation, and scaling to remove the geometric information unrelated to shape variation. Goodall's model of shape [45, 91] describes each of these nonshape components and the residual variation around the mean correspondence configuration. For configuration matrix C_i , the model is given by

$$\mathbf{C}_i = a_i \boldsymbol{\mu} + \mathbf{E}_i \mathbf{R}_i + \mathbf{1} \mathbf{t}_i, \quad (2.2)$$

where a_i is a scalar representing the relative size of specimen i relative to the mean size, \mathbf{E}_i are the residuals from the mean configuration $\boldsymbol{\mu}$, \mathbf{R}_i is a rotation matrix describing the orientation of sample i , $\mathbf{1}$ is a d -dimensional vector of 1s, and \mathbf{t}_i is a translation vector describing the locational information for sample i .

The most common method for estimating $\boldsymbol{\mu}$, and the nonshape components \mathbf{R}_i , \mathbf{t}_i , and a_i is generalized Procrustes analysis (GPA) [68, 47, 9, 34]. When transformed using GPA, correspondences are said to be in *Procrustes space*. Statistical analysis is commonly done in Procrustes space because, for reasonably similar sets of shapes, distance measures between Procrustes coordinates have been shown to be good linear approximations to the geodesic distances in Kendall's shape space [34, 65]. GPA and the statistical analysis of correspondence models is discussed further in Chapter 4.

Correspondence models are distinct from parametric models in that they consist only of a relatively sparse set of correlated points. Recall from Section 1.4.1, that

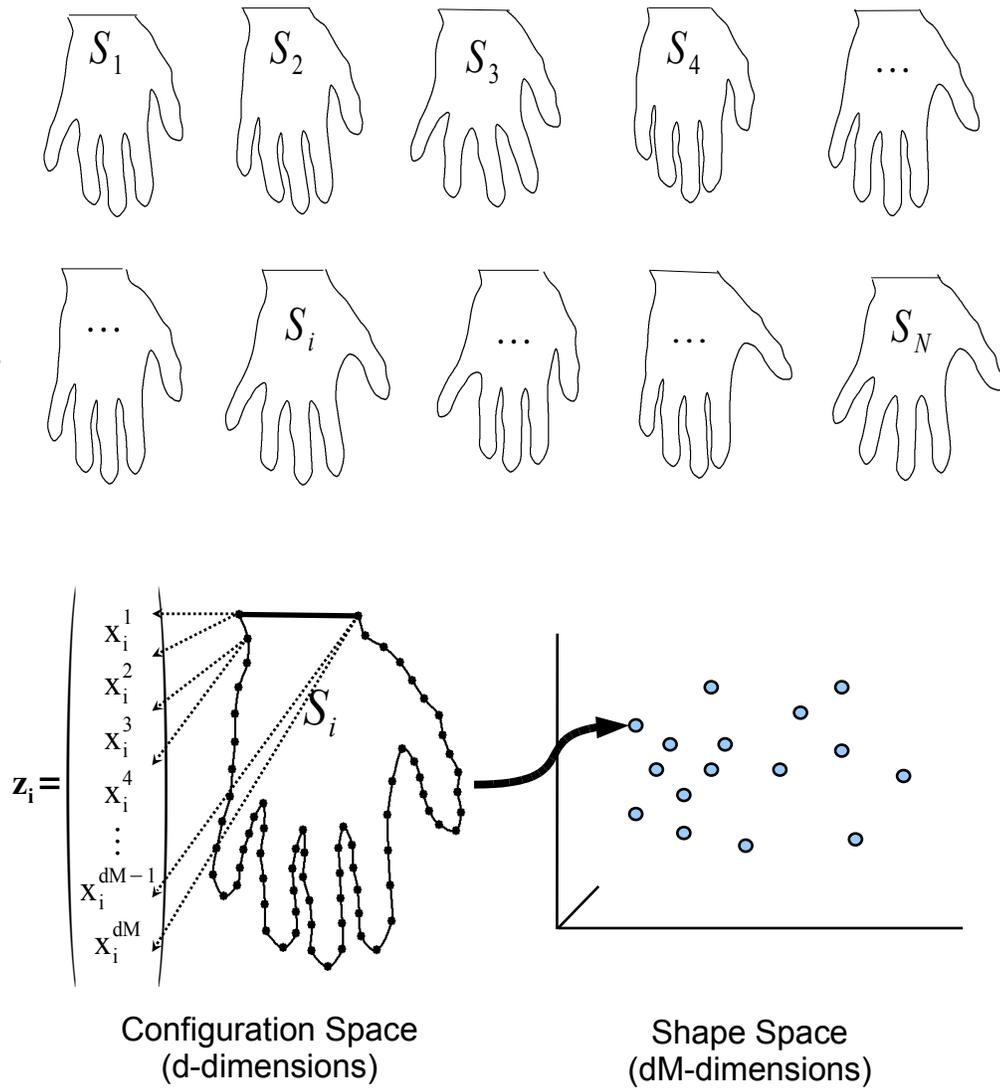


Figure 2.5. An illustration of the relationship between configuration space and shape space.

a parametric surface is explicitly defined using a set of mathematical functions. Parametric shape models allow for continuous mappings between surfaces, but are not concerned with choosing specific correspondences. Correspondence among parametric surfaces is therefore implicitly defined and can be sampled on the mappings from the parametric model. Another process, however, is required in order to determine how to choose optimal sets of correspondences.

2.2.2 Optimized Correspondences

As discussed in the preceding chapter, a correspondence model that is optimized with respect to its information content represents an important advance in shape modeling technology. The goals for optimized model construction are consistent with standard approaches in the statistics literature: to choose a model that efficiently explains the sample data while avoiding overfitting to that data [50, 105]. The result is often a model with fewer parameters and greater statistical power than one chosen by an ad-hoc process [50]. The PBM algorithm proposed in this dissertation is an optimized modeling method. It is inspired by several key ideas from the literature that are discussed in this section.

The problem of how to choose correspondence configurations with minimal information content, so that they efficiently explain the sample data, was first addressed by Kotcheff and Taylor. In [70], they propose the idea of manipulating correspondence configurations through an optimization process on their resulting shape-space distribution. This basic strategy is illustrated in Figure 2.6. From left to right, the top of the figure shows a change in single configuration for a hand contour. A cartoon depiction of the resulting movement of its position in shape space is shown from left to right at the bottom of the figure. The bottom of the figure also shows the hypothetical change in the distribution of the hand shape samples after some optimization process. In this case, the distribution has been modified by the optimization process to have less variance.

Specifically, the method of Kotcheff and Taylor given in [70] applies to shape contours in two dimensions, and minimizes a function of the magnitude of the covariance of the shape space distribution. The method represents a population of shape contours with a set of M correspondences, initialized at equal intervals along the contour. Each shape is then treated as a point in a $2M$ -dimensional shape space, with an associated covariance Σ . Correspondence configurations are then manipulated along the contours so that they minimize a cost function

$$\sum_k \log(\lambda_k + \alpha), \quad (2.3)$$

where λ_k are the eigenvalues of Σ . The parameter α is a regularization parameter that prevents the very thinnest modes (smallest eigenvalues) from dominating the process. This process is the same as minimizing

$$\log |\Sigma + \alpha I|, \quad (2.4)$$

where I is the identity matrix and $|\cdot|$ denotes the matrix determinant. Thus, we can see that this is effectively a minimization of the total variation associated with the model.

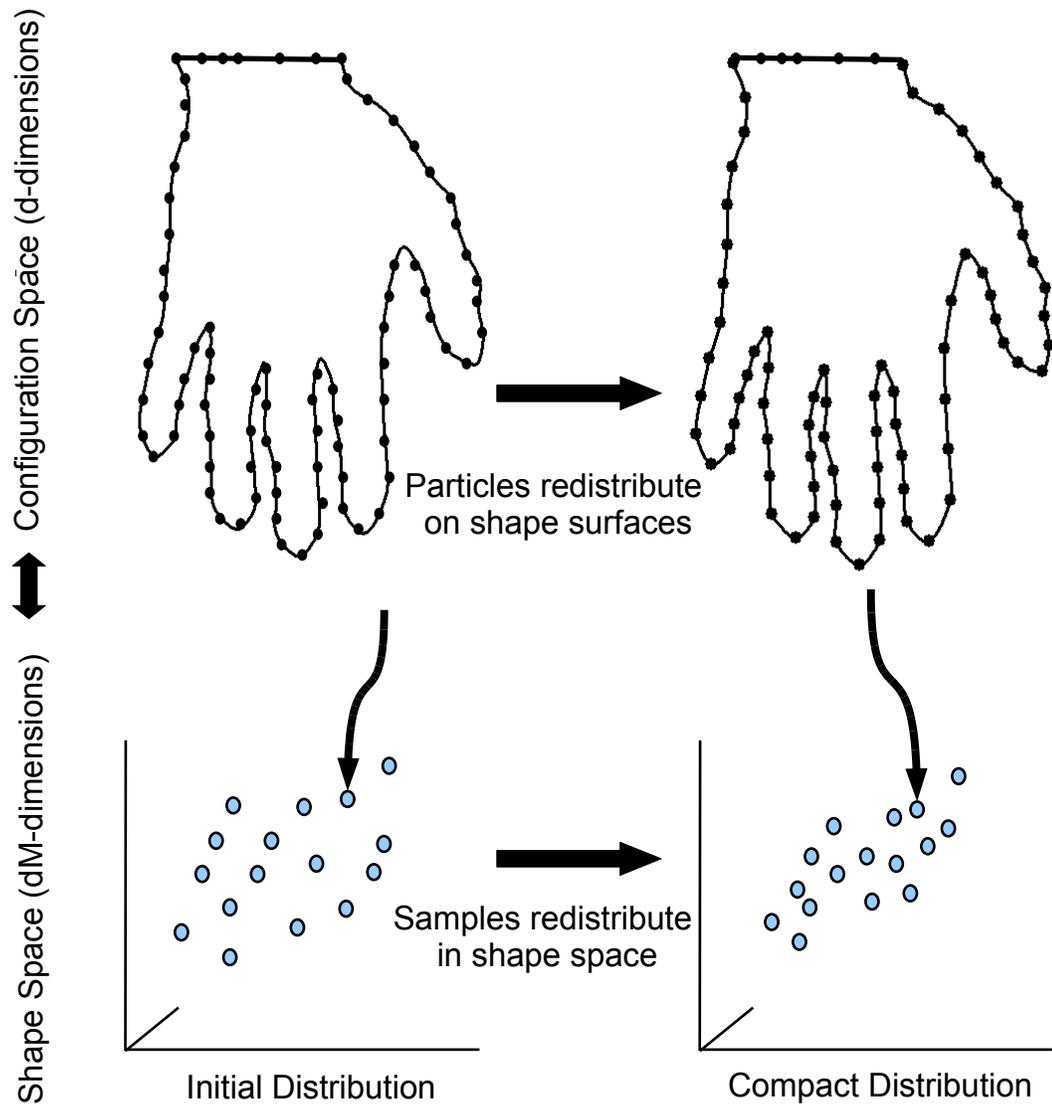


Figure 2.6. An illustration of the basic concepts of a correspondence configuration optimization.

Davies, et al. [30] propose a similar cost function for two-dimensional shapes based on *minimum description length* (MDL). The model is optimized by minimizing the cost of transmitting a principal component model of the correspondences. They use quantization arguments to limit the effects of thin modes, and to determine the optimal number of components that should influence the process. They propose a piecewise linear reparameterization and a hierarchical minimization scheme. In [29] they describe a three-dimensional extension to the MDL method. The extension relies on spherical parameterizations and subdivisions of an octahedral base shape. Correspondence positions are manipulated through the parameterization using smoothed updates represented as Cauchy kernels. The parameterization must be obtained through another process such as [15], which relaxes a spherical parameterization onto each input mesh from the sample set.

In Figure 2.6, the correspondence configurations before and after optimization are depicted as equally valid representations of the hand contour. A difficult problem, however, is how to effectively regularize the optimization so that this condition holds. There is often an inherent tradeoff between the compactness of the model in shape space and the accuracy of the shape representations. This trade-off is illustrated in Figure 2.7, which shows a degenerate, optimized case without regularization. In the figure, two-dimensional contours of a set of star shapes are shown overlaying one another. The correspondences have converged after an optimization process to lie only near the intersection points of the contours. The resulting model is compact, but probably not useful, given that it fails to represent the geometry of the samples. Several regularization strategies to avoid this problem in the basic MDL formulation have been proposed [29, 103], but they entail additional free parameters and assumptions about the quality of the initial parameterizations. One strategy, for example, constrains the solution so that it remains close to an anchor shape [29], which artificially limits the minimization process and biases the solution towards the anchor shape.

Despite concerns with the regularization strategies, however, the MDL approach to constructing correspondence models has been demonstrated to be effective for certain classes of shapes and represents the state-of-the-art in optimized correspondence modeling. Styner et al. [96] describe an empirical study that shows MDL improves correspondences relative to pure geometric regularization and notes that its performance is similar to that of $\min\text{-log} |\Sigma + \alpha I|$.

2.3 Particle-Based Modeling for Shape Correspondence

2.3.1 Motivation

It follows from the discussion in the previous section that there are several important challenges in advancing the state-of-the-art for optimized correspondence models. First, existing methods are specifically designed for closed shapes with spherical topologies. Many applications in biomedicine, however, require analysis of more diverse classes of shapes, such as shapes with boundaries, toroidal topologies, and collections of multiple, disconnected sets of surfaces (see Chapter 3). A second

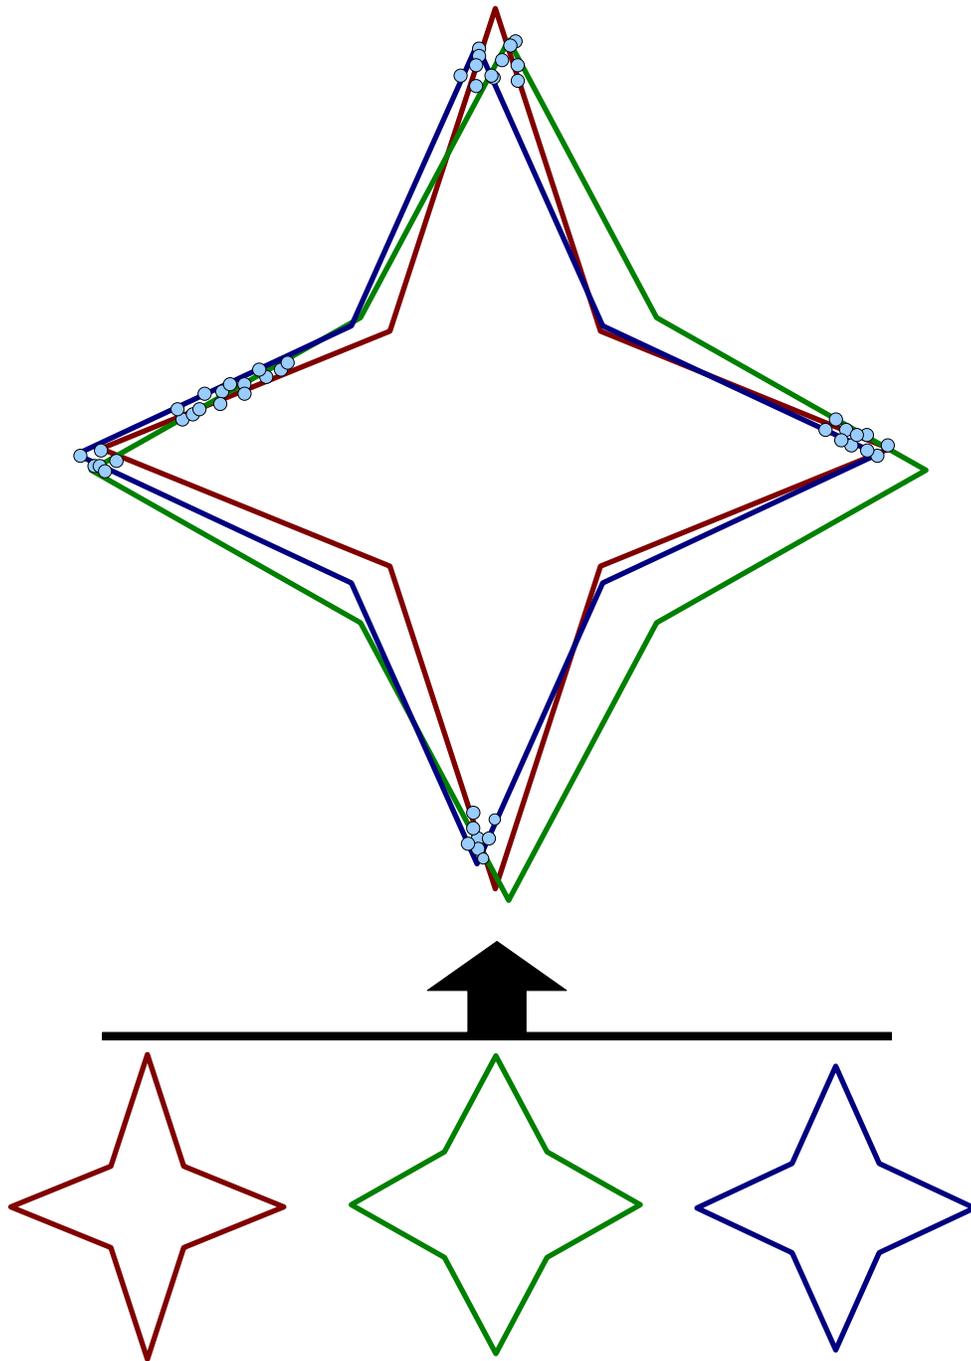


Figure 2.7. An illustration of the tradeoff between model compactness and the quality of shape representation.

challenge is how to limit the numbers of free parameters that must be tuned for modeling. This second consideration is important for the overall ease-of-use of the method and the reproducibility of the results. A third challenge is how to avoid ad-hoc regularization strategies, such as the minimum description length (MDL) approach that constrains solutions to be near an initial template shape [29, 103]. Finally, the overall complexity of the modeling algorithm affects the ease with which it can be implemented and the processing time that it requires to produce results.

The particle-based modeling (PBM) algorithm, which is described in this section, addresses the challenges outlined above through a new formulation of the idea of optimized correspondences. In this new formulation, surfaces are modeled non-parametrically as collections of dynamic particle systems, which are not constrained to a specific topology and can be applied to open surfaces. The particle system approach also avoids the algorithmic complexity and parameter tuning associated with constructing parameterizations. In the PBM approach, correspondences are modeled as freely moving particles in a system, without defining local topology or neighborhood connectivities. Their positions can therefore be *directly* optimized, instead of indirectly manipulated through a parameterization. In order to avoid assumptions about the initial quality of sample surface representations and ad-hoc approaches to regularizing the optimization, the PBM algorithm *explicitly constructs* good representations during the optimization procedure. The surface representations are constructed by an entropy measure on their distributions, which is the same criteria used to construct compact distributions of the samples in shape space.

The remainder of this section describes the PBM optimization process in more detail. The next section presents a brief overview of the algorithm. Sections 2.3.3-2.3.7 present the strategy for constructing surface representations during the optimization procedure. Section 2.3.8 describes the entropy minimization process in shape space that produces a compact statistical model. The section concludes with a discussion of the free parameters in the optimization.

2.3.2 Algorithm Overview

The PBM algorithm models correspondence positions as sets of dynamic particles that are constrained to lie on the surface of the sample set, as in the surface sampling methods described in Section 2.1.4. The optimization is based on the idea of treating correspondence position in configuration space as a random variable, while simultaneously treating correspondence configuration as a random variable. Correspondence positions are optimized by gradient descent on an energy function that balances the negative entropy of the distribution of particles in configuration space with the positive entropy of the distribution of the configurations in shape space.

More specifically, and with reference to Figure 2.5, the method is to consider $\mathbf{z}_k \in \mathfrak{R}^{dM}$, $k = \{1, 2, \dots, N\}$ both as observations on a $dM \times 1$ vector random variable \mathbf{Z} and as N samples of M observations on $N, d \times 1$ vector random variables \mathbf{X}_k . The optimization to establish correspondence minimizes the energy function

$$Q = H(\mathbf{Z}) - \sum_{k=1}^N H(\mathbf{X}_k), \quad (2.5)$$

where H is an estimation of differential entropy. Minimization of the first term in Q produces a compact distribution of samples in shape space, while the second term seeks uniformly-distributed correspondence positions on the shape surfaces for accurate shape representation. Each term is given in commensurate units of entropy, avoiding the need for a separate regularization strategy. Because correspondence points in this formulation are not tied to a specific parameterization, the method operates directly on volumetric data and extends easily to arbitrary shapes, even nonmanifold surfaces. It also avoids the additional complexity and processing time required by the parameterizations in methods such as MDL.

2.3.3 Surface Representation

Consider a single configuration for a shape surface $\mathcal{S} \subset \mathfrak{R}^d$. The configuration consists of a discrete set of M points, which are the correspondence positions. The PBM formulation represents these positions with a set of particles, whose positions are considered a sample on a vector random variable $\mathbf{X} \in \mathfrak{R}^d$, with an associated probability density function describing their distribution. This probability density function $p(\mathbf{X} = \mathbf{x})$ gives the probability of an observation \mathbf{x} on \mathbf{X} , denoted as $p(\mathbf{x})$. In the limit, the amount of information contained in the sample on \mathbf{X} is the differential entropy of $p(\mathbf{X})$,

$$H(\mathbf{X}) = - \int_{\mathcal{S}} p(\mathbf{X}) \log p(\mathbf{X}) dx = -E\{\log p(\mathbf{X})\}, \quad (2.6)$$

where $E\{\cdot\}$ is the expectation. When there are a sufficient number of points sampled from p , the expectation can be approximated by the sample mean [23], which gives

$$H(\mathbf{X}) \approx -\frac{1}{M} \sum_{i=1}^M \log p(\mathbf{x}_i). \quad (2.7)$$

The PBM algorithm manipulates particle positions using a gradient-descent optimization on a cost function C , that is an approximation of negative entropy,

$$C(\mathbf{x}_1, \dots, \mathbf{x}_M) \approx -H(\mathbf{X}). \quad (2.8)$$

The optimization problem is given by

$$\mathbf{z} = \arg \min_{\mathbf{z}} C(\mathbf{x}_1, \dots, \mathbf{x}_M) \text{ s.t. } \mathbf{x}_1, \dots, \mathbf{x}_M \in \mathcal{S}, \quad (2.9)$$

and uses a Gauss-Seidel update with forward differences. Each particle therefore moves with a time parameter and positional update,

$$\mathbf{x}_i \leftarrow \mathbf{x}_i - \gamma \frac{\partial C}{\partial \mathbf{x}_i}, \quad (2.10)$$

where γ is a time step. The partial gradient of C for particle i is

$$\frac{\partial C}{\partial \mathbf{x}_i} = \frac{\partial}{\partial \mathbf{x}_i} \frac{1}{M} \sum_{j=1}^M \log p(\mathbf{x}_j) = \frac{1}{M} \sum_{j=1}^M \frac{\frac{\partial}{\partial \mathbf{x}_i} p(\mathbf{x}_j)}{p(\mathbf{x}_j)}. \quad (2.11)$$

The gradient requires estimates of the probability $p(\mathbf{X} = \mathbf{x}_j)$. For distributions of particles on surfaces, a density function may be quite complex, which suggests a nonparametric, kernel-based approach. The PBM algorithm uses a Parzen windowing density estimation [86] that is based on the particle configurations. The probability of the position of a particle in this formulation is given by the mixture of multivariate Gaussian kernels,

$$p(\mathbf{x}, \sigma) \approx \frac{1}{M} \sum_{j=1}^M G(\mathbf{x} - \mathbf{x}_j, \sigma), \quad (2.12)$$

where $G(\mathbf{x} - \mathbf{x}_j, \sigma)$ is a d -dimensional, isotropic Gaussian with standard deviation σ . When $j = i$ in 2.11, the partial derivative of p with respect to particle position is

$$\frac{\partial}{\partial \mathbf{x}_i} p(\mathbf{x}_i, \sigma_i) = \frac{1}{\sigma_i^2 M} \sum_{j=1}^M G(\mathbf{x}_i - \mathbf{x}_j, \sigma_i) (\mathbf{x}_i - \mathbf{x}_j). \quad (2.13)$$

When $i \neq j$, the derivative is

$$\begin{aligned} \frac{\partial}{\partial \mathbf{x}_i} p(\mathbf{x}_j, \sigma_j) &= \frac{1}{M} \left[\frac{\partial}{\partial \mathbf{x}_i} G(\mathbf{x}_j - \mathbf{x}_1, \sigma_j) + \frac{\partial}{\partial \mathbf{x}_i} G(\mathbf{x}_j - \mathbf{x}_2, \sigma_j) + \dots \right. \\ &\quad \left. + \frac{\partial}{\partial \mathbf{x}_i} G(\mathbf{x}_j - \mathbf{x}_i, \sigma_j) + \dots + \frac{\partial}{\partial \mathbf{x}_i} G(\mathbf{x}_j - \mathbf{x}_M, \sigma_j) \right] \\ &= \frac{1}{M} [0 + 0 + \dots - \sigma_j^{-2} G(\mathbf{x}_i - \mathbf{x}_j, \sigma_j) (\mathbf{x}_i - \mathbf{x}_j) + \dots + 0] \\ &= \frac{1}{\sigma_j^2 M} G(\mathbf{x}_j - \mathbf{x}_i, \sigma_j) (\mathbf{x}_i - \mathbf{x}_j). \end{aligned} \quad (2.14)$$

Substituting 2.13 and 2.14 into 2.11 gives,

$$\frac{\partial C}{\partial \mathbf{x}_i} = \frac{1}{M} \sum_{j=1}^M \frac{G(\mathbf{x}_i - \mathbf{x}_j, \sigma_i) (\mathbf{x}_i - \mathbf{x}_j)}{\sigma_i^2 p(\mathbf{x}_i, \sigma_i)} + \frac{1}{M} \sum_{j=1}^M \frac{G(\mathbf{x}_j - \mathbf{x}_i, \sigma_j) (\mathbf{x}_i - \mathbf{x}_j)}{\sigma_j^2 p(\mathbf{x}_j, \sigma_j)}. \quad (2.15)$$

The computational complexity for Equation 2.15 is $\mathcal{O}(M^2)$, since the entire density function p must be recomputed for each particle update. To simplify the computation, the PBM formulation instead considers p to be fixed for a given particle update: for $j \neq i$ in 2.15, the estimation of the density function at j is allowed to lag behind the update of particle position i . Under this assumption, $\frac{\partial}{\partial \mathbf{x}_i} p(\mathbf{x}_j, \sigma_j) = 0$, and the second term in 2.15 drops out, simplifying the gradient computation to only $\mathcal{O}(M)$.

After dropping the second term, the final approximation to the gradient of particle positional entropy is given by

$$\begin{aligned}
\frac{\partial C}{\partial \mathbf{x}_i} &\approx \frac{1}{M} \sum_{j=1}^M \frac{G(\mathbf{x}_i - \mathbf{x}_j, \sigma_i)(\mathbf{x}_i - \mathbf{x}_j)}{\sigma_i^2 p(x_i, \sigma_i)} \\
&= \frac{1}{M} \sum_{j=1}^M \frac{G(\mathbf{x}_i - \mathbf{x}_j, \sigma_i)}{\sigma_i^2 \frac{1}{M} \sum_{k=1}^M G(\mathbf{x}_i - \mathbf{x}_k, \sigma_i)} (\mathbf{x}_i - \mathbf{x}_j) \\
&= \frac{1}{M} \sum_{j=1}^M w_{ij} (\mathbf{x}_i - \mathbf{x}_j),
\end{aligned} \tag{2.16}$$

where w_{ij} are Gaussian weights based on interparticle distance and $\sum_j w_{ij} = 1$. To minimize C , the particles must move away from each other. Thus, we have a set of particles moving under a repulsive force and constrained to lie on the surface, with $\gamma < \sigma^2$ in 2.10 for stability. The motion of each particle is away from all of the other particles, but interactions are effectively local for sufficiently small σ , where w_{ij} vanishes with increasing interparticle distance.

2.3.4 Adaptive Distributions on Surface Features

The preceding minimization produces a uniform sampling of a surface. For some applications, a strategy that samples adaptively in response to higher order shape information is more effective for several reasons. From a numerical point of view, the minimization strategy relies on a degree of regularity in the tangent planes between adjacent particles, which argues for sampling more densely in high curvature regions. Geometric features such as high curvature are also considered more interesting than flat regions as important landmarks for biological shapes. To this end, the above uniform sampling method can be extended to adaptively sample more densely in high-curvature regions. This is done by modifying the Parzen windowing in Eqn. 2.12 as follows:

$$\tilde{p}(x_i) \approx \frac{1}{M} \sum_{j=1, j \neq i}^M G\left(\frac{\mathbf{x}_i - \mathbf{x}_j}{k_j}, \sigma_i\right) \tag{2.17}$$

where k_j is a scaling term proportional to the curvature magnitude computed at each neighbor particle j . The effect of this scaling is to warp space in response to local curvature. A uniform sampling based on maximum entropy in the warped space translates into an adaptive sampling in unwarped space, where points pack more densely in higher curvature regions. The extension of Eqn 2.16 to incorporate the curvature-adaptive Parzen windowing is straightforward to compute. Since k_j is not a function of x_i , the modified gradient is

$$\frac{\partial C}{\partial \mathbf{x}_i} \approx \frac{1}{M} \sum_{j=1}^M \frac{G((\mathbf{x}_i - \mathbf{x}_j)/k_j, \sigma_i)(\mathbf{x}_i - \mathbf{x}_j)}{\sigma_i^2 k_j p(x_i, \sigma_i)}. \tag{2.18}$$

There are many possible choices for the scaling term k . Meyer, et al. [78] describe an adaptive surface sampling that uses the scaling,

$$k_i = \frac{1 + \rho \kappa_i \left(\frac{s}{2\pi}\right)}{\frac{1}{2}s \cos(\pi/6)}, \quad (2.19)$$

where κ_i is the root sum-of-squares of the principal curvatures at surface location x_i . The user-defined variables s and ρ specify the ideal distance between particles on a planar surface and the ideal density of particles per unit angle on a curved surface, respectively. Note that the scaling term in this formulation could easily be modified to include surface properties other than curvature.

2.3.5 Surface Constraint

The surface constraint in both the uniform and adaptive optimizations is specified by the zero set of a scalar function $F(x)$. This constraint is maintained, as described in several papers [79], by projecting the gradient of the cost function onto the tangent plane of the surface, as prescribed by the method of Lagrange multipliers. The projection operator is given by

$$\mathbf{I} - \mathbf{n} \otimes \mathbf{n}, \quad (2.20)$$

where \mathbf{I} is the identity matrix, \mathbf{n} is the normal to the surface, and \otimes denotes the outer, or tensor, product. The tangent-plane projection is followed by iterative reprojection of the particle onto the nearest root of F by the method of Newton-Raphson. Principal curvatures are computed analytically from the implicit function, as described in [67].

2.3.6 The Kernel Width σ for PDF Estimation

Finally, the kernel width σ of the Parzen windowing estimation of particle density must be chosen at each particle. This is done automatically, before the positional update, using a maximum likelihood optimality criterion. The contribution to C of the i th particle is simply the probability of that particle position. Optimizing that quantity with respect to σ therefore gives a maximum likelihood estimate of σ for the current particle configuration. Using Newton-Raphson, the strategy is to find σ such that

$$\partial p(\mathbf{x}, \sigma) / \partial \sigma = 0, \quad (2.21)$$

which typically converges to machine precision in several iterations. For the adaptive sampling case, we find σ such that

$$\partial \tilde{p}(\mathbf{x}, \sigma) / \partial \sigma = 0, \quad (2.22)$$

so that the optimal σ is scaled locally based on the curvature. The iteration is given by

$$\sigma^{t+1} \leftarrow \sigma^t + \frac{\frac{\partial p}{\partial \sigma}}{\frac{\partial^2 p}{\partial \sigma^2}}, \quad (2.23)$$

and the first derivative of p with respect to σ , from 2.12, is

$$\begin{aligned} \frac{\partial}{\partial \sigma} \sum_{j=1}^M \frac{1}{M} G(\mathbf{x} - \mathbf{x}_j, \sigma) &= \frac{\partial}{\partial \sigma} \sum_{j=1}^M \frac{1}{M(2\pi)^{d/2} \sigma^d} e^{\frac{-r_j}{2\sigma^2}} \\ &= \frac{1}{M(2\pi)^{d/2} \sigma^{d+3}} \sum_{j=1}^M e^{\frac{-r_j}{2\sigma^2}} (r_j - d\sigma^2), \end{aligned} \quad (2.24)$$

where $r_j = (\mathbf{x} - \mathbf{x}_j)^T(\mathbf{x} - \mathbf{x}_j)$ is the distance from x to x_j . The second derivative follows from 2.24, and is given by

$$\begin{aligned} \frac{\partial^2}{\partial \sigma^2} \sum_{j=1}^M \frac{1}{M} G(\mathbf{x} - \mathbf{x}_j, \sigma) &= \frac{1}{M(2\pi)^{d/2} \sigma^{d+6}} \\ &\times \left[\sum_{j=1}^M e^{\frac{-r_j}{2\sigma^2}} (r_j^2 - (3 + 2d)\sigma^2 r_j + d(1 + d)\sigma^4) \right]. \end{aligned} \quad (2.25)$$

2.3.7 Numerical Considerations

There are a few important numerical considerations in computing the particle-based surface representation. First, the Gaussian kernels must be truncated. For all of the results presented in this dissertation, kernels are truncated so that $G(x, \sigma) = 0$ for $|x| > 3\sigma$. This means that each particle has a finite radius of influence, and a spatial binning structure to identify neighboring particles can be used to reduce the computational burden associated with particle interactions. A second consideration is the case where σ for a particle is too small to allow the particle to interact with its neighbors, and updates of σ or position cannot be computed. When σ is small, kernel size is updated using $\sigma \leftarrow 2 \times \sigma$, until σ is large enough for the particle to interact with its neighbors. A final numerical consideration is that the system must include bounds σ_{\min} and σ_{\max} to account for anomalies such as bad initial conditions or too few particles. These are not critical parameters, and as long as they are set to include the minimum and maximum resolutions, the system operates reliably.

One final aspect of the particle formulation to consider is that it computes the Euclidean distance between particles, rather than the geodesic distance on the surface. The PBM algorithm therefore assumes sufficiently dense samples so that nearby particles lie in the tangent planes of the zero sets of F . This is an important consideration; in cases where this assumption is not valid, such as highly convoluted surfaces, the distribution of particles may be affected by neighbors that are outside of the true manifold neighborhood. Limiting the influence of neighbors whose normals differ by some threshold value (e.g., 90 degrees) does limit these effects. The question of particle interactions with more general distance measures remains for future work and may be important for more effectively addressing applications to highly convoluted surfaces, such as the cortical surface of the brain.

2.3.8 Correspondence: Entropy Minimization in Shape Space

A sample set, \mathcal{E} , is a collection of N surfaces, each with their own set of M particles mapped to a single, dM -dimensional vector in shape space, i.e. $\mathcal{E} = \mathbf{z}^1, \dots, \mathbf{z}^N$. The sample set in vector form may be collected into a single matrix $\mathbf{P} = \mathbf{z}_j^k$, with particle positions along the rows and shape samples across the columns. Modeling $\mathbf{z}^k \in \mathfrak{R}^{dM}$ as an instance of random variable \mathbf{Z} , the PBM method for correspondence minimizes the combined sample and shape cost function

$$Q = H(\mathbf{Z}) - \sum_{k=1}^N H(\mathbf{X}^k), \quad (2.26)$$

which favors a compact representation of the sample, and is balanced against a uniform distribution of particles on each surface.

For this discussion we assume that the complexity of each shape is greater than the number of samples, and so normally $dM > N$. Given the low number of examples relative to the dimensionality of the space, the density estimation requires some assumptions. The PBM algorithm therefore assumes a normal distribution and models $p(\mathbf{Z} = \mathbf{z})$ parametrically using an anisotropic Gaussian with covariance Σ . The entropy is then given by

$$H(\mathbf{Z}) \approx \frac{1}{2} \log |\Sigma| = \frac{1}{2} \sum_{j=1}^{dM} \log \lambda_j, \quad (2.27)$$

where $\mathbf{e}_k, \lambda_k, j = 1, \dots, dM$ are the eigenvalues of Σ .

In practice, Σ will not have full rank, in which case the entropy is not finite. The problem must therefore be regularized with the addition of a diagonal matrix $\alpha \mathbf{I}$ to introduce a lower bound on the eigenvalues. The covariance is estimated from the data, and is given by

$$\Sigma = (dMN - 1)^{-1} \mathbf{Y} \mathbf{Y}^T, \quad (2.28)$$

where

$$\mathbf{y}^k = \mathbf{z}^k - \mu, \text{ and } \mu = \frac{1}{N} \sum_{k=1}^N \mathbf{z}^k. \quad (2.29)$$

Thus, \mathbf{Y} denotes the matrix of sample vectors \mathbf{P} minus the sample mean μ , i.e. $\mathbf{Y} = \mathbf{P} - \mu \mathbf{1}^T$, where $\mathbf{1}$ is a $dM \times 1$ vector of ones. Because $N < dM$, the eigenanalysis in 2.27 is done on the dual space of the $N \times N$ covariance matrix $\Sigma^T = (dMN - 1)^{-1} \mathbf{Y}^T \mathbf{Y}$. The nonzero eigenvalues of Σ can be obtained from Σ^T by noting the following relationships (see also [29]). For eigenvalues and eigenvectors

$\{\mathbf{e}_k, \lambda_k\}$ of Σ , $\Sigma \mathbf{e}_k = \lambda_k \mathbf{e}_k$. Similarly, for eigenvalues and eigenvectors $\{\mathbf{e}'_k, \lambda'_k\}$ of Σ^T , $\Sigma^T \mathbf{e}'_k = \lambda'_k \mathbf{e}'_k$. Substituting for Σ^T , we have

$$(dMN - 1)^{-1} \mathbf{Y}^T \mathbf{Y} \mathbf{e}'_k = \lambda'_k \mathbf{e}'_k, \quad (2.30)$$

and premultiplying each side by \mathbf{P} , gives

$$(dMN - 1)^{-1} \mathbf{Y} \mathbf{Y}^T \mathbf{Y} \mathbf{e}'_k = \lambda'_k \mathbf{Y} \mathbf{e}'_k, \quad (2.31)$$

which is equivalent to

$$\Sigma(\mathbf{Y} \mathbf{e}'_k) = \lambda'_k (\mathbf{Y} \mathbf{e}'_k). \quad (2.32)$$

Thus, $\mathbf{e}_k = \mathbf{Y} \mathbf{e}'_k$, and $\lambda_k = \lambda'_k$, for nonzero eigenvectors of Σ . The covariances $|\Sigma|$ and $|\Sigma^T|$ are therefore equivalent (up to a constant factor of α), and the final cost function G associated with the sample entropy is given by

$$G(\mathbf{P}) = \frac{1}{2} \log |\Sigma| = \frac{1}{2} \log \left| \frac{1}{dMN - 1} \mathbf{Y}^T \mathbf{Y} + \alpha \mathbf{I} \right|. \quad (2.33)$$

To compute the gradient of G , we follow a similar logic to that used in the derivation of 2.15, and allow the estimation of the mean μ of the distribution \mathbf{Z} to lag behind the updates $\frac{\partial G}{\partial \mathbf{P}}$. This allows for the simplifying assumption $\frac{\partial G}{\partial \mathbf{P}} \approx \frac{\partial G}{\partial \mathbf{Y}}$. This approximation becomes more accurate as the number of shape samples is increased, and changes in individual particle positions have increasingly less of an effect on the sample mean. The matrix of partial derivatives of G with respect to \mathbf{Y} is derived as follows.

$$\begin{aligned} \frac{\partial G}{\partial \mathbf{Y}} &= \frac{\partial}{\partial \mathbf{Y}} \left(\frac{1}{2} \log \frac{1}{(dMN - 1)^n} + \frac{1}{2} \log |\mathbf{Y}^T \mathbf{Y}| \right) \\ &= 0 + \frac{1}{2} |\mathbf{Y}^T \mathbf{Y}|^{-1} \frac{\partial}{\partial \mathbf{Y}} |\mathbf{Y}^T \mathbf{Y}| \\ &= |\mathbf{Y}^T \mathbf{Y}|^{-1} |\mathbf{Y}^T \mathbf{Y}| \mathbf{Y} (\mathbf{Y}^T \mathbf{Y})^{-1} \\ &= \mathbf{Y} (\mathbf{Y}^T \mathbf{Y})^{-1}. \end{aligned} \quad (2.34)$$

Adding the regularization to the covariance, we have the following equation for the updates

$$\frac{\partial G}{\partial \mathbf{P}} \approx \mathbf{Y} (\mathbf{Y}^T \mathbf{Y} + \alpha \mathbf{I})^{-1}. \quad (2.35)$$

The regularization α on the inverse of $\mathbf{Y}^T \mathbf{Y}$ can now be seen to account for the possibility of a diminishing determinant. The negative gradient $-\partial G / \partial \mathbf{P}$ gives a vector of updates for the entire system, which is recomputed once per system update. This term is added to the shape-based updates described in the previous section to give the update of each particle:

$$\mathbf{z}_j^k \leftarrow \gamma [-\partial G / \partial \mathbf{z}_j^k + \partial E^k / \partial \mathbf{z}_j^k]. \quad (2.36)$$

The stability of this update places an additional restriction on the time steps, requiring γ to be less than the reciprocal of the maximum eigenvalue of $(\mathbf{Y}^T \mathbf{Y} +$

$\alpha I)^{-1}$, which is bounded by α . Thus, we have $\gamma < \alpha$ and note that α has the practical effect of preventing the system from slowing too much as it tries to reduce the thinnest dimensions of the sample distribution. This also suggests an annealing approach for computational efficiency in which α starts off somewhat large (e.g., the size of the shapes) and is incrementally reduced as the system iterates. The annealing approach has been used for all studies described in this dissertation.

The choice of a Gaussian model for $p(\mathbf{Z} = \mathbf{z})$ is not critical for the proposed method. The framework easily incorporates either nonparametric, or alternate parametric models. In this case, the Gaussian model allows us to make direct comparisons with the MDL method, which contains the same assumptions. The Gaussian assumption also allows for an equivalent, maximum-likelihood formulation and subsequent extension to regression modeling (see Section 3.4). Research into alternative models for Z is outside the scope of this dissertation, but remains of interest for future work.

2.3.9 Setting Parameters

The mechanism that has been described in this section is a self tuning system of particles that distribute themselves across the shape surface using repulsive forces to achieve optimal distributions. Particles may also optionally adjust their sampling frequency locally in response to surface curvature. Free parameters of the system are limited to the choice of the number of particles (M), and the parameters s and ρ from 2.19, if adaptive sampling is used. For all of the results presented in this dissertation, the adaptivity parameters are determined empirically based on the data under analysis. To choose the number of particles, the system is initialized using the splitting strategy until the resulting representation is deemed to captured enough detail for the given application (see, for example, Figure 2.4).

In order to explicitly manage the tradeoff between model compactness and the geometric regularization, an additional free parameter β may be introduced into Equation 2.5 as follows:

$$Q = H(\mathbf{Z}) - \beta \sum_{k=1}^N H(\mathbf{X}_k), \quad (2.37)$$

Empirical results, however, suggest that the two terms in this function are already well balanced and $\beta = 1$ represents a good default setting. To generate all of the results presented in this dissertation, for example, β is set to 1 with good results in each case.

2.4 A Shape Modeling Pipeline

This section describes a full process for shape modeling using the PBM algorithm, including preprocessing considerations, initialization of the particle systems, and setting system parameters. A typical data-flow “pipeline” for establishing correspondences from binary image volume inputs is outlined in Figure 2.8. The preprocessing steps in the pipeline establish an initial alignment of the segmentations and generate suitable distance transforms. The optimization phase consists

of initializing the particle system, running the PBM optimization, and iteratively refining the alignment using the Procrustes algorithm to remove residual nonshape data. Iterations of the correspondence optimization are interleaved with alignment steps until convergence. The remainder of this section discusses each of these steps in more detail and then concludes with a description of the software implementation of the PBM framework that was used to generate all of the results in this dissertation.

2.4.1 Preprocessing Binary Segmentations

Any set of implicitly defined surfaces, such as a set of binary segmentations, is appropriate as input to the PBM algorithm. The algorithm, however, can be applied directly to binary segmentation volumes, which are often the output of a manual or automated segmentation process. Binary volumes contain an implicit shape surface at the interface of the labeled pixels and the background. Any suitably accurate distance transform from that interface may be used to form the implicit surface necessary for the particle optimization.

Segmentation data typically requires some processing to remove aliasing artifacts in the binary mask. Aliasing artifacts can adversely affect numerical approximations of surface features and the computations required to maintain the surface constraint in the PBM algorithm. One effective method for antialiasing binary volumes is given by Whitaker in [112], who describes a method for fitting an antialiased, level-set surface to a binary volume through an iterative relaxation process. The process uses curvature flow of the surface, with constraints on the flow dictated by the binary voxel locations of the segmentation. Another effective antialiasing method is the r -tightening algorithm given by Williams et al. [114]. The surface tightening method follows a similar approach to that of Whitaker, but constrains the level-set relaxation process using binary volumes that result from morphological opening and closing of the targeted binary surface. This method has proven to be particularly effective at removing aliasing artifacts without compromising the precision of the segmentation. As a final preprocessing step, the

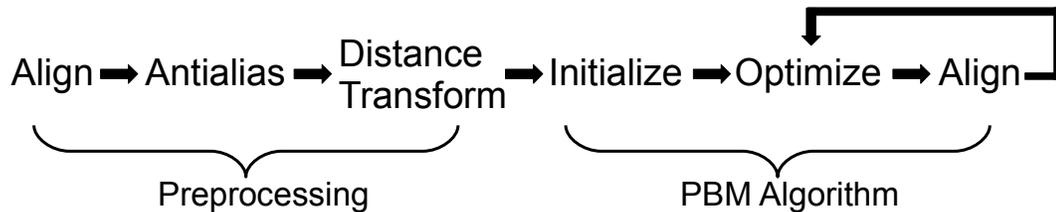


Figure 2.8. A typical shape modeling pipeline of operations using the PBM correspondence optimization algorithm.

distance transform is typically followed by a slight Gaussian blurring to remove the high-frequency artifacts that can occur as a result of numerical approximations.

2.4.2 Initialization and Optimization

There are number of possibilities for initializing the particle systems on the sample shapes, including manual specification of points and regular surface sampling. One particularly effective approach, however, is to use an iterative, *particle splitting* strategy. The splitting strategy proceeds as follows. First, the PBM system is initialized with a single particle on each shape that finds the nearest zero of F . This single particle is then split to produce a new, nearby particle. The two-particle (per shape) system is then optimized for correspondence until a steady state is reached. The splitting process, followed by optimization, is then repeated until a specific number of particles have been produced. Thus, the initialization proceeds simultaneously with the optimization in a multiscale fashion, generating progressively more detailed correspondence models with each split. Figure 2.9 illustrates the splitting process on a single sphere.

For all the experiments that are described in this dissertation, the parameters for the PBM optimization are set automatically as follows. The numerical parameter σ_{\min} is set to machine precision and σ_{\max} is set to the size of the domain. The annealing parameter α starts with a value roughly equal to the diameter of an average shape and is reduced to machine precision over several hundred iterations. Particles are initialized on each shape using the splitting procedure described above. These default settings have been found to produce reliably good results that are very robust to the initialization.

Processing time for the PBM algorithm on a 2GHz desktop machine averages around 1/8000 second/particle per iteration. This translates to full optimization times that scale linearly with the number of particles in the system and are on the order of minutes for small systems of only a few thousand particles to several hours for larger systems of tens-of-thousands of particles. Optimizations of very large systems of hundreds-of-thousands of particles may take processing times of several dozen hours.

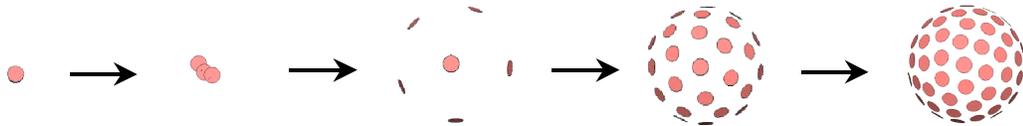


Figure 2.9. A system of 100 particles on a sphere, produced by a particle splitting algorithm.

2.4.3 Alignment of Shape Surfaces

A collection of shape segmentations must often be aligned in a common coordinate frame for modeling and analysis. Where no information exists to specify a correct alignment, one approach is to first align segmentations with respect to their centers of mass and the orientation of their first principal eigenvectors. Then, during the optimization, the PBM method may optionally further align shapes with respect to rotation, translation, and scale using generalized Procrustes analysis (GPA) [45]. The GPA alignment is applied at regular intervals after particle updates in order to remove any residual, nonshape information from the model. GPA alignment during the optimization process is only enabled once the full set of M particles have been initialized on all surfaces. Where the true shape alignments are known, however, the GPA iterations may be omitted. A subset of the GPA alignment parameters may also be applied, such as only the rotational and translational components, leaving the scale unaffected.

2.4.4 Software Implementation

All of the results for experiments described in this dissertation were computed using a C++ software implementation of the PBM algorithm. Because the PBM algorithm is completely generalizable to higher dimensions, shapes in two and three dimensions were processed using the same software, templated on dimension. The software consists of a core library of code for managing particle systems, associated libraries for the numerical optimizations, and a set of user interfaces for interactively optimizing and visualizing correspondences. The code was constructed in part from components of the Insight Toolkit (ITK), an open-source image processing toolkit for medical imaging applications [60]. Visualization components were constructed using code from the open-source Visualization Toolkit (VTK) [93]. All of the PBM algorithm code produced for this dissertation is freely available for download as the ShapeWorks software package, which is released through the SCI Institute at the University of Utah [17].

2.5 Preliminary Evaluation

This section details several experiments designed as a preliminary evaluation of the PBM correspondence optimization method. Further studies on clinical data are described in Chapter 5. First, models generated using the particle method are compared with models generated using MDL for two synthetic, two-dimensional datasets. Next, a simple experiment on tori illustrates the applicability of the method to nonspherical topologies and illustrates how the method is capable of finding underlying, true parameters in a synthetic system. Finally, the PBM method is applied to a full statistical shape analysis of several three-dimensional neuroanatomical structures from published clinical datasets.

2.5.1 Comparison with MDL

This experiment compares results from the PBM algorithm on closed curves in a two-dimensional plane with results from the two-dimensional open-source Minimum

Description Length (MDL) implementation given by Thodberg [103]. In the first half of the experiment, we used the proposed, particle method to optimize 100 particles per shape under uniform sampling on 24 “box-bump” shapes, similar to those described in [103]. Each shape was constructed as a fast-marching distance transform from a set of boundary points using cubic b-splines with the same rectangle of control, but with a bump added at a random location along the top of its curve. This example is interesting because we would, in principle, expect a correspondence algorithm that is minimizing information content to discover this single mode of variability in the sample set.

MDL correspondences were computed using 128 nodes and *mode 2* of Thodberg’s Matlab software, with all other parameters set to their defaults (see [103] for details). Principal component analysis (PCA) identified a single dominant mode of variation for each method, but with different degrees of leakage into orthogonal modes. MDL lost 0.34% of the total variation from the single mode, while the proposed method lost only 0.0015%. Figure 2.10 illustrates the mean and three standard deviations of the first mode of the two different models. Shapes from the particle method remain more faithful to those described by the original training set, even out to three standard deviations where the MDL description breaks down. A striking observation from this experiment is how the relatively small amount of variation remaining in the minor modes of the MDL case produce such a significant effect on the results of shape deformations along the major mode.

The second part of the MDL comparison was conducted on the set of 18 hand shape contours described in [30], again applying both the particle method and MDL using the same parameters as described above. Distance transforms from spline-based contour models again form the inputs. In this case, results were also compared to a set of manually selected correspondences, which introduce some anatomical knowledge of the digits by selecting landmarks at known homologous locations, such as the tips of the fingers. Figure 2.11 compares the three resulting models in the top three modes of variation to ± 3 standard deviations. A detailed analysis of the principal components shows that the proposed particle method and the manually selected points both produce very similar models, while MDL differs significantly, particularly in first three modes. This observation is supported by Figure 2.12, which compares the cumulative percentage of total variance for the first 10 modes. It is clear that the particle method discovers a set of correspondences for this data set that conform more closely to the anatomical relationships.

2.5.2 Nonspherical Topologies

Existing three-dimensional MDL implementations rely on spherical parameterizations and are therefore only capable of analyzing shapes that are topologically equivalent to a sphere. The particle-based method does not have this limitation. This section presents an experiment illustrating the use of the PBM algorithm on a sample set of 40 randomly chosen tori. The tori were parameterized by the small radius r and the large radius R (see also the example given in Figure 2.1) and were chosen from a distribution with mean $r = 1, R = 2$ and $\sigma_r = 0.30, \sigma_R = 0.15$. A rejection policy was used to exclude invalid tori (e.g., $r > R$). Correspondences were optimized using 1024 particles per shape, and a uniform sampling (no adaptivity).

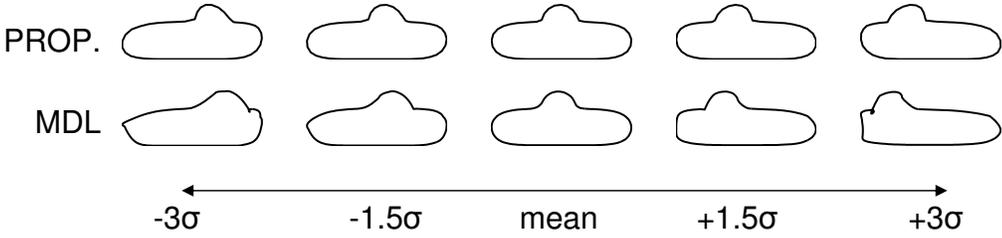


Figure 2.10. The “box-bump” experiment.

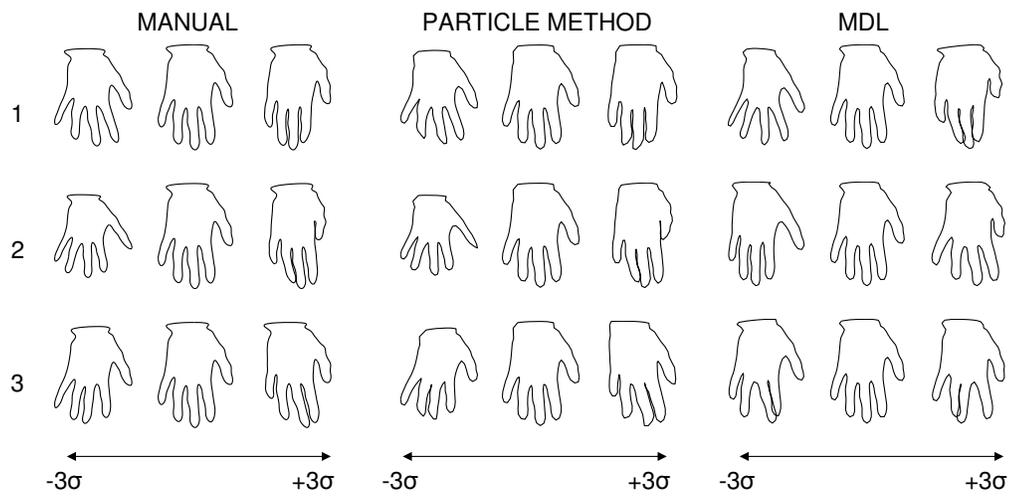


Figure 2.11. The mean and ± 3 standard deviations of the top 3 modes of the hand models.

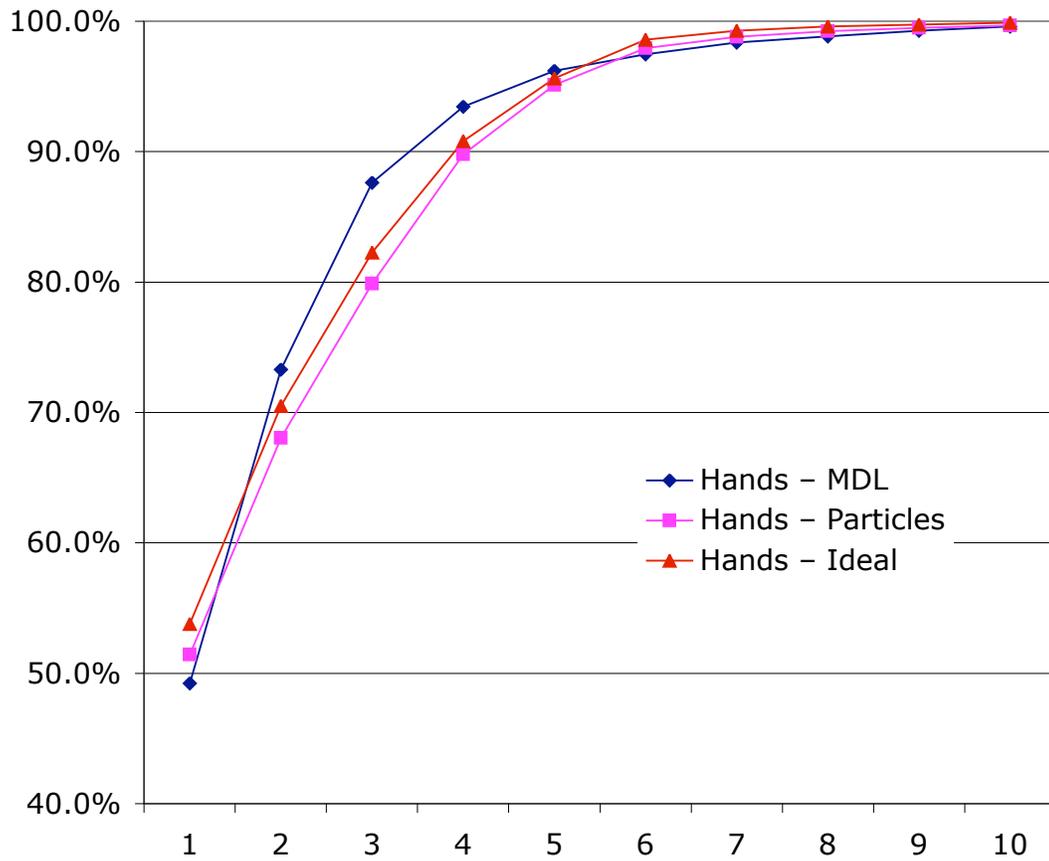


Figure 2.12. Comparison of the percentage of total variance in PCA modes of the hand experiment.



Figure 2.13. An illustration of particle correspondences on several tori from a random distribution on r and R . Corresponding particle positions across the three shapes are indicated by glyphs with matching colors.

Figure 2.13 shows the particle system distribution across several of the torus shapes in the sample set with 1024 correspondences. Correspondence positions are indicated by spherical glyphs and correspondence across shapes is indicated by the color of the glyph. Surface reconstructions for each sample were done using the correspondence positions and the algorithm given by Hoppe, et al. for collections of unorganized points [54]. A principal component analysis (PCA) of the resulting correspondence positions indicates that the particle system method discovered two pure modes of variation. Only the top two modes of the PCA exhibit significant variation. PCA 1 contains 69.7870% of total variation and PCA 2 contains 30.2076% of total variation. Less than 0.006% of total variation is accounted for in the smaller, “error” modes.

Empirical observation of the top two PCA modes suggests that they correspond well to variation in r and R from the parametric model that was used to generate the sample data. Figure 2.14 shows the mean correspondence positions from the model moved along each of the top two PCA modes. Torus shapes along each mode are reconstructed from -3 to $+3$ standard deviations from the mean. The top row illustrates variation in PCA 1, which corresponds to r , and the bottom row indicates variation in PCA 2, which corresponds to R . In this experiment, the PBM method appears to have estimated the true orthogonal modes of variation of the shape sample.

2.5.3 Shape Analysis of Neuroanatomical Structures

As a further validation of the PBM algorithm, this section describes hypothesis testing of group shape differences on data from two published clinical studies in neuroanatomy. The first study addresses the shape of the hippocampus in patients with schizophrenia. The data consist of left and right hippocampus shapes from 56 male adult patients versus 26 healthy adult male controls, segmented from MRI using a template-based semi-automated method [100]. The second study addresses the shape of the caudate in males with schizo-typal personality disorder. The data consist of left and right caudate shapes from 15 patients and 14 matched, healthy controls that were manually segmented from MRI brain scans of the study subjects by clinical experts [74]. In each study, variation in overall subject size is controlled for by normalizing the segmentation volumes with respect to intercranial volume.

The raw binary segmentations from each study were aligned and processed as described in Section 2.4, including Procrustes registration. Correspondence models were optimized using the PBM method, with 1024 correspondence points per shape and the curvature-adaptive sampling strategy. Separate models were created for left and right structures using the combined data from patient and normal populations. Models were generated without knowledge of the shape classifications so as not to bias the correspondences to one class or the other, an important consideration for statistical analysis that is discussed further in Chapter 3. On inspection, all of the resulting models appear to be of good quality; each major mode of variation describes some plausible pattern of variation observed in the training data. As an example, Figure 2.15 shows several surface meshes of shapes generated from the correspondence point sets of the right hippocampus model.

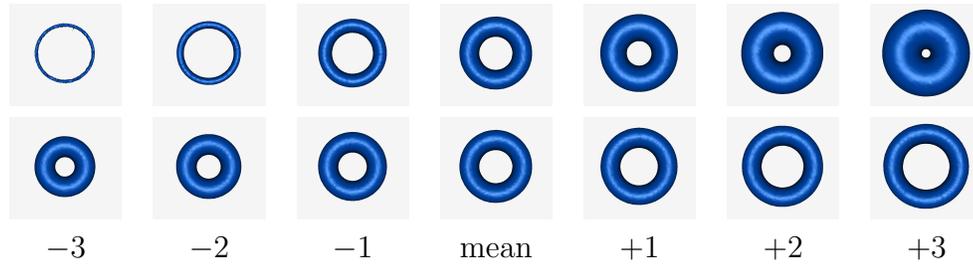


Figure 2.14. The top two PCA modes of variation for the torus shape model, shown in the first and second rows, respectively. Surface reconstructions are from the correspondence positions moved along the axis of each principal component at positions -3 to $+3$ standard deviations from the mean shape. PCA 1 (top row) corresponds to r and PCA 2 (bottom row) correspondence to R .

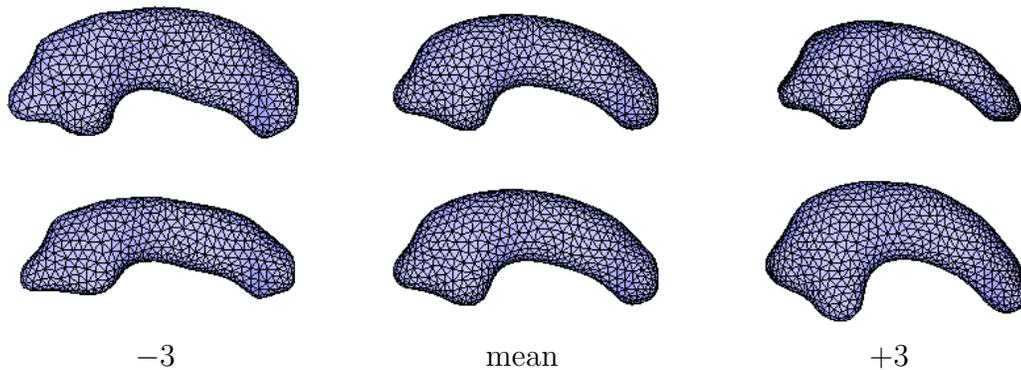


Figure 2.15. Right hippocampus model mean and ± 3 standard deviations in the top two PCA modes (top and bottom rows, respectively).

After computing the models, separate statistical tests for differences in correspondence positions were applied to every correspondence index. The specific test used is a nonparametric permutation test of the Hotelling T^2 metric, with false-discovery-rate (FDR) correction for multiple comparisons. This is a method described by Styner, et al. [98] that is useful for identifying and visualizing localized regions of significant shape differences. The statistical results were computed using an open-source implementation of the algorithm [98], with 20,000 permutations among groups and an FDR bound set to 5%. The null hypothesis for these tests is that the distributions of the locations of corresponding sample points are the same regardless of group. The point-wise hypothesis test was chosen for this analysis, because it has also been applied to several other published studies of this data [100, 81] and thus provides a basis for comparison of the results.

Figure 2.16 shows the raw and FDR-corrected p-values for the left and right hippocampi from the schizophrenia study. Areas of significant group differences ($p \leq 0.05$) are shown in red. Areas with insignificant group differences ($p > 0.05$) are shown in blue. The right hippocampus shows significant differences in the mid-region and the tail, even after FDR-correction. The left hippocampus appears to exhibit few group differences, with none detected after FDR correction. Differences in the tail, especially on the right structure were also observed by Styner et al. in [96]. These results also correlate with those reported for the spherical harmonics method (SPHARM) [100] and spherical wavelet analysis [81].

Raw p-values for the caudate analysis are shown at the bottom of Figure 2.16. No significant differences on either shape were found after FDR correction. The raw p-values, however, suggest that both structures may exhibit group differences in the tail and that the right caudate contains more group differences than the left, an observation that agrees with results given in [98], [74], [100], and [81].

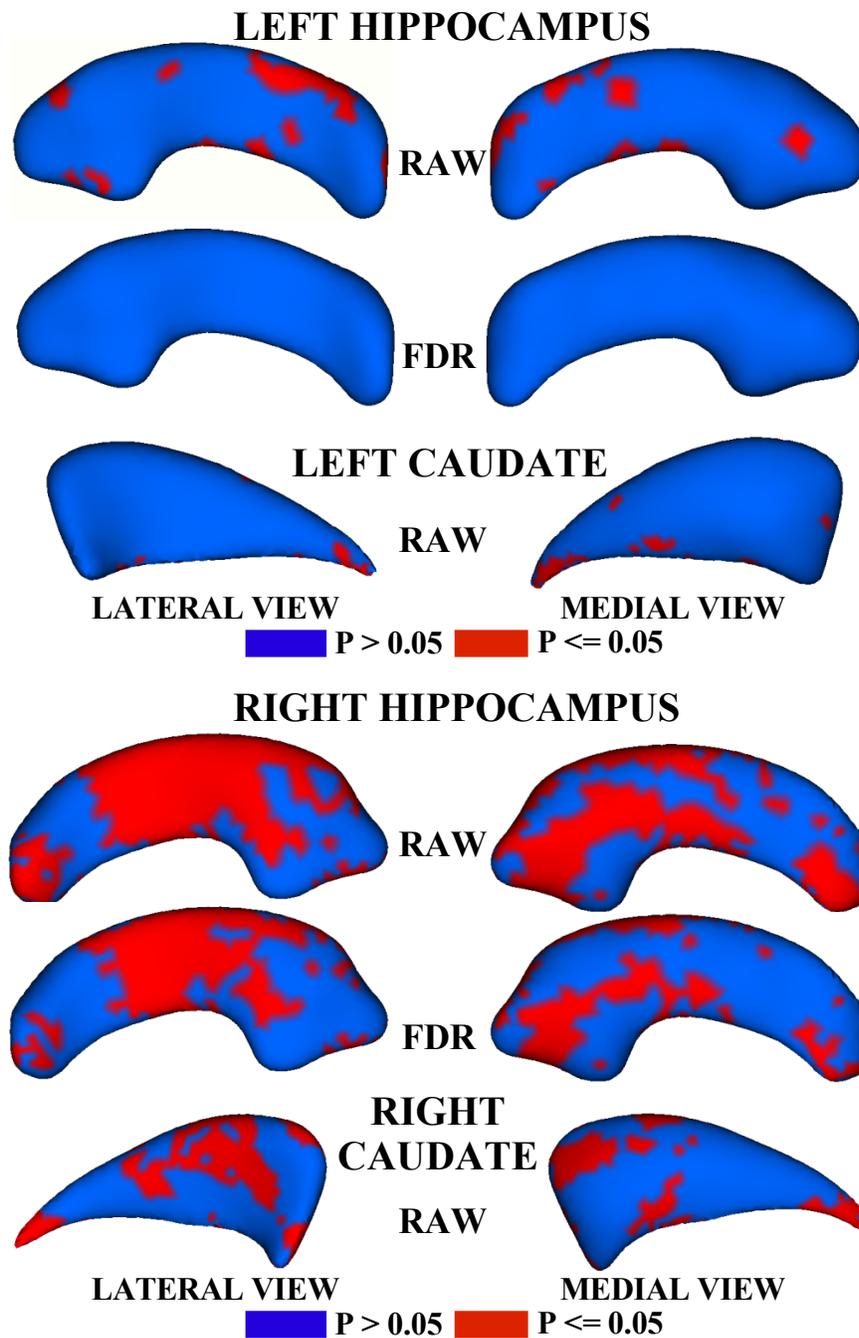


Figure 2.16. P-value maps for the hippocampus and caudate shape analyses, shown on the mean shape. Red (dark in grayscale) indicates significant group differences ($p \leq 0.05$), and blue indicates no significant group differences ($p > 0.05$).

CHAPTER 3

MODELING EXTENSIONS

This chapter describes several mathematical extensions of the particle-based modeling (PBM) algorithm that are designed to make it more robust to realistically complex shape analysis problems. The work in this chapter is motivated by the needs of the biomedical research community for tools to model more complicated anatomical shapes and statistical designs than the basic examples presented in the previous chapter. The first set of extensions allows the PBM algorithm to operate on more complicated surface geometries. A second set of extensions provide for more flexible notions of correspondence by generalizing the underlying statistical optimization. These tools represent an important contribution to the shape analysis literature because, to date, proposed correspondence optimization methods have mainly focused on closed, manifold surfaces with spherical topologies. Statistical optimization criteria has also received a fairly narrow treatment, with most correspondence methods considering only structural information in the absence of confounding or explanatory variables.

Many anatomical structures of interest to clinical and biological research are not well represented as closed, manifold surfaces. Some examples include regions of the brain and the heads of bones. A region of the brain may be defined more by functional information than by structural considerations and therefore requires a boundary specification. Similarly, the extent of the bone that comprises its head must also be specified by a boundary. The head of a bone may also be part of a larger *complex* of structures, one that includes a socket, for example, as in a knee or elbow joint. Such a complex of structures is a shape that consists of a collection of multiple, disconnected surfaces. Other examples of shape complexes are sets of bones in the hands and feet and collections of brain structures.

In many cases, the geometric features of an anatomical object are not sufficient to properly establish correspondence. Some anatomy is highly variable across subjects and additional information, such as functional data, is helpful in determining how surface regions correspond. One such example, that I will examine in more detail in Chapter 5, is the cortical surface. Because of the highly variable folding patterns on the cortex, neuroanatomists typically rely on other information, such as sulcal depths and vascular connectivity, for correspondence [84]. Another common situation that requires additional information for proper correspondence is when the shape of anatomy is correlated with one or more independent variables. Most clinical studies, for example, control for subject age, and require a regression model that properly accounts for variability explained by age.

The remainder of this chapter presents modifications to the PBM algorithm to manage the scenarios outlined above. This chapter is organized as follows. In

Section 3.1, I remove the closed surface requirement for the PBM algorithm, and present an extension for modeling open surfaces, or surfaces with boundaries. The open surface method allows an arbitrary boundary to be defined as the intersection of a closed surface with a set of shape primitives. Section 3.2 describes how the PBM algorithm can be used to optimize correspondence among shapes that consist of *multiple* anatomical objects. This capability is important in the study of shape covariance among anatomy that is functionally or structurally correlated. Section 3.3 describes a generalization of the PBM optimization criteria to correspondence in arbitrary, multivariate functions of position, rather than only considering positional information. This approach is useful for problems where there is data other than geometric information that indicates correspondence, such as multimodal imaging studies, and studies with functional imaging data. Finally, Section 3.4 describes a methodology for including a regression model on independent variables into the PBM correspondence optimization. Shape regression modeling is important for examining, or controlling for, correlations between shape and certain factors such as age or clinical variables. This chapter discusses the motivation, theory, and implementation of each of the extensions. Examples of their application to scientific studies are given in Chapter 5.

3.1 Modeling Shape with Open Surfaces

As discussed in Section 2.1.2, there are many important anatomical structures that require an open surface representation for shape analysis. Some examples include flat structures such as the diaphragm, functional subregions of the cortical surface, and the head of a bone (see Figure 2.2). In Chapter 5, I will present an application of the PBM algorithm to neonatal head shape data, which requires a specification of the boundary of the cranium. This section describes an extension to the basic PBM algorithm to optimize correspondence on surfaces with open boundaries.

Conceptually, there are two ways to handle the surface boundary when establishing correspondences. The first approach is to explicitly represent and model the boundary. To properly model the boundary, correspondences must be allowed to lie on the boundary, and the optimization must track particle movement on and off of the boundary. This approach is appropriate for applications where the boundary shape is of specific interest to the problem. In many cases, however, it is not important, or even desirable, to model the variation in the shape of the boundary. A segmentation, for example, may contain noise in the boundary shape due to ambiguities in its specification during the segmentation process. In this situation, where the boundary is considered noisy, it can simply be treated as a constraint on the particle optimization, which is the approach developed in this section.

The proposed algorithm for correspondence on open surfaces represents the surface boundary as the intersection of a closed surface (e.g., \mathcal{S} in 2.3.3) with a set of geometric primitives, such as cutting planes and spheres. The boundary representation is then used to influence the entropy maximization of the PBM algorithm particle position (Section 2.3.3), so that it indirectly constrains the positions of particles to lie within the surface boundary. The goal of the open surface modeling algorithm is to formulate particle interactions with the boundaries so that

the positions of the constraints have as little influence as possible on the statistical shape model. This approach is consistent with the idea that the boundary shape may contain noise, and we wish to minimize the influence of this noise on the model.

The algorithm proceeds as follows. For each geometric primitive in the surface boundary representation, the algorithm constructs a *virtual* particle distribution that consists of all of the closest points on its surface to the particles with positions \mathbf{x}_j on \mathcal{S} . During the gradient descent optimization, particles \mathbf{x}_j interact with the virtual particles, and are therefore effectively repelled from the geometric primitives, and thus from the open surface boundary. The virtual distributions are updated after each iteration, as the particles on \mathcal{S} redistribute under the optimization. Because the virtual particles are allowed to factor into the Parzen windowing kernel size estimation (Equation 2.12), particles \mathbf{x}_i maintain a distance from the boundary proportional to their density on the surface \mathcal{S} . In this way, features near the boundary may be sampled, but particles are never allowed to lie on the boundary itself, limiting the effect of errors in the boundary specification on the configuration. Note that the virtual particle distributions are also not used in the correspondence optimization term (the sample entropy from Equation 2.27) and therefore do not directly affect the distribution of samples in shape space.

Figure 3.1 illustrates a particle configuration using the method outlined above for open surfaces, and shows the effect of increasing the number of particles. In the figure, the open surface boundary is defined by the intersection of an implicit bone surface, a cutting plane, and a sphere. As the number of particles is increased, the distribution samples regions of the bone closer and closer to the surface boundary. Note, however, that the particle distribution never touches or crosses the boundaries. For another example of PBM optimizations on open surfaces, see the experiment for head shape regression that is given in Section 5.4.

3.2 Modeling Shape Complexes

This section presents an extension of the PBM algorithm for shape modeling of multiple, disconnected anatomical surfaces, or shape *complexes*. Like the open surface method from the previous section, this extension allows the PBM to operate on more diverse classes of shape. In many areas of clinical psychiatric and neurological research, the joint analysis of different anatomical structures is of increasing interest. Certain spectrum disorders, such as autism, for example, are thought to represent a confluence of several underlying abnormalities, impacting the relationships between brain regions [20]. Statistical models that capture those relationships are important to gain insight into the mechanisms and development of such diseases. Shape models of anatomical complexes are also important tools for geneticists and developmental biologists, who study the covariation of shape for complexes of anatomical structures in small animal models (e.g., [31]). Another important set of applications are biomechanical studies of bone and joint complexes, where shape correlations can provide insight into the origin and progression of diseases such as hip dysplasia (e.g., [107]).

A *multiobject complex* is defined as a set of solid shapes, each representing a single, connected biological structure. The complex of structures are assembled into a scene within a common coordinate frame. Each structure in a multiobject

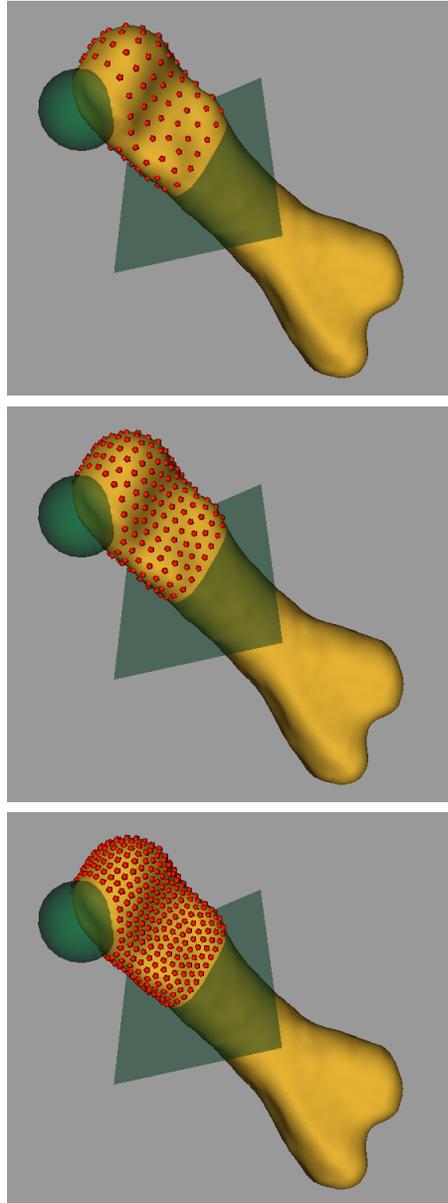


Figure 3.1. An example of distributing increasing numbers of particles on an open surface, where the boundary is defined by the intersection of an implicit bone surface with a sphere and a cutting plane.

complex contains shape, pose, scale, and positional information. Some examples include the segmentations of multiple brain structures from a single MRI of a patient and sets of bones segmented from a CT scan. The proposed correspondence method for establishing correspondence on multiobject complexes is novel in that it optimizes correspondence positions in the full, joint shape space of the object complex. Researchers have previously only considered the correspondence problem separately for each structure, thus ignoring the interstructural shape correlations in the optimization process. By explicitly modeling the correlations among variabilities, however, optimization in the joint space may produce more compact distributions for correspondences, resulting in fewer model parameters and greater statistical power.

State-of-the-art methods for correspondence optimization that rely on parameterized surface representations and assume a spherical or toroidal topology are not suitable for multiobject complexes, which consist of disconnected sets of discrete surfaces. Some parametric methods have been applied to shape complexes by finding correspondences for each structure independently and then treating those correspondences as the marginal distributions of the full shape complex model [46]. This particular approach however ignores any covariation in shape across structures during the optimization, which is typically the variation of special interest for the model. Parametric approaches in general may be limited for shape complex modeling because the reliance on individual anchor shapes for regularization would restrict the degree to which parameterizations of different objects in the sample set could interact.

The particle-based correspond method described in Chapter 2 can be directly applied to multiobject complexes by treating all of the objects in the complex as defining a single surface. However, if the objects themselves have distinct identities (i.e., object-level correspondence is known a priori), we can assign each particle to a specific object, decouple the spatial interactions between particles on different shapes, and constrain each particle to its associated object. In this way, each correspondence is guaranteed to stay on a particular anatomical structure, and the surface sampling is not influenced by regions where structures in the complex happen to be near to one another. The shape-space statistics using this method, however, remain coupled, and the covariance Σ (Eqn. 2.27) includes all particle positions across the entire complex, so that optimization takes place on the joint, multiobject model.

As with the single-object framework, any set of implicitly defined surfaces is appropriate as input to the multiobject framework, with similar preprocessing considerations as those discussed in Section 2.4. In the case of binary segmentations, the input is now a set of N segmentations of K -object complexes, which contains $N \times K$ distinct, volumetric label masks. The optimization can be initialized using the splitting strategy described in Section 2.4, starting with a single particle on each structure in the complex, and splitting until a specified number of particles per structure is reached. Additional particles can then be placed on individual structures within the complex as necessary, using the splitting strategy for individual objects, but under the joint correspondence optimization.

An example application of the multiobject PBM algorithm to a study of clinical brain structure data is given in Section 5.2. In the study, I demonstrate group

differences between autism subjects and typical controls in the shape of a complex of ten subcortical brain structures. I also demonstrate how the joint optimization produces a model with better group discrimination than a collection of ten separate models.

3.3 Correspondence Based on Functions of Position

The previous two sections describe methods that extend the PBM algorithm to more complex geometries. As I discussed in the introduction to this chapter, however, the geometric features of an anatomical object are often not sufficient to properly establish correspondence. The basic PBM algorithm (Chapter 2) only considers particle position, which is the geometric, or structural, information of the shape surface. Thus, we need a way to incorporate additional information about correspondence into the optimization process.

This section describes an extension to the basic PBM algorithm to establish correspondence by minimization of the entropy of arbitrary, vector-valued functions of position. This more general method is useful in cases where the notion of correspondence is not well defined by the surface geometry, but can be described by other metrics. Correspondence for cortical surfaces of the brain, for example, is difficult to establish structurally because of the high variability of cortical folding patterns across subjects. Neuroanatomists therefore often rely on other information for correspondence, such as cortical thickness and functional imaging data [36]. Many other biomedical studies involve multiple imaging modalities, such as CT imaging combined with MRI, or SPECT. Because it incorporates more information about the targeted anatomy, a correspondence model based on data from multiple modalities may be more accurate than a model using only structural information.

The extension to the PBM algorithm to incorporate functional data, which I will refer to as the *generalized* PBM algorithm, is straightforward. It consists of substituting the entropy estimation of the matrix of particle positions with an entropy estimation on an arbitrary, vector-valued function of the particle position. From Section 2.3, the energy term for the basic PBM optimization is given by

$$Q = H(\mathbf{Z}) - \beta \sum_{k=1}^N H(\mathbf{X}_k), \quad (3.1)$$

where H is an estimation of entropy, X_k is a vector random variable with the distribution of particle configuration k , and Z is the vector random variable with the distribution of the shape samples in the dM -dimensional shape space. The extension to the generalized PBM algorithm only modifies the correspondence term $H(\mathbf{Z})$. The entropy associated with individual correspondence configurations, $H(\mathbf{X}_k)$, is not modified, and still operates on positional information. In other words, particles are still constrained to lie on the surface of the shape and distribute themselves across shape surfaces using the maximization of positional entropy, but their correspondence is established using a function of positional information. Note that a function of position could be designed to also include particle position, so that both structural and functional data influence the correspondences.

Recall from 2.3.8, that the entropy estimation of the sample distribution in shape space is given by

$$H(\mathbf{Z}) \approx \frac{1}{2} \log |\Sigma|, \text{ and } \Sigma = (dMN - 1)^{-1} \mathbf{Y} \mathbf{Y}^T, \quad (3.2)$$

where Σ is the covariance matrix, and \mathbf{Y} is the $dM \times N$ data matrix P of sample vectors $\mathbf{z}^k, k = \{1, \dots, N\}$, minus the sample mean μ , and each vector \mathbf{z}^k consists of the positional information from M particles on shape surface k . In the case of computing entropy of vector-valued functions of the correspondence positions, the extension to functional data considers the more general case where columns of the data matrix are instead given by

$$\tilde{\mathbf{p}}^k = \begin{bmatrix} f(\mathbf{x}_0^k) \\ f(\mathbf{x}_1^k) \\ \vdots \\ f(\mathbf{x}_j^k) \\ \vdots \\ f(\mathbf{x}_{M-1}^k) \end{bmatrix}, \quad (3.3)$$

where \mathbf{x}_j^k is the positional information of particle j for shape k , and $f : \mathbb{R}^d \rightarrow \mathbb{R}^q$.

The matrix \mathbf{Y} now becomes a matrix $\tilde{\mathbf{Y}}$ of the function values at the particle points, minus the means of those functions at the points. Columns of $\tilde{\mathbf{Y}}$ are given by

$$\tilde{\mathbf{y}}^k = \begin{bmatrix} f(\mathbf{x}_0^k) - \frac{1}{N} \sum_{i=1}^N f(\mathbf{x}_0^i) \\ f(\mathbf{x}_1^k) - \frac{1}{N} \sum_{i=1}^N f(\mathbf{x}_1^i) \\ \vdots \\ f(\mathbf{x}_{M-1}^k) - \frac{1}{N} \sum_{i=1}^N f(\mathbf{x}_{M-1}^i) \end{bmatrix}. \quad (3.4)$$

The new cost function \tilde{G} is the estimation of entropy of the samples $\tilde{\mathbf{y}}^k$. With the same assumption of a Gaussian distribution in shape space, by the same logic as the derivation of the cost function G in 2.33, we have

$$\tilde{G}(\tilde{\mathbf{z}}) = \log \left| c \tilde{\mathbf{Y}}^T \tilde{\mathbf{Y}} \right|, \quad (3.5)$$

with c a constant.

Let $Q = (\tilde{\mathbf{Y}}^T \tilde{\mathbf{Y}} + \alpha \mathbf{I})^{-1}$. By the chain rule, the partial derivative of \tilde{G} with respect to the data \mathbf{y}^k becomes

$$-\frac{\partial \tilde{G}}{\partial \tilde{\mathbf{P}}^k} = \mathbf{J}_k^T \mathbf{Q}^k, \quad (3.6)$$

where \mathbf{J}_k is the Jacobian of the functional data for shape k . The matrix \mathbf{J}_k has the structure of a block diagonal matrix with $M \times M$ blocks, with diagonal blocks

the $q \times d$ submatrices of the function gradients at particle j . Specifically, for each shape k , we have function data

$$\mathbf{y}^k = [f_1^0, f_2^0, \dots, f_q^0, f_1^1, f_2^1, \dots, f_q^1, \dots, f_1^{M-1}, f_2^{M-1}, \dots, f_q^{M-1}]^T, \quad (3.7)$$

and a diagonal submatrix block of the Jacobian $\mathbf{J}_k = \nabla_{\mathbf{z}^k} \mathbf{y}^k$ has the structure

$$\begin{bmatrix} \partial f_1^j / \partial x_{dj+1} & \partial f_1^j / \partial x_{dj+2} & \dots & \partial f_1^j / \partial x_{dj+d} \\ \partial f_2^j / \partial x_{dj+1} & \partial f_2^j / \partial x_{dj+2} & \dots & \partial f_2^j / \partial x_{dj+d} \\ \vdots & \vdots & \ddots & \vdots \\ \partial f_q^j / \partial x_{dj+1} & \partial f_q^j / \partial x_{dj+2} & \dots & \partial f_q^j / \partial x_{dj+d} \end{bmatrix}, \quad (3.8)$$

where $j = \{0, 1, 2, \dots, M-1\}$ is the block number, which corresponds to a single particle, and $\{x_1, x_2, \dots, x_{dM}\}$ are the directional components of the full set of M particles. The correspondence optimization proceeds by gradient descent, as described in Section 2.3, with the substitution of the gradient of the new cost function \tilde{G} for the original cost function G in Eqn. 2.36.

In summary, the generalized PBM algorithm replaces the entropy of positional information with entropy of an arbitrary function of positional information. This modification offers a much more generalized framework for optimizing the statistical properties of an ensemble of shapes. Note that the standard PBM algorithm from Chapter 2 is a special case of the generalized PBM algorithm, where $f(\mathbf{z}) = \mathbf{z}$. In Section 5.3, I present an application of the generalized PBM algorithm to the problem of correspondence on cortical surfaces.

3.4 Correspondence with Regression Against Explanatory Variables

In general, the design of a scientific study in biology or medicine cannot control for all confounding variables. The variability in shape due to such factors as age, differential growth rates, or clinical variables, for example, must be accounted for during the analysis phase. In other cases, this variability is the specific focus of the study, and researchers want to examine the correlation of an explanatory variable with shape. A typical experiment, for example, might examine the correlation of disease progression with the shape of anatomical structures or the change in the shape of anatomy with age. If such correlations can be established, they may lead to new diagnostic protocols or interventional planning.

This section extends the PBM algorithm to the problem of establishing correspondence in the presence of confounding variables and examining the correlation of shape with explanatory variables. Like the previous section, this method allows for a more general notion of correspondence that takes into account additional information about the data under study. The algorithm works by expanding the point-based correspondence model from Section 2.2 to include a regression against the independent variables. The optimization of correspondence position is then done on the *residual* to the regression model. The theory behind this approach is

the same principle of parsimony behind the basic PBM formulation: to minimize model parameters and maximize statistical power.

Under the assumption of a Gaussian distribution for the random variable \mathbf{Z} from Equation 2.37, which is the distribution of shape samples in shape space, we can write the generative statistical model

$$\mathbf{z} = \mu + \epsilon, \epsilon \sim \mathcal{N}(\mathbf{0}, \Sigma) \quad (3.9)$$

for particle correspondence positions, where ϵ is normally-distributed error. Replacing μ in this model with a function of an explanatory variable t gives the more general, regression model

$$\mathbf{z} = f(t) + \hat{\epsilon}, \hat{\epsilon} \sim \mathcal{N}(\mathbf{0}, \hat{\Sigma}). \quad (3.10)$$

The optimization described for the basic PBM algorithm minimizes the entropy associated with ϵ , which is the difference from the mean. In this section, the goal is to optimize correspondences under the regression model in Eqn. 3.10 by instead minimizing entropy associated with $\hat{\epsilon}$, the residual from the regression model. For the simple case where particle correspondence is a linear function of t , given as $f(t) = \mathbf{a} + \mathbf{b}t$, parameters \mathbf{a} and \mathbf{b} must be estimated to compute $\hat{\epsilon}$. These parameters are estimated with a least-squares fit to the correspondence data,

$$\arg \min_{\mathbf{a}, \mathbf{b}} E(\mathbf{a}, \mathbf{b}) = \frac{1}{2} \sum_k [(\mathbf{a} + \mathbf{b}t_k) - \mathbf{z}_k]^T \Sigma^{-1} [(\mathbf{a} + \mathbf{b}t_k) - \mathbf{z}_k]. \quad (3.11)$$

Setting $\frac{\delta E}{\delta \mathbf{a}} = \frac{\delta E}{\delta \mathbf{b}} = 0$ and solving for \mathbf{a} and \mathbf{b} , we have

$$\mathbf{a} = \frac{1}{n} \left(\sum_k \mathbf{z}_k - \sum_k \mathbf{b}t_k \right), \quad (3.12)$$

and

$$\mathbf{b} = \left(\sum_k t_k \mathbf{z}_k - \sum_k \mathbf{z}_k \sum_k t_k \right) / \left(\sum_k t_k^2 - \left(\sum_k t_k \right)^2 \right). \quad (3.13)$$

The proposed regression model optimization algorithm proceeds as follows. Correspondences are first optimized under the nonregression model (Eqn 3.9) to minimize the entropy associated with the total error ϵ . This process also establishes an initial estimate for \mathbf{a} and \mathbf{b} . The next step is to optimize under the regression model, which proceeds by gradient descent on $H(\mathbf{Z}) \approx \frac{1}{2} \log |\hat{\Sigma}| + H(P^k)$. In other words, the method follows the same optimization procedure as the basic PBM framework (Section 2.3), but replaces the covariance of the model with the covariance of the underlying residual, relative to the generative model. The two estimation problems are interwoven: the parameters \mathbf{a} and \mathbf{b} are re-estimated after each iteration of the gradient descent on the particle positions.

As an example of a correspondence optimization, the method was applied to a set of $N = 40$ tori. To generate each torus, the large radius R was randomly drawn from a Gaussian distribution $R \sim \mathcal{N}(35, 3)$, and small radius r randomly drawn

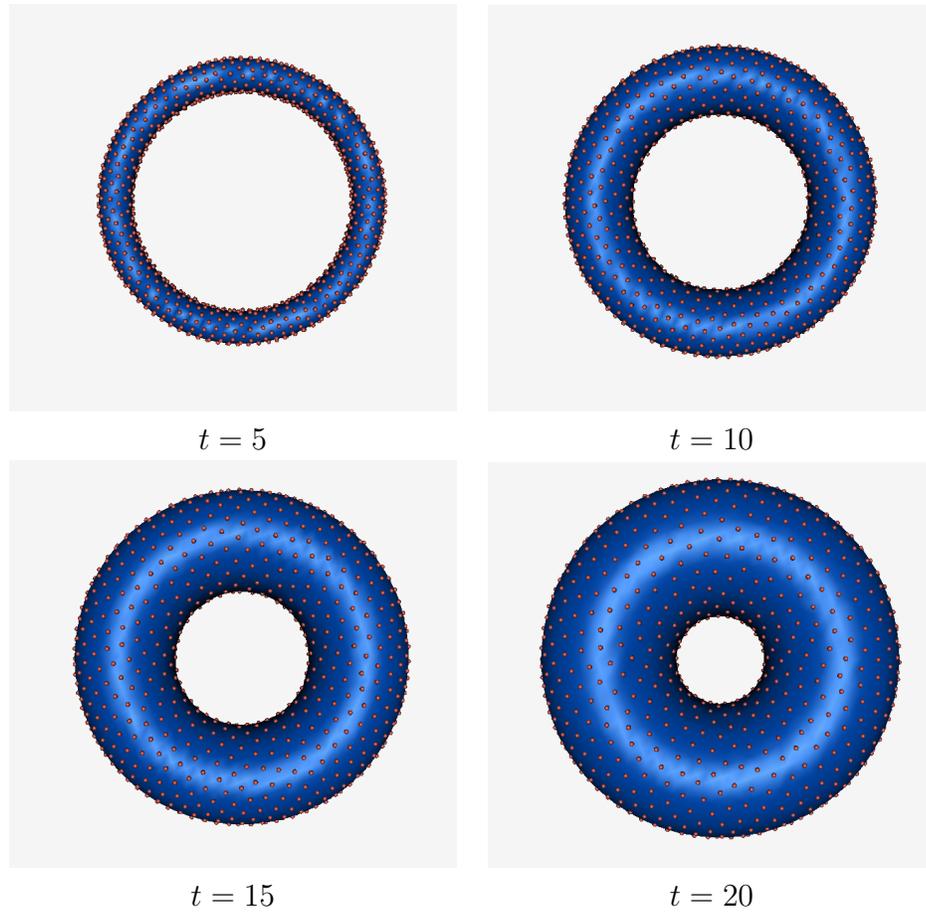


Figure 3.2. Shape regression on an explanatory variable t correlated with the small radius of a set of tori with $r \sim \mathcal{U}(5, 20)$.

from a uniform distribution $r \sim \mathcal{U}(5, 20)$. An explanatory variable $t_i = r_i + \epsilon$, with $\epsilon \sim \mathcal{N}(0, .3)$, was assigned to each shape sample $i \in \{1, \dots, 40\}$ to establish a good correlation with variation in the small torus radius. Correspondences were optimized using 1024 particles per shape and the PBM regression algorithm outlined above. In the resulting correspondences, variation in the residuals to the regression line exhibits one major mode that empirically corresponds to r . Empirical observation of the regression line, which is shown in Figure 3.2 suggests good correlation with R .

A further example of the PBM regression algorithm is given in Section 5.4, where I apply the method to compute a model of the change in neonatal head shape with age. The analysis of the statistical significance of regression models from this method is discussed in Section 4.5.

CHAPTER 4

STATISTICS AND INTERPRETATION

This section describes methods for statistical analysis of correspondence positions. These methods are tools developed specifically for use with the PBM framework and its extensions, but are generally applicable to any method for dense landmark or correspondence-based shape modeling. There are two major challenges for statistical analysis of shape models that I address in this chapter. The first challenge is the very high dimensionality of the shape space, coupled with relatively low numbers of shape samples. The second challenge is how to visualize the significant geometric variation identified by statistical tests. This latter consideration is important in order for researchers to interpret statistical results in the context of biological hypotheses.

To date, research in computationally-derived shape modeling has focused mainly on the construction of representative models, and the field has not reached a consensus regarding the best approach to statistical analysis. For this dissertation, I therefore propose to adapt well-established multivariate methods from the statistics literature and the geometric morphometrics community to correspondence models. These methods require careful application and interpretation in the context of the high dimensionality of the correspondences in order to produce valid results. New methods for visualizing the statistical significance of the results are also required.

The chapter is organized as follows. First, I describe computation of the sample mean and median, and the rationale behind statistical analysis using distance among Procrustes-aligned correspondences. Next, I discuss group-wise hypothesis testing with correspondences and propose a principled approach that avoids the difficulties inherent in high-dimensional, low-sample size models. This approach is based on multivariate methods from the statistics literature that are commonly used and understood by biomedical researchers, an important consideration for its practicality in real scientific studies. I also present ideas for visualizing statistically significant shape variation. Next, I discuss more complex analysis scenarios, such as the analysis of variance, and special considerations for significance in optimized regression analysis. Finally, I describe the analysis of object pose and scale, which are metrics related to shape that are often examined in studies of anatomical variability.

4.1 Summary Statistics

Recall from Section 2.2 that correspondence-based modeling methods, such as the PBM algorithm, represent shape by sampling each shape surface in a consistently ordered fashion so as to define homologous object surface points. Thus, we

have a statistical sample of N surface representations drawn from a population of surfaces, and the model of shape variation is the variation in a set of M , d -dimensional points on each of the N surfaces. The set of positions for all M correspondences on a shape can be mapped to a dM -dimensional shape vector. The positions of the vectors in the dM -dimensional shape space give rise to the statistical analysis.

In order to assess statistical differences among positions in shape space, it is necessary to first define a metric for the distance between positions. The pioneering work of David Kendall and others has shown that the true distance between shapes in shape space is nonlinear because shapes actually lie on a complex manifold within the dM -dimensional, Euclidean shape space [64]. Because the true, nonlinear distances are difficult to compute in practice, however, and statistics using nonlinear metrics is not generally well understood, researchers in the geometric morphometrics community have developed linear distance approximations that are currently in widespread use. One standard approach to landmark-based shape statistics, for example, has been described by Bookstein, Goodall, Kent, and others and uses linear distance in a local tangent plane of the Kendall shape space manifold, which is computed around the mean sample shape [57, 66, 45]. The mean sample shape is computed as the *Procrustes mean* of the landmarks, obtained by Generalized Procrustes Analysis (GPA) [47, 9, 45], which places the mean landmark configuration on the Kendall shape space manifold [45, 34].

In this dissertation, I propose to use the Euclidean distance between Procrustes-aligned points in dM -dimensional shape space as an approximate distance measure. The rationale behind this choice is as follows. The Procrustes mean has been shown to be equivalent to the Euclidean mean of the Procrustes-aligned landmark points [66, 45, 34]. When the variation among shapes in a sample is small, statistical metrics that rely on Euclidean distances to the mean are therefore reasonable approximations to the tangent-space approaches proposed by Bookstein, et al., with the error in the approximation limited to the residual distance from the tangent plane (see also [34]). Furthermore, since the statistical methods developed in this dissertation are linear, other distance approximations can be easily substituted, provided that they are also linear. Thus, the choice of a specific distance metric is not critical to the research discussed in this dissertation and only depends on the linearity of the metric.

From the arguments above, it follows that a good choice for the sample mean of the correspondence positions is the Euclidean mean of the Procrustes-aligned points, which is given by

$$\mu = \frac{1}{N} \sum_{k=1}^N \mathbf{z}^k, \quad (4.1)$$

where \mathbf{z}^k are the mappings of the correspondence positions to the dM -dimensional shape space, as described in Section 2.2. The median shape is computed as the L_1 multivariate median (e.g., [111]), with all distances between shapes weighted equally. The problem equates to finding the sample with the minimum sum of its Euclidean distances to all other samples. Thus, we have a cost function and a minimization problem given by

$$C(\mathbf{z}) = \sum_{k=1}^N \|\mathbf{z} - \mathbf{z}^k\|, \text{ and } \mathbf{z}_{\text{median}} = \arg \min C(\mathbf{z}) : \mathbf{z} \in \{\mathbf{z}^1, \mathbf{z}^2, \dots, \mathbf{z}^N\}. \quad (4.2)$$

For the applications described in this dissertation, the sample size N is small enough to find $\mathbf{z}_{\text{median}}$ by “brute force”, that is, simply computing C for every sample and choosing the sample with the smallest C .

4.2 Principal Component Analysis

Generalized Procrustes analysis, as discussed in the previous section, describes a mean shape and a space in which linear approximations to the nonlinear shape-space distances are defined. Another important consideration for shape analysis, however, is to characterize the *variability* of shape. One widely used approach is principal component analysis (PCA) [87]. PCA identifies a set of no more than $N - 1$ orthogonal basis vectors that describe the distribution of the N samples in shape space. Each basis vector is chosen to minimize the sum-of-squares residual error in the data; thus, the bases are optimal in the least-squares sense. For a correspondence model, these basis vectors account for the correlations among correspondence positions and thus describe the independent *modes* of variation of shape. Shape space coordinates projected onto individual PCA modes, or *PCA loadings*, are useful for visualizing variation and for dimensionality reduction in statistical analysis (see Chapter 4).

PCA has been widely used in landmark-based shape analysis for both visualization and dimensionality reduction. Kent, for example, describes PCA in the Procrustes tangent space [66], and Cootes, et al. discuss PCA for correspondence models [21]. A generalization of PCA to nonlinear distances on manifolds, called principal geodesic analysis (PGA), has been developed by Fletcher, et al. [39]. While PGA offers a much more accurate approach by accounting for the true nonlinear distances between shape samples, multivariate statistical methods on PGA loadings are not as well understood by researchers, and standard PCA remains the most common approach.

For this dissertation research, I propose to use standard PCA on the correspondence positions of Procrustes-aligned points to describe shape variability. This approach follows from the choice of the linear distance approximation described in the previous section. I will also use PCA for dimensionality reduction in statistical analysis, as described in Section 4.3.1. Note that the error in this approximation increases with the distance from the Procrustes mean. As an example of the error, consider the “box-bump” shapes in Figure 2.10. Shapes at larger distances from the mean (e.g., 3σ) are noticeably distorted, relative to the shapes in the training set. The distortion is due, at least in part, to the error in the linear distance approximation.

Principal component analysis is a linear transformation of data into a new coordinate space, in which each coordinate axis, or *component*, represents a decreasing amount of variability in the data. The first component is the axis describing the greatest amount of variability in the data, the second component describes the second greatest amount of variability, and so on. PCA can be computed by an

eigenanalysis of the covariance of the data (see, for example, [21]). Thus, a principal component analysis for PBM correspondence models is given by the eigenvectors and eigenvalues of the matrix

$$\Sigma = (dMN - 1)^{-1} \mathbf{Y}\mathbf{Y}^T, \quad (4.3)$$

which is the covariance of the correspondence positions from Equation 2.27. Because $N < dM$, the eigenanalysis is computed in the dual space of Σ^T , as described in Section 2.3.8. The eigenvectors represent the principal component axes, and the magnitudes of the eigenvalues indicate the relative amounts of variation in each component. Examples of visualization and statistical analysis of PCA loadings are given in Sections 2.5.1 and throughout Chapter 5.

4.3 Hypothesis Testing

In the introductory material for this dissertation, I discussed how statistical shape models can be powerful tools for describing anatomical structures and are increasingly being used in a wide variety of clinical and biological contexts. One of the promising applications of this technology is the quantification of anatomical shape differences between control and study populations and the visualization of those group differences. The geometric morphometrics community has established some standard approaches to statistical analysis for small collections of landmarks that rely on traditional multivariate statistical methods [34, 1]. The statistical analysis of higher-dimensional, computationally-derived models, however, is more difficult due the larger numbers of degrees of freedom and the challenge of obtaining sufficient numbers of subjects to ensure statistical power. This problem is known in the statistics literature as the high-dimension, low sample-size (HDLSS) problem [2]. Another challenging aspect of HDLSS models is how to visualize the statistical significance of geometric variation. In other words, biomedical researchers need to understand not only whether there are significant group differences, but also what those differences look like. This section discusses these issues in the context of HDLSS correspondence models and proposes a hypothesis testing and visualization framework based on standard approaches in multivariate statistics.

4.3.1 Challenges

In the context of point-based models, or surface samplings obtained from parameterized models, one approach to shape statistics is point-wise analysis of correspondences, which are elements of R^3 , for shapes in three dimensions. These data are drawn from the correspondence configurations (Section 2.2) and are marginal distributions of the full shape space. The mean shape is the Euclidean averages of correspondence positions, as described in the previous section, and point-wise differences in correspondence positions define local shape variation [98]. Hypothesis tests, in this case, reveal regions of significant differences between groups, which can be directly visualized as p -value maps on the mean shapes. Styner [98], for example, proposes a statistical analysis of the correspondence positions using a nonparametric permutation test with the Hotelling T^2 metric, followed by a false-discovery rate (FDR) correction for the multiple-comparison problem inherent in the analysis.

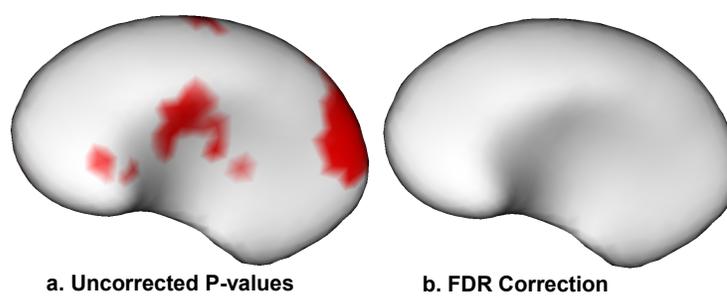


Figure 4.1. Point-wise hypothesis test results for the putamen in a pediatric autism study.

Figure 4.1 is an illustration of point-wise hypothesis testing as proposed in [98] on a model of the right putamen from normal control subjects and autism patients. The data are taken from an ongoing longitudinal pediatric autism study [51]. For the test, there are 10 autism structures and 15 matched normal controls. The PBM algorithm was used to optimize 1024 correspondence points, and the analysis included 20,000 test permutations. The correspondences were computed without knowledge of group by including both autism and normal controls in a single model optimization. The p -values before FDR correction that indicate significance at the 5% level are colored in red on the mean normal putamen surface in Figure 4.1a and suggest several distinct areas of shape differences. Figure 4.1b shows that in this case, however, which is not uncommon in neurological shape analysis, no significant p -values remain after FDR correction (using a 5% bound). This example illustrates a major difficulty encountered in point-wise analysis: the large number of comparisons results in a very conservative correction of the hypothesis test results, significantly reducing the statistical power of the test.

To avoid the multiple-comparisons problem, the shape model can be analyzed directly in the dM -dimensional shape space. The analysis in this case, however, is also difficult because traditional statistical metrics no longer apply with small sample sizes [2]. At issue is the fact that the convergence of any estimator in very high dimensional space is prohibitively slow with the respect to the number of samples. We would need many more samples than are practical for clinical studies in order to reliably estimate distances between, for example, population means. A common solution is to employ dimensionality reduction by choosing a subspace of the dM -dimensional space in which to project the data for traditional multivariate analysis. If this subspace is linear, this entails the selection of a set of basis vectors, and we can easily evaluate the loss of information by projecting the population onto the k -dimensional subspace and evaluating the residual.

Principal components analysis (PCA) [105] is an attractive choice for dimensionality reduction. The basis vectors derived by PCA are orthogonal and determined solely from the data. Thus, they provide an effective means of visualizing the major modes of shape variation. PCA in the Procrustes tangent space, for example, has been widely used by biomedical researchers for landmark-based shape analysis [34, 66]. With PCA, we can find no more than $N - 1$ modes that have nonzero variance, meaning that the problem is reduced to $k < N$ without loss of information. Other basis functions such as wavelets [81] have also been used for dimensionality reduction, with the difference being that they impose an a-priori choice of how the space of the model should be decomposed.

In a suitably low-dimensional shape space, such as basis vectors from PCA, we can apply traditional statistical methods such as the Hotelling T^2 test [105]. There are two major challenges for dimensionality reduction in shape analysis, however. First, is the the model selection problem of *how many* basis vectors to choose, which can be hard to resolve when the choice of different numbers of components leads to different statistical results. The second challenge is how to visualize group differences, which is important for researchers in order to relate the outcomes of statistical tests to scientific hypotheses regarding the growth and form of anatomy. Visualization is discussed further in Section 4.3.3. The remainder of this section describes the dimensionality reduction problem.

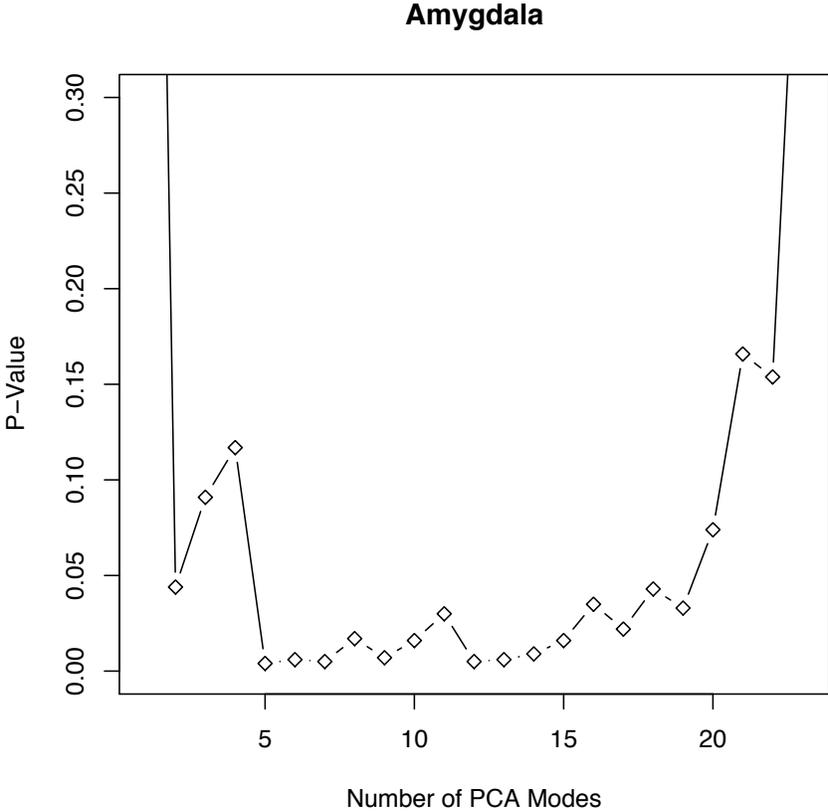


Figure 4.2. Hotelling T^2 test results with increasing numbers of PCA modes for the amygdala brain structure from a pediatric autism study.

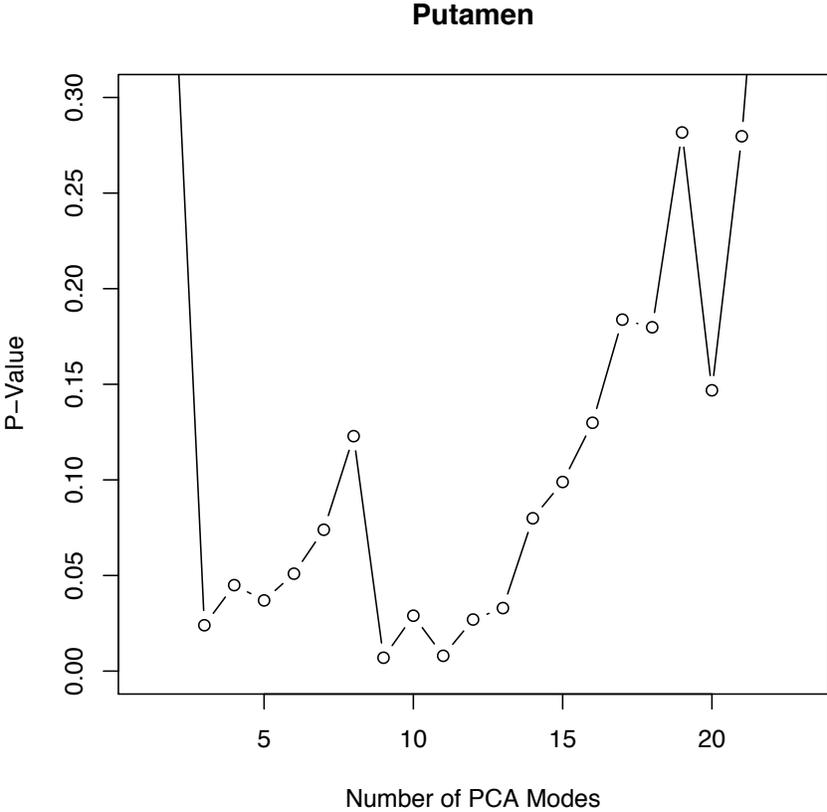


Figure 4.3. Hotelling T^2 test results with increasing numbers of PCA modes for the putamen brain structure from a pediatric autism study.

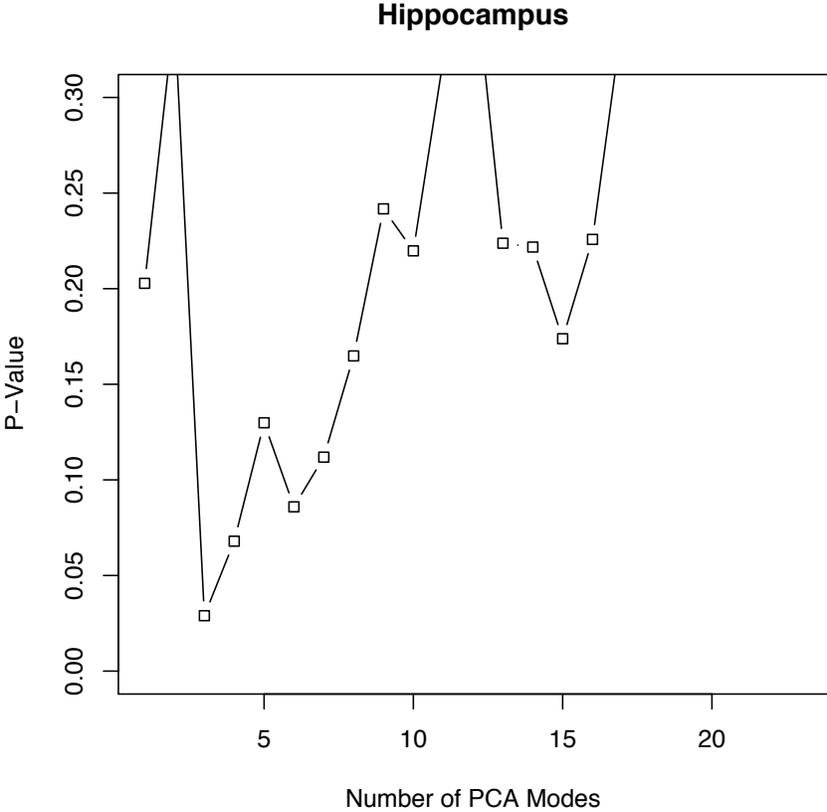


Figure 4.4. Hotelling T^2 test results with increasing numbers of PCA modes for the hippocampus brain structure from a pediatric autism study.

To illustrate the challenge in choosing an appropriate subspace for analysis, Figures 4.2–4.4 show the p -value results of Hotelling T^2 tests using increasing numbers of PCA modes on three brain structures from the pediatric autism study referenced in Figure 4.1 above. Several trends can be observed that pose a challenge for the analysis. First, is the trend at higher numbers of modes towards increasing p -values, which is likely due to the cumulative effects of noise in these lower-variance modes of the PCA. The second trend is that the p -value curves are not monotonic, that is, they do not smoothly decrease to a minimum value. It is therefore possible to choose a variety of inconsistent interpretations of significance simply by choosing different numbers of modes. The challenge for this type of analysis is to choose as many modes as possible that contain meaningful variation, i.e. variation that is distinguishable from noise, with the caveat that too few modes may result in the loss of information that is useful for group discrimination. I will refer to this challenge as the *number-of-bases* problem.

Many methodologies have been proposed to address the number-of-bases problem and good reviews, such as [35], are available. Methodologies range from simple conventions, such as choosing only PCA modes that account for at least 5% of the total variance, to more systematic approaches that attempt to model the problem based on its dimensionality and sample sizes. Notable among these latter approaches, *parallel analysis* is commonly recommended as the best method for determining modes with variance that is distinguishable from noise [43].

4.3.2 Hypothesis Testing Framework

The previous section analyzed several important challenges for hypothesis testing with correspondence shape models and suggested that a systematic approach to dimensionality reduction, followed by standard multivariate tests, will avoid the HDLSS problem. In this section, I outline a specific approach in more detail. The remaining challenge of how to present an intuitive visualization of hypothesis test results is left for the next section.

The first step for hypothesis testing is to compute a single correspondence model using the combined data from all of the groups. I will refer to this model as the *combined* group model. It is important that the model is constructed without knowledge of the group classifications, or the results of the test will be biased. If the analysis compares two separate correspondence models, they will typically show a very high probability of group differences, even if the data was originally drawn from the same shape distribution. Consider two models for the same set of spherical shapes, for example. An infinite number of equally-optimal correspondence models can be found by simply rotating all of the correspondences around the sphere centers. Thus, it is possible to find group differences where none exist, unless group statistics are measured within a single, unified model.

For a correspondence point model of shapes embedded in a d -dimensional space, we have a $dM \times N$ shape matrix \mathbf{P} , where columns of \mathbf{P} are the shape vectors of correspondence points for the set of all samples. After the optimization of the combined model, dimensionality is reduced by projecting \mathbf{P} onto the basis vectors determined by PCA analysis, i.e., $\tilde{\mathbf{P}} = \mathbf{E}\mathbf{P}$, where columns of \mathbf{E} are the eigenvectors

of the covariance matrix of correspondences, in decreasing order of the magnitude of their eigenvalues.

Following the PCA, the next step is to choose the number of PCA bases to retain using parallel analysis. In the context of principal components analysis of N , vector-valued data samples of dimensionality dM , the goal of parallel analysis is to identify the subset of the components that contain variation distinguishable from the expected variation resulting from noise, where noise is modeled by an isotropic, multivariate unit Gaussian, i.e. a random dM -vector $\mathbf{X} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. To make this distinction, we need an estimator \mathbf{E} for the expected values of the variances in the ordered PCA modes of random samplings on \mathbf{X} , given the fixed sample size N . Because PCA reorders modes based on their percentage of the total variation, there is no obvious closed-form expression for \mathbf{E} , so it is estimated using Monte Carlo simulation. Many random sample sets of size N are independently drawn from \mathbf{X} , followed by PCA on each sample set and ordering of the associated eigenvalues. The ordered eigenvalues are then averaged to produce an estimate of the Gaussian noise variance profile across modes. Note that the eigenvalues in this case quantify variance, and the percentage of total variance for a PCA mode is equivalent to the ratio of its eigenvalue to the sum of all eigenvalues.

In order to determine the number of modes to use from parallel analysis, the percent-total-variance profiles, or *scree plots*, from the Monte Carlo simulation and the PCA of the true data are compared. Only the modes where the percent-total-variance in the true data is greater than the simulation data are retained. Figure 4.5, for example, is a scree plot of the percent-variances associated with shape data of a putamen brain structure ($N = 25$, $dM = 3000$) [18] (solid line) and the variances from the Monte Carlo noise variance simulation (dashed line). The two lines intersect just before mode 6, and so we would consider only modes 1-5 in the analysis.

Once we have chosen the set of k PCA basis vectors by parallel analysis, the next step for hypothesis testing is to project the correspondence data onto the bases to give a set of k , principal component *loadings*. Standard multivariate statistics are now possible on the PCA loadings. A widely used approach to multivariate hypothesis testing is the parametric Hotelling T^2 test [105], which is the method used for all results reported in this dissertation. The null hypothesis for the test is that the two groups under comparison are drawn from the same distribution. A low p -value therefore provides evidence for rejection of the null hypothesis. The Hotelling T^2 , two-sample metric is given by

$$T^2 = \frac{(n_a n_b)(n_a + n_b - 2)}{n_a + n_b} (\mu_{\mathbf{a}} - \mu_{\mathbf{b}})^T \mathbf{w}, \quad (4.4)$$

$$\mathbf{w} = (\Sigma_a + \Sigma_b)^{-1} (\mu_{\mathbf{a}} - \mu_{\mathbf{b}}),$$

where $\mu_{\mathbf{a}}$ and $\mu_{\mathbf{b}}$ are the means, Σ_a and Σ_b are the covariances, and n_a and n_b are the sample sizes of the two groups, respectively.

To illustrate the application of proposed hypothesis testing method, I present an analysis of the brain structure models for the pediatric autism study introduced in the previous section. The hypothesis test p -value results for the amygdala, putamen, and hippocampus models from Figures 4.2–4.4, respectively, are 0.003, 0.046, and

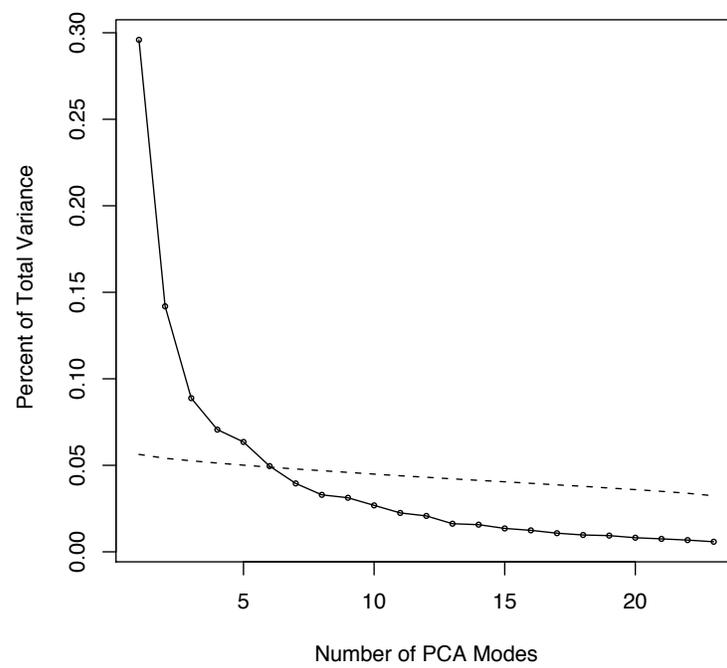


Figure 4.5. Parallel analysis for the putamen data

0.100, with the number of PCA modes chosen as 5, 6, and 5. Of particular interest is the fact that the result for the putamen indicates group differences at the 5% significance level, while the point-wise hypothesis tests shown in Figure 4.1 indicate no significant differences. This difference illustrates the increased statistical power of the proposed testing method, which avoids the multiple-comparisons problem.

4.3.3 Visualizing Significant Group Difference

In contrast to the point-wise statistical method illustrated in Figure 4.1, a significant drawback of hypothesis testing in a PCA subspace is that the group differences in this space are not necessarily easy to visualize and interpret from an anatomical perspective. The hypothesis test poses the question of whether there is evidence for significant group differences. The next logical question that researchers may ask is what those group differences look like. This section proposes two methods for visualizing group differences. The first approach is a direct visualization of the linear discriminant implicit in the Hotelling T^2 hypothesis test metric. I also discuss how this approach differs from the group mean differences, which does not confer any information regarding statistical significance. Secondly, I propose an approach to visualizing mean differences using strain tensors from a deformation between the group means that is analogous a standard method used for landmark data by the geometric morphometrics community.

To gain some insight into the geometric differences that are most statistically significant between groups, one approach is to transform the group differences measured in the PCA space back into the full shape space where they can be visualized on the mean shape surfaces. Implicit in the Hotelling T^2 metric is a linear discriminant vector which indicates the direction in the PCA space along which the maximum group difference is observed. The discriminant vector \mathbf{w} (Equation 4.4) is also known as Fisher’s linear discriminant and is the line along which the between-group variance is maximized with respect to within-group variance [105]. The Hotelling T^2 metric is therefore a scaled projection of the group difference onto the discriminant line. The most significant group differences are therefore seen by transforming \mathbf{w} back from PCA space into the full-dimensional shape space, i.e. $\hat{\mathbf{w}} = \mathbf{E}^{-1}\tilde{\mathbf{w}}$, where $\tilde{\mathbf{w}}$ is \mathbf{w} padded to n -dimensions with $n-k$ zeros. The components of the dM -vector $\hat{\mathbf{w}}$ can then be mapped onto a visualization of mean correspondence point positions. One possible visualization is a surface mesh reconstruction from the correspondence points of group a . The resulting linear discriminant analysis (LDA) visualization indicates group differences in shape that the test metric identified as the most effective for discrimination. The differences are shown on the mean shape for a , in the direction of the mean shape for b .

Figure 4.6 is an example of an LDA visualization for the putamen data introduced in Section 4.3.1. The discriminant vector is shown on the the mean control population correspondence points, with arrows indicating the direction from patient to the normal control populations. The visualization indicates a shortening of the anterior and posterior regions of the putamen, with a thickening in the medial region.

One common approach to *empirical* evaluation of group differences is to visualize the point-wise difference vectors between group mean shapes. This visualization,

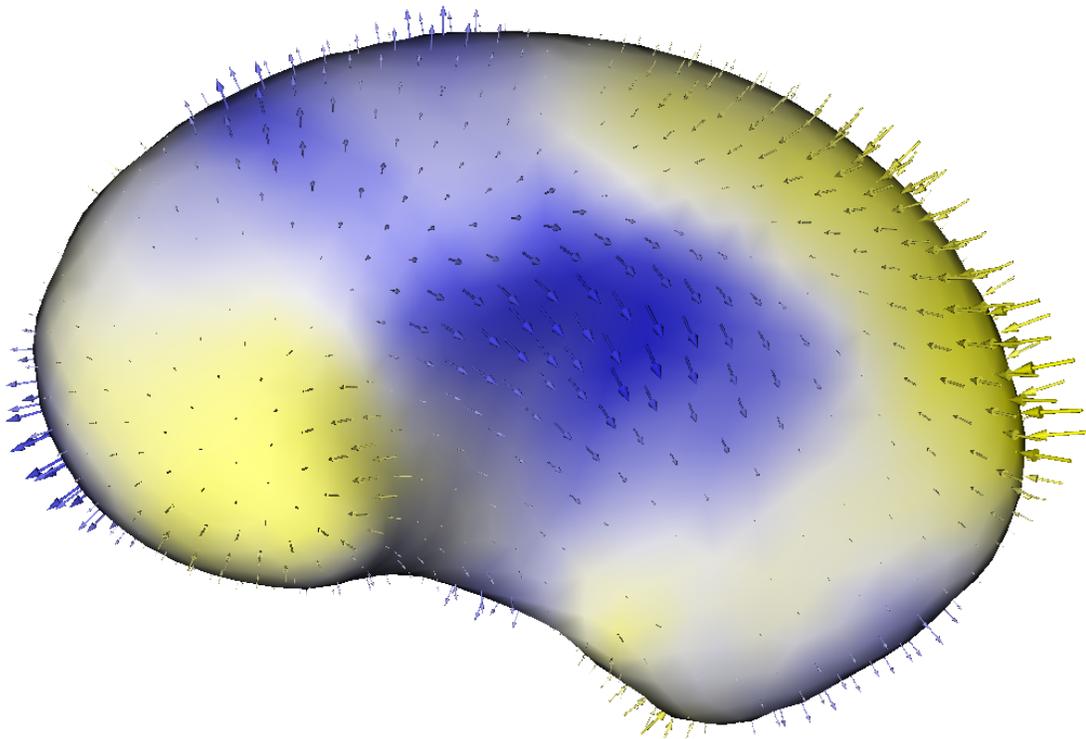


Figure 4.6. LDA visualization of the right putamen from an autism study

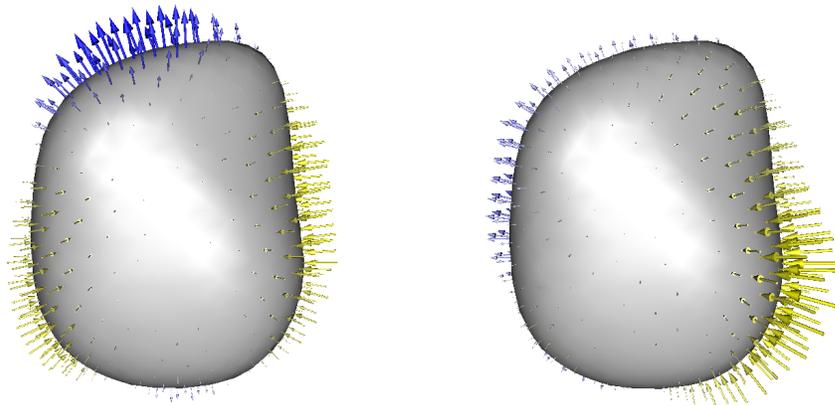


Figure 4.7. A comparison of the linear discriminant visualization (left) with a visualization of the mean group difference (right) on a PBM correspondence model of the left putamen in a shape study of autism versus normal control.

however, does not give any indication of statistical significance unless it is accompanied by a point-wise hypothesis test method (see Section 4.3.1). It is possible, of course, that the group mean difference vector may be similar to the discriminant vector, especially in cases where group shape differences are pronounced. In many other cases, however, the LDA visualization can be very different. Figure 4.7 compares the two approaches for a PBM model of the left putamen from the autism study described in Section 4.3.1. Note that the group mean differences in this case are very different from the linear discriminant.

Though group mean differences do not indicate statistical significance, they are still a traditional approach to empirical evaluation of shape data. The most common method for empirical evaluation of group differences in two-dimensional landmark models are visualizations of a thin-plate spline deformation between the group means [11]. Visualizations of deformations indicate regions of expansion and contraction of the surface, which better reflect a model of biological growth than simple directional vectors. Note that this type of analysis is very similar to the voxel-based morphometry approaches reviewed in Section 1.2.

A similar approach to the thin-plate spline methods for landmark models can be applied in the context of correspondence models of three-dimensional shapes. To visualize deformations between the group mean shapes, we can compute metrics on the displacement field describing the mapping from points \mathbf{x} on one group mean to corresponding points \mathbf{x}' on the another. Using the set of correspondence points, a smooth transformation $T(\mathbf{x}) = \mathbf{x}'$ may be computed using a thin-plate spline interpolation. Details for computing $T(\mathbf{x})$ are omitted here, as good descriptions can be readily obtained elsewhere (e.g., [11, 113]). To describe mean differences, I propose a visualization of *strain*, which is a measure on the Jacobian \mathbf{J} of the deformation field $\mathbf{x} - T(\mathbf{x})$, and describes the local stretching and compression caused by the deformation. The Lagrangian strain tensor is a symmetric, second order tensor given by

$$\mathbf{E} = \frac{1}{2}(\mathbf{J} + \mathbf{J}^T + \mathbf{J}^T \mathbf{J}). \quad (4.5)$$

The eigenvectors of \mathbf{E} indicate the principal directions of strain, and the eigenvalues of \mathbf{E} indicate the unit elongations in those directions. An effective visualization for the strain tensor is an ellipsoid with principal axes given by the eigenvalues and oriented along the eigenvector directions.

Figure 4.8 is a visualization of the strain tensors computed from a thin-plate spline deformation of the mean patient shape to the mean normal control shape for the putamen data. In the figure, the three principal axes of each ellipsoid are scaled by the three principal eigenvalues of the strain tensor at each correspondence position and oriented according to their corresponding eigenvectors. Ellipsoids and the surrounding surface are colored according to the value of the first principal eigenvector (the longest axis), with yellow indicating negative (compression) and blue indicating positive (stretching). While a clinical interpretation of this result is beyond the scope of this analysis, this visualization may offer a more detailed insight into how groups differ than a mean difference visualization. Note, however, that this visualization also gives no indication of the statistical significance of these differences.

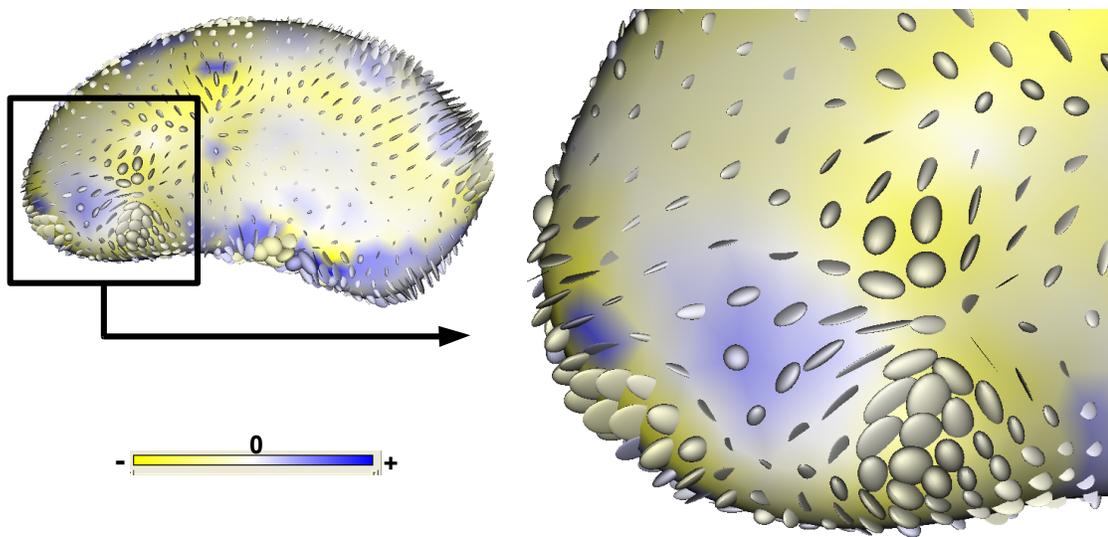


Figure 4.8. Strain tensors for the right putamen from an autism study. Tensor scale is exaggerated for visualization purposes.

4.4 Linear Models and Regression

Many studies in biomedicine must account for variables, such as age, that are independent of shape, but that are somehow correlated with shape, either directly or indirectly through a common factor. These variables are called *explanatory* variables and often cannot be completely controlled for by the experimental design. This section illustrates how standard multivariate methods may be used on the PCA loadings from the previous section to apply the PBM correspondence framework to statistical analysis in the presence of explanatory variables.

A common approach to controlling for explanatory variables is to model the outcome variables (e.g., shape) as a linear function of the explanatory variables and analyze the residual variation from this model. A model of outcome as a function of explanatory variables is known as a regression model, and it is important to assess whether the regression is a significantly better estimate of the data than the sample mean. Significance is commonly assessed by comparing the ratio of the residual variation in the regression model to the total variation in the data [105]. A common test statistic for significance is

$$T = \frac{R^2/(p-1)}{(1-R^2)/(N-p)}, \quad (4.6)$$

where R^2 is Pearson's coefficient of regression, generally defined as $R^2 = 1 - \frac{SS_{\text{err}}}{SS_{\text{tot}}}$. The constant p is the number of parameters in the regression model, N is the number of data points, SS_{err} is the sum-squared residual error in the regression model, and SS_{tot} is the total variance in the data. In general, R^2 can be related to the unexplained variance of the generated model and is also commonly used to measure the goodness-of-fit for the regression model. When the residuals of the linear model are independently and identically-distributed Gaussian, the statistic T follows an F -distribution with $N_1 - 1$ and $N_2 - N_1$ degrees of freedom under the null hypothesis [105]. The significance of the correlation is therefore estimated as the probability of T under the F -distribution.

For univariate shape parameters, such as single PCA loadings, we can perform a simple regression analysis with significance assessed as described above. In the context of hypothesis testing, a group comparison with the residuals to the regression line can be done using standard *analysis of covariance* (ANCOVA) methods. ANCOVA introduces the group label as an explanatory variable and tests the significance of its influence on the variance of the model. When including a group factor, it is also important to establish that there is no significant correlation, or *interaction*, between the group variable and the explanatory variables, otherwise the simple regression model cannot be considered valid.

As an example of an ANCOVA analysis, consider the case of a single explanatory variable $a \in \mathfrak{R}$ and a two-group comparison with group variable $g \in \{0, 1\}$. The interaction model is given by

$$y = \beta_1 a + \beta_2 g + \beta_3 a g + \epsilon, \quad (4.7)$$

where ϵ is the residual, and y is the response variable. The response variable y might represent, for example, a single shape PCA loading, and the variable a might

indicate test subject age. To test whether the interaction term ag is significant, we can use a t -statistic to test the hypothesis that $\beta_3 = 0$, following standard approaches from the statistics literature (for example, [105]). If the interaction term β_3 is not found to be significant, then we drop the term ag and test the significance of β_2 in the regression model

$$y = \beta_1 a + \beta_2 g + \epsilon, \quad (4.8)$$

under the assumption that the response y in both groups varies in the same way with a .

The linear models in 4.7 and 4.8 can be extended to vector-valued response variables (e.g., multiple correspondence PCA loadings), and significance of vector-valued parameters β_i assessed following a similar logic to the ANCOVA analysis. The extension of the ANCOVA model to multiple response variables is known as multiple analysis of covariance (MANCOVA), which determines significance by comparing metrics on the covariance of vector-valued residuals to the fitted models. For the results presented in this dissertation, the above understanding of the statistical analysis of linear models is sufficient. For more detail, however, the reader is referred to the many comprehensive texts on the subject (for example, [53]).

To visualize shape variation in linear regression models, I propose to compute predicted response variables from estimated model parameters using the linear regression equation

$$\hat{\mathbf{y}} = \hat{\beta}_1 a + \hat{\beta}_2 g. \quad (4.9)$$

The fitted k -vectors $\hat{\mathbf{y}}$ represent PCA loadings of the mean shapes for each group at a given value of a . The PCA loadings can be rotated back to the full dM -dimensional shape space, as described for the linear discriminant vector w in Section 4.3.2, and added to the Euclidean mean correspondences for visualization purposes. Some examples of this type of visualization are given in Section 5.1. Visualizing reconstructed shapes along the regression line allows for empirical investigation of both shape change with explanatory variables and how group differences change with explanatory variables.

4.5 Statistical Significance of Regression

While regression analysis can be done in the PCA space using standard linear regression and multivariate metrics, as outlined in the preceding section, the PBM algorithm extension described in Section 3.4 proposes a method for computing regression in the full shape space and an optimization of the correspondence positions with respect to the regression model. In order to understand the statistical significance of the optimized model, we need a systematic, unbiased method for testing the correlations. This section describes such an approach and is joint research with Manasi Datar, at the University of Utah.

In this case where the outcome variables are correspondence-optimized shape parameters, the underlying assumptions of the parametric F -test of the statistic T in Equation 4.6 may not hold. Furthermore, optimization with knowledge of

the underlying parameter could lead to optimistic estimates of significance, because we are explicitly minimizing the residual. To overcome this, we propose a nonparametric permutation test for significance. A permutation test repeatedly permutes the values of the explanatory variables, systematically constructing a histogram of the resulting test statistics. The histogram is thus an estimation of the distribution of the test statistic under the null hypothesis that the explanatory variable has no relationship to the dependent variable. Given data (\mathbf{z}_i, t_i) , we generate the k th permuted data set as $(\mathbf{z}_i, t_{\pi_k(i)})$, where π_k is a permutation of $1, \dots, n$. For each permutation we compute a test statistic T_k using Equation 4.6. Then comparing our unpermuted test statistic T to the distribution of T_k , we can compute the p -value as the percentage of T_k that are greater than T . In the case of regression-optimized correspondences, described in Section 3.4, the correspondence optimization is performed on each permutation separately, and thus the results of the permutation test are not biased by the correspondence method.

To illustrate the regression analysis, this section concludes with an experiment on a several sets of 40 synthetically generated tori, parameterized by the small radius r and the large radius R . The values for the shape parameters were chosen as independent functions of a uniformly distributed explanatory variable t . The definition of R^2 in the test statistic described in the preceding section (4.4 is extended to include the two independent variables for this experiment:

$$R^2 = 1 - \frac{(SS_{err})_r + (SS_{err})_R}{(SS_{tot})_r + (SS_{tot})_R}. \quad (4.10)$$

Two sets of time-dependent shapes with p -values $\{0.01, 0.1\}$ were generated in order to examine the performance of the system with and without significant correlation. To construct these example data sets, a known value for the statistic T was used to generate a target R^2 . The values of r and R were chosen so that R^2 for the generated set was approximately equal to the target R^2 for that experiment. Along with explicit correspondences generated from the standard torus parameterization, we optimized 256 correspondences using the PBM regression method from Section. 3.4, and also standard PBM optimization without regression for comparison. An analysis of the resulting models showed that both sets of correspondences exhibited two, essentially pure modes of variation corresponding empirically to variation in r and R .

For the analysis of the test of statistical significance of the correlations, we constructed a histogram of T -statistic values using 1000 permutations of the explanatory variable t . Test statistics were computed using the two dominant PCA modes as the r and R values in 4.10. The p -values were then estimated from the permutation test histogram. Table 4.1 shows the results of the two permutation tests for the explicit correspondences and correspondences generated using the proposed methods. A comparison of the parametric p -value with the p -values obtained by the permutation tests confirms that the proposed methods preserve the relationship between the explanatory variable and the dependent variables. The correspondence-based approaches, particularly with the regression model, show greater significance than the parametric case. This might be an inherent property of the statistic, or it could be an artifact due to the limited number of example datasets

and the limited number of permutations. Future work is warranted to investigate these questions, including simulations with more datasets, more permutations, and a bootstrapping procedure to analyze variability of the p -values computed by the various methods.

4.6 Analysis of Nonshape Components: Pose and Scale

In Section 3.4, the shape of an object was defined as all geometric information distinct from translation, orientation, and overall size. This is the common definition introduced by Kendall, Bookstein, and other pioneers in shape analysis [65, 11, 45, 34]. Typically, however, a biomedical study is often also interested in the nonshape components that describe geometric variability in an object. Thus, it is important to establish a definition of object size and briefly discuss how the size and the orientation, or pose, of biological structures can be examined.

For this dissertation, the size of an object is given by the normalizing scale computed by Generalized Procrustes Analysis (GPA) and the pose of the object by the registration transforms computed by GPA. These are natural definitions that are compatible with the proposed use of the Procrustes mean and distances among Procrustes-aligned correspondences (see Section 4.1). Procrustes scale for a given dM -dimensional shape vector z^k is determined by an iterative process that alternates groupwise least-squares, rigid registrations of all N correspondence configurations, with a scale normalization step determined by

$$s_k = \left(\frac{\sum_{i=1}^N \|\hat{\mathbf{z}}_i\|^2}{\|\hat{\mathbf{z}}_k\|} \right)^{\frac{1}{2}}, \quad (4.11)$$

where $\hat{\mathbf{z}}_i$ represent the partially-optimized configurations. After convergence of the GPA algorithm, the final scale normalization represents Procrustes scale and the final rigid registration transforms represent the pose parameters. For a more complete description of Generalized Procrustes Analysis, the reader is referred to the original papers [9, 47] and the many subsequent review works [45, 34].

Hypothesis testing on Procrustes scale parameters is done using a standard, two-tailed parametric t -test. The group differences in relative position, which are the translational components of the rigid Procrustes alignment parameters, can

Table 4.1. Results of permutation tests (1000 permutations)

p -value (theory)	p -value (parametric)	Correspondence Type		
		Explicit	Min. Entropy	Regression-based
0.01	0.011	0.011	0.007	0.004
0.1	0.095	0.095	0.067	0.066

be analyzed with a parametric Hotelling T^2 test. The analysis of rotational pose parameters requires a general, nonparametric hypothesis test for metric spaces [49], which relies only on pairwise distances between the data and geodesic distances in the rotation group.

In a multiobject complex setting (Section 3.2), where each of the N shape samples consists of $K > 1$ structures, we must decide the level of granularity at which to align shapes in order to analyze pose and scale. Previous work [46, 42] employs a hierarchical strategy, with a *global* coordinate frame for the entire complex, followed by a set of *local* coordinate frames for each object. The global frame is established by alignment of the entire complex, resulting in N sets of global pose parameters. Remaining pose discrepancies among the individual objects constitute the local coordinate frames, which are determined through GPA on the set of N individual structures within the complex, resulting in a set of K local pose parameters for each of the N complexes.

Finally, it should be noted that the role size and pose play in shape analysis may vary with the study design and the data under consideration. To be consistent with established definitions of shape, distinct from pose and orientation, I have proposed an iterative application of GPA during the PBM optimization process (Section 2.4). The PBM algorithm can, however, also accommodate a definition of shape that includes scale. This is accomplished in one of two ways: either by omitting the GPA entirely during the optimization or by re-introducing the Procrustes scale into the Procrustes-space correspondence positions after optimization with GPA. As a general rule, I propose the latter approach because it produces correspondences that are more interpretable with respect to established definition of shape.

CHAPTER 5

APPLICATIONS

This chapter presents applications of the particle-based modeling (PBM) algorithm, its extensions, and the statistical analysis tools that are developed in the preceding chapters. The methods are applied to six different studies in neuroanatomy and genetics. These studies represent validation of the specific contributions of the dissertation, which are to develop a more robust and unified framework for addressing the wide variety of challenging shape analysis problems in biomedicine (Section 1.5). Using the PBM tools, the studies in this chapter provide new scientific insights into current biomedical problems and suggest new hypotheses for further investigation.

The application areas in this chapter are neurobiology, with a focus on clinical psychiatry, and genetic phenotyping of mouse models of human disease and development. As discussed in Section 1.3, these two fields of basic research increasingly rely on statistical models of shape to characterize normal shape variation and to quantify shape differences between populations. The neurobiology studies presented here focus on basic insights into the process of neurodevelopment and investigate the possibility of new diagnostic measures of individuals against normative models. The mouse phenotyping experiments described in this chapter are among the first applications of correspondence models to phenotyping, for which the state-of-the-art is still empirical evaluation and landmarking methods. The phenotyping experiments are designed to investigate gene function in the normal patterning and development of skeletal structures and provide insights into normal human growth and development, with implications for important questions in anthropology.

Neurobiology and phenotyping each pose special challenges for shape analysis, which are among the topics addressed in this chapter and the motivation for much of the work in the preceding chapters. Brain structures often exhibit subtle variation in shape and require methods with sufficient statistical power to examine differences among populations. By contrast, anatomical features in genetically modified mice can change dramatically when mutated. Features are also typically irregularly shaped. Correspondence in mutated specimens is therefore even more difficult to assign and the statistical challenge is how to detect subtle group differences in the presence of such large anatomical variation. As discussed in the preceding chapters, many problems in biomedicine also require assessment of the *covariation* in shape among discrete structures. Some brain structures, such as the cortical surface, exhibit a great deal of geometric variability among normal subjects and require additional information to reliably assign correspondence.

In light of the variety of challenges for shape analysis in biomedicine, the results in this chapter are a demonstration of the efficacy of the PBM algorithm in ad-

addressing diverse shape analysis problems with a uniform statistical approach and a minimum of parameter adjustment and preprocessing considerations. For all of the results presented in this chapter, I have followed the same shape analysis pipeline procedures described in Section 2.4 to prepare and process the segmentation data, with any modifications in the the PBM optimization given by the relevant algorithm extensions described in Chapter 3. Only the curvature-adaptivity parameters required empirical adjustment for different datasets; these parameters were set according to the criteria outlined in Section 2.4. No other free parameters are required for the PBM correspondence optimization.

A few notes on the software for statistical analysis and visualization are warranted for the presentation to follow. For statistical analysis, I have used the freely available GNU R project software, which includes implementations of standard methods for univariate and multivariate hypothesis testing, estimation of linear model parameters, and the analysis of model variance [90]. To reconstruct surfaces from correspondence positions, I have used the method of Hoppe, et al. [54] for surface reconstruction from a collection of unorganized points, which operates directly on the set of correspondence point position, and is implemented in the open source Visualization Toolkit (VTK) [93]. I have also used VTK to construct most of the visualizations presented in this chapter.

The remainder of this chapter is organized as follows. The first section presents a group comparison of the caudate between populations of normal and autistic children, with age as a covariate. The second section describes a shape analysis of a complex of ten subcortical brain structures and a groupwise comparison between autism and normal controls. The third section presents a test of the effectiveness of the PBM algorithm for establishing correspondence among cortical surfaces, using sulcal depth to drive the correspondence. The fourth experiment is an examination of the change in pediatric head shape with age. The chapter concludes with two mouse phenotyping studies: The first study examines the patterning of forelimb bones in a *Hoxd11*-knockout mouse. The second mouse phenotyping study examines shape changes in the cranial base of a *Pax7* mutant mouse.

5.1 Group Comparison of the Caudate in Autism

5.1.1 Introduction

In vivo imaging of brain anatomy, particularly magnetic resonance imaging (MRI), has made shape analysis a feasible tool for the study of many problems in neuroanatomy. In clinical psychiatry, one important new application for shape analysis is the study of autism and its relationship to the morphology of subcortical brain structures. The study presented in this section focuses specifically on the shape of the caudate. The overall size of the caudate has previously been observed to correlate with autism diagnoses [20, 72]. It is therefore reasonable to hypothesize that significant *shape* differences might also be observed. If significant relationships can be observed among brain structure shape and autism diagnosis,

they may suggest new diagnostic procedures and standards for quantification of disease progression.

The study in this section is a preliminary examination of caudate shape and its relationship to autism diagnosis. This research was conducted in collaboration with Martin Styner and Heather Hazlett from the University of North Carolina (UNC), Chapel Hill. Using the PBM correspondence algorithm, I examine brain structure data taken from an ongoing study in pediatric autism that is currently underway at UNC [51]. The results of this study provide some preliminary evidence for patterns of correlation with shape and suggest that further, more exhaustive investigations are warranted. The caudate study also provides a nice illustration of the application of the PBM algorithm and statistical methods to a realistically complex clinical study design.

5.1.2 Methods

The data for this study consist of segmentations of the left and right caudates from MRI brain scans of autistic subjects (AUT), typically-developing controls (TC), developmentally-delayed subjects (DD), and pervasive developmentally delayed subjects (PDD). The latter two categories are present in the study to control for the effects of developmental disorders unrelated to autism. The segmentations were done by trained experts using semi-automated procedures [46]. The UNC data includes subjects at both two and four years of age and includes multiple MRI scans for some subjects that were acquired at different ages. For this study, where multiple MRI scans exist for a chosen subject, only one was selected, and the choice of scans was made in an effort to balance the overall distribution of subject age.

A total of 51 subjects were selected for the study. Subject age is clustered around two timepoints with means $T_1 = 33.55$ months and $T_2 = 45.66$ months, with standard deviations $\sigma_1 = 5.86$ and $\sigma_2 = 5.62$, respectively. The four group classifications are pooled into two categories for the statistical analysis. The first group (TC+DD) consists of 19 typical controls and 4 developmentally-delayed subjects. The second group (AUT+PDD) consists of 22 autism subjects and 6 pervasive developmentally-delayed subjects. In the TC+DD group, there are 10 subjects clustered around time T_1 and 9 subjects clustered around time T_2 . In the AUT+PDD group, there are 10 subjects clustered around time T_1 and 12 subjects clustered around T_2 .

Correspondence models of both the right and left caudate structures were computed using the PBM algorithm to optimize 512 correspondence points per shape simultaneously on all $N = 51$ subjects, without knowledge of group classification. The optimization procedure employed a curvature-adaptive sampling (Section 2.3.3) and iterative alignment and scaling by the generalized Procrustes analysis.

5.1.3 Results and Discussion

Hypothesis testing was done using a MANCOVA analysis (Section 4.4) of the top 5 principal component loadings from a PCA of the correspondence data. The left caudate and the right caudate models were each analyzed separately. The number of components was chosen using parallel analysis, as described in Section 4.3.2. The

null hypothesis for the test is that caudate shapes for the AUT+PDD and TC+DD groups are drawn from the same distribution. Including age as an explanatory variable, the p -values for the group \times age interaction term of the MANCOVA models are $p = 0.076$ for the right caudate and $p = 0.445$ for the left caudate. Subsequent analysis of the MANCOVA models without the interaction terms gives $F = 2.64$ and $p = 0.041$ for the left caudate and $F = 1.37$ and $p = 0.260$ for the right caudate. Thus, there is evidence at the 5% significance level for group differences in shape in the left caudate, but there is no evidence for group difference in the right caudate. The relatively low p -value on the interaction term, however, suggests that there may be some interaction between age and group in the right caudate.

In order to gain further insight into the possible interaction between group and age in the right caudate, I performed a second analysis of the data for this structure at each time point T_1 and T_2 separately. Using the same PCA component loadings as above, MANCOVA analysis at T_1 indicates no significant interaction with group ($p = 0.956$) and a significant difference in the shape components, with $F = 3.373$ and $p = 0.0358$. At T_2 there is also no evidence for group interaction ($p = 0.856$), but no evidence for significant differences in group shape, with $F = 0.250$ and $p = 0.933$.

Figure 5.1 summarizes the hypothesized relationship between shape, age, and group in the right caudate. The figure shows a *cartoon* depiction of the sample distributions in the five-dimensional (5-D) shape space. The actual regression in the 5-D space would be difficult to visualize, and the cartoon regression lines in the figure are used instead to illustrate the main hypotheses resulting from the MANCOVA analysis. The group levels in the analysis are $g = 0$ for TC+DD and $g = 1$ for AUT+PDD, thus the regression line for the AUT+PDD level is shown above in red and the regression line for TC+DD is shown below in blue. The samples are clustered around time points T_1 and T_2 . A visualization of the mean group difference at T_1 is shown on the left. Group difference at T_2 is shown on the right. The group difference is computed as the Euclidean difference between shapes in the 5-D space along the true regression at for the average times T_1 and T_2 and visualized as arrow glyphs in the normal direction to the mean TC+DD surface. From the figure, it is clear that the group differences are larger at T_1 than T_2 , which may account for the better group discrimination at T_1 .

Figure 5.2 illustrates the change in shape with respect to age *within* the two groups. The caudate shapes illustrate the change in shape from $T_1 = 33.55$ to $T_2 = 45.66$ along the each of the TC+DD (left image) and AUT+PDD (right image) regressions. In both cases, the change in shape from T_1 to T_2 is in a similar direction in shape space, but the magnitude of the change is much larger in the TC+DD group. This analysis suggests the hypothesis that the right caudate in the AUT+PDD group may develop faster than the TC+DD, with the group differences at T_1 in part due to the difference in developmental rates.

A similar analysis to that given for the right caudate in Figures 5.1–5.2 is illustrated for the left caudate in Figures 5.3–5.4. In contrast to the right caudate, MANCOVA of the left caudate indicates significant group difference when both T_1 and T_2 are considered together, but no significant differences when analyzed separately at T_1 ($p = 0.122$) and T_2 ($p = 0.723$). This suggests that there may be systematic group differences in shape regardless of age, but that a larger sample

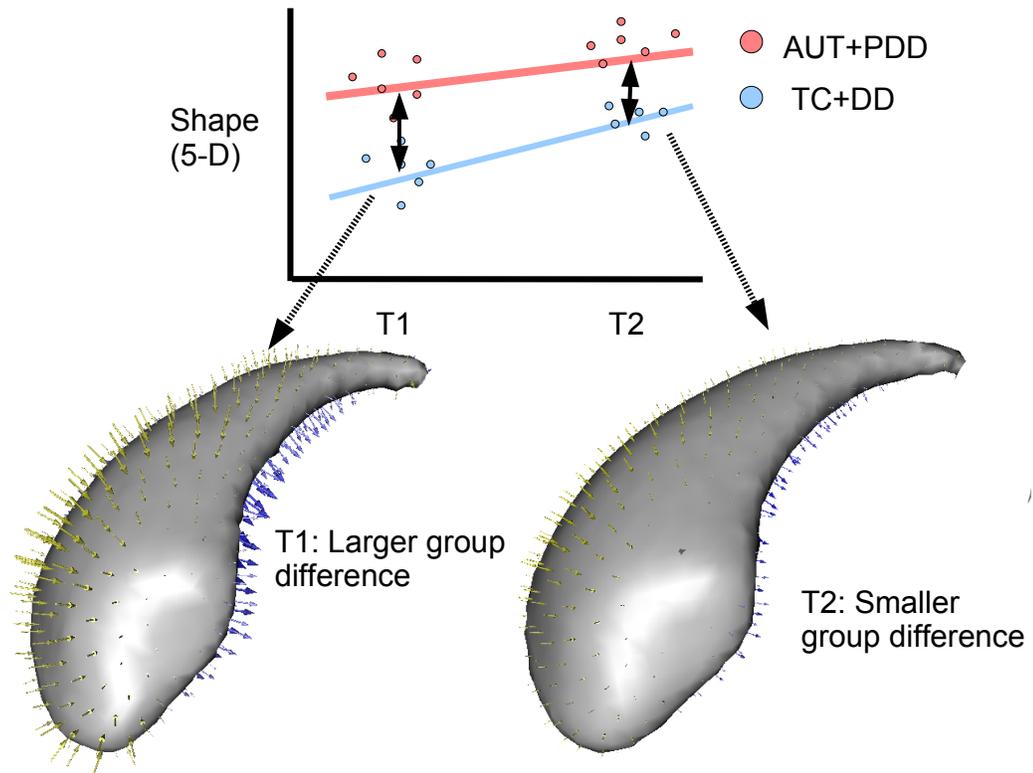


Figure 5.1. Hypothesized relationships among shape, age, and group in the right caudate.

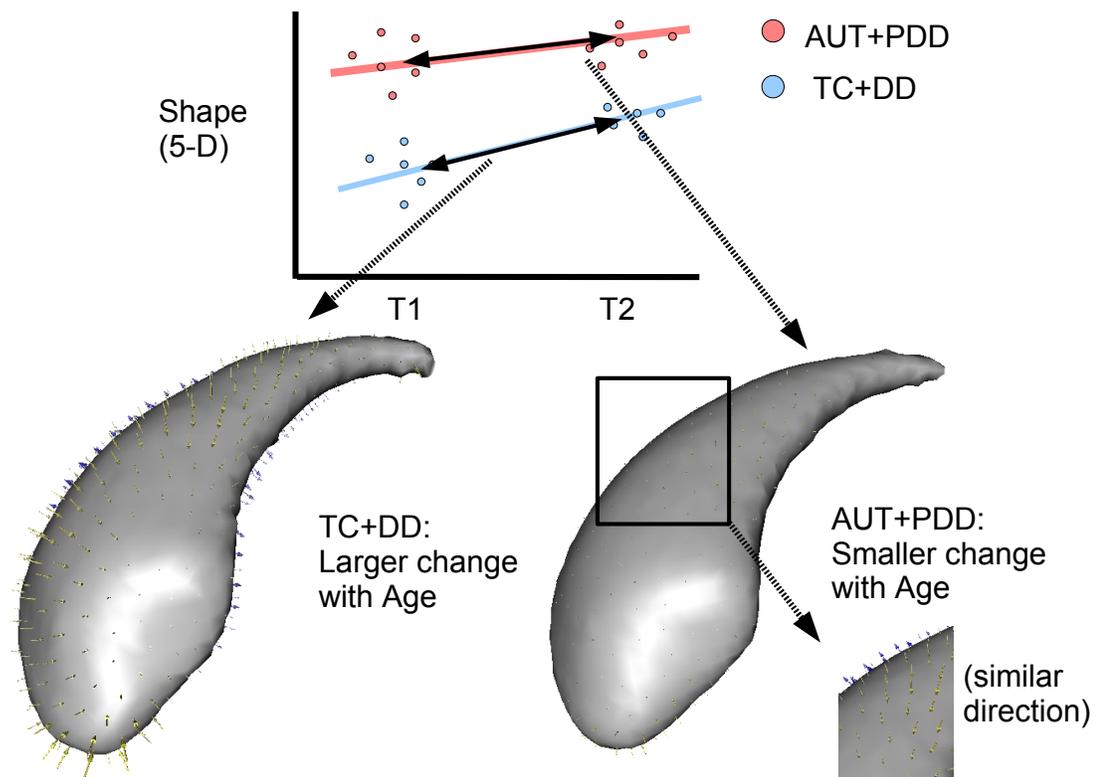


Figure 5.2. Shape change with respect to age within TC+DD group and the AUT+PDD group for the right caudate.

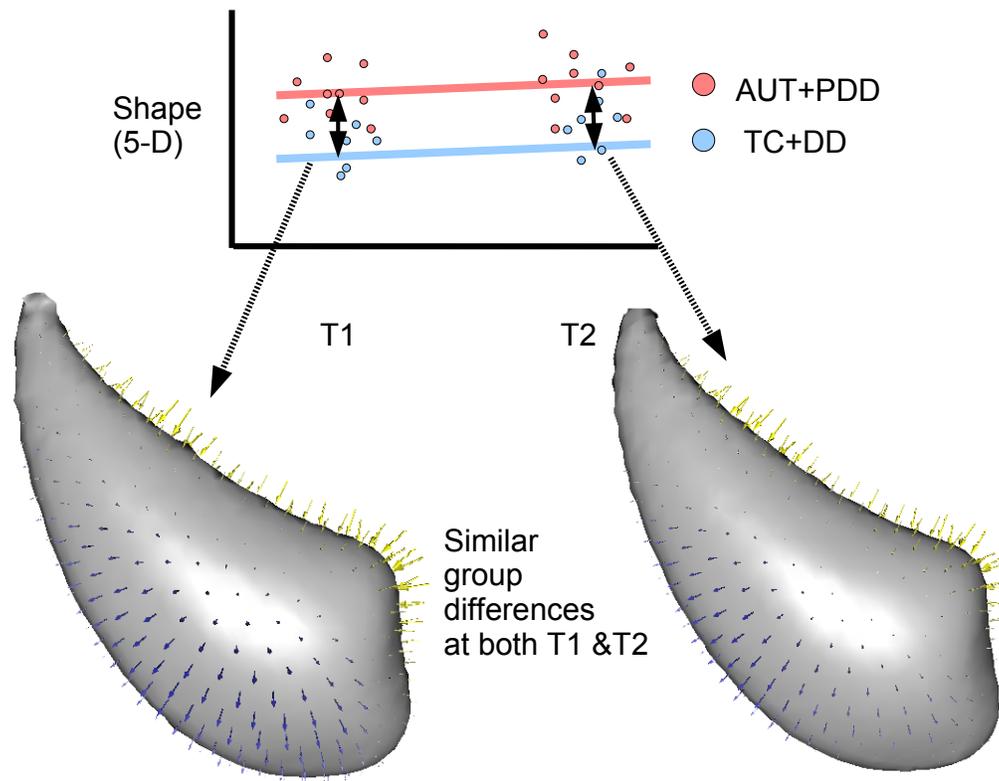


Figure 5.3. Hypothesized relationships among shape, age, and group in the left caudate.

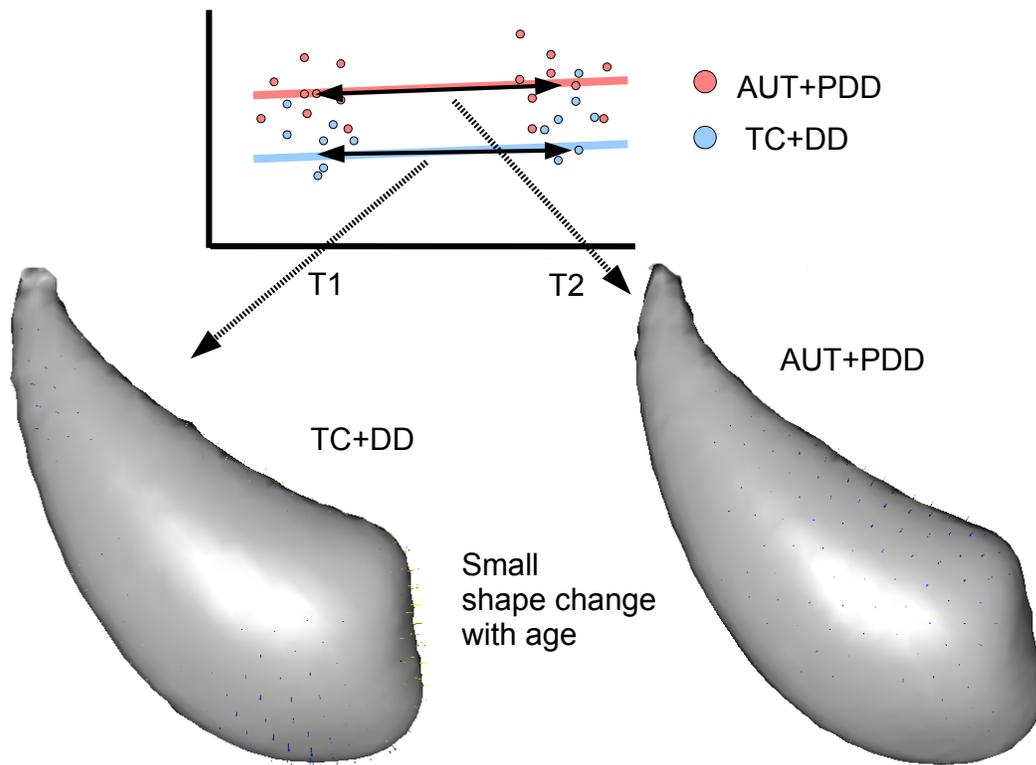


Figure 5.4. Shape change with respect to age within the TC+DD group and the AUT+PDD group for the left caudate.

size than is available at any single timepoint is required for sufficient statistical power to reveal those differences. Note that Figure 5.4 also shows how there is only a relatively small change in shape with age in both groups.

In conclusion, the PBM analysis of the caudate reveals group shape differences between autism and control groups in the left caudate, regardless of subject age. A more detailed analysis of the right caudate suggests that there may be significant differences in younger subjects that decrease with age. These results represent only a preliminary investigation and are not sufficient to draw clinical conclusions. They do, however, suggest interesting new hypotheses for investigation. Further research on this data that includes a much larger sample size is currently underway.

5.2 Multiobject Analysis of 10 Subcortical Structures in Autism

5.2.1 Introduction

This section presents results from a PBM analysis of a multiobject shape complex. Recall from Section 3.2 that a *multiobject complex* is defined as a set of solid shapes, each representing a single, connected biological structure. The shape complexes for this study are segmentations of 10 subcortical brain structures, taken from the clinical study of pediatric autism at the University of North Carolina, Chapel Hill that was introduced in the previous section [51]. For autism, the joint analysis of multiobject complexes is of special interest because the disorder is thought to represent a confluence of several underlying abnormalities and thus it potentially impacts the *relationships* between brain structures[20]. The PBM analysis presented in this section finds significant group differences in shape between normal and patient populations that have not been seen previously in the UNC data and demonstrates how an optimized joint model can yield results with a higher statistical power than a model constructed by simply optimizing the correspondences on individual structures. These results are a proof-of-concept example of the applicability of the multiobject PBM extension (Section 3.2) to a non-trivial and compelling clinical application.

5.2.2 Methods

The data in this study consist of the multiobject segmentations taken from the UNC autism study, which includes MRI brain scans of autistic subjects and typically-developing controls at time points of 2 and 4 years of age. Using a semi-automated procedure, trained experts produced binary segmentations of 10 subcortical structures, which consist of the right and left caudate, globus pallidus, amygdala, putamen, and hippocampus [46]. For this analysis, we used 10, 2-year old male typical normal (TC) controls and chose 15 age-matched autism (AUT) subjects for comparison. Multiobject correspondences were computed from the segmentations as described in Section 3.2 to produce a combined model of both groups. We sampled each complex of segmentations with 10,240 correspondence points, using 1024 particles per structure. For comparison, we also computed point-

correspondence models for each of the 10 structures separately and concatenated their correspondences together to form a *marginally optimized* joint model.

5.2.3 Results and Discussion

Scale and pose were calculated for structures in the complex as described in Section 4.6. A hypothesis test indicates significant group differences in scale only for the right and left amygdala, with p -values of 0.0017 and 0.018, respectively. These results are consistent with previous observations in the literature of volume differences in structures from autistic and normal brains (e.g., [46]), although the analysis does not indicate significant scale differences for other structures that were reported in that work. Hotelling T^2 tests on relative structure positions do not suggest any differences between the groups, with $p > 0.05$ for all structures. Similarly, group differences in pose are not indicated by the statistical pose analysis, with $p > 0.05$ for all structures. This result for pose is also consistent with results given in [46], who similarly did not observe significant pose differences in this data.

The hypothesis test method outlined in Section 4.3.2 results in a highly significant p -value of 0.0087, with 8 PCA modes chosen by parallel analysis. This result is the first evidence shown for these data for group differences in shape alone. Gorcowski [46] reports group differences when scale is included with shape, but reports insignificant shape discrimination between groups when the shapes are normalized to the same size, as in this study. Parallel analysis of the marginally-optimized model indicates that the first six modes should be used, which gives a p -value of 0.0480. While the test still suggests group differences at the 5% significance level, note that the result is an order of magnitude lower in statistical power.

To illustrate the morphological differences that are driving the shape result, the direction and magnitude of the linear discriminant vector in the Hotelling T^2 test (\mathbf{w} in Equation 4.4) is visualized by the colormap in Figure 5.5. The length in the surface normal direction of each of the point-wise discriminant vector components is given by the colormap. Yellow indicates a negative (inward) direction and blue indicates a positive (outward) direction. The right amygdala for the normals is shown in Figure 5.6 as a more detailed example, with the vectors depicted as arrows. This type of visualization gives an indication of the group differences in shape that the hypothesis test statistic identified as the most effective for group discrimination. Note the clear trend towards a shortening of the anterior end of the amygdala in the autistic versus the normal population.

Previous work in shape analysis has suggested hypothesis testing on individual correspondence point positions instead of working in the full dimensional shape space [98]. For comparison, I ran statistical tests at every correspondence point location, using an open-source implementation of the nonparametric Hotelling T^2 method described in [98], with 20,000 permutations among groups and an FDR bound set to 5%. Uncorrected p -values show widespread differences, but no significance remains after FDR correction. This result is in contrast to the global shape result and underscores one of the difficulties with point-based shape analysis at a local feature scale: the unavoidable reduction in statistical power due to the necessary correction for multiple comparisons (see also Section 4.3.1).

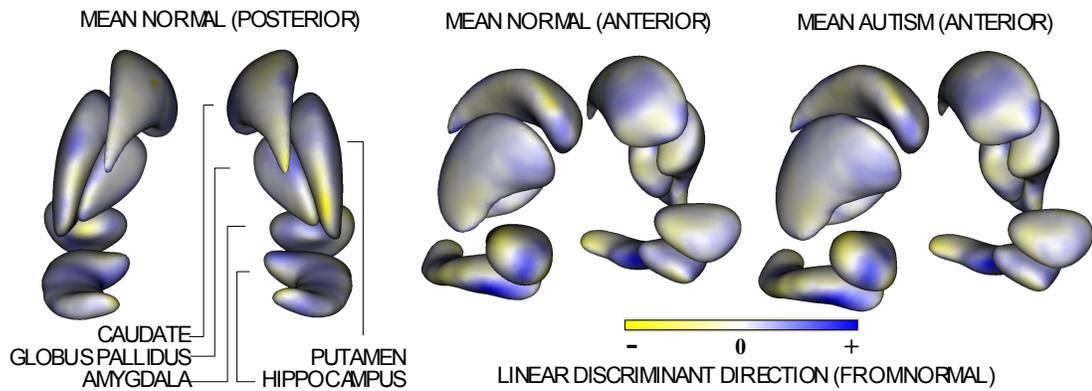


Figure 5.5. Mean brain structure complexes with average pose. Colormap indicates the magnitude and direction of the linear discriminant.

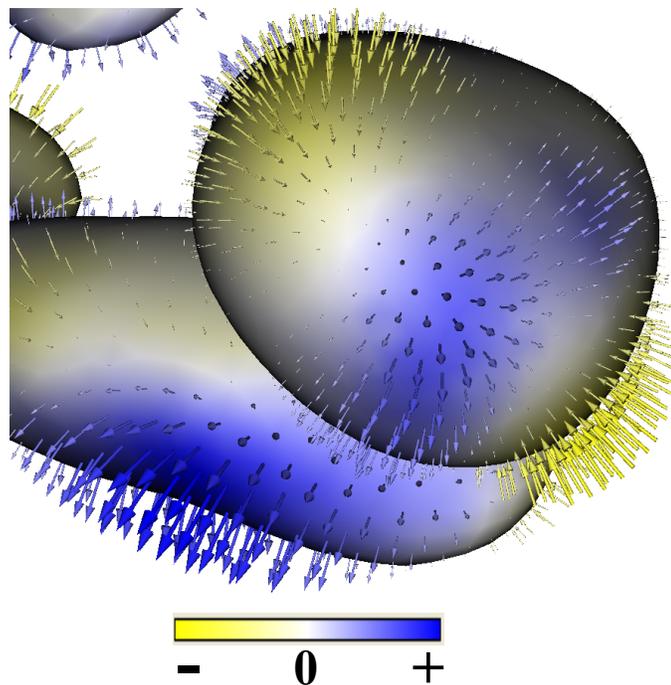


Figure 5.6. Detail of the right amygdala from Figure 5.5, with the linear discriminant visualized as vectors.

Figure 5.5 shows the mean shape surfaces for the normal and autistic groups, as reconstructed from the Euclidean averages of the correspondence points. Each structure is displayed in its mean orientation, position, and scale in the global coordinate frame. The average orientation for each structure was computed using methods for averaging in curved spaces [39]. I used the arithmetic mean of position and the geometric mean of scale. Mean pose differences between the two groups appear small, as might be expected from their statistical analyses.

In summary, our results suggest that the multiobject PBM algorithm and the proposed statistical analysis framework can effectively analyze group differences in the autism data that have not been seen with other methods. Furthermore, greater statistical power is obtained in this case by optimization in the joint space than by optimization in the marginal space. This analysis, however, is only a proof-of-concept example of how the particle method may be applied to multiobject data. A more rigorous study, which remains for future work, is required in order to draw clinical conclusions from the autism data.

5.3 Cortical Thickness Analysis with Correspondence Based on Sulcal Depth

5.3.1 Introduction

This section explores the feasibility of constructing correspondence models of the human cortical surface using the *generalized* PBM algorithm, which is the extension developed in Section 3.3 for establishing correspondence based on functions of position. The correspondence computation on the cortex is a very challenging problem due to the convoluted geometry of the brain and the high variability of folding patterns across subjects. Correspondence based solely on the surface geometry is therefore generally considered unreliable. The generalized PBM algorithm allows data that are coregistered with the surface geometry to drive the correspondence. Incorporating data such as sulcal depth maps, diffusion tensor imaging connectivity maps, or vascular information from magnetic resonance angiography images, may, in theory, improve the quality of the correspondence.

The experiment described in this section was conducted by Ipek Oguz and Martin Styner at the University of North Carolina, Chapel Hill, with the author as a collaborator. The experiment is a proof-of-concept demonstration of the generalized PBM framework to a difficult correspondence problem. This section presents a summary of their results. In the study, sulcal depth maps are used to drive the cortical correspondence optimization and evaluate the results by comparing the variance of sulcal depth and cortical thickness measures at correspondence positions. The complete study description and results are published in Oguz, et al. [84]. Further, related cortical correspondence work from Oguz and Styner that uses the generalized PBM method and probabilistic fiber connectivity data is published in Oguz, et al. [85].

5.3.2 Methods

The data for this study consist of 9 cortical surface representations, derived from T1 magnetic resonance image (MRI) scans of healthy subjects. Surface representations were constructed from the MRI data using the FreeSurfer system, which provides an entire framework for the segmentation, surface reconstruction, topology correction, inflation and spherical parameterization of the cortex [94, 25, 36]. Correspondences were optimized on the cortical surfaces using the generalized PBM algorithm, with sulcal depth as the functional data.

The cortical surface presents a significant challenge for the PBM algorithm due to the complexity of the folding patterns and their variability across subjects. Because of the tangent plane assumption in the formulation of spatial interactions among particles, good distributions over highly convoluted surfaces are difficult to achieve with the current PBM implementation. In addition, the high variability in geometric information leads to many possible local minima in the optimization. Ideas for extending the PBM algorithm to more complicated surfaces like the cortex are discussed further in Chapter 6. To resolve this issue for the cortical thickness study, a cortex “inflation” technique was used to smooth the convolutions. The particles were then easily distributed on the smoothed surface.

Two additional sets of correspondence positions were constructed for comparison with the correspondences constructed via sulcal depth. The first was constructed using the standard PBM optimization and was therefore based only on spatial position. The second set of correspondences was constructed by the FreeSurfer system, which uses a pair-wise, semirigid alignment to an average parameterized cortex shape, according to the algorithm given in [36]. Because there is no ground truth to assess the quality of the optimized correspondence positions, the results were evaluated based on how well they reduced the local variability of both sulcal depth and cortical thickness. The local variances of sulcal depth and cortical thickness were measured around correspondence positions before and after optimization, with a lower variability among the optimized positions considered to be more desirable.

5.3.3 Results

The mean sample variances for both the cortical thickness and sulcal depth measurements are summarized in Table 5.1. For sulcal depth measurements, the generalized PBM algorithm reduces variance almost 75-fold over initial data and almost 25-fold over the FreeSurfer results. For cortical thickness, the generalized PBM method has considerable improvement over initial data, but a slightly higher average variance than FreeSurfer. An inspection of the distribution of this mean variance over the surface is shown in Figure 5.7. The colormap in the figure shows increasing (normalized) variance from green = 0 to red = 1. The figure reveals that FreeSurfer has a higher variance across the *entire* surface, whereas the PBM method has only a localized high variance around the temporal lobe (which was not perfectly reconstructed by FreeSurfer due to input image noise) and performs much better in other areas of the cortical surface. It should be noted that both the PBM method and FreeSurfer use the sulcal depth information as part of the correspondence optimization process and therefore the sulcal depth evaluation is biased. Also note that the higher degree of remaining variability in cortical thickness

Table 5.1. Mean sample variance for cortical thickness and sulcal depth.

<i>Correspondence Method</i>	<i>Sulcal Depth</i>	<i>Cortical Thickness</i>
<i>Initial Data</i>	0.2276	0.3349
<i>PBM (positional entropy)</i>	0.2196	0.3417
<i>generalized PBM (sulcal depth entropy)</i>	0.0035	0.3108
<i>FreeSurfer</i>	0.0756	0.3034

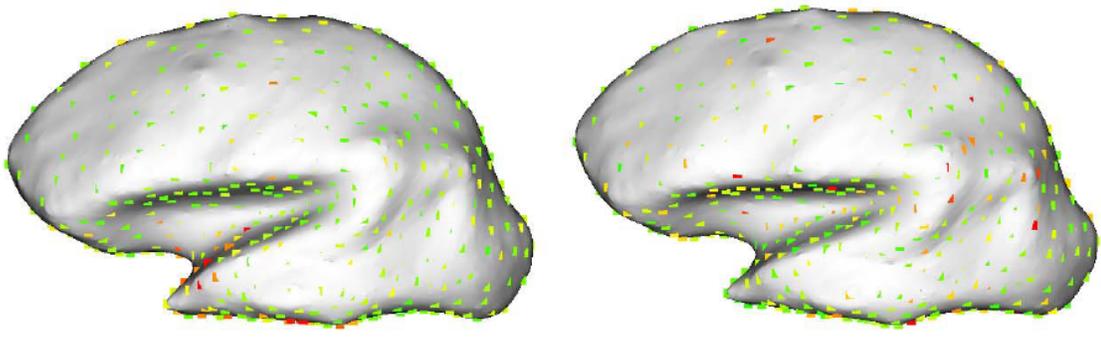


Figure 5.7. Comparison of the variance of cortical thickness at correspondences from the generalized PBM method (left) and to FreeSurfer (right). Figure by Ipek Oguz.

can be largely attributed to intersubject variability, as cortical thickness patterns tends to vary largely among individuals.

5.4 A Regression Model of Pediatric Head Shape with Age

5.4.1 Introduction

The analysis of shape as a function of explanatory variables, or shape *regression*, promises to be an important tool to study the relationship between anatomy and underlying clinical or biological parameters, such as age. In developmental analyses, such as pediatric neurodevelopment, shape regression analysis formulates aggregate models of growth, with variability. These models give not only basic insights into the process of development, but also allow comparisons of individuals against normative models.

Quantitative magnetic resonance imaging (MRI) has significantly advanced our understanding of brain development during childhood and adolescence. Courchesne, et al. [22], for example, describe an imaging study that looks at differences in growth patterns in autism compared to controls in children over four years in age. Data measured in infants from birth to 4 years, however, are mostly *volumetric measurements*, such as intracranial volume and volumes of brain lobes and subcortical structures [69], and there has been little work to date examining the change in *shape* with age in early childhood brain development.

This section describes an application of the PBM shape regression algorithm from Section 3.4 to the problem of correlating head shape with age in children. The work in this section was conducted with Manasi Datar, at the University of Utah. In this study, the PBM regression algorithm is used to model the shape variability in a set of neonatal head shape boundaries extracted from MRI images. The children in this study are between 2 and 5 years of age, which is a time of rapid brain growth and development. Current clinical practice characterizes growth by head circumference, a metric correlated with the overall size of the head. In this work we examine whether overall head shape, independent of size, is a significant indicator of developmental progress and whether the PBM regression algorithm is suitable for examining that relationship. We also test the effectiveness of the open surface modeling extensions from Section 3.1.

5.4.2 Methods

The PBM regression optimization was applied to structural MRI data obtained from clinical studies spanning the age range from neonate to five year olds. The 40 clinical cases include 1.5T, T1-weighted MRI scans with resolutions of $1mm \times 1mm \times 1mm$ and $0.4mm \times 0.4mm \times 3.6mm$. The scans were segmented using a semiautomated procedure to obtain binary volume representations of the head surfaces. These segmentations were preprocessed according to the methods outlined in Section 2.4 to produce distance transform inputs to the PBM optimization process. The PBM regression optimization was run on the segmentations using 500

particles per shape sample. The optimization took roughly 40 minutes to complete on a standard 2GHz desktop machine.

Manually placed landmarks on the bridge of the nose and the openings of the left and right ear canals specify a cutting plane and a pair of spheres, which define the boundary of the surface, as in Section. 3.1. During the optimization, the boundary repels the correspondences in order to restrict the analysis to the cranium, which is the most interesting from a neurological point of view. Shape in the neck, face, and ears is not hypothesized to correspond with brain growth and adding the variability of these structures to the model may reduce statistical power and complicate the analysis. Figure 5.8 is an illustration of the correspondences on one of the head shapes after optimization.

5.4.3 Results and Discussion

Head size, measured in volume or circumference is well known to correlate with age and is confirmed by our own study of the Procrustes scale of the correspondence models. A linear regression plot of Procrustes scale versus \log of age, with $p < 2 \times 10^{-16}$, is shown in Figure 5.9. The correlation is highly significant.

Changes in head shape versus \log of age are shown in Figure 5.10. The figure is a set of reconstructed surfaces from correspondences along the linear regression line in shape space, with shape versus \log of age (see Section 3.4). Note the relative lengthening of the head and the narrowing at the temples with increasing age. These shape changes are consistent with clinical observations that neonatal brain growth proceeds more rapidly in the forebrain. These results tie head shape to age in the paediatric setting. A test for significance of the correlation was done using the method from Section 4.5, with 1000 permutations run in multiple small batches of parallel processes. None of the permutations gave a better correlation than the input data, resulting in an effective p -value of zero. While this result is not conclusive, it does give strong evidence for significance. Future work that includes more permutations is probably warranted to more accurately evaluate the significance. At 40 minutes per optimization, however, increasing numbers of permutations represent a significant amount of processing time.

Results from this study of head shape growth indicate that the PBM regression optimization can be effectively applied to quantitative characterization of the relationship between age and head shape in young children. Expanded analyses thus offer potential for understanding of growth beyond the currently established standard of the head circumference index. Moreover, this type of analysis will generate normative data as a continuous growth model of shape, which can be useful in building optimal MRI head coils for young infants. The continuous shape model could also find use in population studies where two groups are compared with respect to growth trajectory, rather than differences at individual time points.

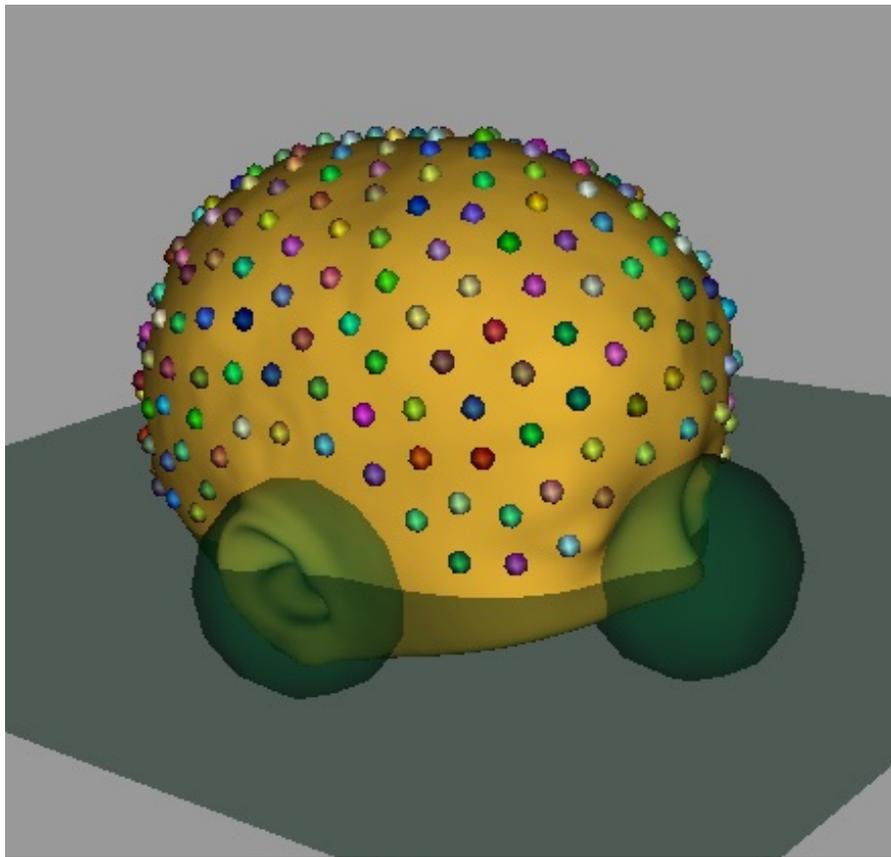


Figure 5.8. An illustration of correspondences on the cranial surface for the neonatal head shape study.

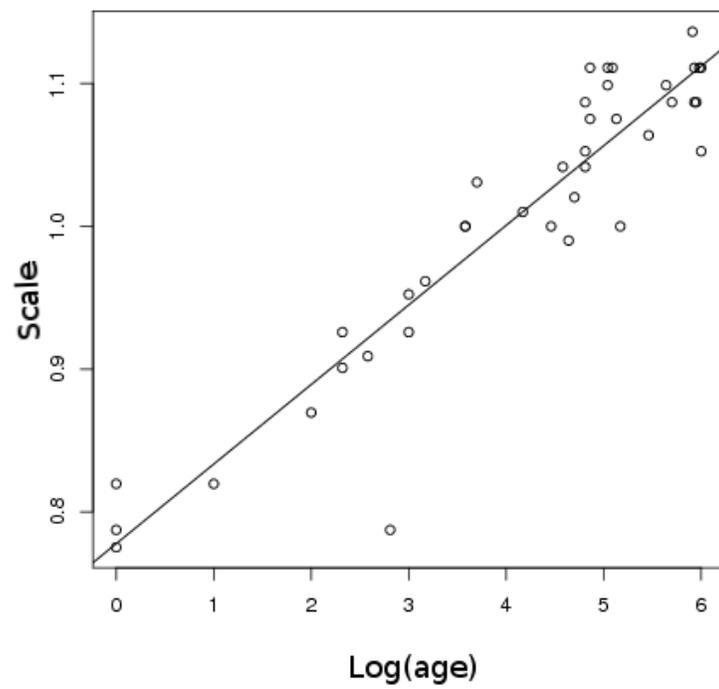


Figure 5.9. Changes in head scale with *log* of age.

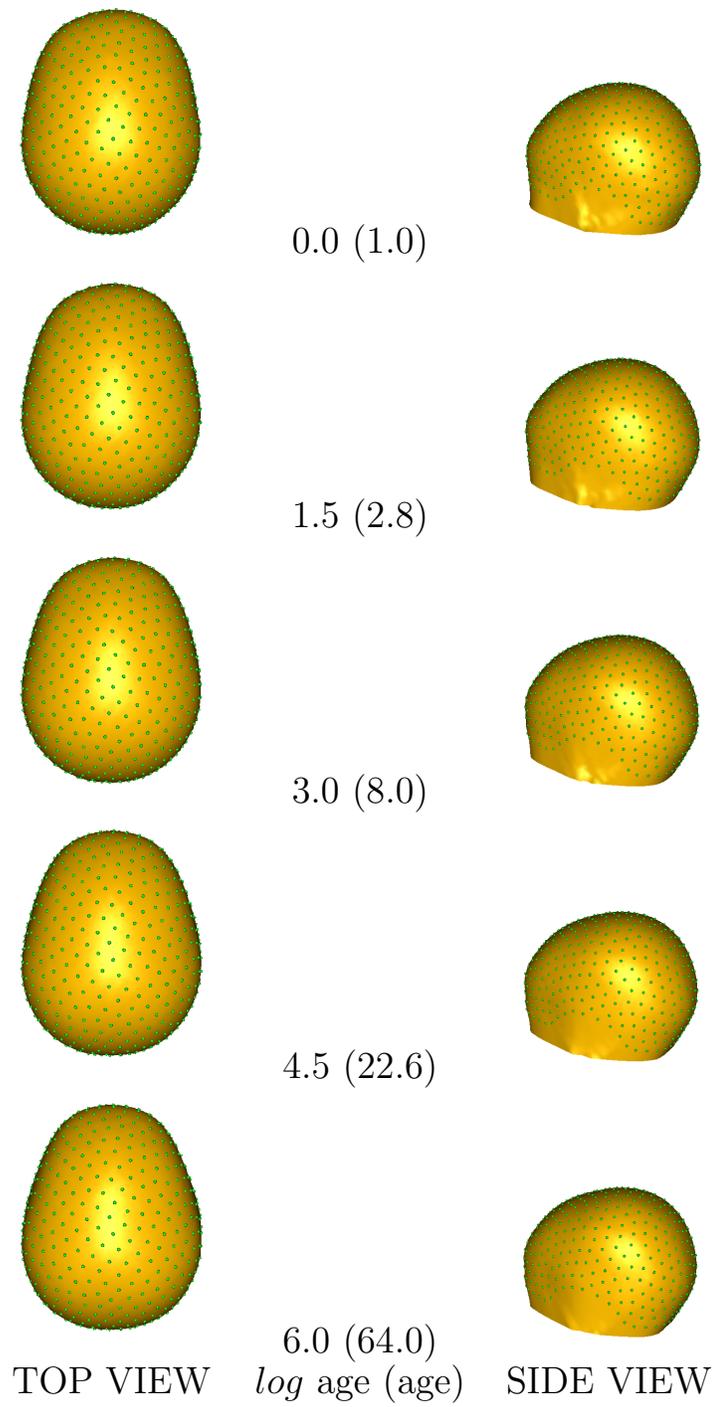


Figure 5.10. Overview of head shape regression: Changes in head shape with *log* age. Age is given in months.

5.5 Phenotyping in *Hoxd11*-deficient Mice

5.5.1 Introduction

In cooperation with researchers from the Mario Capecchi lab at the University of Utah, we applied the PBM shape modeling framework to a phenotypic study of several of the forelimb bones of mice with a targeted disruption of the *Hoxd11* gene. The hypothesis for this study is that the underexpression of *Hoxd11* affects the growth of forelimb bones, resulting in changes in their shape. Research has previously suggested a systematic shortening of several of these bones in mutant mice, but a characterization of overall shape differences is necessary in order to characterize more subtle effects of *Hoxd11* deficiency. In this work, the PBM algorithm is shown to be an effective tool for characterizing those shape differences. This study also is among the first to use high-dimensional, point-based shape modeling and analysis for small animal phenotyping, which is traditionally the domain of landmarking methods and simple, univariate metrics such as length and volume.

The *Hox* complex of genes is known to play an important role in the proper development of the mouse and a better understanding of the function of these genes may offer important insights into the cause of certain human birth defects. Through a series of gene targeting experiments, Boulet, Davis and Capecchi have previously shown that the *Hoxd11* gene in particular is especially important for the normal development and patterning of the appendicular skeleton [31, 13]. In [31], for example, Davis prepared murine newborn pup skeletons with alizarin red stain and measured the lengths of the autopod bones under digital light microscopy, a method of phenotype characterization that is the current standard for phenotyping studies [76]. From these length measures, Davis, et al. hypothesize that mice deficient in the *Hoxd11* gene show significant differences in the average length of specific bones of the mouse forelimb. In our study, we are interested in assessing what shape changes are present that account for these length differences and whether there are more subtle variations that can be seen.

5.5.2 Methodology

Twenty male mice homozygous for targeted disruption of the *Hoxd11* gene were bred from an existing colony of *Hoxd11*^{-/-} female and *Hoxd11*^{+/-} male mice. (Male *Hoxd11*^{-/-} mice are infertile for unknown reasons.) Details on creation of *Hoxd11* mice by construction of a targeting vector and establishing a targeted ES cell line are described in [31]. *Hoxd11*^{-/-} was confirmed by PCR analysis and gel electrophoresis using tail DNA from each of the mice. Twenty *C57BL/6* (wildtype) male mice were used as the control group phenotype. Mice were sacrificed by carbon dioxide asphyxiation at 10 weeks and immobilized in a 50 mL centrifuge tube before scanning in order to minimize motion artifact and provide consistency in mouse position and placement for all scans. The eXplore Locus Small Animal MicroCT Scanner (GE Healthcare), which utilizes a 3500 x 1750 CCD detector and Feldkamp cone-beam reconstruction, was used to perform high resolution, 360 degree volumetric CT of each mouse. All mouse specimens were bred and scanned at the Eccles Institute for Human Genetics at the University of Utah.

Following the breeding and imaging of the mice, three forelimb bones were segmented from the CT volumes by experts: the metacarpal (MC), the first phalange (P1), and the second phalange (P2) of digit 2 of the right forepaw of the mouse. These bones were chosen because they were identified in [31] as exhibiting significantly different average lengths in mutant and normal mice. The right humerus bone was also segmented as a variable to use in controlling for development rate, as indicated by [31]. Several experts in the relevant anatomy performed volumetric segmentations of the bones from the micro-CT images using the Seg3D volume segmentation software [33]. The segmentation procedure consisted of a region-growing segmentation of the entire complex of forepaw bones, followed by a manual delineation of the boundaries between specific bones of interest.

The collection of bone segmentations required an initial alignment into a common coordinate frame for modeling and analysis. The segmentations were first automatically aligned with respect to their centers of mass and the orientation of their first principal eigenvectors. The automatic alignment was followed by a manual adjustment, using a tool constructed in VTK for this purpose that visualizes overlays of surface renderings, allowing the user to adjust their relative orientations. Using this tool, orientation was adjusted around the principal axis until characteristic features were in rough alignment with a reference bone shape. During the correspondence computation phase, this rough alignment was automatically refined with respect to rotation and translation using the Procrustes algorithm, as described in Section 2.4. Combined group models for each of the three bones of interest were then computed by the PBM algorithm, with 1024 correspondence points per shape. Each optimization took approximately two hours on a standard 2GHz desktop machine.

5.5.3 Results and Discussion

For the mouse data, we have a clear biological hypothesis as to the group differences between mutant and normal mouse strains. Specifically, the work of Boulet, Davis, and Capecchi has suggested length differences in the P1, P2, and MC bones of the right forepaw. For this study, I analyzed the PBM correspondence models for those same three bones using the multivariate hypothesis testing procedure from Section 4.3.2. Parallel analysis indicates using 2, 2, and 3 modes for the P1, P2, and MC bones, respectively. Hotelling T^2 tests on the PCA loadings show highly significant p -values in each case ($p \ll 0.01$), indicating significant shape differences probably do exist between the two groups. Figure 5.11 depicts the linear discriminant from the Hotelling T^2 test for the P1 bone. The discriminant is shown on the mean wildtype shape in the direction of the mutant shape. The linear discriminant visualization reveals two clear morphological differences: a reduction in bone “length” and an increase in bone “thickness” in the *Hoxd11*-deficient population. Empirical analyses of the P2 and MC bones each show very similar patterns.

The relationship between dominant shape variable (PC1) and the morphometry-based shape variable (bone length) from the literature [31] is of key interest in this study, because length differences were the main hypotheses coming from the original phenotyping studies of the *Hoxd11* mouse. This relationship can be analyzed

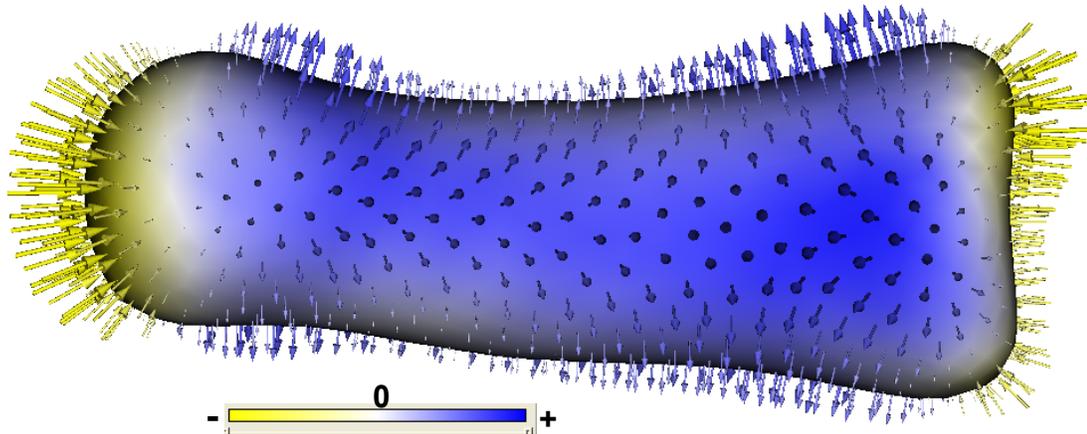


Figure 5.11. LDA visualization of the P1 bone for the *HoxD11* phenotype study

through linear regression of PC1 on bone length and an analysis of the correlation between the two measures. Correlation between bone lengths and the PC1 measures is high for each of the bones and statistically significant $p \ll 0.01$, results that support the hypothesis presented in [31] that reduced bone length is the dominant morphological effect of *Hoxd11*-deficiency. A univariate hypothesis test was done as a further analysis of group differences in PC1. To account for variability in the rates of development of each mouse, we tested group difference with an analysis of covariance (ANCOVA) model that included the length of the humerus bone as a covariate, a bone not hypothesized to be affected by the *Hoxd11* gene. For the univariate hypothesis tests on PC1, we found no significant interaction between group classification and humerus length and significant differences in the group means for all bones, with $p \ll 0.01$ in each case. Figure 5.12, for example, shows this analysis for the P2 bone. Similar results were obtained for the MC and P1 bones.

In conclusion, the *Hoxd11* PBM analysis has quantified and statistically validated one of the major conclusions from empirical studies of the *Hoxd11* gene, as well as revealing a new significant phenotypic effect. These results are significant in that they represent one of the first applications of computationally-derived point-based models to genetic phenotyping. Additionally, in contrast to the more exploratory nature of the neuroanatomy shape studies from previous sections, this study poses a specific biological hypothesis for which the PBM analysis has provided new evidence and further insight. Thus, it represents a strong validation of the usefulness of the PBM algorithm in scientific work.

5.6 Analysis of the Cranial Base in *Pax7* Mutant Mice

5.6.1 Introduction

This study presents the results of a second phenotyping analysis of a genetically engineered mouse. It applies the PBM correspondence method to the analysis of shape differences in the cranial base of newborn wildtype mice and mice genetically deficient for *Pax7* [77]. Specifically, the analysis focuses on the basioccipital bone. The basioccipital bone varies subtly for *Pax7* mutant mice in a way that is analogous to the changes in shape that occur within the hominin fossil record in connection with the adoption of an upright posture [82]. In our analysis, we detected a statistically significant antero-inferior inflection of the posterior edge of the basioccipital, among other shape differences. This study was conducted in collaboration with Lisa Nevell, Charles Keller, and Suresh Prajapati from the Greehey Children’s Cancer Research Institute, University of Texas Health Science Center.

Craniofacial development is of increasing interest to computational biologists and developmental and evolutionary biologists alike [73]. The cranial base is especially interesting because its shape defines physiological parameters characteristic of species taxonomies [56]. The basioccipital bone is part of the midline of the cranial base and forms the anterior and lateral aspects of the foramen magnum, where

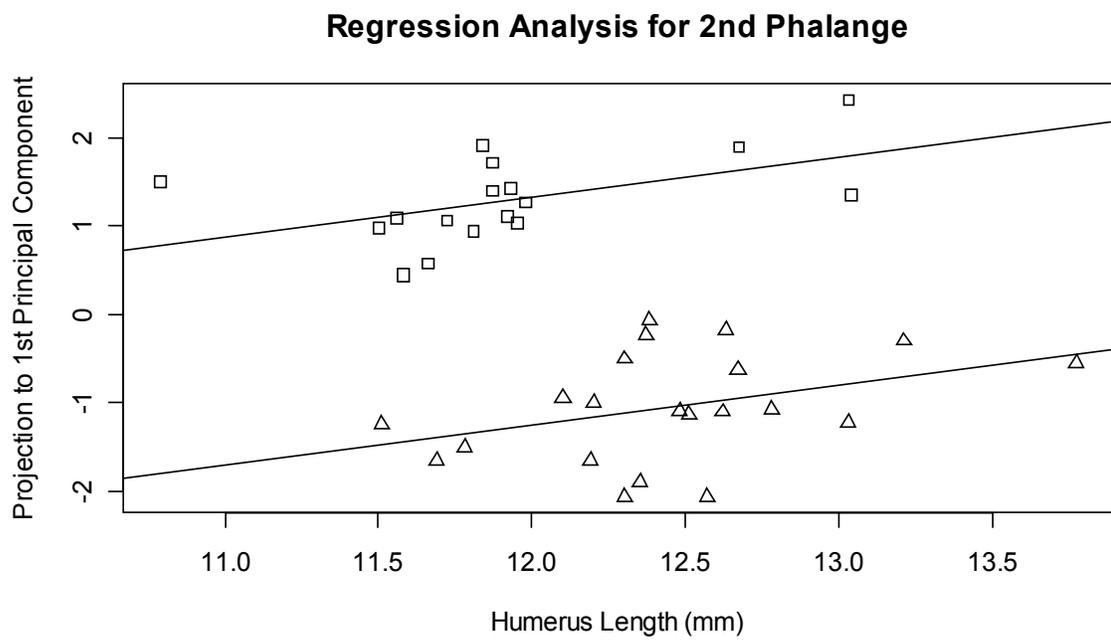


Figure 5.12. ANCOVA for the PC1 measurements on the P2 bone. Wildtypes are shown as squares and mutants as triangles.

the spinal cord exits the cranium. In hominin evolution, the forward (anterior) placement of the foramen magnum is a key difference between upright hominins and primates [82]. Therefore, the methods described here are applicable not only to phenotyping genetically-engineered mice, but also to the study of the evolution of the hominin cranial base.

5.6.2 Methodology

Data for this experiment were generated at the Greehey Children’s Cancer Research Institute and consist of micro computed tomography (micro-CT) scans of mice with a LacZ gene insertion interrupting the native Pax7 gene. Twenty-four mutant mice pups at postnatal day 0 were euthanized and imaged at $27\ \mu\text{m}$ isometric resolution. For comparison, 24 postnatal day 0 wildtype *C57BL/6* pups were also scanned using the same protocol. In the mutant mouse sample set, genotyping revealed 18 mice to be heterozygous for the LacZ mutation *Pax7(LacZ/WT)* (WT=wildtype) and 6 mice as homozygous for the mutation *Pax7(LacZ/LacZ)*.

We performed volumetric segmentations of each the basioccipital bones from the CT image volumes. Segmentations were done by several expert observers using manual delineation. The software used for the segmentations was the freely available open source Seg3D image processing tool [33]. Antialiased distance transforms were produced for input to the PBM algorithm using the methods outlined in Section 2.4.

We computed three correspondence models on the segmented data using the PBM algorithm: a *homozygous* model with the combined data from the wildtype and *Pax7(LacZ/LacZ)*, a *heterozygous* model with the combined data from the *Pax7(LacZ/WT)* and wildtype groups, and a *combined mutant* model with the combined data from the *Pax7(LacZ/LacZ)* and *Pax7(LacZ/WT)* groups. Each model was computed with 1024 correspondence positions and curvature-adaptive sampling. Run times for each optimization using a C++ implementation on standard 2GHz PC desktop hardware were approximately one hour.

For comparison with the PBM analysis, we also conducted a simple morphometric analysis of the basioccipital segmentations using four metrics: the maximum widths in each of the sagittal, coronal, and axial directions and the volume. To compute the maximum widths, we first coregistered the segmentations into the same coordinate frame using the rotation transforms from the generalized Procrustes analysis (without scaling). Following alignment, the width along a desired direction was automatically measured at all two-dimensional coordinates in the remaining two directions and the maximum of these widths recorded. Width was measured by counting the number of voxels inside the segmented object along the desired direction and multiplying by the voxel spacing in that direction. Volume was measured by iterating over the entire image volume, counting the number of voxels inside the segmented object, and scaling that number by the cubic volume of an image voxel.

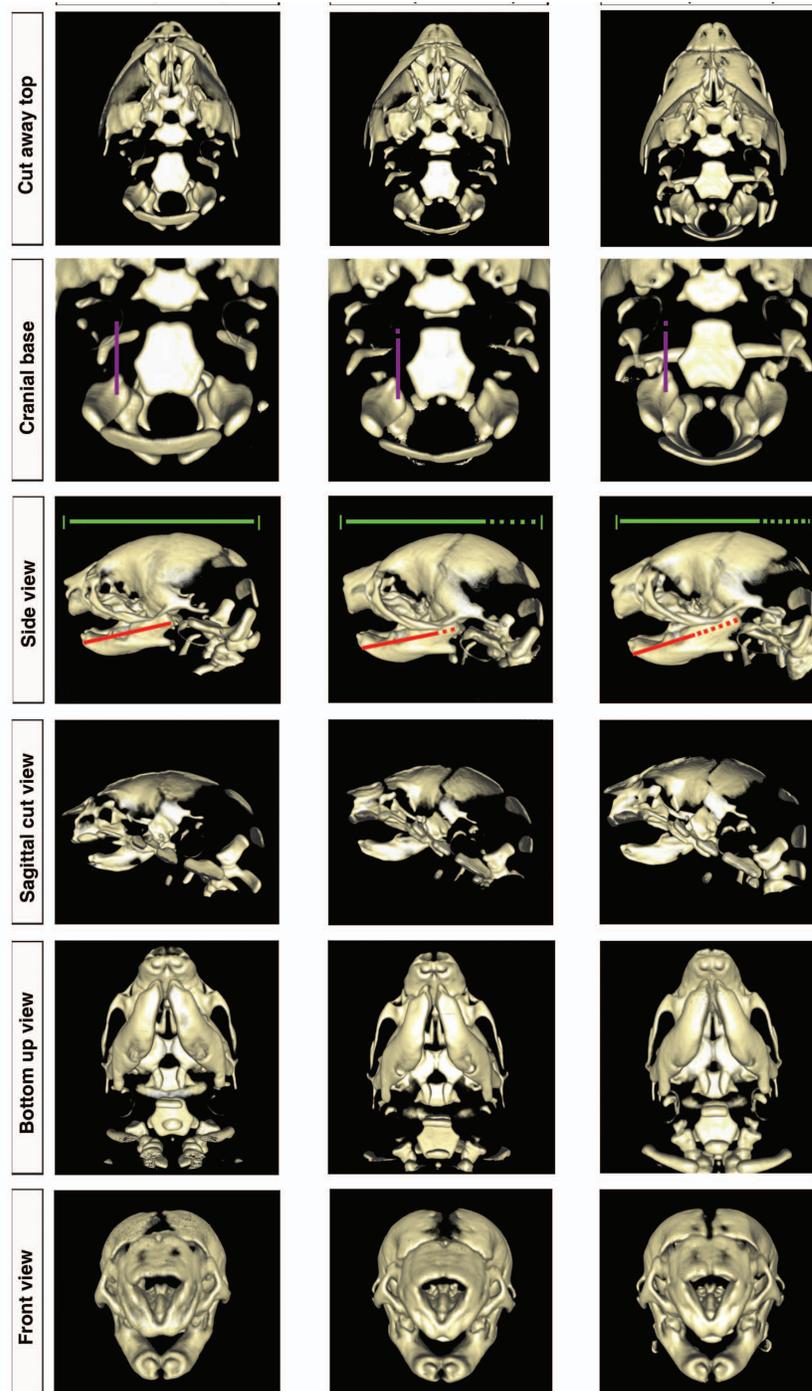


Figure 5.13. Volume renderings of micro-CT scans of wildtype Pax7(WT/WT), homozygous Pax7(LacZ/LacZ), and heterozygous Pax7(LacZ/WT) median specimens. Images created by Suresh Prajapati.

5.6.3 Results and Discussion

Volume rendering visualizations of the micro-CT scans for representative specimens from the wildtype Pax7(WT/WT), homozygous Pax7(LacZ/LacZ) and heterozygous Pax7(LacZ/WT) groups are shown in Figure 5.13. The figure depicts the neonate animal with the median shape for each group, where the median is computed as the L_1 multivariate median of the correspondence positions (Section 4.1). Row 1 depicts the neonatal skullbase, with the temporo-parietal bones cropped away. Row 2 is a closer view of the cranial base and the basioccipital bone. Row 3 compares the overall jaw and skull lengths. Side-view images for the sagittal hemi-sections of each sample are shown in Row 4 and images looking on to lower jaws from bottom are given in Row 5. Row 6 compares craniofacial features of the front side of the skull. These images were created by Suresh Prajapati using the freely available open source ImageVis3D volume rendering software (ImageVis3D, <http://www.sci.utah.edu/cibc/software>).

Key phenotypic differences by inspection in Figure 5.13 include lengthening of the mandible and an increase in the overall anteroposterior length (Row 3). Row 2 indicates gross structural changes are present in the basioccipital. While more overt changes to the mandible and anteroposterior dimension are present, we focus on the basioccipital bone of the cranial base because of the evolutionary interest in the basioccipital in relationship to the placement of the foramen magnum.

Figure 5.14 shows superior views of the group mean shapes, reconstructed from the mean PBM correspondence positions. Note the similarity between the homozygous and heterozygous mean shapes. The mean correspondence positions are shown as red sphere glyphs. Table 5.2 summarizes the results of the multivariate, Hotelling T^2 statistical test for differences in the mean shapes from Figure 5.14, along with the number of PCA loadings determined by parallel analysis. For the homozygous and heterozygous models, the results indicate highly significant shape differences from the wildtype group ($p \ll 0.001$). No significant differences between homozygous and heterozygous groups are indicated. We also conducted Welch two-sample t-tests for mean differences in Procrustes scale for each cohort in the analysis, which gave no indication of any significant differences.

Figure 5.15 shows the linear discriminant vector from the Hotelling T^2 test between homozygous and wildtype cohorts and the heterozygous and wildtype cohorts. In each case, the discriminant vector is shown on the mean wildtype shape, in the direction of the mutant shape. Qualitatively these images give an indication of how the wildtype shapes would change if they were to increase their probability of being part of the mutant cohort; thus, these visualizations represent the local shape differences that are statistically significant. Larger arrows indicate local regions that exhibit more shape change between groups and smaller arrows indicate relatively smaller shape differences. Note that the linear discriminants are very similar for the homozygous and heterozygous cases. In these representations, the differences in basioccipital shape include antero-inferior inflection of the posterior edge (which forms the anterior boundary of the foramen magnum), lateral-superior outgrowth of the lateral prominences, anterior elongation in the midline, and flattening of the posterior $\frac{1}{3}$ inferior surface.

In addition to the multivariate statistical analysis, univariate statistical tests (Welch 2-sample t-test for mean difference) for each of the first five principal com-

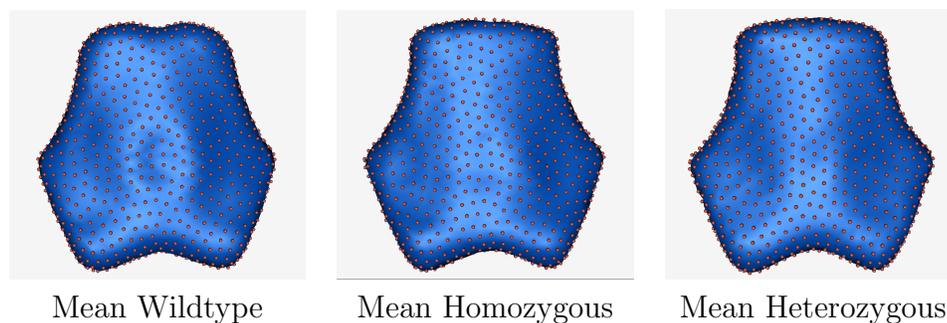
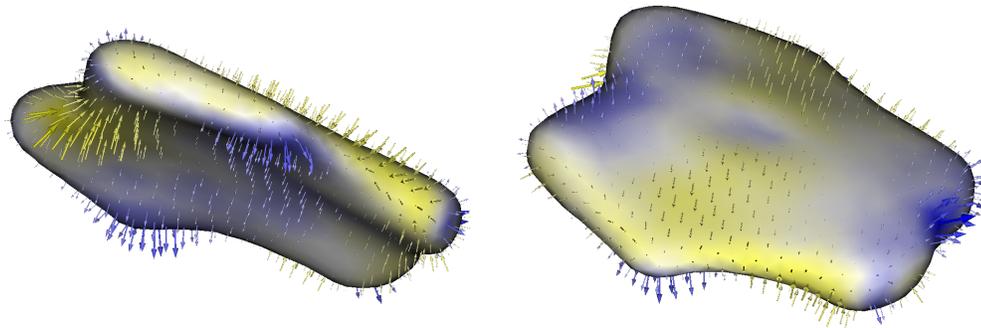


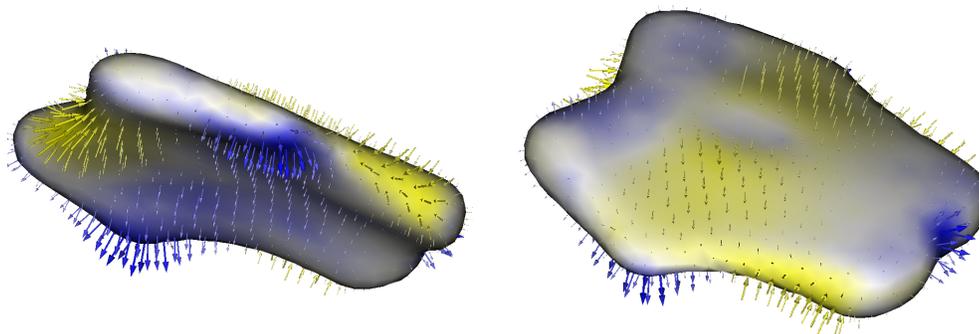
Figure 5.14. A visual comparison of the group mean shapes. Shape surfaces are reconstructed from the Euclidean means of the correspondence positions for each group. Mean correspondence positions are indicated by spheres.

Table 5.2. Multivariate hypothesis test results for group differences in mean shape for PacZ(LacZ/LacZ), PacZ(LacZ/WT), and combined mutant groups.

<i>Correspondence model</i>	<i>Hotelling T^2 test (# PC modes)</i>
Homozygous (LacZ/LacZ vs WT)	$T^2 = 95.33, p \ll 0.001$ (5)
Heterozygous (LacZ/WT vs WT)	$T^2 = 150.16, p \ll 0.001$ (4)
Combined mutant (LacZ/LacZ vs LacZ/WT)	$T^2 = 1.54, p = 0.852$ (4)



TOP: Homozygous vs. wildtype linear discriminant



BOTTOM: Heterozygous vs. wildtype linear discriminant

Figure 5.15. Visualization of the linear discriminant from the Hotelling T^2 test for the homozygous shape model (TOP ROW) and the heterozygous shape model (BOTTOM ROW). Arrows indicate direction from the wildtype distribution to the mutant distribution.

ponent (PC) loadings help quantify and visualize local shape difference. Table 5.3 shows the results of these univariate t-tests for each of the five principal components. Table 5.4 shows the percentage of the total shape variation represented by each of the principal component modes. The univariate tests for individual modes show significant differences among groups, with the difference between means in the third principal component (PC3) as the most significant. Figure 5.16 illustrates the shape variation described by each of the principal component modes in the homozygous model at ± 1 , 2 and 3 standard deviations from the mean. Variation in the heterozygous model modes appears similar.

The results of the morphometric analysis are given in Tables 5.5-5.6. Table 5.5 shows the mean widths and volumes, along with their standard deviations, for each study group. Table 5.6 presents a groupwise comparison of the morphometrics. The results consist of Welch two-sample t-tests for group differences in the means for each of the univariate measures of width and volume and a multivariate Hotelling T^2 test using all four metrics together. Each group is compared with every other group. The results suggest some evidence (at the 5% significance level) for group differences in width in the sagittal direction between the Pax7(LacZ/WT) and the Pax7(WT/WT) groups. No other evidence of significant differences is indicated, however.

The results of the simple morphometric analysis of width and volume, which indicate little or no significant differences in morphometric features between groups, are in stark contrast to the highly significant shape differences revealed by the the PBM correspondence analysis. In general, given sufficient sample size, we would expect a shape analysis method that better represents the geometry of the anatomy to have more statistical power. In this case, the correspondence analysis has revealed shape differences that are too subtle for the more traditional univariate shape features to capture.

In summary, the PBM analysis has demonstrated highly significant group differences in the basioccipital bone due to the *Pax7* mutation. Significant changes include an antero-inferior inflection of the posterior edge of the basioccipital, which is the anterior boundary of the foramen magnum, a landmark whose relative placement in hominins has been associated with bipedalism. By contrast, a more traditional morphometric analysis, which uses a collection of univariate metrics, found only weak significance in one length measure of the data. This analysis thus not only serves as further validation of increased statistical and descriptive power of the PBM method over traditional methodologies, but also strengthens the interplay of reverse genetics in model organisms to paleoanthropology.

Table 5.3. Welch two-sample t-tests for mean shape differences in the first 5 principal component loadings for *Pax7* cohorts.

PCA Loading	Homozygous	Heterozygous	Combined mutant
<i>PC 1</i>	$t = -0.20, p = 0.847$	$t = -1.37, p = 0.182$	$t = -1.25, p = 0.236$
<i>PC 2</i>	$t = -1.29, p = 0.195$	$t = 2.53, p = 0.018$	$t = 0.29, p = 0.775$
<i>PC 3</i>	$t = 5.83, p \ll 0.001$	$t = 5.81, p \ll 0.001$	$t = 0.27, p = 0.793$
<i>PC 4</i>	$t = -0.81, p = 0.446$	$t = 0.33, p = 0.747$	$t = -0.37, p = 0.720$
<i>PC 5</i>	$t = -0.68, p = 0.516$	$t = 1.77, p = 0.088$	$t = 1.32, p = 0.230$

Table 5.4. The percentage of total shape variation represented by each of the first 5 principal modes of variation.

<i>Correspondence model</i>	<i>PC 1</i>	<i>PC 2</i>	<i>PC 3</i>	<i>PC 4</i>	<i>PC 5</i>
Homozygous	42%	19%	12%	11%	6%
Heterozygous	40%	28%	8%	8%	7%
Combined mutant	37%	24%	20%	7%	3%

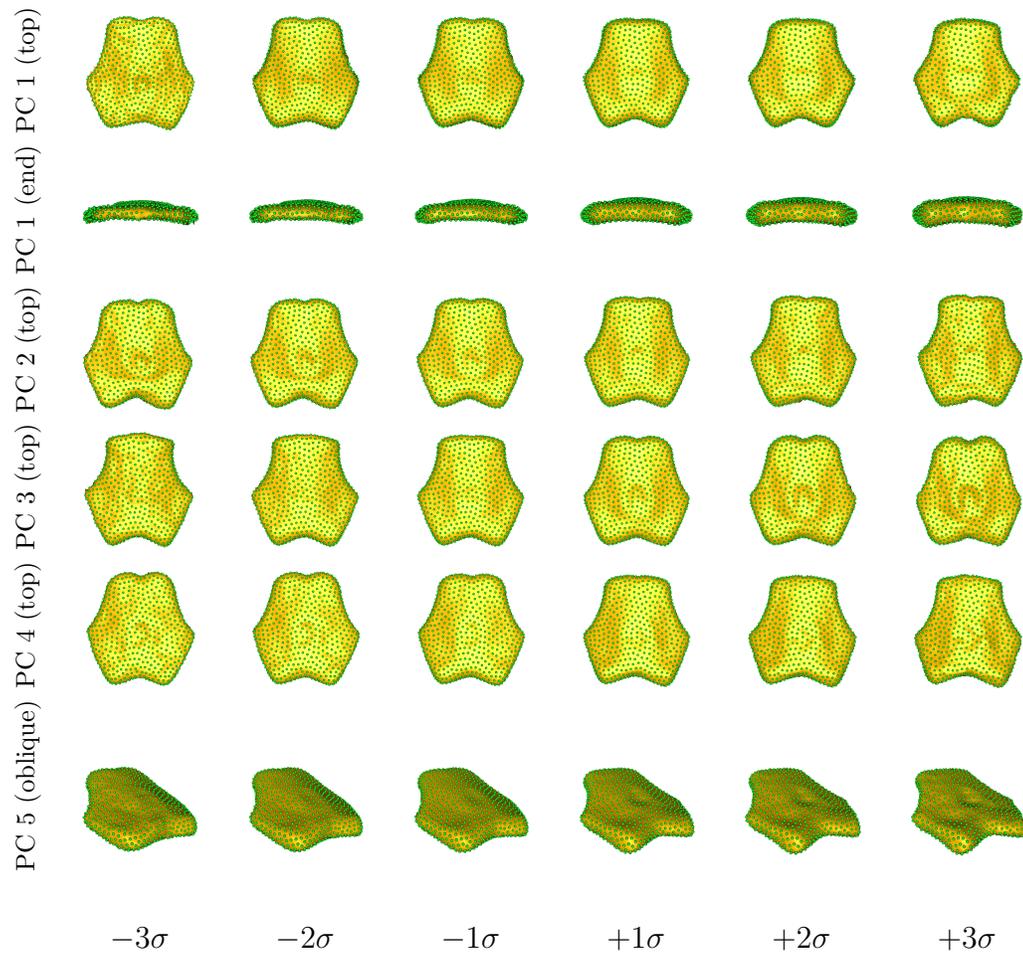


Figure 5.16. Variation in each of the top 5 principal components of the homozygous basioccipital model. Shape (yellow surface) is reconstructed from correspondence positions (green points) at $\pm 1, 2$, and 3 standard deviations (σ) from the mean shape.

Table 5.5. The mean (μ) and standard deviation (σ) in millimeters of the maximum widths for each group of basioccipital bone segmentations.

Metric	<i>LacZ/LacZ</i>		<i>LacZ/WT</i>		<i>WT/WT</i>	
	μ	σ	μ	σ	μ	σ
Sagittal Width	1.9907	0.0423	2.0401	0.0686	1.9769	0.0595
Coronal Width	1.9941	0.1084	2.0217	0.1268	1.9544	0.0621
Axial Width	0.3761	0.0530	0.3726	0.0482	0.3623	0.0428
Volume	0.8205	0.1822	0.8144	0.2019	0.7861	0.1343

Table 5.6. Hypothesis test results for the groupwise comparison of the simple morphometrics. The metrics used are the maximum widths of the segmentations in the sagittal, coronal, and axial directions, along with the segmentation volume. The multivariate Hotelling T^2 test was conducted using all four metrics together.

Metric	Homozygous		Heterozygous		Combined Mutant	
	t	p	t	p	t	p
Sagittal Width	0.5667	0.5801	2.6813	0.0126	2.0922	0.0546
Coronal Width	0.8304	0.4349	1.9304	0.0645	0.5165	0.6167
Axial Width	0.5535	0.5944	0.6170	0.5427	0.1411	0.8913
Volume	0.4101	0.6928	0.4613	0.6481	0.0688	0.9466
Hotelling T^2	1.7338	0.8379	16.5015	0.0172	8.0020	0.1855

CHAPTER 6

CONCLUSION AND FUTURE WORK

This chapter concludes the presentation of my dissertation research with a brief summary of its contributions and early scientific impact. I also discuss some technical limitations and possible future research directions. The main contribution of this work is the particle-based modeling (PBM) framework and its extensions, which I summarize in Section 6.1. The PBM framework a novel, well-validated framework for statistical shape modeling and analysis that is more robust to a wider variety of applications than existing technologies. The reception of the PBM framework in the biomedical research community has thus far been very encouraging, and in Section 6.2, I describe several of the research projects for which it is currently in use. While the PBM algorithm offers some significant advantages, there are still limitations to this work. I discuss some of these limitations in Section 6.3, along with possible solutions and other ideas for future work.

6.1 Research Contributions

In this dissertation, I have described a new computational framework called particle-based modeling (PBM) for constructing statistical representations of sets of similar shapes. The PBM algorithm produces a distribution of a large set of surface-point correspondences using a fast, entropy-based optimization process. The optimization balances the simplicity of the model (compactness) with the accuracy of the shape representations by using two commensurate entropy metrics and a minimum of free parameters. The work in this dissertation significantly advances the state-of-the-art in shape analysis by overcoming many of the limitations of existing methods, including the reliance on spherical representations and ad-hoc regularization strategies.

The PBM algorithm has proven to be a robust analysis tool for a wide variety of shapes and applications. The method works directly on binary segmentation volumes, requires very little parameter tuning, and generalizes easily to accommodate alternate sampling strategies such as curvature adaptivity. I have presented extensions of the basic algorithm to handle open surfaces and complexes of surfaces, which are common to many important biomedical research applications. Additionally, I have extended the PBM framework to accommodate more general notions of correspondence, including correspondence based on functions of surface location and correspondence in the presence of explanatory variables (regression analysis). The PBM algorithm and extensions have been shown to produce results that compare favorably to the state-of-the-art, and statistical analysis of several clinical datasets yields results consistent with those seen in the literature. Additionally, I

have demonstrated results on data that could not be analyzed with existing shape analysis methods, due to its topological complexity.

In this dissertation, I have also presented an extensive evaluation of the PBM framework in the context of six, novel scientific studies in the areas of neurobiology and genetic phenotyping. For the statistical analysis and visualization of the results, I adapted existing multivariate methods from the geometric morphometrics literature to correspondence modeling. The results of these studies have underscored existing scientific hypotheses and produced new hypotheses for future study. These results therefore represent a strong validation of the engineering contributions in this dissertation work.

6.2 Scientific Impact

The PBM shape modeling framework is currently impacting several areas of biomedical research, including neurobiology, cardiology, orthopedics, and evolutionary biology. Several preliminary studies using PBM are either planned or ongoing in several university and clinical research laboratories. By maintaining and distributing ShapeWorks, the open-source distribution of the PBM implementation [17], the Scientific Computing and Imaging Institute is facilitating the use of these tools by the community. This section lists some of the ongoing studies that are using the PBM algorithm, which are evidence of some early scientific impact of this dissertation work.

In the area of neurobiology, the work of Oguz, et al., from the University of North Carolina, has already produced additional results for improved cortical surface correspondence using the generalized PBM method. In [85], they incorporate probabilistic connectivity data into the optimization, which reflects the fiber structures between the cortical surface and various regions of interest, such as the corpus callosum and the internal capsule.

In cardiology, several research efforts are beginning to make use of the PBM algorithm. The Comprehensive Arrhythmia Research and Management Center (CARMA), at the University of Utah, is an interdisciplinary team of scientists and cardiologists who are pioneering new treatments for cardiac arrhythmias. The CARMA group is investigating the feasibility of using the PBM algorithm to create population models of the left and right atria. Because the shapes of the atria are known to vary in patients with atrial fibrillation, shape is one of the promising metrics for image-based assessment of disease progression.

In collaboration with Raimond Winslow and Laurent Younes from the Johns Hopkins University, Manasi Datar and Ross Whitaker at the University of Utah are investigating shape changes in the ventricles of the heart in patients with ischemic conditions. They are currently comparing PBM results with voxel-based morphometry analysis (e.g., [4]). While the concave shape and thin walls of the ventricle are challenging features to model, Figure 6.1 (courtesy of Manasi Datar) illustrates some preliminary results from this study that suggest PBM algorithm will be an effective tool. The figure illustrates several views of a mean left ventricle shape.

At the University of Utah Department of Orthopedics, researchers are investigating the role of the shapes of the femur and hip joint in hip dysplasia. This project



Figure 6.1. Preliminary results from a PBM study of ventricle shapes. Images courtesy of Manasi Datar.

is led by Jeff Weiss and Andy Anderson, also in collaboration with Manasi Datar and Ross Whitaker. The PBM algorithm for open surfaces has already proven effective in modeling the femoral head, as shown in Figure 6.2, which illustrates the mean shape of a normal control population. Further research will model femur covariation with the hip joint shape, in order to understand the relationship between these two structures in both normal and patient populations.

6.3 Limitations and Future Work

While the PBM algorithm and its extensions are proving to be a robust approach to statistical shape modeling for a variety of biomedical problems, there are some technical issues that impose limitations on its use. Two issues that limit its application to more complex surface geometries, for example, are the tangent plane assumption in the particle distribution formulation (Section 2.3.7) and the need for a more robust initialization strategy. Another issue, with implications for hypothesis testing and statistics, is the Gaussian assumption for the sample distributions in shape space (Section 2.3.8). In this section, I briefly describe these issues and suggest some future research that may help address them.

Complicated surfaces that consist of multiple ridges and valleys remain a challenge for the PBM optimization. Recall from Section 2.3.7 that the numerical implementation for distributing particles across shape surfaces relies on a degree of regularity in the tangent planes between adjacent particles. For particles on ridges and in valleys, this assumption is not satisfied, and distributions can become “stuck” in these regions, having achieved a false local minimum by effectively losing track of their true neighbors. A variable sampling density that is proportional to surface curvature (Section 2.3.4) can somewhat alleviate this problem, but is not sufficient in general. Conceptually, the most straightforward way to handle this problem is to remove the tangent plane approximations and use geodesic distances between particles. Unfortunately, this approach is very computationally expensive and would severely impact the speed of the algorithm. As an alternative to geodesics on the implicit surface representation, however, the algorithm could be modified so that the surface constraint is maintained on a mesh-based representation. The geodesic distances between vertices on the meshes could then be *precomputed*, and

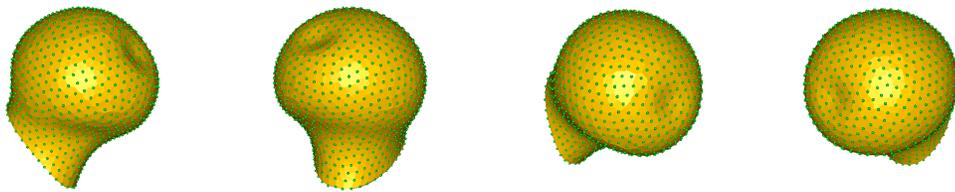


Figure 6.2. Preliminary results from a PBM study of femur shapes. Images courtesy of Manasi Datar.

distances between particle positions quickly interpolated from those values during the optimization.

For all of the examples and applications of the PBM algorithm in this dissertation, I have successfully used the splitting strategy described in Section 2.4.2. The splitting method, however, is susceptible in some cases to producing poor configurations from which the optimization cannot recover. Specifically, the problem is that corresponding particles sometimes find their way onto opposite sides of ridges or opposite sides of flat shapes (i.e. the basioccipital bone). While a geodesic distance formulation of PBM would likely alleviate this problem, there are probably simpler solutions to be found. One idea for a more robust initialization strategy is to relax the surface constraint during the initialization, so that particles are allowed to move off of the surfaces. The constraint can then be gradually phased back in, so that particles once again find the surface. I have done some preliminary tests of this strategy for the basioccipital bone data (Section 5.6) with good results. These tests suggest that an added benefit may be a significant speedup of the initialization process. This approach also suggests a new optimization metric that adjusts a particle’s affinity for the surface based on the likelihood of its position, given the positions of all of its corresponding particles. Under such an approach, “outlier” particles would be allowed to move with increased degrees of freedom towards better correspondence, while movement of particles already in good correspondence would remain constrained on the surface.

Recall from Section 2.3.8 that, due to the limited sample size, the current implementation of the PBM algorithm assumes a Gaussian distribution for the samples in shape space when estimating entropy. It is possible that this assumption may bias the results of group-wise hypothesis tests by favoring optimized configurations that resemble a single Gaussian configuration, rather than a multimodal distribution. Because the Gaussian assumption is not critical to the general correspondence method, a nonparametric estimation of sample density, such as a kernel-based method, could be used instead to remove this bias. As with other challenges for statistical analysis of correspondences (Section 4.3.1), however, the difficulty in estimating the distribution lies in the high dimensionality of the model and the low numbers of samples. One solution may be to estimate the sample entropy in a lower-dimensional subspace of the model and, perhaps, also employ a bootstrapping method to compute confidence intervals for estimates under this model.

Many other expansions and optimizations of the PBM framework can be also be considered. The algorithm is, for example, very easy to parallelize by partitioning the problem into one surface sample per processor, with the covariance calculation in Equation 2.33 being the only point of synchronization among the processes. Image segmentation using PBM correspondences is another interesting area for possible future research. Segmentation could be done, for example, using an approach analogous to active shape models [21]. Finally, because the flexibility of the PBM framework makes it relatively easy to adapt to new classes of shapes and optimization criteria, I look forward to seeing future applications to problems in fields beyond biomedicine, such as industrial engineering and computer graphics.

REFERENCES

- [1] D. Adams, F. J. Rolf, and D. Slice. Geometric morphometrics: Ten years of progress following the “revolution”. *Italian Journal of Zoology*, 71:5–16, 2004.
- [2] J. Ahn, J. S. Marron, K. M. Muller, and Y. Chi. The high-dimension, low-sample-size geometric representation holds under mild conditions. *Biometrika*, 94(3):760–766, 2007.
- [3] M. Alexa, J. Behr, D. Cohen-Or, S. Fleishman, D. Levin, and C. T. Silva. Computing and rendering point set surfaces. *IEEE Transactions on Visualization and Computer Graphics*, 9(1):3–15, January 2003.
- [4] J. Ashburner and K. J. Friston. Voxel-based morphometry — the methods. *Neuroimage*, 11:805–821, 2000.
- [5] J. Ashburner, C. Hutton, R. Frackowiak, I. Johnsrude, C. Price, and K. Friston. Identifying global anatomical differences: Deformation-based morphometry. *Human Brain Mapping*, 6(5):348–357, 1998.
- [6] M. Audette, F. Ferrie, and T. Peters. An algorithmic overview of surface registration techniques for medical imaging. *Medical Image Analysis*, 4:201–217, 2000.
- [7] R. Bajcsy and S. Kovačič. Multiresolution elastic matching. *Comput. Vision Graph. Image Process.*, 46(1):1–21, 1989.
- [8] M. F. Beg, M. I. Miller, A. Trounev, and L. Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Int. J. Comput. Vision*, 61(2):139–157, 2005.
- [9] J. T. Berge. Orthogonal procrustes rotation for two or more matrices. *Psychometrika*, 42:267–276, 1977.
- [10] J.-D. Boissonnat and S. Oudot. Provably good sampling and meshing of surfaces. *Graphical Models*, 67:405–451, 2005.
- [11] F. Bookstein. Principal warps: Thin plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(6), 1989.
- [12] F. Bookstein. Biometrics, biomathematics, and the morphometric synthesis. *Bulletin of Mathematical Biology*, 58:313–365, 1996.

- [13] A. M. Boulet and M. R. Capecchi. Duplication of the *hoxd11* gene causes alterations in the axial and appendicular skeleton of the mouse. *Developmental Biology*, 249:96–107, 2002.
- [14] F. O. Bower. Size and form in plants. *Nature*, 126(3175):355–361, 1930.
- [15] C. Brechbuhler, G. Gerig, and O. Kubler. Parametrization of closed surfaces for 3-d shape description. *Computer Vision Image Understanding CVIU*, 61:154–170, 1995.
- [16] M. R. Capecchi. Generating mice with targeted mutations. *Nat Med*, 7(10):1086–1090, October 2001.
- [17] J. Cates, 2009. ShapeWorks [Computer Software]. Scientific Computing and Imaging Institute (SCI).
- [18] J. Cates, P. T. Fletcher, M. Styner, H. C. Hazlett, and R. Whitaker. Particle-based shape analysis of multi-object complexes. In *MICCAI '08: Proceedings of the 11th international conference on Medical Image Computing and Computer-Assisted Intervention - Part I*, pages 477–485, Berlin, Heidelberg, 2008. Springer-Verlag.
- [19] X. J. Chen, N. Kovacevic, N. J. Lobaugh, J. G. Sled, R. M. Henkelman, and J. T. Henderson. Neuroanatomical differences between mouse strains as shown by high-resolution 3d mri. *NeuroImage*, 29:99–105, 2005.
- [20] R. Chultz and D. Robins. *Handbook of autism and pervasive developmental disorders (3rd Ed.)*, chapter Functional neuroimaging studies of autism spectrum disorders, pages 515–533. J Wiley and Sons, 2005.
- [21] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models-their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, January 1995.
- [22] E. Courchesne, C. M. Karns, H. R. Davis, R. Ziccardi, R. A. Carper, Z. D. Tigue, H. J. Chisum, P. Moses, and K. Pierce. Unusual brain growth patterns in early life in patients with autistic disorder: An mri study. *Neurology*, 57:245–254, 2001.
- [23] T. Cover and J. Thomas. *Elements of Information Theory*. Wiley and Sons, 1991.
- [24] P. Dalal, B. C. Munsell, S. Wang, J. Tang, K. Oliver, H. Ninomiya, X. Zhou, and H. Fujita. A fast 3d correspondence method for statistical shape modeling. *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, 0:1–8, 2007.
- [25] A. M. Dale, B. Fischl, and M. I. Sereno. Cortical surface-based analysis: I. segmentation and surface reconstruction. *NeuroImage*, 9(2):179 – 194, 1999.

- [26] C. Darwin. *The Origin of Species*. Gramercy, 1995.
- [27] R. Davies, C. Twining, P. Allen, T. Cootes, and C. Taylor. Shape discrimination in the hippocampus using an MDL model. In *Information Processing in Medical Imaging*, pages 38–50, 2003.
- [28] R. H. Davies, C. J. Twining, P. D. Allen, T. F. Cootes, and C. J. Taylor. Shape discrimination in the hippocampus using an mdl model. In *IPMI*, pages 38–50, 2003.
- [29] R. H. Davies, C. J. Twining, T. F. Cootes, J. C. Waterton, and C. J. Taylor. 3d statistical shape models using direct optimisation of description length. In *ECCV (3)*, pages 3–20, 2002.
- [30] R. H. Davies, C. J. Twining, T. F. Cootes, J. C. Waterton, and C. J. Taylor. A minimum description length approach to statistical shape modeling. *IEEE Trans. Med. Imaging*, 21(5):525–537, 2002.
- [31] A. Davis and M. Capecchi. Axial homeosis and appendicular skeleton defects in mice with targeted disruption of *hoxd-11*. *Development*, 120:2187–2198, 1995.
- [32] B. C. Davis, P. T. Fletcher, E. Bullitt, and S. C. Joshi. Population shape regression from random design data. In *Proceedings of the 10th International Conference on Computer Vision (ICCV 2007)*, 2007.
- [33] M. Davis, 2009. Seg3D: Volumetric Image Segmentation and Visualization [Computer Software]. Scientific Computing and Imaging Institute (SCI).
- [34] I. Dryden and K. Mardia. *Statistical Shape Analysis*. John Wiley and Sons, 1998.
- [35] L. R. Fabrigar, D. T. Wegener, R. C. MacCallum, and E. J. Strahan. Evaluating the use of exploratory factor analysis in psychological research. *Psychological Methods*, 4:272–299, 1999.
- [36] B. Fischl, M. I. Sereno, and A. M. Dale. Cortical surface-based analysis: I: Inflation, flattening, and a surface-based coordinate system. *NeuroImage*, 9(2):195 – 207, 1999.
- [37] R. A. Fisher. *The Design of Experiments*. Oliver and Boyd, 1935.
- [38] S. Fleishman, D. Cohen-Or, M. Alexa, and C. T. Silva. Progressive point set surfaces. *ACM Trans. Graph.*, 22(4):997–1011, 2003.
- [39] P. Fletcher, C. Lu, S. Pizer, and S. Joshi. Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Trans. Med. Imaging*, 23(8):995–1005, 2004.
- [40] G. Galilei. *Dialogues Concerning Two New Sciences*. Dover, 1960.

- [41] J. Gayon. History of the concept of allometry. *American Zoologist*, 40(5):748–758, 2000.
- [42] G. Gerig, S. Joshi, T. Fletcher, K. Gorczowski, S. Xu, S. Pizer, and M. Styner. Statistics of populations of images and its embedded objects: Driving applications in neuroimaging. In *IEEE Symp. on Biomed. Imaging ISBI*, pages 1120–1123, 2006.
- [43] L. W. Glorfeld. An improvement on horn’s parallel analysis methodology for selecting the correct number of factors to retain. *Educational and Psychological Measurement*, 55:377–393, 1995.
- [44] P. Golland, W. Grimson, M. Shenton, and R. Kikinis. Detection and analysis of statistical differences in anatomical shape. *Medical Image Analysis*, 9:69–86, 2005.
- [45] C. Goodall. Procrustes methods in the statistical analysis of shape. *J. R. Statistical Society B*, 53:285–339, 1991.
- [46] K. Gorczowski, M. Styner, J. Jeong, J. Marron, J. Piven, H. Hazlett, S. Pizer, and G. Gerig. Statistical shape analysis of multi-object complexes. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2007.
- [47] J. Gower. Generalized procrustes analysis. *Psychometrika*, 40:33–50, 1975.
- [48] J. S. Gregory, D. Testi, A. Stewart, P. E. Undrill, D. M. Reid, and R. M. Aspden. A method for assessment of the shape of the proximal femur and its relationship to osteoporotic hip fracture. *Osteoporosis Int.*, 15(1):5–11, 2004.
- [49] P. Hall and N. Tajvidi. Permutation tests for equality of distributions in high-dimensional settings. *Biometrika*, 89(2):359–374, 2002.
- [50] M. Hansen and B. Yu. Model selection and the principle of minimum description length. *Journal of the American Statistical Association*, 96(454):746–774, June 2001.
- [51] H. Hazlett, M. Poe, G. Gerig, R. Smith, J. Provenzale, A. Ross, J. Gilmore, and J. Piven. Magnetic resonance imaging and head circumference study of brain size in autism: Birth through age 2 years. *Arch Gen Psych*, 62:1366–1376, 2005.
- [52] P. A. Helm, L. Younes, M. F. Beg, D. B. Ennis, C. Leclercq, O. P. Faris, E. McVeigh, D. Kass, M. Miller, and R. L. Winslow. Evidence of structural remodeling in the dyssynchronous failing heart. *Circulation Research*, 98(1):125–132, 2006.
- [53] K. Hinkelmann and O. Kempthorne. *The Design and Analysis of Experiments*. Wiley, 2008.

- [54] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. In *ACM SIGGRAPH 1992 Conference Proceedings*, pages 71–78, 1992.
- [55] H. Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of Educational Psychology*, 24:417–441, 1933.
- [56] D. A. Hoyte. A critical analysis of the growth in length of the cranial base. *Birth Defects Orig Artic Ser*, 11:255–282, 1975.
- [57] J. R. Hurley and R. B. Cattell. The procrustes program: producing direct rotation to test a hypothesised factor structure. *Behavioural Science*, 7:258–262, 1962.
- [58] J. Huxley. *Problems of relative growth*. New York, L. MacVeagh, The Dial Press,, 1932.
- [59] J. S. Huxley and G. Teissier. Terminology of relative growth. *Nature*, 137:780–781, 1936.
- [60] L. Ibanez, W. Schroeder, L. Ng, and J. Cates. *The ITK Software Guide*. Insight Consortium, 2003.
- [61] P. Jolicoeur. The generalization of the allometry equation. *Biometrics*, 19:497–499, 1963.
- [62] S. Joshi, B. Davis, M. Jomier, and G. Gerig. Unbiased diffeomorphic atlas construction for computational anatomy. *Neuroimage*, 23:S150–S160, 2004.
- [63] D. G. Kendall. The diffusion of shape. *Adv. Appl. Probab.*, 9:428–430, 1977.
- [64] D. G. Kendall. Shape-manifolds, procrustean metrics and complex projective spaces. *Bulletin of the London Mathematical Society*, 16:81–121, 1984.
- [65] D. G. Kendall. *Shape and Shape Theory*. Wiley, 1999.
- [66] J. T. Kent. The complex bingham distribution and shape analysis. *Journal of the Royal Statistical Society, Series B*, 56:285–299, 1994.
- [67] G. Kindlmann, R. Whitaker, T. Tasdizen, and T. Moller. Curvature-based transfer functions for direct volume rendering. In *Proceedings of IEEE Visualization 2003*, pages 512–520, 2003.
- [68] C. P. Klingenberg. Morphometrics and the role of the phenotype in studies of the evolution of developmental mechanisms. *Gene*, 287:3–10, 2002.
- [69] R. C. Knickmeyer, S. Gouttard, C. Kang, D. Evans, K. Wilber, J. K. Smith, R. M. Hamer, W. Lin, G. Gerig, and J. H. Gilmore. A structural mri study of human brain development from birth to 2 years. *J. Neurosci.*, 28(47):12176–12182, November 2008.

- [70] A. Kotcheff and C. Taylor. Automatic Construction of Eigenshape Models by Direct Optimization. *Medical Image Analysis*, 2:303–314, 1998.
- [71] F. P. Kuhl and C. R. Giardina. Elliptic fourier features of a closed contour. *Computer Graphics and Image Processing*, 18(3):236–258, 1982.
- [72] M. Langen, S. Durston, W. Staal, S. Palmen, and H. van Engeland. Caudate nucleus is enlarged in high-functioning medication-naive subjects with autism. *Biological Psychiatry*, 62:262–266, 2007.
- [73] S. M. Leach, H. Tipney, W. Feng, W. A. Baumgartner, Jr., P. Kasliwal, R. P. Schuyler, T. Williams, R. A. Spritz, and L. Hunter. Biomedical discovery acceleration, with applications to craniofacial development. *PLoS Comput Biol*, 5(3):e1000215, 03 2009.
- [74] J. Levitt, C.-F. Westin, P. Nestor, R. Estepar, C. Dickey, M. Voglmaier, L. Seidman, R. Kikinis, F. Jolesz, R. McCarley, and M. Shenton. Shape of caudate nucleus and its cognitive correlates in neuroleptic-naive schizotypal personality disorder. *Biol. Psychiatry*, 55:177–184, 2004.
- [75] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. *Computer Graphics*, 21(4):163–169, July 1987.
- [76] S. L. Mansour, K. R. Thomas, and M. R. Capecchi. Disruption of the proto-oncogene int-2 in mouse embryo derived stem cells: a general strategy for targeting mutations to non-selectable genes. *Nature*, 336:348–352, 2004.
- [77] A. Mansouri, A. Stoykova, M. Torres, and P. Gruss. Dysgenesis of cephalic neural crest derivatives in pax7^{-/-} mutant mice. *Development*, 122:831–838, 1996.
- [78] M. Meyer, B. Nelson, R. Kirby, and R. Whitaker. Particle systems for efficient and accurate high-order finite element visualization. *Visualization and Computer Graphics, IEEE Transactions on*, 13(5):1015–1026, Sept.-Oct. 2007.
- [79] M. D. Meyer, P. Georgel, and R. T. Whitaker. Robust particle systems for curvature dependent sampling of implicit surfaces. In *Proceedings of the International Conference on Shape Modeling and Applications*, pages 124–133, June 2005.
- [80] B. C. Munsell, P. Dalal, and S. Wang. Evaluating shape correspondence for statistical shape analysis: A benchmark study. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(11):2023–2039, 2008.
- [81] D. Nain, M. Styner, M. Niethammer, J. Levitt, M. Shenton, G. Gerig, A. Bobick, and A. Tannenbaum. Statistical shape analysis of brain structures using spherical wavelets. In *Biomedical Imaging: From Nano to Macro, 2007. ISBI 2007. 4th IEEE International Symposium on*, pages 209–212, April 2007.

- [82] L. Nevell and B. Wood. Cranial base evolution within the hominin clade. *Journal of Anatomy*, 212(4):455–468, 2008.
- [83] B. J. Nieman, A. M. Flenniken, S. L. Adamson, R. M. Henkelman, and J. G. Sled. Anatomical phenotyping in the brain and skull of a mutant mouse by magnetic resonance imaging and computed tomography. *Physiological Genomics*, 24:154–162, 2006.
- [84] I. Oguz, J. Cates, P. T. Fletcher, Z. Warnock, R. Whitaker, D. Cool, S. Aylward, and M. Styner. Cortical correspondence using entropy-based particle systems and local features. In *Proc. 5th IEEE International Symposium on Biomedical Imaging (ISBI '08)*, pages 1637–1640, 2008.
- [85] I. Oguz, M. Neithammer, J. Cates, R. Whitaker, P. T. Fletcher, C. Vachet, and M. Styner. Cortical correspondence with probabilistic fiber connectivity. In *Information Processing in Medical Imaging (IPMI 2009)*, LNCS 5636, pages 651–663, 2009.
- [86] E. Parzen. On estimation of a probability density function and mode. *The Annals of Mathematical Statistics*, 33(3):1065–1076, 1962.
- [87] K. Pearson. On lines and planes of closest fit to systems of points in space. *Philosophical Magazine*, 2(6):559–572, 1901.
- [88] E. Persoon and K. S. Fu. Shape discrimination using fourier descriptors. *IEEE Trans. Systems, Man and Cybernetics SMC*, 7(3):388–397, 1977.
- [89] S. M. Pizer, J.-Y. Jeong, C. Lu, K. E. Muller, and S. C. Joshi. Estimating the statistics of multi-object anatomic geometry using inter-object relationships. In *Deep Structure, Singularities, and Computer Vision*, volume 3753 of LNCS, pages 60–71, 2005.
- [90] R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2008.
- [91] F. J. Rohlf. Bias and error in estimates of mean shape in geometric morphometrics. *Journal of Human Evolution*, 44:665–683, 2003.
- [92] M. Rousson, N. Paragios, and R. Deriche. Implicit active shape models for 3d segmentation in mr imaging. In *Medical Image Computing and Computer-Assisted Intervention MICCAI 2004*, pages 209–216, 2004.
- [93] W. Schroeder, K. Martin, and B. Lorensen. *The Visualization Toolkit, An Object Oriented Approach to 3D Graphics*. Prentice-Hall, 1998.
- [94] F. Segonne, A. M. Dale, E. Busa, M. Glessner, D. Salat, H. K. Hahn, and B. Fischl. A hybrid approach to the skull stripping problem in mri. *NeuroImage*, 22(3):1060 – 1075, 2004.
- [95] J. Sethian. *Level Set Methods and Fast Marching Methods*. Cambridge University Press, 1996.

- [96] M. Styner, J. Lieberman, and G. Gerig. Boundary and medial shape analysis of the hippocampus in schizophrenia. In *MICCAI*, pages 464–471, 2003.
- [97] M. Styner, J. A. Lieberman, D. Pantazis, and G. Gerig. Boundary and medial shape analysis of the hippocampus in schizophrenia. *Medical Image Analysis*, 2004.
- [98] M. Styner, I. Oguz, S. Xu, C. Brechbuhler, D. Pantazis, J. Levitt, M. Shenton, and G. Gerig. Framework for the statistical shape analysis of brain structures using SPHARM-PDM. *The Insight Journal*, 2006.
- [99] M. Styner, K. T. Rajamani, L.-P. Nolte, G. Zsemlye, G. Székely, C. J. Taylor, and R. H. Davies. Evaluation of 3d correspondence methods for model building. In *IPMI*, pages 63–75, 2003.
- [100] M. Styner, S. Xu, M. El-Sayed, and G. Gerig. Correspondence evaluation in local shape analysis and structural subdivision. In *Biomedical Imaging: From Nano to Macro, 2007. ISBI 2007. 4th IEEE International Symposium on*, pages 1192–1195, April 2007.
- [101] G. Székely, A. Kelemen, C. Brechbuhler, and G. Gerig. Segmentation of 2-d and 3-d objects from mri volume data using constrained elastic deformations of flexible fourier contour and surface models. *Medical Image Analysis*, 1(1):19–34, 1996.
- [102] T. Terriberry, S. Joshi, and G. Gerig. Hypothesis testing with nonlinear shape models. In *IPMI'05*, pages 15–26, 2005.
- [103] H. H. Thodberg. Minimum description length shape and appearance models. In *IPMI*, pages 51–62, 2003.
- [104] D. W. Thompson. *On growth and form*. Cambridge University Press, 1945.
- [105] N. H. Timm. *Applied Multivariate Analysis*. Springer-Verlag, 2002.
- [106] M. W. Tocheri, C. M. Orr, S. G. Larsen, T. Sutikna, Jatmiko, E. W. Saptomo, R. A. Due, T. Djubiontono, M. J. Morwood, and W. L. Jungers. The primitive wrist of *homo floresiensis* and its implications for hominim evolution. *Science*, 317(5845):1743 – 1745, 2007.
- [107] M. Umer, A. Thambyah, W. T. Tan, and S. Das De. Acetabular morphometry for determining hip dysplasia in the singaporean population. *J Orthop Surg (Hong Kong)*, 14(1):27–31, Apr 2006.
- [108] M. Vaillant, M. Miller, L. Younes, and C. Ceritoglu. Initial value diffeomorphic landmark matching and its application shape statistics. In *IEEE Workshop on Statistical Signal Processing*, pages 306–320, 2003.
- [109] M. Vaillant, M. Miller, L. Younes, and A. Trouv. Statistics on diffeomorphisms via tangent space representations. *NeuroImage*, 23(Supplement 1):S161 – S169, 2004. Mathematics in Brain Imaging.

- [110] B. van Ginneken, A. F. Frangi, R. F. Frangi, J. J. Staal, B. M. T. H. Romeny, and M. A. Viergever. Active shape model segmentation with optimal features. *IEEE Transactions on Medical Imaging*, 21:924–933, 2002.
- [111] Y. Vardi and C.-H. Zhang. The multivariate l_1 -median and associated data depth. *Proceedings of the National Academy of Sciences*, 97(4):1423–1426, 2000.
- [112] R. Whitaker. Reducing aliasing artifacts in iso-surfaces of binary volumes. In *IEEE Volume Visualization and Graphics Symposium*, pages 22–23, 2000.
- [113] M. Whitbeck and H. Guo. Multiple landmark warping using thin-plate splines. In *IPCV*, pages 256–263, 2006.
- [114] J. Williams and J. Rossignac. Tightening: curvature-limiting morphological simplification. In *Proc. Ninth ACM Symposium on Solid and Physical Modeling*, pages 107–112, 2005.
- [115] A. P. Witkin and P. S. Heckbert. Using particles to sample and control implicit surfaces. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 269–277. ACM Press, 1994.
- [116] I. Wright, P. McGuire, J. Poline, J. Travere, R. Murray, C. Frith, R. Frackowiak, and K. Friston. A voxel-based method for the statistical analysis of gray and white matter density applied to schizophrenia. *NeuroImage*, 2:244–252, 1995.