

Groupwise Shape Correspondence with Local Features

İpek Oğuz

A dissertation submitted to the faculty of the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Computer Science.

Chapel Hill
2009

Approved by:

Martin A. Styner, Advisor

Marc Niethammer, Co-principal Reader

Stephen M. Pizer, Reader

J. Stephen Marron, Reader

Ross Whitaker, Reader

© 2009
İpek Oğuz
ALL RIGHTS RESERVED

Abstract

İpek Oğuz: Groupwise Shape Correspondence with Local Features.
(Under the direction of Martin A. Styner)

Statistical shape analysis of anatomical structures plays an important role in many medical image analysis applications such as understanding the structural changes in anatomy in various stages of growth or disease. Establishing accurate correspondence across object populations is essential for such statistical shape analysis studies. However, anatomical correspondence is rarely a direct result of spatial proximity of sample points but rather depends on many other features such as local curvature, position with respect to blood vessels, or connectivity to other parts of the anatomy.

This dissertation presents a novel method for computing point-based correspondence among populations of surfaces by combining spatial location of the sample points with non-spatial local features. A framework for optimizing correspondence using arbitrary local features is developed. The performance of the correspondence algorithm is objectively assessed using a set of evaluation metrics.

The main focus of this research is on correspondence across human cortical surfaces. Statistical analysis of cortical thickness, which is key to many neurological research problems, is the driving problem. I show that incorporating geometric (sulcal depth) and non-geometric (DTI connectivity) knowledge about the cortex significantly improves cortical correspondence compared to existing techniques. Furthermore, I present a framework that is the first to allow the white matter fiber connectivity to be used for improving cortical correspondence.

To İsmail Hakkı Birler (1927-2009)

Acknowledgments

I'd like to express my gratitude to those who have contributed to this work. First, I would like to thank Martin Styner for being an outstanding advisor *and* mentor during these 5 years. It makes me happy to know that our collaboration is not coming to an end.

I am also very grateful to my committee members. In particular, Marc Niethammer has been extremely helpful with the preparation of this document with his detailed (and extremely fast!) feedback as well as his insights throughout the project. It has been a privilege to learn from Steve Pizer, whose immense knowledge and understanding of the field as well as his enthusiasm for his work is a constant source of inspiration. I am thankful to Steve Marron for lending his statistical expertise and to Ross Whitaker for a very productive collaboration throughout the project.

I'd like to thank Josh Cates who generously shared his knowledge, insights, code and time throughout this project. Tom Fletcher has been a great mentor and friend who gave me advice on what classes to take during my first semester, what to expect at job interviews during my last semester, and everything in between (including bird watching advice that I completely ignored, sorry). The MDL-based portion of this work wouldn't have been possible without the very patient and friendly Tobias Heimann. The MIDAG team at UNC has been a constant source of knowledge, ideas and support.

I am thankful to Randy Gollub and the MIND Clinical Imaging Consortium for the main cortical dataset, and to Stephen Aylward and Derek Cool for the remaining cortical data. The caudate and striatum segmentations were provided by James Levitt and Martha Shenton, the femur segmentations were provided by Markus Fleute, and the lateral ventricle segmentations were provided by Daniel Weinberger and Douglas Jones.

This work was supported by the National Alliance for Medical Image Computing (NA-MIC), funded by the NIH. I have been incredibly fortunate to be part of NA-MIC; I am absolutely grateful for the opportunity to meet and collaborate with many great researchers.

Finally, I would like to thank my friends, my dear family, and Jeff.

Table of Contents

List of Figures	x
List of Abbreviations	xii
1 Introduction	1
1.1 Correspondence of local features	2
1.2 Framework for correspondence of local features without parametrical mapping	3
1.2.1 Application to the human brain: dealing with the geometric challenges of the cortex	5
1.3 Using white matter fiber structure for cortical correspondence	6
1.3.1 Probabilistic connectivity	7
1.3.2 Mapping the connectivity to the cortical surface	7
1.4 Evaluation of correspondence quality	9
1.5 Thesis and contributions	10
1.6 Overview of chapters	11
2 Shape Correspondence	13
2.1 Pairwise Correspondence	14
2.1.1 Surface-based pairwise correspondence	14
2.1.2 Volume-based pairwise correspondence	21
2.2 Groupwise Correspondence	23
2.2.1 Determinant of the Covariance Matrix	24

2.2.2	Minimum Description Length	24
2.3	Correspondence Evaluation	25
2.3.1	Distance to Manual Landmarks	25
2.3.2	Jaccard Coefficient Difference	26
2.3.3	Generalization, Specificity and Compactness	27
2.3.4	Goodness of Prediction ρ^2	29
3	Parameterization-Based Group Correspondence	30
3.1	Traditional MDL	30
3.1.1	Original Algorithm Formulation	31
3.1.2	Gradient Descent Optimization of the MDL Function	33
3.1.3	Shape Image Based MDL Optimization Schemes	35
3.2	Generalized MDL Correspondence	35
3.2.1	Using Local Features for Generalizing MDL	36
3.3	Geometric Features	37
3.3.1	Principal Curvatures (κ_1, κ_2)	37
3.3.2	Mean Curvature and Gaussian Curvature (H, K)	38
3.3.3	Curvedness and Shape Index (C, S)	38
3.4	Results and Discussion	40
3.4.1	Lateral Ventricle	43
3.4.2	Caudate	44
3.4.3	Femur	44
3.4.4	Striatum	44
3.4.5	Discussion	45
4	Parameterization-Free Group Correspondence and Application to Highly Folded Surfaces	51

4.1	Entropy-based Particle Framework	52
4.1.1	Surface entropy	52
4.1.2	Ensemble entropy	55
4.2	Generalized ensemble entropy	57
4.3	Application to cortex	58
4.4	Surface interpolation from particles	60
4.5	Results	63
5	DTI-based Connectivity for Cortical Correspondence	69
5.1	Diffusion Tractography for Connectivity	71
5.1.1	Streamline tractography	71
5.1.2	Stochastic tractography	72
5.1.3	Optimal path methods	74
5.2	Mapping Connectivity to the Cortical Surface	75
5.2.1	Kernel-based averaging	75
5.2.2	Brain deflation	76
5.3	Results	82
6	Discussion	87
6.1	Summary of contributions and thesis	87
6.2	Future work and discussion	94
	Appendix A: Spherical Harmonics	99
A-1	Surface representation	99
A-2	Partial derivatives	100
A-3	Surface normal	103
A-4	First and second fundamental forms	103
A-5	Curvature metrics	103

Bibliography	105
------------------------	-----

List of Figures

1.1	Pipeline overview	8
2.1	Example of parameterization-based correspondence	16
3.1	Issues with using spatial proximity as a correspondence indicator . . .	36
3.2	Shape index S	39
3.3	Curvedness C	40
3.4	The relative locations of the various brain structures used in experiments	41
3.5	The results of the SPHARM, traditional MDL and CombinationMDL demonstrated on two striata and two femoral heads	42
3.6	Comparison of correspondence results for lateral ventricle population	47
3.7	Comparison of correspondence results for caudate population	48
3.8	Comparison of correspondence results for femur population	49
3.9	Comparison of correspondence results for striatum population	50
4.1	Sulcal depth on the white matter (WM) surface and the inflated WM surface	60
4.2	Surface interpolation from particles	61
4.3	Example of surface interpolation	62
4.4	Correspondence results comparison for the first cortical dataset . . .	64
4.5	Comparison of the distribution of variance across the cortical surfaces for entropy method and FreeSurfer	65
4.6	Correspondence results comparison for the second cortical dataset . .	66
4.7	Cortical thickness generalization, $G(M)$	67
4.8	Cortical thickness specificity, $S(M)$	68

5.1	Connectivity features pipeline overview	70
5.2	Impact of brain deflation algorithm on surface connectivity values . .	76
5.3	Brain deflation progress for one subject	79
5.4	Motivation for normalizing connectivity maps	80
5.5	Comparison of connectivity maps with linear normalization and his- togram equalization	81
5.6	Correspondence results comparison using average variances of cortical thickness and sulcal depth	84
5.7	Cortical thickness generalization and specificity comparison	85

List of Abbreviations

<i>MRI</i>	Magnetic Resonance Imaging, introduced chapter 1
<i>MRA</i>	Magnetic Resonance Angiography, introduced chapter 1
<i>MDL</i>	Minimum Description Length, introduced section 1.1
<i>WM</i>	White Matter, introduced section 1.2
<i>GM</i>	Gray Matter, introduced section 1.2
<i>DTI</i>	Diffusion Tensor Imaging, introduced section 1.3
<i>DWI</i>	Diffusion Weighted Imaging, introduced section 1.3
<i>ROI</i>	Region of Interest, introduced section 1.3
<i>SPHARM</i>	Spherical Harmonics, introduced section 2.1.1
<i>SPM</i>	Statistical Parametric Mapping, introduced section 2.1.2
<i>FMRIB</i>	Oxford Center for Functional MRI of the Brain, introduced section 2.1.2
<i>FSL</i>	FMRIB Software Library, introduced section 2.1.2
<i>PCA</i>	Principal Components Analysis, introduced section 2.2
<i>ASM</i>	Active Shape Models, introduced section 2.2
<i>SD</i>	Sulcal Depth, introduced section 4.3

Chapter 1

Introduction

The variability of anatomical structures among individuals is large within anatomical populations. This variability makes it necessary to use statistical modeling techniques to study shape similarities and to assess deviations from the healthy range of variability. For instance, studying the local cortical thickness measurements of the brain is a common tool for studying many medical conditions such as autism and Alzheimer’s disease in humans. Statistical analysis of anatomical objects is therefore becoming increasingly important in segmentation, analysis and interpretation of medical datasets.

The construction of such statistical models requires the ability to compute local shape differences among similar objects, which introduces the problem of finding corresponding points across the population. Consistent computation of corresponding points on 3D anatomical surfaces is a difficult task, since manually choosing landmark points not only is cumbersome, but also does not yield a satisfyingly dense correspondence map. It should also be noted that no generic “ground truth” definition of dense correspondence exists across different anatomical surfaces. The choice of particular correspondence metric must, therefore, be application-driven.

In this dissertation, I present a framework for finding corresponding points on populations of anatomical surfaces based on segmented medical images. Particular

emphasis is given to the human cortical surface using Magnetic Resonance Imaging (MRI) scans. The correspondence computation on the cortex is a highly challenging problem due to the convoluted geometry of the brain and the high variability of folding patterns across subjects. Using mere spatial locations of surface points produces a weak and inadequate correspondence map. My work allows the use of additional local information, called ‘features’ throughout this manuscript, for computing correspondence. These features can be structural, such as sulcal depth measurements, as well as nonstructural, such as connectivity maps computed from Diffusion Tensor Imaging (DTI) images, or vessel structure extracted from Magnetic Resonance Angiography (MRA) images. The particular choice of features should be determined by the target applications. This dissertation in particular explores various features that can be extracted from DTI scans.

1.1 Correspondence of local features

Previous population-based correspondence computation methods, such as Minimum Description Length (MDL) [1, 2], optimize correspondence by minimizing the covariance of the sample locations across the population. These methods work well for objects of simple enough geometry, such as caudates, but are inadequate with objects of complex geometry with rapidly changing curvature values across the surface, such as femoral heads or striata.

I propose to solve this problem by introducing local curvature measurements into the objective function being optimized. The curvature function variation contains local shape information that is lacking in the spatial positions of point samples; therefore, adding the local curvature to the objective function stabilizes the optimization process. I propose to extend the traditional MDL algorithm to allow for the inclusion of an arbitrary number of local features in the objective function. This is done by

replacing the data matrix in MDL, which contains the spatial location of each sample point of each object in the population, by a matrix that contains the desired local feature value at each sample point of each object.

Using this framework, I show that using a version of MDL based on solely curvature values leads to poor correspondence, since the curvature values are not unique across the object surfaces. However, using a combination of spatial position and local curvature improves the correspondence quality significantly for objects of complex geometry, and it gives similar results to traditional (location-only) MDL for objects of simple geometry. In this context, correspondence quality is measured by using the generalization and specificity properties of the statistical model of the population that results from the correspondence optimization, as discussed later in detail.

These results show that there is room for improvement of the correspondence quality by exploring local features other than spatial location. My framework allows for the use of any local feature whose absolute difference defines a metric in the feature space, as long as the feature values at each sample point are provided as input to the system.

The choice of the particular feature set to be used should be made based on the application context; one of the novel contributions of this dissertation is the exploration of possible features to be used for optimizing the correspondence of the human cortex. The ideal feature set would provide enough variability across the object surface and across the population.

1.2 Framework for correspondence of local features without parametrical mapping

Even though using a parametric approach like MDL has the advantage of relative ease of imposing surface constraints (such as no foldings), the MDL algorithm has

limited applicability to human cortex. The available implementations of the algorithm rely on a spherical parameterization of the surfaces, which is computationally expensive to obtain for the cortical surface (defined as the white matter (WM) - gray matter (GM) boundary) given the complex geometry. It also is computationally very expensive for high resolution meshes that are necessary for representing the cortical surface. Therefore, the current formulation of MDL is not suitable for the cortical correspondence optimization. However, it is a well-known fact in information theory that MDL is, in general, equivalent to minimum entropy ($\min \log|\Sigma + \alpha I|$, where Σ is the covariance matrix of the population and αI is a diagonal regularization matrix that introduces a lower bound α to the eigenvalues of the covariance matrix). Therefore, I chose to use an entropy-based dynamic particle framework introduced by Cates et al.[3, 4] as my starting point. The entropy-based correspondence scheme does not require a spherical topology and is much more computationally efficient. I propose a solution to the problems caused by the geometric complexity of the cortex in Section 1.2.1.

The main idea for the entropy-based correspondence method is to construct a point-based sampling of the shape ensemble that simultaneously maximizes a combination of the geometric accuracy and the statistical simplicity of the model. Surface point samples, which also define the shape-to-shape correspondences, are modeled as sets of dynamic particles that are constrained to lie on a set of implicit surfaces. Sample positions are optimized by gradient descent on an energy function that balances the negative entropy of the points' distribution on each shape, which represents an even sampling of the individual surfaces, with the positive entropy of the ensemble of shapes, which represents a high similarity of corresponding points across the population.

I extend this method to incorporate local feature measurements on the surface to provide a general correspondence definition. This is done by introducing the local

feature values in the ensemble entropy term, in lieu of the spatial locations. The surface entropy term remains the same, since this term ensures an even sampling of the surfaces regardless of the local feature values. The incorporation of the local features in the particle framework requires a modification to the ensemble entropy term. Consequently, the associated gradient has to change, which can be accomplished via the chain rule.

1.2.1 Application to the human brain: dealing with the geometric challenges of the cortex

One of the main constraints of the entropy-based correspondence method is that it assumes the particles exist on local tangent planes of the surface. This assumption makes it possible to avoid the costly computation of geodesic distances on the surface. However, the assumption becomes problematic for surfaces of a highly convoluted geometry, such as the human cortex. I overcome this problem by first ‘inflating’ the cortex to obtain a less convoluted surface. The particles live and interact on this blob-like surface; however, the local feature values, such as curvature and sulcal depth, are still associated with the original surface. A one-to-one correspondence between the original cortex surface and the inflated surface is therefore needed.

A set of automated tools, distributed as part of the FreeSurfer package [5, 6, 7, 8, 9, 10], are used to inflate the cortical surface, as well as to perform surface reconstruction. However, any other surface reconstruction and inflation pipeline could easily be used to replace FreeSurfer, since my algorithm does not depend on the specific FreeSurfer methodology.

The FreeSurfer inflation process is such that points that lie in convex regions move inwards while points in concave regions move outwards over time. Therefore, the average convexity/concavity of the surface over a region, also referred to as sulcal depth, can be computed as the integral of the normal movement of a point during

inflation. It should be noted that sulcal depth captures the high level foldings of the cortical surface but is relatively insensitive to the smaller folds; this property makes sulcal depth an attractive correspondence metric since it is relatively stable across individuals, in addition to being available across the whole cortical surface.

1.3 Using white matter fiber structure for cortical correspondence

The choice of suitable local features is central to the quality of correspondence results. For the cortical correspondence problem, I propose to use a feature set derived from Diffusion Tensor Imaging (DTI) scans of the subjects in order to incorporate available knowledge about the white matter (WM) fiber tracts of the brain in addition to the structural features discussed earlier. Structural MRI scans show white matter homogeneously, such that it is impossible to infer the orientation of the fiber tracts within each voxel. The understanding of the white matter structure, however, can be largely improved by additional information on fiber tracts that can be fully automatically extracted from DTI data.

The main challenge is to find a suitable mapping of the fiber tract structure to the cortical surface. Probabilistic connectivity maps, which represent for each voxel on the cortical surface the probability of its being connected via fiber tracts to a given region of interest (ROI), is the proposed solution to this problem.

Note that both low-resolution features that vary on a larger scale and spatial features that are useful for identifying higher-resolution structures are necessary. Sulcal depth, a by-product of the cortical inflation algorithm, as well as the proposed DTI-derived metrics are used as low-resolution features, whereas spatial location and curvature are used as higher-resolution features.

1.3.1 Probabilistic connectivity

The probabilistic connectivity maps are obtained via stochastic tractography. For every voxel in the white matter, the connectivity probability to various ROI's is computed via a Monte Carlo approach. In this dissertation, I use a stochastic tractography implementation, described by Ngo [11], based on a modification of Friman's algorithm [12, 13]. In this approach, fiber tracts are modeled as sequences of unit vectors whose orientation is determined by sampling a posterior probability distribution. The posterior distribution is proportional to a prior of the fiber orientation multiplied by the likelihood of the orientation given the Diffusion Weighted Imaging (DWI) data. The tracking stops when the tract either reaches a voxel with a low probability of belonging to the white matter, or it exceeds a predetermined maximum length, or it makes an improbably sharp turn. In order to estimate the probabilistic connectivity with appropriate accuracy, a high number of sample fibers need to be tracked from each voxel included in the input ROI. The probabilistic connectivity of a voxel to the ROI is then defined as the ratio of fiber samples that travel through a voxel to the total number of samples.

The connectivity to each separate ROI is represented as a separate feature in the particle framework. It should be noted that the spatial features used in the correspondence (such as the sulcal depth) have to be normalized to match the range of values of the connectivity probabilities, in order to avoid assigning a heavier weight to the spatial features.

1.3.2 Mapping the connectivity to the cortical surface

The stochastic tractography algorithm provides connectivity values, but the tracking often stops short of the white matter boundary, due to a low degree of anisotropy near the surface as well as noise in the DWI signal. For this reason, a method of mapping the connectivity probability to the white matter surface is necessary




Figure 1.1: Pipeline overview. I use T1 images to generate white matter (WM) surfaces and inflated cortical surfaces, as well as local sulcal depth. Selected ROI's and the DWI image are input to the stochastic tractography (ST) algorithm. WM surface deflated using proposed algorithm is used to construct connectivity maps on the surface from ST results. Inflated cortical surfaces and the connectivity maps are used to optimize correspondence. Note that there are three different surfaces representing the cortex. The original surface is used for the computation of the geometrical features, such as sulcal depth. The inflated surface is used for the computation of the particle inter-distances. The deflated surface is used for evaluating the connectivity probabilities.

after the stochastic tractography algorithm is executed. The connectivity probability values assigned by the stochastic tractography algorithm strongly correlate with sulcal depth. Directly using these probabilities for correspondence optimization is therefore not appropriate.

I propose a surface evolution algorithm for this purpose. In this method, I compute a new, deflated surface inwards from the white matter boundary. This gives a deflated surface voxel corresponding to each cortical surface voxel. Then, I define the projected connectivity probability at each cortical voxel as the connectivity probability value at the corresponding inner-curve voxel. Figure 1.1 summarizes the full pipeline.

1.4 Evaluation of correspondence quality

In order to evaluate the proposed framework, I apply my proposed techniques to a small set of clinical studies and compare the results with other existing algorithms. This is for evaluation purposes only, in order to compare the performance of my algorithm with others. Thorough validation studies are outside the scope of this dissertation.

Metrics for assessing correspondence quality are necessary to compare the different correspondence algorithms. For this purpose, I analyze the generalization and specificity properties of the resulting shape models. Lower variability as well as better generalization and specificity properties point to an improved correspondence across the population. I use the cortical thickness measurements to compute the generalization and specificity for the cortical datasets, instead of the surface sample locations. Location-based analysis is considered biased since both my technique and many other algorithms use the surface sample locations for the optimization; the cortical thickness provides an unbiased measurement more suitable for evaluation. I also present visual assessments of correspondence quality where suitable.

Using my proposed methodology, one should expect to find improved correspondence over certain regions and little improvement in other regions where no relevant additional local information is used. Furthermore, it would be up to each individual application to define what regions are important for the application’s context, and what additional data can be used to improve the correspondence in these critical regions. For example, the regions that are strongly identifiable by fiber tract connections to subcortical regions are expected to have significantly improved correspondence if the connectivity probability is designated as a feature.

1.5 Thesis and contributions

Thesis: Statistical shape analysis of anatomical structures, which is essential to understanding the structural changes in anatomy in various stages of growth or disease, requires establishing accurate correspondence across object populations. However, anatomical correspondence is rarely a direct result of spatial proximity of sample points on the surface. A generalized correspondence framework that incorporates the similarity of non-spatial local features provides a more accurate correspondence of sample points across populations of surfaces. In particular, incorporating features based on cortical geometry as well as the fiber connectivity of the white matter significantly improves correspondence of the human cortical surfaces.

The contributions of this dissertation are as follows:

1. I demonstrate that the use of an approach allowing for the incorporation of arbitrary local features into the similarity metric to be used for correspondence optimization enhances correspondence, as measured by objective evaluation criteria.
2. I present a novel parametric groupwise correspondence optimization method that allows using arbitrary local features for establishing correspondence.
3. I demonstrate that using geometric information, such as local curvature measures, as additional local features improves correspondence quality when the objects in the population exhibit complex geometry.
4. I present a novel nonparametric groupwise correspondence optimization method that allows using arbitrary local features for establishing correspondence.
5. I show that this nonparametric groupwise correspondence technique can be applied to the human cortex despite the geometric challenges presented by the convoluted surface; inflation of the surfaces as a preprocessing step solves this

problem by producing surfaces smooth enough to avoid these challenges. Furthermore, I show that cortical correspondence significantly improves when sulcal depth is used as an additional local feature.

6. I present a novel framework for integrating white matter fiber connectivity information into cortical correspondence, the first such method that uses fiber connectivity patterns to establish structural correspondence. To this end, I compute probabilistic connectivity maps from diffusion weighted images via a stochastic tractography algorithm. I project these connectivity values to the cortical surface by a new cortical deflation algorithm. I present empirical evidence showing that using connectivity features enhances cortical correspondence.
7. I develop open-source software that implements all the above techniques, as well as a visualization tool that allows qualitative examination of the surfaces, the local features associated with them and the surface samples used in the correspondence algorithm.

1.6 Overview of chapters

The remainder of this dissertation is organized as follows.

Chapter 2 contains an overview of concepts and existing methodology regarding the correspondence problem as well as techniques for evaluating correspondence quality.

Chapter 3 describes a novel methodology for integrating local features into the optimization of parameter-based shape correspondence and demonstrates, using applications to a variety of clinical datasets that using local features can strongly improve correspondence. The suggested methodology is an extension to the traditional Minimum Description Length algorithm.

Chapter 4 describes the entropy-based particle correspondence algorithm, with

and without local features. This methodology is no longer dependent on a particular parameterization of the surface. This chapter also discusses some issues regarding the application of this technique to the human cortical surface, and it presents results that demonstrate that the use of even simple geometric local features (such as sulcal depth) is beneficial to correspondence quality.

Chapter 5 discusses the proposed methodology for integrating DTI-based connectivity information into cortical correspondence. This technique computes connectivity probabilities via stochastic tractography and applies a surface evolution algorithm to deflate the cortical surface in order to map the connectivity probabilities to the cortical surface. It presents experimental results that show, via an evaluation based on cortical thickness, that the additional knowledge of white matter structure significantly improves correspondence quality.

Chapter 6 presents a discussion of the contributions of this thesis and future research directions.

Chapter 2

Shape Correspondence

This chapter provides a summary of various existing techniques to solve the correspondence problem. These techniques fall in two main categories: Given a population of objects, pairwise correspondence methods establish the correspondence between each object and an atlas; the correspondence across multiple objects follows by transitivity. Groupwise correspondence methods, on the other hand, consider the entire population at once in an attempt to capture the variability in the population. Pairwise correspondence methods, unlike groupwise approaches, fail to incorporate information from the entire population and treat each surface separately, which can lead to suboptimal correspondence results. In both cases, the correspondence computation is typically formulated as an optimization problem with an objective function that incorporates a similarity measure between the objects and often regularization terms.

In order to compare various correspondence optimization schemes, one needs a suitable correspondence evaluation method. For this purpose, I review existing algorithms for objectively evaluating correspondence quality in the second part of this chapter.

2.1 Pairwise Correspondence

Pairwise correspondence methods aim to optimize the correspondence between each object in a given population and either a labeled atlas or one of the objects in the population chosen to serve as a template. Correspondence optimization can be done based on either surface representations or volumes. Surface-based correspondence methods typically lend more weight to geometrical properties of the objects, whereas volume-based methods focus on image intensities. Further classification can be done based on whether the algorithm aims for an exact match or an approximate match enforced by a soft penalty.

2.1.1 Surface-based pairwise correspondence

In most surface-based schemes, correspondence is defined through a parameterization of both objects, such that points in each object with the same parameter space coordinates correspond (see Fig. 2.1). Therefore, it is necessary to compute a one-to-one mapping between each object and a standard parameter space.

During the optimization stage, there are two possible ways of manipulating the correspondence: either the vertices on the surface can be moved around while keeping the parameterization fixed, or the parameterizations can be manipulated while keeping the surface vertices fixed. However, manipulating the surface vertices directly is a difficult task, since it would then be necessary not only to confine the vertices to the surface as they move in $\mathbb{R}^3$ but also to construct a mapping of the surface onto itself at each iteration, a task far from trivial for arbitrary 3D surfaces. Most algorithms therefore choose to manipulate the parameterizations rather than the surfaces themselves.

A number of approaches use various shape-based information as additional cues in the correspondence optimization. This is similar to the approach I suggest for using

local features for improving correspondence; however, the existing techniques focus on a preselected feature or set of features, whereas I propose a generalized framework where the user can determine what features are relevant for the particular application context.

Spherical harmonics

The spherical harmonics (*SPHARM*) description, introduced in [14], is commonly used as a parameterization-based correspondence scheme. Here, a continuous one-to-one mapping from each surface to the unit sphere is computed. The parameterization is defined such that it is area preserving and distortion minimizing, using a constrained optimization method. Each object is then described as a weighted sum of spherical harmonics basis functions (see Appendix A for details). The correspondence is established by rotating the parameter mesh such that the axes of their first-order spherical harmonics, which are ellipsoidal, coincide with the coordinate axes in the parameter space. Since each object can thus be parameterized without any knowledge of the others in the dataset, data sets that have significant shape variability become problematic, because the SPHARM method does not have a proper means of optimizing shape similarity but rather focuses on area preservation and minimized distortion during the parameterization.

Another major shortcoming of the SPHARM correspondence is the poor handling of objects that are rotationally symmetric around the major axes with respect to the first-order ellipsoid. This often happens if the second and third axis have similar sizes. Brechbühler argues that using information from higher-order harmonics can disambiguate such cases.

These problems can be at least partially overcome by introducing a ‘flip template’, an additional object in SPHARM representation that can be used to disambiguate the orientation of the parameterization [15]. This is achieved by testing all possi-




Figure 2.1: Example of parameterization-based correspondence. A parameterization is computed for objects A and B individually. The correspondence is defined via the parameter space, such that points with the same parameterization correspond together when the correspondence optimization is completed. Color map shows the ϕ parameter, the longitude in the spherical representation.

ble flips of the parameterization and selecting the one whose reconstruction has the minimum distance to the flip template. The use of the flip template puts SPHARM correspondence in the pairwise category. An additional limitation of this method is that parameterization-based correspondence schemes in general are restricted to objects of a given topology (e.g. spherical in the case of SPHARM).

Shape-based nonrigid correspondence

One of the earlier works that use geometrical measures to align two shapes is the shape-based nonrigid correspondence scheme of Tagare et al. [16]. In this simple 2D correspondence algorithm, the local curvature is used as a similarity metric, and the total squared difference of local curvature is minimized along with an arc length-based regularization term. Although this method is intuitive and seems to work well for simple 2D populations such as hand outlines and the endocardium of the heart, it relies on a toroidal parameterization of the two objects in 2D and therefore the extension to 3D is not straightforward.

Geometry driven multispectral optical flow

Tosun et al. [17] propose to use geometrical measures to align an atlas and a subject surface in a more sophisticated framework in 3D. They propose to use two curvature-related measures introduced by Koenderink [18] as the similarity metric for the correspondence. These measures, the shape index S and the curvedness C , decouple the shape and the size of curved surfaces and will be further discussed in Section 3.3.3. Tosun computes these curvature measures on a partially inflated cortical surface in order to capture only the geometry of the most prominent anatomical features to allow meaningful comparison among different individuals. This partially inflated cortical surface is obtained by a mean curvature based evolution algorithm with a preselected curvature threshold as the stopping condition. Once the shape measure maps are computed for the subject and the atlas brain, a multispectral optical flow algorithm is used to warp the subject cortical surface into the atlas using spherical parameterizations of the surfaces.

Tosun uses a surface-based iterated closest point (*ICP*) matching scheme [19] as an initialization point to her correspondence algorithm. This ensures that key anatomical landmarks such as major sulci are mapped to reasonably close locations on the parameter space and therefore provides a good initialization. The variational problem to estimate the optical flow field is solved using a Euclidean framework, and a gradient-descent optimization is applied.

Even though satisfactory correspondence results can be achieved by this algorithm, there are several shortcomings. The method completely ignores spatial location information as a correspondence metric and therefore relies on a very good initialization to avoid false correspondences on a global scale. Additionally, as all pairwise correspondence methods, the optical flow method fails to incorporate information from the entire population, which can lead to suboptimal correspondence results.

Parameter space warping

Meier et al. [20] propose a method that represents a pair of objects using SPHARM description. However, unlike the original SPHARM correspondence, they warp the parameter space to optimize the correspondence between the two objects instead of relying on the first-order ellipsoid alignment. The objective function is a similarity metric based on Euclidean distances, normal directions and Koenderink’s shape index S [18]. The SPHARM representation allows both a hierarchical optimization approach and the robust computation of differential features since no additional discretization error is introduced.

As with other spherical parameterization-based techniques, the method is limited to objects of spherical topology. Additionally, both the initial SPHARM computation as well as the warping procedure itself are computationally expensive and limit the practical applicability of the algorithm. For the cortical surface dataset reported in the manuscript, 24 harmonics were used, with 9126 vertices on the whole surface. While this may be an adequate representation of the surface for $5mm$ -thick MRI slices used in the study, it is far from satisfactory for higher resolution data commonly available nowadays.

FreeSurfer

FreeSurfer [5, 6, 7, 8, 9, 10] is an image analysis suite for brain studies. Aside from cortical reconstruction and inflation tools, FreeSurfer also offers a spherical surface-based coordinate system based on a cortical correspondence optimization method. This method, described in detail in [7], is based on nonrigidly aligning each individual’s folding patterns with an atlas. The folding pattern in this context is quantified by the sulcal depth of the surfaces.

The algorithm is initialized by first mapping the reconstructed cortical surface onto a sphere using a maximally isometric transformation. Then, the spherical parameter-

izations for the individual and the atlas are morphed into register by a combination of sulcal depth alignment and isometry-preserving forces. The sulcal depth alignment is achieved by minimizing the mean squared difference between the average sulcal depth computed from a fixed training set and that of the individual modulated by the variance across the training set.

Fischl further argues that simply maximizing the sulcal depth alignment is not enough to prevent folds and distortions. He therefore introduces two additional energy terms to the objective function of the optimization. The first term is for the preservation of local distances, which gives the surface some local stiffness and discourages excessive shear. The second term is for area preservation and aims to prevent excessive compression or expansion. These terms are weighted against the sulcal depth alignment term by free parameters.

Like the geometrically driven optical flow method described above, the FreeSurfer correspondence method ignores the spatial location of the vertices in its similarity metric definition and therefore has to resort to a nontrivial initialization procedure as well as the additional regularization terms described above. Additionally, it can produce suboptimal correspondence results since it doesn't capture the variability of the whole population.

BrainVoyager

A multiscale correspondence method based on cortical curvature is used as part of the BrainVoyager QX software [21]. The curvature of the cortex typically reflects the gyral/sulcal folding pattern of the brain, similar to the sulcal depth metric used by FreeSurfer. The algorithm first computes a spherical parameterization of the cortical surface to simplify vertex location manipulation during the optimization stage. Then, a gradient-descent optimization is used to minimize the mean squared differences between the curvatures of the source and target brains. The target typically is an

atlas or a designated template object in the population; however, BrainVoyager allows the use of a ‘moving target’, defined as the population’s average, updated after each iteration. The optimization is performed by an iterative coarse-to-fine procedure by means of different levels of smoothing of the curvature maps.

Like FreeSurfer, BrainVoyager bases its correspondence definition entirely on the curvature, which can lead to suboptimal correspondence results if the initialization is not very good. However, the ‘moving target’ approach makes it a pseudo-groupwise technique, which is an advantage over typical pairwise algorithms.

BrainVisa

Cachier et al. [22] propose an intensity and geometric feature based registration algorithm. Given a brain image, the sulci are first automatically extracted and labelled using a neural network trained on a manually labelled set. The sulcal border and the sulcal bottom (the edge of the sulcus deep in the brain) are extracted from these sulci. The sulcal bottoms with the same label on two brains are then matched with a nonparametric approach based on an objective function that has similarity terms for the image intensities at landmark locations as well as location on the sulcus. The image intensities are used to overcome problems such as sulci of different topologies across different brains, as well as to increase the robustness of the automatic sulcus labeling. Although this approach produces satisfying correspondence results, it is strongly dependent on the initialization step, which is the automatic extraction and labeling of the sulci. Furthermore, it is not clear how to extend the correspondence outside the sulcal bottoms.

Shape-Based Correspondence Using Geodesics and Local Geometry

Wang et al. [23] propose a pairwise correspondence method that uses geodesic distance and surface curvature to capture the complex geometry of the cortical surface.

It is necessary to define a sparse set of anatomical landmark points on the atlas manually to use the algorithm. The points corresponding to these on the subject brain are generated by a nonparametric shape-based matching procedure via an objective function based on Euclidean distance, Koenderink’s curvedness metric, and a surface normal match measure. Next, geodesic interpolation of these initial points is used to obtain a dense set of corresponding points between the subject brain and the atlas.

This method is unique in its usage of geodesic interpolation of sparse set of corresponding points. However, it requires the manual selection of the initial landmark points on the atlas (69 landmarks were necessary for the study reported in [23]), which can be tedious. It is also unclear how well the finer details of the folding pattern are matched with this method, since these are not explicitly identified by the manual landmarking.

2.1.2 Volume-based pairwise correspondence

A fundamentally different approach to correspondence optimization is via the registration of an image volume to an atlas. The correspondence on the surface can be obtained by applying the warp field implied by the volume registration. Applied to brain images, the main advantage of these methods is that both the cortical surface as well as subcortical structures can be treated in a unified framework. A full discussion of volumetric registration is beyond the scope of this dissertation; therefore, this section is limited to a review of a representative selection of the methods used in neuroimaging.

Talairach registration procedure is a classical volume-based correspondence method [24]. The Talairach coordinate system is defined by ensuring that the anterior commissure and the posterior commissure lie on a horizontal line. Since these two landmarks lie on the midsagittal plane, the coordinate system is completely defined by requiring this plane to be vertical. The registration procedure therefore only requires a few

parameters to represent the entire transformation, which can lead to poor anatomical accuracy in the cortex. In fact, several studies have shown that the between-subject variability of landmarks after Talairach registration can be on the order of several centimeters [25, 26]. This means that the location estimations of many small cortical areas can be severely mistaken.

Another volume-based correspondence approach is the Statistical Parametric Mapping (*SPM*) spatial normalization method, which is part of the SPM software toolkit for analysis of functional imaging data [27]. The registration method used by SPM5 (and subsequent SPM versions) aims to match each skull-stripped image to the skull-stripped reference or atlas image. The registration involves minimizing the mean squared difference between the images that have been presmoothed with a Gaussian kernel. The first step of the registration is the estimation of a 12-parameter affine transformation. Excessive zooms and shears are penalized via a regularization term. The second step involves nonlinear registration which targets shape differences between the two brains, which the affine transformation cannot account for. The warping is modeled by a linear combination of low-frequency cosine transform basis functions. Regularization is obtained by minimizing the membrane energy of the warp [27].

Another volume-based correspondence method is proposed by the Oxford Centre for Functional MRI of the Brain, FMRIB. Their software package FSL (FMRIB Software Library) provides two different methods: a linear multi-resolution global optimization method [28], as well as a more recently introduced nonlinear registration method [29]. The nonlinear registration method is similar to the SPM registration but uses cubic b-splines rather than cosine transforms for parameterizing the deformation.

Although the SPM and FSL registrations provide a more accurate correspondence than the Talairach registration procedure, they nonetheless have similar limitations. Because the deformations are defined by only a few hundred parameters, deformations

for higher levels of detail cannot be modeled in these frameworks.

Christensen et al. [30] propose a much higher-dimensional deformation field for the registration to accomodate shape differences between the atlas and images of other brains. This is accomplished by defining probabilistic transformations on the atlas coordinate system modeled by the physical properties of viscous fluids. Although a near-perfect match can be obtained in the image intensities of different brains with a high-dimensional deformable registration algorithm, the alignment of the sulcal and gyral patterns is not ensured, because the cortical geometry is ignored (as in all the other volumetric methods discussed so far). Moreover, since gyral and sulcal landmarks are typically accurate predictors of the location of functional areas of the brain, it seems appropriate to use these folding patterns as features to drive the registration of the cortical surfaces, rather than image intensities.

The HAMMER (Hierarchical Attribute Matching Mechanism for Elastic Registration) method [31] addresses the issue of taking into account the underlying anatomy rather than simply matching image intensities across volumes. For this purpose, HAMMER uses an attribute vector associated with each voxel to drive the elastic registration. These vectors contain geometric information of different scales and they can help differentiate between different parts of the anatomy that might otherwise be indistinguishable if only image intensities are considered. However, this technique still has the limitations of pairwise correspondence methods; furthermore, in the current formulation, the user is not allowed to define new attributes but rather is forced to use a predetermined set of 13 attributes.

2.2 Groupwise Correspondence

Since their introduction by Cootes and Taylor [32], active shape models (*ASM*) have become very popular in medical imaging. This method is in fact one of the earliest to

introduce the problem of choosing correspondence points for a population, along with Bookstein’s work [33]. The solution proposed by Cootes and Taylor is to manually choose landmarks and to perform a generalized full Procrustes alignment on the entire population to align them with each other. Generalized full Procrustes alignment is the alignment of n objects using translations, rotations and scaling such that the sum of distances between all pairs of objects is minimized. The ASM method therefore consists of manually choosing landmarks, aligning them by considering the entire population, and finally performing a Principal Components Analysis (*PCA*) on the landmark locations. This can be viewed as a first iteration of a correspondence method, where the only additional step would be the optimization of the landmark positions. This seminal work thus paved the way for the groupwise correspondence approaches that were subsequently proposed.

2.2.1 Determinant of the Covariance Matrix

Kotcheff et al. [34] propose to automatically find correspondence points by optimizing an objective function that leads to compact and specific models. They argue that the appropriate objective function is the determinant of the covariance matrix of the landmark locations, and they optimize this objective function via a genetic algorithm that manipulates the parameterization and pose of the objects in the parameterization. This leads to better correspondence than some of the earlier pairwise algorithms. However, as Davies et al. later pointed out [2], the choice of the determinant of the covariance matrix as the objective function is not clearly justified and is therefore solely based on intuition.

2.2.2 Minimum Description Length

The MDL method, introduced by Davies et al. [35], is an information theoretic approach to the correspondence problem. The main idea is that the *simplest* description

of a population is the best; simplicity is measured in terms of the length of the code to transmit the data as well as the model parameters. Ward et al. [36] extend the method to medial object representations. Chapter 3 presents a novel approach to MDL formulation, and detailed reviews of both the original algorithm and a variety of techniques attempting to improve it will be provided in Section 3.1.1. Cates et al. [3, 4] propose an entropy-based formulation that can be shown to be equivalent to MDL. This nonparametric approach forms the basis for the framework presented in Chapter 4, and a detailed review of the technique proposed by Cates will be presented in Section 4.1.

2.3 Correspondence Evaluation

Objective methods for evaluating correspondence quality are necessary in order to compare various correspondence optimization schemes. This section reviews major methods of assessing correspondence quality.

2.3.1 Distance to Manual Landmarks

Since the ‘ground truth’ for correspondence of biological objects is not known, it is common to use a small set of anatomical landmarks selected manually by a human expert on each object for evaluation purposes. The human expert knowledge is thus considered the ‘gold standard’ [37]. Typically, the mean distance between the manual landmarks and the landmark points in a template is used to quantify the performance of each correspondence algorithm. Moreover, reproducibility error of the manual landmark selection, typically in the range of a few millimeters [37], should be taken into account to judge the results.

2.3.2 Jaccard Coefficient Difference

Munsell et al. [38] propose to overcome the lack of ground truth knowledge in a benchmark study. Given an arbitrary statistical shape model, they generate a large set of new shape instances. This new data set can be input to different shape correspondence algorithms after resampling and the addition of random affine transformations. The correspondence performance can thus be objectively evaluated since the ‘ground truth’ for this data set can be defined via the underlying shape model.

Munsell proposes two metrics to evaluate correspondence quality based on the Jaccard-coefficient difference. The Jaccard-coefficient difference of two shape contours is defined as one minus the ratio of the area enclosed by their intersection and the area enclosed by their union.

The first evaluation metric is the bipartite-matching difference, which is the total Jaccard-coefficient difference between each shape S_i in the original set of shape contours and the shape contour with the minimum Jaccard-coefficient difference to S_i in the contour set obtained via the correspondence algorithm. A small value indicates that the shapes are closely similar for the whole population and therefore that the correspondence method used is satisfactory.

The second metric is a statistical test applied to the minimum spanning tree of the fully connected graph of all shape contours, where the edge weights are defined by the Jaccard-coefficient difference of the contours represented by the neighboring vertices. This provides an estimate of the probability that the two sets of continuous shape contours are from the same shape space.

To apply Munsell’s metrics to real anatomical structures, one can in theory use a very large number of real shape contours instead of the input statistical shape model. However, obtaining such a large dataset is difficult and labor-intensive, especially in 3D. Furthermore, there is no guarantee that the entire shape space will be adequately sampled with this approach.

2.3.3 Generalization, Specificity and Compactness

Davies [35] proposes three evaluation metrics that measure the properties an optimal statistical shape model should have: generalization, specificity and compactness. These three associated metrics are designed to be used for comparing different correspondence methods applied to similarly sized datasets, since their values are dependent on the number of surfaces.

Generalization describes a model’s capability to represent unseen instances of the class of objects being studied. This is a useful metric since it penalizes models that have been over-fitted to the training set. In practice, the generalization $G(M)$ of a model can be computed by a leave-one-out algorithm. For each object in the training set, a model is constructed by Principal Components Analysis (PCA) using the remaining $n - 1$ objects. The model is used to reconstruct the left-out object using M principal modes of variation, and the reconstruction error is computed. This process is repeated for all n objects and the reconstruction error is averaged. Formally,

$$G(M) = \frac{1}{n} \sum_{i=1}^n |x_i - x'_i(M)|^2, \quad (2.1)$$

$$\text{with standard deviation} \quad \sigma_{G(M)} = \frac{\sigma}{\sqrt{n-1}}, \quad (2.2)$$

where x_i is the location matrix for the i^{th} object, $x'_i(M)$ is the reconstructed object using M modes of variation as described above, and σ is the sample standard deviation of $G(M)$. Since the generalization is a measure of average error, lower values of generalization are desirable. The standard deviation is necessary in order to reason about the significance of differences in $G(M)$ for different correspondence methods.

Specificity captures the fact that a good model should only generate object instances similar to those in the training set. In practice, a PCA model is computed based on the training set and a large number N of random instances are generated from this model using M modes of variation. For each new object, its distance to the

closest object in the training set is computed; this distance is averaged for all new objects to compute the specificity $S(M)$. Formally,

$$S(M) = \frac{1}{N} \sum_{i=1}^N |\text{new}_i(M) - \text{nearest}_i|^2, \quad (2.3)$$

$$\text{with standard deviation} \quad \sigma_{S(M)} = \frac{\sigma}{\sqrt{N-1}}, \quad (2.4)$$

where new_i is the location matrix of the i^{th} new random object, nearest_i is the location matrix of the object in the training set with the minimum distance to new_i , and σ is the sample standard deviation of $S(M)$. As in the case of the generalization, specificity is an average error measure and therefore low values are desirable.

Note that the case $M = 0$ corresponds to the population average. $G(0)$, therefore, measures how far, on the average, the individual shapes are from the population average. $S(0)$, on the other hand, measures the distance between the population average and the individual shape that is closest to that average.

The third evaluation metric Davies proposes is compactness. A compact model is one that has as little variance as possible and requires as few parameters as possible to represent an object. This property is captured by the compactness $C(M)$, defined as a cumulative variance:

$$C(M) = \sum_{i=1}^M \lambda^i, \quad (2.5)$$

$$\text{with standard deviation} \quad \sigma_{C(M)} = \sum_{i=1}^M \sqrt{\frac{2}{n}} \lambda^i, \quad (2.6)$$

where λ^i is the i^{th} eigenvalue in the PCA model, and n is the number of objects in the training set. Based on the definition of compactness, it is clear that a low value of $C(M)$ is desirable.

2.3.4 Goodness of Prediction ρ^2

Jeong [39] proposes a metric that captures the goodness of fit of a shape model to a data set. This fitness is evaluated by the squared correlation ρ^2 . Jeong demonstrates that this correlation reduces to the following formula:

$$\rho^2 = \frac{\sum_{i=1}^N d(\hat{m}_{i,\text{test}}, \bar{m}_{\text{train}})^2}{\sum_{i=1}^N d(m_{i,\text{test}}, \bar{m}_{\text{train}})^2} \quad (2.7)$$

where $\hat{m}_{i,\text{test}}$ is the projection of the test model $m_{i,\text{test}}$ on the shape space, $\bar{m}_{\text{train}}$ is the Frechet mean of the training set, and $d()$ is an appropriate distance metric (e.g. Euclidean distance for Cartesian space, geodesic distance for manifolds). ρ^2 can be interpreted as the amount of variation of a test set explained by the retained principal directions estimated by a training set. This can be used as a correspondence evaluation tool by performing N leave-one-out experiments. Thus, higher values of ρ^2 indicate better correspondence models.

Chapter 3

Parameterization-Based Group Correspondence

A natural way to establish correspondence across a population represented in a parameterized form is to manipulate the parameterization to optimize an objective function. In this chapter, I first review in detail the Minimum Description Length (MDL) approach to correspondence. As discussed in 2.2.2, MDL is a groupwise correspondence approach that uses ideas from information theory. I then demonstrate how this method can be extended to incorporate additional local features to effectively provide a generalized parameter-based correspondence method. Finally, I discuss some geometric metrics as local feature candidates and report results from several anatomical data sets.

3.1 Traditional MDL

The goal of the MDL technique is to build optimal statistical shape models from boundary representations of 2D or 3D object populations. The key idea is that the ‘best’ model will minimize the description length, or the length in bits of a description of the entire dataset, including both the model parameters and the encoded data. This leads to an objective function comprised of two terms, one that aims to minimize the

model complexity (description length of model parameters), and one that aims to ensure the quality of the fit between the model and the data (description length of encoded data).

3.1.1 Original Algorithm Formulation

The original MDL method for shape correspondence was introduced by Davies et al. [2] for 2D. Each shape is individually parameterized such that points on different surfaces that have the same parameterization correspond. The correspondence optimization task thus becomes a problem of finding the optimal parameterization.

Davies proposed to represent the shapes using the shape parameters from a Principal Components Analysis (PCA) of the population. Then, new shapes can be generated by choosing random values for the shape parameters in the range of the training set. If the correspondence is appropriate, this will provide an efficient statistical model of the data set; otherwise, illegal shapes can occur.

The objective function is the description length, which is the sum of the description length for the parameters of the PCA model ($DL_{parameters}$) and the description length of the data (DL_{data}). The training data is modeled with a 1D Gaussian distribution along each principal direction. Therefore, the parameters that need to be transmitted are the mean μ and the variance σ of the Gaussian distribution. The data is assumed to be bounded and quantized, with an upper bound R and a quantization parameter Δ ; the data description length is therefore the entropy associated with a bounded, quantized 1D Gaussian. Furthermore, the quantization parameter for the variance, δ , must be included in $DL_{parameters}$. Finally, the n principal direction vectors must be transmitted; however, the description length of these is constant for a given training set with a fixed number of sample points per object, and therefore it does not need to be considered in the objective function. Similarly, R and Δ can be omitted for a

given training set. Therefore, we have

$$DL_{parameters} = DL_{\mu} + DL_{\sigma} + DL_{\delta} \quad (3.1)$$

$$DL = DL_{parameters} + DL_{data} = \sum_{m=1}^n DL_m \quad (3.2)$$

with

$$DL_m = \begin{cases} f(R, \Delta, n) + (n-2)\ln(\sigma_m), & \text{if } \sigma_m \geq \sigma_{cut} \\ f(R, \Delta, n) + (n-2)\ln(\sigma_{cut}) + \frac{n+3}{2}((\frac{\sigma_m}{\sigma_{cut}})^2 - 1), & \text{if } \sigma_m < \sigma_{cut} \end{cases} \quad (3.3)$$

where σ_{cut} is the lower bound on the variances along each principal direction (σ_m), and f is a function that depends only on R, Δ and n (see [2] for the derivation of Eq. 3.3). In the limit as $\Delta \rightarrow 0$, σ_m approaches $\sqrt{n_p \lambda_m}$, where n_p is the number of samples per object and λ_m are the eigenvalues of the covariance matrix associated with the sample locations. It can be seen that the description length is closely related to the determinant of the covariance matrix.

In this original formulation, Davies uses a parameterization that is piecewise linear to ensure one-to-one and monotonic mappings and optimizes the MDL function by using a stochastic algorithm. To avoid the whole system to collapse to a trivial global minimum, he proposes to fix the parameterization of one object in the population.

Davies further extended [1] this basic formulation to 3D. An equal-area mapping to the sphere (as described by Brechbuhler [14]) is used to parameterize the shapes, which is the only major change from the 2D formulation. This parameterization is obtained by a nonlinear constrained optimization aiming for minimal edge length distortion via manipulation of the spherical coordinates of the vertices. An initial spherical parameterization based on a heat conduction model is used as initialization point to this optimization. Each face of the object mesh is constrained to have the same area. Davies claimed that any other spherical mapping method that does not

introduce surface folding or tearing would perform equally as well for this purpose, since MDL further optimizes the parameterization and the final result should not significantly depend on the initialization, at least in theory.

In this 3D formulation, a multi-resolution optimization is performed using Cauchy kernels to create symmetric θ transformations in order to manipulate the parameterization, by initially using only a few big Cauchy kernels and then using additional, smaller kernels. This reparameterization strategy causes the points near the kernel center to be spread over the sphere, while landmarks in other regions of the surface are compressed. Using a large number of kernels at different locations, the parameterization can be arbitrarily manipulated.

3.1.2 Gradient Descent Optimization of the MDL Function

The stochastic optimization method used in the original algorithm description is very time consuming. Heimann et al.[40] propose a gradient descent optimization scheme to resolve this issue. The PCA is computed using singular value decomposition on the data matrix $\mathbf{A}$, defined as $\mathbf{A} = \frac{1}{\sqrt{N-1}}(\mathbf{L} - \bar{\mathbf{L}})$, where $\mathbf{L}$ is the matrix that encodes the vertex locations for each object along successive columns, $\bar{\mathbf{L}}$ is a matrix with all columns set to the mean shape μ , and N is the number of objects in the population. A simplified version of the MDL objective function is then used, as proposed by Thodberg et al. [41]:

$$F = \sum_{m=1}^n F_m,$$

with

$$F_m = \begin{cases} 1 + \log(\lambda_m / \lambda_{cut}), & \text{if } \lambda_m \geq \lambda_{cut} \\ \lambda_m / \lambda_{cut}, & \text{if } \lambda_m < \lambda_{cut} \end{cases} \quad (3.4)$$

where λ_m are the eigenvalues of the covariance matrix associated with the sample locations and λ_{cut} is a free parameter that corresponds to σ_{cut} from the original MDL formulation.

Heimann also claims that the Cauchy kernels used by Davies to manipulate the parameterization are inefficient because adding one new kernel modifies all vertex positions. However, it is desirable to keep established correspondences stable. He proposes to confine the kernels to be strictly local instead by truncating them at a predefined distance from the kernel center. By decreasing the threshold distance as the optimization progresses, a hierarchical optimization effect is achieved, such that larger regions are handled first and finer details are handled last. In addition, the parameterization meshes are randomly rotated throughout the optimization to ensure that all regions of the sphere are treated equally by the Cauchy kernels, whose locations remain fixed in the parameter space.

This series of randomized parameterization rotations, combined with the hierarchical optimization approach, makes convergence to local minima unlikely. Consequently, the gradient descent approach produces significantly better correspondences than the original MDL formulation, as evidenced by better generalization and specificity properties as well as lower convergence values for the MDL cost function, which indeed suggest an optimization more robust to local minima. Furthermore, Heimann reports the gradient descent optimization convergences up to 5000 times faster than the original approach. However, the method nonetheless takes reportedly up to 20 hours for relatively small and simple datasets (10-20 subjects with less than 5000 vertices per surface); the computational cost therefore becomes prohibitive for highly detailed datasets such as human cortical surfaces.

3.1.3 Shape Image Based MDL Optimization Schemes

The computation bulk of the MDL method lies in the reparameterization step, which requires the interpolation of the mesh on the parameter sphere at each iteration. Twining et al. [42] argue that moving this interpolation from the spherical domain to a plane would significantly improve the correspondence computation time. For this purpose, the sphere is first mapped to an icosahedron, which is then cut open and flattened, as proposed by Praun et al. [43]. Twining calls these flattened images of the parameter sphere *shape images* and proposes to store spatial location information in them. An elastic registration of the shape images is performed to optimize the MDL function.

Although this method is significantly faster than traditional MDL implementations, it does nonetheless require a spherical parameterization as the initial input, which can be a limiting factor. The entropy-based particle correspondence method discussed in the next chapter offers fast computation times in addition to being parameterization-free, therefore allowing surfaces of arbitrary topologies.

3.2 Generalized MDL Correspondence

Although traditional location-based MDL (as discussed in Section 3.1.1) has been demonstrated to be a powerful approach for solving the groupwise correspondence problem in an automated fashion, it typically performs rather poorly on objects with complex geometry, when the surfaces in question are convoluted. For such surfaces, spatial proximity of vertices does not provide a distinctive enough correspondence objective function and can easily result in false correspondence results. A typical example is the case of vertices of a brain mesh lying on opposing sides of a sulcus, as illustrated in Figure 3.1.




Figure 3.1: Spatial proximity can be a false indicator of correspondence, as illustrated by points lying on opposing sides of a sulcus of the brain. Point A on the left brain is closer to point B on the right brain than it is to point C ($\overline{AB} < \overline{AC}$). However, it is intuitively clear that A is much more likely to correspond to C than to B, as C is located on the opposite bank of the sulcus. A’s position is replicated on the right brain for ease of comparison.

3.2.1 Using Local Features for Generalizing MDL

A natural solution to these issues is to make use of geometric information that goes beyond spatial location, such as local curvature. The MDL objective function can be extended to incorporate such local features, as we presented in [44, 45]. The extension I propose is applicable to not only geometric measures such as curvature, but to any local features whose absolute difference defines a metric in the feature space, as long as the feature values are available at each sample location. Depending on the application domain, appropriate local features might be proximity to major blood vessels, connectivity to other anatomical structures, image intensities, etc.

To use local features in the MDL computation rather than vertex locations, it is

necessary to substitute the data matrix L encoding the spatial locations in traditional MDL. In this alternative matrix the columns are the local feature values, such as curvature measurements of the object, instead of spatial locations. Therefore, the cost function for MDL must be modified to be based on the eigenvalues of the curvature matrix instead of the eigenvalues of the location matrix.

With this technique, it is possible to use local features that are high dimensional, such as different curvature measurements, or to even include the spatial location itself as a feature dimension, in addition to the other application-driven features.

3.3 Geometric Features

The choice of specific features for representing local shape is one of the issues in using geometrical features to improve MDL. In this section, I review some candidates for local geometrical features. The feature selection will be made based on some of the correspondence evaluation techniques reviewed in Section 2.3, in particular, the generalization and specificity metrics.

3.3.1 Principal Curvatures (κ_1, κ_2)

In differential geometry, the second fundamental form II is a quadratic form on the tangent plane of a smooth surface in the 3-D Euclidean space. The second fundamental form is a symmetric bilinear map that captures the local shape of the surface [18]. In this context, shape is how the surface normal direction changes while moving along the surface in arbitrary directions. The principal directions of the surface can be computed as the eigenvectors of II . The associated eigenvalues of the second fundamental form are called the principal curvatures κ_1 and κ_2 (chosen such that $\kappa_1 \geq \kappa_2$). These values measure the maximum and minimum values of bending of a surface, and they bound the local values of the normal curvature.

3.3.2 Mean Curvature and Gaussian Curvature (H , K)

The trace and determinant of the second fundamental form capture geometric invariants regarding the shape of the surface. The mean curvature is defined as $H = \frac{1}{2}(\kappa_1 + \kappa_2)$ — half the trace of II . Koenderink [18] describes it as “the nosedive averaged over all directions”, where the ‘nosedive’ refers to the amount of twist-free turning of the principal frame field. Mean curvature is an extrinsic measure of curvature.

The Gaussian curvature is the determinant of II , defined as $K = \kappa_1\kappa_2$. The Gaussian curvature can be interpreted extrinsically as the measure of the spread of surface normals per unit surface area. This corresponds to the area magnification of the Gauss map, hence the name “Gaussian curvature”. However, K is an intrinsic measure since it is invariant under local isometries and its value can be computed from measurements on the surface itself, regardless of the way the surface is situated in 3D space. This is the result of Gauss’s famous Theorema Egregium [18].

H and K are both algebraic invariants, meaning that they do not change depending on the choice of frame field for the surface, and geometric invariants, meaning that they do not change when the surface is rotated or translated. However, these measures are not scale-invariant.

3.3.3 Curvedness and Shape Index (C , S)

Koenderink [18] points out that all spheres intuitively have the same *shape* even though they may have different *sizes*. All of the above curvature metrics fail to capture this intuitive property. Koenderink therefore proposes two new metrics to reflect these properties, the curvedness C and the shape index S .

The shape index S describes the local shape in terms of concavity and convexity.




Figure 3.2: Shape index S . The shape index is a curvature-based metric that intuitively captures local shape. S can take on values from $[-1..1]$. Shapes with opposite values of S have the relationship of an object and its *mold*. Figure reprinted from [18].

It is formally defined as:

$$S = -\frac{2}{\pi} \tan^{-1} \frac{\kappa_1 + \kappa_2}{\kappa_1 - \kappa_2} \quad (3.5)$$

Figure 3.2 illustrates surfaces corresponding to various values of S . S takes values in the interval $[-1..1]$, with the endpoints corresponding to the concave and convex umbilics (i.e., points where $\kappa_1 = \kappa_2$).

The curvedness C represents how curved the surface is. It is formally defined as:

$$C = \frac{2}{\pi} \ln \sqrt{\frac{\kappa_1^2 + \kappa_2^2}{2}} \quad (3.6)$$

Figure 3.3 illustrates surfaces corresponding to various values of C . C can take any value in $(-\infty..\infty)$, with $-\infty$ corresponding to a flat point and $+\infty$ corresponding to a singular point.

C and S basically correspond to a polar representation of the principal curvatures, in an attempt to decouple measurements for the size and shape of the curved surface. In particular, S is a very useful property when comparing objects of different sizes (and thus different curvature ranges), since trying to match variables with different ranges could adversely affect the optimization process. As discussed previously in




Figure 3.3: Curvedness C . Curvedness captures the *size* of a surface regardless of its shape. The unit sphere has a C value of 0. Positive values of C correspond to increasingly sharp points, and negative values correspond to increasingly flat surfaces. Figure reprinted from [18].

Section 2.1.1, Tosun et al. [17] use C and S as part of their cortical correspondence scheme.

3.4 Results and Discussion

In this section, I present a practical comparison of four correspondence methods discussed so far: SPHARM correspondence (Sec. 2.1.1), Heimann’s implementation of position-based MDL (Sec. 3.1.2), generalized MDL correspondence (Sec. 3.2) using a pair of curvature measures (denoted *CurvMDL*) and generalized MDL correspondence using position and a pair of curvature measures (denoted *CombinationMDL*). Furthermore, I have applied the CurvMDL and CombinationMDL methods separately using each pair of metrics discussed in Sec. 3.3.

I have applied these correspondence methods to four different object populations: lateral ventricles, caudates, striata (which is the union of caudate, putamen and nucleus accumbens), and left femoral heads. Figure 3.4 illustrates the relative positions




Figure 3.4: The relative locations of the various brain structures used in experiments. Representative axial (A) and coronal (B) slices are shown, as well as a 3D rendering (C). The caudate nuclei, shown in blue, are located near the centre of the brain. The putamen, shown in green, is the outermost part of the basal ganglia. The striatum is the structure comprised of the caudate and the putamen, which are separated by the internal capsule fiber tract, and the nucleus accumbens (shown in yellow). The pallidum, shown in purple in the 2D views, is the other nucleus among the rostral basal ganglia. The lateral ventricles, shown in red, are roughly wrapped around the rostral basal ganglia.

of the brain structures (lateral ventricles, caudates and striata). All populations of brain structures included healthy subjects as well as patients with various disorders. This setting leads to a higher variability in the populations. The results of the SPHARM correspondence have been used as initialization to all MDL-based methods, since these require a spherical parameterization for the surface. All curvature measures have been analytically computed from the SPHARM coefficients (see Appendix A for details).

For all populations, traditional MDL performed better than SPHARM, which agrees with findings reported by Styner et al. [37], as can be seen in Figures 3.6 - 3.9. SPHARM correspondence does not handle well rotational variations in the populations, since it is based on first-order ellipsoid alignment. It is surprising, however, that SPHARM nonetheless performed better than pure curvature MDL (CurvMDL) for all populations. This is due to the non-unique nature of curvature measurements; without the spatial location information, curvature metrics on their own are not enough




Figure 3.5: The results of the SPHARM, traditional MDL and CombinationMDL demonstrated on two striata and two femoral heads. The coloring shows the ϕ coordinates on the spherical parameterization of the objects, with each (x, y, z) point on the object surface mapped to a (ϕ, θ) point on the unit sphere. Locations with same coloring on the two subjects correspond to each other for a given method. As shown in (c), the line separating blue and red regions corresponds to a longitude line on the unit sphere. Note that the SPHARM correspondence for the striata is very poor, and the CombinationMDL method gives the best visual correspondence for both populations.

for a stable correspondence computation. The natural solution is to combine location and curvature for MDL computation (CombinationMDL), so that the spatial location will be used for low-resolution correspondence and curvature for high-resolution. My framework enables the seamless integration of these features of different resolutions in a single objective function. The performance of CombinationMDL is satisfactory for all populations. However, for some populations, there seems to be little reason to use CombinationMDL instead of the traditional MDL. This finding is expected since the traditional MDL is quite adequate for simpler shapes like caudates. However, for objects with complex curvature patterns, such as the striata and the femur illustrated

in Figure 3.5, the incorporation of curvature improves the correspondence.

For all populations, the choice of particular pair of curvature measures proved not to be critical, since any of the 3 pairs of measures I used resulted in performances in the same range. This result is not surprising, since given any pair of measures I presented, one can easily compute the other 2 pairs: each pair encodes the same information about the surface shape. To improve graph readability, I only show results using the (C, S) pair of metrics on most of the figures in this chapter.

In the generalized MDL methods (CurvMDL and CombinationMDL), I have scaled the feature values appropriately. This is accomplished by using the inverse of the feature variance averaged over the whole surface as the scaling factor. This forces all feature channels to have a variance of 1 and to therefore have equal weight in the correspondence optimization.

The MDL implementation that I used requires a parameterization mesh as an additional input to determine the spatial resolution at which correspondence optimization is performed. In all computations I have used a mesh of the same resolution as the one used for SPHARM computation to obtain full-resolution correspondence. However, one might want to use a lower-resolution parameterization mesh to speed up the computation and interpolate to obtain a denser correspondence. This would lead to a faster but less accurate correspondence optimization.

3.4.1 Lateral Ventricle

The first population consists of 116 lateral ventricle segmentations. These are large, C-shaped structures located in the cerebrum. Shape analysis of the lateral ventricles is important for the study of schizophrenia and Parkinson’s disease [46, 47].

For this population, the traditional MDL performs very well, and introducing additional curvature information does not have any benefits. In fact, since curvature measurements are more prone to sampling noise than locations, using the curvature

data is similar to adding noise to the dataset. However, the performances of the two methods are still in the same range. Figure 3.6 summarizes these results.

3.4.2 Caudate

The next dataset includes 56 caudate segmentations. The caudate is one of the basal ganglia nuclei, with a vaguely C-shaped geometry with a wide head and a tapering tail. The traditional MDL and the CombinationMDL perform similarly well on this population as expected, given the simple shape of the objects in the population. In fact, Figure 3.7 demonstrates that even SPHARM correspondence performs reasonably well on this dataset, since there are no twists in the structure.

3.4.3 Femur

The third dataset I used consists of 18 truncated femur bones. The femur is a large bone in the upper leg, and its analysis is important to computer assisted surgery as well as prosthesis design [48]. The superior part of the femur has three major protuberances: the head, the greater trochanter and the lesser trochanter. The lower part, including the lower extremity, was truncated and excluded for this study.

For this dataset, due to the complex structure of the objects, the correspondence improves when curvature information is included. Note that, for the generalization metric, traditional MDL catches up with the performance of CombinationMDL when a high number of shape eigenmodes (M) are used. The CombinationMDL has superior specificity values independent of M . Figure 3.8 captures these results.

3.4.4 Striatum

Finally, I have compared the methods on a population of 20 striata, which consist of the caudate nucleus, the putamen and the nucleus accumbens. The striatum best

demonstrates the type of object where using curvature for establishing correspondence is beneficial. A typical striatum has both highly convex and highly concave surface patches, and the curvature pattern quickly changes along the surface. Traditional MDL does not perform optimally on these convoluted objects because the vertices with similar spatial locations across subjects do not necessarily correspond together. In this case, the local geometry becomes varied enough that it provides a means of identifying corresponding points, and thus it helps to use curvature as a stabilizer on striatum correspondence. The results are shown in Figure 3.9.

3.4.5 Discussion

In this chapter, I presented a framework for integrating local features into MDL computation. When using geometrical measurements as local features, such as various curvature measures, my method invariably produces results at least comparable to traditional MDL, and it handles objects with highly complex curvature structure much better than traditional MDL.

It is clear that incorporating local measurement into correspondence computation has the potential of significantly improving correspondence. Any other measurements can also be used to improve correspondence besides geometric information such as curvature.

In human cortical surfaces, which is the main target application of this dissertation, geometric measurements such as sulcal depth (see Section 1.2.1) and connectivity measures that can be derived from DTI scans are some of the intuitive features that can be used to improve correspondence. Chapter 4 discusses a parameterization-free correspondence method and its application to the human cortex using sulcal depth as a feature. Chapter 5 discusses connectivity measures that can be used as additional features for human cortical surfaces.

Although MDL is a powerful method for optimizing correspondence, it has the

disadvantage of requiring parameterized surfaces, which limits its applicability. For the human cortex, for instance, it is a nontrivial task to produce a surface with spherical topology from MR images. The level of folding of the cortical surface further complicates the problem, since the surface meshes necessarily will have a very large number of vertices, which leads to unacceptably long MDL run times, even with faster implementations such as Heimann's. The nonparametric correspondence framework introduced in the next chapter not only typically converges significantly faster than MDL even on very large meshes, but also does not require spherical topology.




Figure 3.6: Top, the shape and SPHARM parameterization of a typical lateral ventricle surface. Middle and bottom, generalization ($G(M)$) and specificity ($S(M)$) comparison on a population of lateral ventricles. In all figures, *XYZ* refers to the traditional MDL method; *CS*, *HK* and *Kappa* refer to CurvMDL method using the specified curvature metrics; *XYZCS*, *XYZHK* and *XYZKappa* refer to CombinationMDL method. Note that the choice of particular curvature metric has very little effect on the results for both CurvMDL and CombinationMDL methods. Therefore, only results using the C-S metrics will be shown in the subsequent figures to improve graph readability. For the lateral ventricle population, SPHARM and CurvMDL both perform poorly. The performance of traditional MDL and CombinationMDL is within the same range.




Figure 3.7: Top, the shape and SPHARM parameterization of a typical caudate nucleus. Middle and bottom, the generalization ($G(M)$) and specificity ($S(M)$) comparison on a population of caudates. Given the very simple structure of the objects, there is no extra benefit in adding curvature information to MDL computation. Even SPHARM performs well on this dataset, since the first order ellipsoid alignment is satisfactory. The pure curvature method (CurvMDL) performs poorly.




Figure 3.8: The generalization ($G(M)$) and specificity ($S(M)$) comparison on a population of femurs. CombinationMDL provides an improved correspondence compared to traditional MDL, even though the improvement is negligible when a higher number of shape eigenmodes (M) is used. SPHARM and CurvMDL both perform poorly.




Figure 3.9: The generalization ($G(M)$) and specificity ($S(M)$) comparison on a population of striata. CombinationMDL clearly improves correspondence. SPHARM and CurvMDL both perform poorly.

Chapter 4

Parameterization-Free Group Correspondence and Application to Highly Folded Surfaces

MDL implementations rely on parameterizations, which must be obtained through a preprocessing stage. This is a computationally expensive step at best and becomes prohibitive for 3D surfaces of non-spherical topology. Furthermore, as discussed in Chapter 3, MDL optimization itself is a slow process, especially with a large number of vertices per surface, due to the reparameterization step in the algorithm. This chapter presents a nonparameterized, topology-independent and computationally efficient algorithm suitable for correspondence optimization on human cortical surfaces.

Styner et al. [37] describe an empirical study which shows that ensemble-based statistics improve correspondences relative to pure geometric regularization and that MDL performance is virtually the same as that of $\min \log |\Sigma + \alpha I|$ (where Σ is the covariance matrix of the sample positions across the population and αI is a diagonal regularization matrix that introduces a lower bound α to the eigenvalues of the covariance matrix). This last observation is consistent with the well-known result from information theory: MDL is, in general, equivalent to minimum entropy [49]. Cates et al. [3, 4] propose a system exploring this property; this entropy-based

correspondence algorithm is the underlying technique for the methodology presented here.

The entropy-based correspondence method uses a point-based surface sampling to optimize surface correspondence in a groupwise manner. Like-numbered samples, named particles, define correspondence across the population. The optimization consists of moving the particles along the surfaces in the direction of the gradient of an energy functional that strikes a balance between an even sampling of each surface (characterized by shape entropy) and a high spatial similarity of the corresponding samples across the population (ensemble entropy). The novel contribution of this chapter is to extend this method by allowing local measurements on the object surfaces to be incorporated into the ensemble entropy term to provide a generalized correspondence definition.

4.1 Entropy-based Particle Framework

4.1.1 Surface entropy

In this work, as presented in [4], a surface $\mathcal{S} \subset \mathbb{R}^d$ is sampled using a discrete set of N surface points, $Z = (X_1, X_2, \dots, X_N)$. These surface points, called *particles*, are considered to be random variables drawn from a probability density function (PDF), $p(X)$. A realization of this PDF is denoted with lowercase; thus, a particle set is represented by $z = (x_1, x_2, \dots, x_N)$, where $z \in \mathcal{S}^N$. The probability of a realization x is $p(X = x)$, which is denoted simply as $p(x)$.

Cates et al. describe in [4] a nonparametric Parzen windowing method to estimate $p(x_i)$ such that

$$p(x_i) \approx \frac{1}{N(N-1)} \sum_{j=1, j \neq i}^N G(x_i - x_j, \sigma_i). \quad (4.1)$$

where $G(x_i - x_j, \sigma_i)$ represents a d -dimensional isotropic Gaussian with standard deviation σ_i . The value for σ_i is computed using Newton-Raphson method such that $\partial p(x_i, \sigma_i)/\partial \sigma_i = 0$, which optimizes the probability of the particle i being at the position indicated by the current configuration.

The amount of information contained in such a random sampling is the differential entropy of the PDF in the limit, which is $H[X] = -\int_{\mathcal{S}} p(x) \log p(x) dx = -E\{\log p(X)\}$, where $E\{\cdot\}$ denotes expected value. The cost function C is the negative of this expected value, which can be approximated by the sample mean for a sufficiently large sample size. The optimization problem is thus given by:

$$\hat{z} = \arg \min_z C(z) \quad \text{s.t. } x_1, \dots, x_N \in \mathcal{S} \quad (4.2)$$

$$C(z) = -H[X] \approx \sum_{i=1}^N \log \frac{1}{N(N-1)} \sum_{j=1, j \neq i}^N G(x_i - x_j, \sigma_j) \quad (4.3)$$

$$\begin{aligned} -\frac{\partial C}{\partial x_i} &= \frac{1}{\sigma^2} \frac{\sum_{j=1, j \neq i}^N (x_i - x_j) G(x_i - x_j, \sigma_i)}{\sum_{j=1, j \neq i}^N G(x_i - x_j, \sigma_i)} \\ &= \sigma_i^{-2} \sum_{j=1, j \neq i}^N (x_i - x_j) w_{ij} \end{aligned} \quad (4.4)$$

where w_{ij} are weights such that $\sum_{j=1}^N w_{ij} = 1$. This can be interpreted as the particles moving away from each other under a repulsive force while constrained to lie on the surface. The motion of each particle is away from all of the other particles, but the forces are weighted by a Gaussian function of inter-particle distance. Interactions are therefore local for sufficiently small σ . Furthermore, the Gaussian kernels are truncated at 3σ to ensure that each particle has a finite radius of influence.

Note that this particle formulation computes Euclidean distances between particles rather than the geodesic distances on the surface. Thus, a sufficiently dense sampling is assumed, so that nearby particles lie in the tangent planes of the zero sets of a scalar function F which provides the implicit cortical surface. This is an

important consideration; in cases where this assumption is not valid, such as highly curved surfaces, the distribution of particles may be affected by neighbors that are outside of the true manifold neighborhood.

The preceding minimization produces a *uniform* sampling of a surface. A strategy that samples adaptively in response to higher order shape information can be more effective for some applications. In particular, sampling high-curvature regions more densely can be desirable to ensure the validity of the assumption that tangent planes vary smoothly between neighboring particles. This can be achieved by modifying Eq. 4.1 by introducing a scaling factor k_j proportional to the root sum-of-squares of the principal curvatures at each particle j :

$$p'(x_i) \approx \frac{1}{N(N-1)} \sum_{j=1, j \neq i}^N G\left(\frac{1}{k_j}(x_i - x_j), \sigma_i\right). \quad (4.5)$$

The human cortex is a prime example of a highly curved surface. In fact, due to the high level of convolution, even a strong degree of adaptivity does not produce a satisfactory sampling of these surfaces such that nearby particles can be assumed to lie on the local tangent planes, unless a very high number of particles are used. However, using such a high number of particles is undesirable due to computational cost. Therefore, in this work, I overcome this problem by inflating the cortical surface prior to optimizing correspondence instead of adaptive sampling. The particles therefore live in the tangent planes of the inflated surface; they are only pulled back to the original cortical surface for correspondence evaluation purposes.

For both uniform and adaptive sampling, the constraint of the particles to the object surface is achieved via an implicit representation where the surface is defined by the zero-set of a signed distance function $F(x)$. This constraint is maintained at each iteration of the optimization, as described by Meyer et al. [50], by projecting the gradient of the cost function onto the tangent plane of the surface, moving the

particle along this tangent vector, and projecting the particle to the closest root of F by the Newton-Raphson method. The sampling system is initialized with a single particle at a root of F that is recursively split in two to produce the desired number of particles, as described by Cates [4].

4.1.2 Ensemble entropy

An ensemble $\mathcal{E}$ is a collection of M surfaces each with their own set of particles, i.e., $\mathcal{E} = z^1, \dots, z^M$. The ordering of the particles on each shape implies a correspondence among shapes and the entire population can be represented in a matrix of particle positions $P = x_j^k$ with particle positions along the rows and shapes across the columns. Cates et al. [3] model $z^k \in \mathbb{R}^{Nd}$ as an instance of a random variable Z , and they propose to minimize the combined ensemble and shape cost function

$$Q = H(Z) - \sum_{k=1}^M H(P^k), \quad (4.6)$$

which favors a compact ensemble representation balanced against a uniform (or adaptive) distribution of particles on each surface.

For this discussion, the complexity of each object is assumed to be greater than the number of objects, and therefore it is typical to have $N > M$. Given the low number of examples relative to the dimensionality of the space, some conditions must be imposed in order to perform the density estimation. Cates assumes a normal distribution and models the distribution of Z parametrically using a Gaussian with covariance Σ . The entropy of a g -dimensional Gaussian with covariance Σ can be computed as

$$H = - \int \int \dots \int f(x) \ln(f(x)) dx \quad (4.7)$$

$$= \frac{1}{2}(g + g \ln(2\pi) + \ln|\Sigma|) \quad (4.8)$$

The ensemble entropy can therefore be expressed as

$$H(Z) \approx \frac{1}{2} \log |\Sigma| = \frac{1}{2} \sum_{j=1}^{Nd} \log \lambda_j, \quad (4.9)$$

where $\lambda_1, \dots, \lambda_{Nd}$ are the eigenvalues of Σ .

In practice Σ will not have full rank, in which case the entropy is not finite. It is therefore necessary to regularize the problem with the addition of a diagonal matrix αI to introduce a lower bound on the eigenvalues. Let $Y = P - \bar{P}$, where P is the matrix that encodes the vertex locations for each object along successive columns, $\bar{P}$ is a matrix with all columns set to the mean shape μ . The covariance can then be estimated from the data, with $\Sigma = (1/(M-1))YY^T$. Because $N > M$, the computations are performed on the dual space (dimension M rather than N), knowing that the determinant is the same up to a constant factor of α . Thus, the cost function G associated with the ensemble entropy is defined as:

$$\log |\Sigma| \approx G(P) = \log \left| \frac{1}{M-1} Y^T Y \right| \quad (4.10)$$

and

$$-\frac{\partial G}{\partial P} = Y(Y^T Y + \alpha I)^{-1}. \quad (4.11)$$

It is now clear that α is a regularization on the inverse of $Y^T Y$ to account for the possibility of a diminishing determinant. Starting with a large α and incrementally reducing it as the optimization converges yields an annealing approach which improves computational efficiency; this has the effect of preventing the system from attempting to reduce the thinnest dimensions of the ensemble distribution too early in the process.

The negative gradient $-\partial G/\partial P$ 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:

$$x_j^k \leftarrow \gamma \left[-\partial G / \partial x_j^k + \partial C^k / \partial x_j^k \right]. \quad (4.12)$$

4.2 Generalized ensemble entropy

As demonstrated in Chapter 3, incorporating local features into the similarity metric can significantly improve correspondence quality as compared to approaches that only use spatial proximity. In particular, surfaces with complex geometry have been shown to benefit the most from this approach. Furthermore, there is no clear evidence that spatial proximity alone defines anatomical correspondence in medical datasets, where non-spatial measurements can be available along with the surface geometry. In these cases, it is desirable to incorporate such local information into the similarity metric to achieve an optimal correspondence across the population.

The entropy-based particle framework lends itself nicely to this generalized correspondence definition, as we presented in [51]. The ensemble entropy term is modified to reflect the similarity of the local features instead of the spatial locations. The features are represented as a function of location, $f(x_j^k)$, with $f : \mathbb{R}^d \rightarrow \mathbb{R}^q$. The function f is thus vector-valued (with vector dimension q) to allow multiple features to be used at once. The surface entropy term remains unchanged, since it is still desirable to sample the surfaces uniformly (or adaptively, if so chosen).

When computing the ensemble entropy of vector-valued functions of the correspondence positions P , the generalized case can be represented by $\tilde{P} = f(x_j^k)$. $\tilde{Y}$ becomes a matrix of the function values at the particle points minus the means of those functions at the points, and the general cost function can be computed simply as

$$\tilde{G}(\tilde{P}) = \log \left| \frac{1}{M-1} \tilde{Y}^T \tilde{Y} \right|. \quad (4.13)$$

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

$$-\frac{\partial \tilde{G}}{\partial P^k} = J_k^T Q^k, \quad (4.14)$$

where J_k is the Jacobian of the functional data for shape k . Note that each J_k is in the form of a block diagonal matrix with $(q \times N) \times (q \times N)$ blocks, with diagonal blocks the $q \times d$ submatrices of the function gradients at particle j .

When multiple features are present ($q > 1$), each feature channel is scaled separately, as described in Sec. 3.4 for the case of generalized MDL applications. In particular, the inverse of the feature variance averaged over the whole surface is used as the scaling factor. This forces all feature channels to have a variance of 1 and to therefore have equal weight in the correspondence optimization. However, it is also possible to scale the features differently if the application warrants assigning a heavier weight to some of the features in the similarity metric computation.

4.3 Application to cortex

One of the main challenges of using the entropy-based particle framework on the cortical surface is that it assumes the particles to be existing on local tangent planes, which presents a problem for the cortex given the highly convoluted surface geometry as discussed in Sec. 4.1.1. This problem can be overcome by incorporating geodesic distances instead of Euclidean distances in the surface entropy term. However, this would be computationally expensive and thus it would eliminate one of the strengths of the entropy-based correspondence algorithm, making it unsuitable for application to cortical datasets.

I overcome this difficulty by instead ‘inflating’ the cortex surface. This way, I obtain a less convoluted, blob-like surface for the particles to live on, while the local features continue to live on the original cortical surface. Inflation is a commonplace

procedure for cortical visualization, since the extensive sulcal and gyral convolutions make visual inspection difficult. Several cortical unfolding procedures that expose the buried folds of the cortical surface have been proposed [52, 53, 54]. In this work, I use FreeSurfer for inflation [6] since it provides a smooth surface while minimizing metric distortions, without introducing a cut to the surface. However, the technique I propose is not specifically dependent on the FreeSurfer algorithm. Therefore, FreeSurfer inflation could be replaced by any module that delivers similar inflated surfaces and a one-to-one correspondence between the original cortex surface and the inflated surface, since the data to be used for correspondence, such as the sulcal depth, lives on the original cortex surface.

FreeSurfer inflation is achieved via the optimization of an energy functional consisting of the weighted sum of a spring force that works towards ‘inflating’ the surface and a metric preservation term that ensures that as little metric distortion as possible is introduced in the process. The inflation process is such that points that lie in convex regions move inwards while points in concave regions move outwards over time. Therefore, the average convexity/concavity of the surface over a region, also referred to as sulcal depth, can be computed as the integral of the normal movement of a point during inflation. Specifically, the sulcal depth $SD(x_k(0))$ at position x_k is defined as:

$$SD(x_k(0)) = \int v_k(t) \bullet n_k dt \quad (4.15)$$

where n_k is the unit normal vector at position x_k and $v_k(t)$ is the direction in which the vertex x_k moves at time step t , which is equivalent to the partial derivative of the energy functional driving the inflation. It should be noted that SD captures the high level foldings of the cortical surface but is relatively insensitive to the smaller folds; this property makes SD an attractive correspondence metric, since it is relatively




Figure 4.1: The sulcal depth pictured as a color map on the white matter surface and its inflation surface. The blue and red regions correspond to highly negative and highly positive values of sulcal depth, respectively.

stable across individuals. Fig. 4.1 shows the sulcal depth as a color map on both the original white matter surface and the inflated surface.

4.4 Surface interpolation from particles

The entropy-based correspondence algorithm described above produces N particles along each surface such that like-numbered particles across the population correspond. For some applications this can be satisfactory, such as in the case of surfaces with a low level of geometric complexity that can be captured adequately with a few thousand particles or less. However, for other applications such as in the case of the human cortical surface, a much denser sampling of the correspondence map is desired. Directly obtaining this dense sampling from the correspondence algorithm by using a very high number of particles is not appealing, since such an attempt would significantly increase the computation time. Furthermore, it is not possible to manually choose the locations of the particles, which can make the correspondence information for certain anatomical landmarks elusive. Finally, it is desirable to retain geometrical information associated with the surfaces, such as normal vectors and curvature. Therefore, it is advantageous to interpolate the relatively sparse correspondence maps obtained from the entropy-based correspondence algorithm using the original dense surface meshes and a set of desired sample locations. In the experiments presented in




Figure 4.2: Surface interpolation from particles. Given the desired surface sample set, the closest particles A_i, B_i, C_i to each sample point v_i are identified and the barycentric coordinates of v_i in this triangle are computed. Then, the corresponding sample location v'_i on a different surface can be obtained by interpolating using these barycentric coordinates in the triangle $\triangle A'_i B'_i C'_i$, where A'_i, B'_i, C'_i are the particles corresponding to A_i, B_i, C_i respectively.

this chapter, the vertices for one of the original reconstructed surfaces are used as the sample set, typically producing densely interpolated surface meshes of about 150,000 vertices per object, in contrast to $N = 1024$ or 2048 particles.

The interpolation scheme used to produce these dense surface meshes from the particles is a 3-step algorithm. First, interpolation coefficients are computed using the sample set and the associated particles. Then, for each surface, the corresponding sample points are computed by interpolating the corresponding particles using the weights computed in the first step. Finally, the interpolated sample points are projected onto the original surface for each object to compensate for the discrepancy introduced in the first step, as explained below. Figure 4.2 illustrates the process.

The interpolation coefficients for a sample point v_i are obtained by locating the 3 closest particles A_i, B_i, C_i to v_i and computing v_i 's barycentric coordinates in the triangle $\triangle A_i B_i C_i$ formed by these 3 points. Note that these barycentric coordinates represent v_i 's projection on $\triangle A_i B_i C_i$ rather than v_i itself; the discrepancy can become significant in high-curvature regions if the particle sampling is too sparse. Therefore, it is advisable to use the adaptive sampling scheme described in Sec. 4.1.1 for surfaces that have high-curvature patches.




Figure 4.3: Example of surface interpolation. Given M surfaces and the particle locations on each that are in correspondence, the surface interpolation method produces M dense meshes (with 144,827 vertices per mesh for this particular example) such that like-numbered vertices are in correspondence. The rightmost image shows a zoomed-in section of one of the interpolated surfaces with the particle locations overlaid with the wireframe mesh.

For computing the sample locations on each surface S , for each sample point v_i , the particles A'_i, B'_i, C'_i corresponding to the particles A_i, B_i, C_i are interpolated using the barycentric coordinates computed in the first step. The sample point v'_i thus computed belongs to the triangle $\triangle A'_i B'_i C'_i$ and corresponds to the projection of v_i onto the triangle $\triangle A_i B_i C_i$.

To compute the sample point that corresponds to v_i , one needs to project v'_i back to the surface S . This can be achieved by finding the closest point on the original mesh for S to v'_i , whether this is a vertex of the mesh or a point within an edge or a triangle. This step compensates for any discrepancies that might have been introduced in the first step and ensures that all sample points created by the algorithm are indeed on the surface rather than being slightly inside or outside the surface.

An example of the surface interpolation results is shown in Fig. 4.3 for two surfaces. The vertices of one of the original meshes have been used as the sample point set for the interpolation. Even though only a modest number of particles per surface ($N = 1024$) has been used to produce the surfaces in this figure, the quality

of the densely interpolated mesh is satisfactory, as demonstrated by a mean distance of $0.0088mm$ between the Subject 2 original surface and the associated interpolated surface.

4.5 Results

I evaluate my method on two different datasets of human cortical surfaces. In order to evaluate the quality of my results, I analyze the local variability reduction of both the features that are being used for the correspondence computation (sulcal depth in this case), and, more importantly, of a different local measurement, namely, cortical thickness. Additionally, I use the generalization and specificity measurements, discussed in detail in Sec. 2.3.3, to further analyze the correspondence quality. I compare these performance measures with results obtained from FreeSurfer, which is one of the most commonly used methods for cortical correspondence.

As discussed in Sec. 2.1.1, the FreeSurfer correspondence method is essentially a two-step correspondence computation: it is initialized with spatial correspondence (while computing a spherical parameterization of the surfaces) and then it optimizes the sulcal depth correspondence. This method is also fundamentally different from the entropy-based techniques in that it focuses on a pairwise correspondence (subject to average), whereas the entropy-based method emphasizes a groupwise approach.

For both the sulcal depth and cortical thickness measurements, I compute the local sample variance of the measurement before and after correspondence optimization; the desirable result is a lowering of variability across corresponding points. It should be noted that both sulcal depth-based entropy methods and FreeSurfer use the sulcal depth information as part of the correspondence optimization process; therefore, the sulcal depth evaluation is biased. Also note that the higher degree of remaining variability in cortical thickness can be largely attributed to inter-subject variability,

	Mean of Sulcal Depth Variance	Mean of Cortical Thickness Variance
Initial Data	0.228 (0.007)	0.338 (0.012)
XYZ-entropy	0.220 (0.006)	0.342 (0.013)
SD-entropy	0.003 (4e-4)	0.312 (0.012)
FreeSurfer	0.076 (0.004)	0.307 (0.010)

Figure 4.4: Mean sample variances of sulcal depth and cortical thickness measurements across the cortical surface for the first dataset, given different correspondence maps. The values in parentheses show the standard error associated with each mean.

since cortical thickness patterns tend to vary largely among individuals.

The first dataset consists of 10 cortical surfaces from healthy patients reconstructed via FreeSurfer from T1-weighted images; 1 had to be removed due to failure of successful surface reconstruction with FreeSurfer, caused by a high level of noise in the MRI scan (therefore, the final study used $M = 9$ subjects). All reconstructions showed some degree of error in the temporal lobe; no manual interventions were made.

As a first experiment, I simply use the local sulcal depth as the feature f for the generalized entropy method. So the methods compared are the initial data, the sulcal depth-based entropy (denoted *SD-entropy* from here on for brevity), the standard location-only entropy-based particle system (denoted *XYZ-entropy*) as well as FreeSurfer. Mean sample variances for both the cortical thickness and sulcal depth measurements are summarized in Fig. 4.4.

For sulcal depth measurements, SD-entropy reduces variance almost 75-fold over initial data, and almost 25-fold over FreeSurfer results. For cortical thickness, SD-entropy has considerable improvement over initial data and has a slightly higher average variance than FreeSurfer does, but the difference is not significant since it is




Figure 4.5: Comparison of the distribution of variance across the cortical surfaces for SD-entropy method (left) and FreeSurfer (right). The coloring of the particles is linearly growing from values of (blue=0) to (red=1). It can be seen that the SD-entropy results in a localized high variance near the temporal lobe (conceivably due to reconstruction noise), but has low variance elsewhere; FreeSurfer results in relatively high variance across the entire cortical surface.

within the standard error values. Furthermore, an inspection of the distribution of this mean variance over the surface, as shown in Fig. 4.5, reveals that FreeSurfer has a higher variance across the entire surface, whereas SD-entropy has only a localized high variance around the temporal lobe (which is, as noted above, not perfectly reconstructed due to input image noise) and performs much better in other areas of the cortical surface. It is therefore possible that the SD-entropy performance would improve with a better temporal lobe reconstruction. Furthermore, the SD-entropy method completely discards the spatial proximity from the similarity metric. As illustrated in the next example, using spatial proximity in addition to SD in the generalized entropy can also improve the correspondence quality. Note that both SD-entropy and FreeSurfer perform considerably better than the XYZ-entropy method.

The second dataset consists of 9 cortical surfaces also from healthy subjects. Since both T1 and T2 images are available for this population, a bias-corrected T1-weighted image is produced via an atlas-based tissue segmentation tool. This tool, itkEMS, implements an Expectation-Maximization approach to segment the major brain tissue classes and correct for intensity inhomogeneity using both T1 and T2 weighted images

	Mean of Sulcal Depth Variance	Mean of Cortical Thickness Variance
XYZ-entropy	0.109 (0.003)	0.262 (0.006)
SD-entropy	0.009 (8e-5)	0.253 (0.008)
XYZ-SD-entropy	0.109 (0.003)	0.251 (0.006)
FreeSurfer	0.039 (0.003)	0.277 (0.010)

Figure 4.6: Mean sample variances of sulcal depth and cortical thickness measurements across the cortical surface for the second dataset, given different correspondence maps. The values in parentheses show the standard error associated with each mean.

[55]. This bias-corrected T1-weighted image is used as input to the FreeSurfer pipeline to increase the robustness of the reconstruction.

For this dataset, I use a different generalized entropy application, such that $f(x, y, z) = \{x, y, z, sd(x, y, z)\}$. In addition to this method (denoted *XYZ-SD-entropy*), XYZ-entropy, SD-entropy and FreeSurfer are also tested as for the first dataset. The sulcal depth and cortical thickness variance analysis is summarized in Figure 4.6. As in the first dataset, the SD-entropy method outperforms FreeSurfer in sulcal depth variability reduction. However, the XYZ-SD-entropy method provides the best cortical thickness variability reduction of all 4 methods tested. The combination of spatial proximity and sulcal depth similarity outperforms either measurement on its own, as well as FreeSurfer which is a somewhat hybrid method. These results suggest that FreeSurfer and SD-entropy methods might be overfitting the model by a too strong focus on the sulcal depth similarity; however further evaluation is needed since the differences in cortical thickness variation are within the standard error.

To provide further correspondence quality comparison, I provide a comparison of the generalization and specificity measures (see Sec. 2.3.3) between FreeSurfer, XYZ-entropy and XYZ-SD-entropy. It should be noted that these measurements




Figure 4.7: Cortical thickness generalization, $G(M)$. M denotes the number of principal modes of variation used in the surface reconstruction. Note that a lower $G(M)$ value represents a better correspondence. The XYZ-SD-entropy method clearly has better generalization ability than both FreeSurfer and XYZ-entropy.

are computed based on cortical thickness rather than spatial location both to be consistent with the mean variance measure described above and to benefit from the unbiased nature of the cortical thickness. Also note that both generalization and specificity are estimates of error (measured in millimeters) and therefore a low value represents a good correspondence of the population. Figures 4.7 and 4.8 show the results of the analysis on this dataset. The XYZ-SD-entropy clearly has better generalization ability than FreeSurfer, and FreeSurfer performs better than XYZ-entropy. For specificity, the difference between the methods is minor; however, both entropy methods perform slightly better than FreeSurfer. XYZ-SD-entropy also outperforms XYZ-entropy by a small margin. Note that specificity is a measurement designed for statistical modeling purposes rather than statistical analysis; therefore, it is less relevant for cortical thickness or functional analysis purposes as compared to generalization and variability reduction.

This chapter presented a novel method of computing correspondence across datasets of human cortical surfaces using functions of spatial locations in an entropy mini-




Figure 4.8: Cortical thickness specificity, $S(M)$. M denotes the number of principal modes of variation used in the surface reconstruction. Note that a lower $S(M)$ value represents a better correspondence. Both entropy methods have slightly better specificity than FreeSurfer.

mizing particle framework. I showed that this generalized entropy method provides correspondence maps that result in tighter distributions of sulcal depth and cortical thickness compared to other commonly used methods such as FreeSurfer. Using sulcal depth in conjuncture with the spatial location gives the best results for these cortical datasets. For other applications, curvature and spatial location could be a useful combination of geometric features.

The power of the generalized entropy-based correspondence method becomes more clear when one considers non-geometric measurements as part of the feature set. In the next chapter, I will present a framework for using white matter connectivity for improving the cortical correspondence even further. Image intensities, blood vessel proximity (as extracted from MRA images) or functional measurements are other examples of non-geometric features that can improve correspondence in other applications.

Chapter 5

DTI-based Connectivity for Cortical Correspondence

The previous chapter presented an entropy-based framework that incorporates local functions of position to establish correspondence among a surface population. This extension of the particle correspondence framework is critical for the application of the algorithm to populations of cortical surfaces, since additional information sources can have significant impact on correspondence quality. Structural features such as sulcal depth, as presented in the previous chapter, provide additional information about the geometry of the brain. This chapter introduces DWI-based fiber connectivity features to provide augmented knowledge about the white matter structure, as we presented in [56].

Structural MRI scans show white matter (WM) as a mostly homogeneous region, such that it is impossible to infer information about WM fiber tracts. The understanding of the WM structure, however, can be significantly improved via information on fiber tracts that can be extracted from diffusion weighted imaging (DWI) scans. One of the main contributions of this dissertation is a suitable mapping of such fiber tract information to the cortical surface. Cortical connectivity maps, which represent the extent to which each voxel on the cortical surface is connected via fiber tracts to a given region of interest (ROI), is proposed as a possible solution to this problem.




Figure 5.1: Pipeline overview. (A), selected ROI's and the DWI image are input to the stochastic tractography (ST) algorithm. WM surface is deflated using proposed algorithm (B) to construct connectivity maps on the surface from ST results (C). These connectivity maps are normalized via histogram equalization (D) and used to optimize correspondence (E).

I employ a stochastic tractography algorithm, described in Sec. 5.1.2, to generate connectivity maps that represent the probability of each voxel being connected to each one of various designated ROI's. A separate feature channel is used for each ROI's connectivity map in the correspondence optimization.

There is, however, a major obstacle to using these connectivity maps for cortical correspondence: the connectivity probabilities typically decrease drastically near the WM/GM boundary, since the diffusion gets too isotropic and noisy near the surface. This effect is more emphasized at the ridges of the gyri (as opposed to the valleys of the sulci). Thus, the connectivity values sampled directly at the cortical surface voxels reflect the local sulcal depth rather than actual WM connectivity. Sec. 5.2.2 presents a method of computing the connectivity probability at the cortical surface using a surface deflation algorithm. The surface deflation results in a smoother surface that

follows the sulcal cortical boundary closely while leaving out the gyri. Then, the connectivity probability at each cortical voxel is defined as the connectivity probability value at the corresponding deflated surface voxel.

An overview of the pipeline for computing connectivity features is shown in Fig. 5.1. A stochastic tractography algorithm (Section 5.1.2) is used to compute probabilistic connectivity values to selected ROI's based on the DWI image. To project these maps to the cortical surface, a deflated surface is computed (Section 5.2.2). Finally, to normalize these connectivity maps with respect to ROI and brain size, histogram equalization is applied (Section 5.2.2). The normalized connectivity maps are used as features in the particle correspondence framework, as discussed in Chapter 4.

5.1 Diffusion Tractography for Connectivity

The goal of tractography is the extraction of WM fiber paths from DWI or DTI images in order to determine intervoxel connectivity on the basis of the anisotropic diffusion of water [57]. Fiber tracking algorithms can be divided into three main groups: deterministic (streamline) tractography, probabilistic (stochastic) tractography, and optimal path approaches. This section provides a brief review of a variety of tracking algorithms, with a focus on the technique used for generating the connectivity features used in the experiments to be presented in Section 5.3.

5.1.1 Streamline tractography

Deterministic tractography typically generates fiber tracts by following the direction of maximal diffusion. In this approach, fiber trajectories, or 'streamlines' follow the primary eigenvector of the diffusion tensor assuming that this is parallel to the dominant direction of axonal tracts. The streamlines are computed by forward integration of the field of vectors defined by the local principal eigenvectors. The twofold ambigu-

ity of the eigenvector direction is resolved by the continuity of the path. Thresholds for the maximum turning angle of the streamline between integration steps as well as for the minimum FA at a given tract location can be established to constrain the fiber tracts to regions of the brain where the diffusion tensor model realistically represents the white matter pathways [57].

While these methods have low computation costs and simplify the visualization of the extracted fiber tracts, there are some issues with using them for quantitative analysis [58]. They are not robust against noise in the input images since the error in the integration accumulates. They are also prone to partial voluming effects caused by fiber crossings, a problem which renders the primary eigenvector direction ambiguous and thus misleads the streamline. Furthermore, since streamlining methods do not account for the uncertainty in fiber orientation in highly isotropic regions, they have limited applicability to regions containing isotropic voxels. These limitations motivate the use of a stochastic tractography algorithms for computing the cortical connectivity maps to be used for correspondence.

5.1.2 Stochastic tractography

In contrast to deterministic approaches, stochastic tractography methods take the uncertainty of fiber orientations into account and therefore yield results that are more suitable for this dissertation’s purposes.

The bootstrap method is a nonparametric statistical procedure for determining the uncertainty in a given statistic. As shown by Lazar et al. [59] and Jones et al. [60], this method can be applied to fiber tracking in order to compute connectivity between different ROI’s based on estimates of variance in the original DTI data. The method estimates the variance of the diffusion data by taking several (redundant) sets of DWI scans and using these to create new data through random recombinations of the original data. A large number of such mixed DWI sets (called bootstrap samples)

are each converted to DTI images on which streamline tractography is performed with the same start points. The connectivity probability of each voxel in the volume is then computed via the percentage of bootstrap samples where the fiber tract passed that particular voxel.

On the other hand, Bayesian stochastic techniques, such as the method presented by Behrens et al. [61], perform tractography in a probabilistic framework by generating a posterior probability distribution of fiber directions from the observed diffusion data. A stochastic integration algorithm is then applied by choosing the flow vector at each step to be a random sample from this distribution defined around the tensor’s principal eigenvector orientation at the current voxel. Since the fiber orientation distribution is analytically intractable, these techniques are typically combined with Monte-Carlo simulations, which may include tens of thousands of paths from a single seed, of which only a small fraction will typically reach the target. The connectivity probability of each voxel is defined via the percentage of samples that traverse that particular voxel. The explicit modeling and propagation of uncertainty allow stochastic methods to generate tracts in regions of low anisotropy.

In this work, I use an open-source implementation by Ngo et al. [11] of Friman’s stochastic tractography algorithm [12, 13]. Because Monte-Carlo simulations tend to be computationally expensive, Friman et al. introduce a stochastic method that avoids Monte-Carlo integration. In this approach, as in Behrens’ work above, fiber tracts are modeled as sequences of unit vectors whose orientation is determined by sampling a posterior probability distribution. The posterior distribution is given by a prior likelihood of the fiber orientation multiplied by the likelihood of the orientation given the DWI data. Friman uses a tensor model constrained to be linearly anisotropic to lower the computational cost of the algorithm; deviations from this distribution are modeled as uncertainty in the fiber orientation. At each step, the orientation of the previous vector in the sequence affects the prior, ensuring no backtracking occurs.

The tracking stops when the tract reaches a voxel with a low posterior probability of belonging to the white matter.

A high number of sample fibers are tracked from each voxel included in the input ROI; the probabilistic connectivity of a voxel to the ROI is defined as the ratio of fiber samples that travel through a voxel over the total number of samples. As described in the next section, the connectivity values on the WM/GM boundary are not appropriate for usage in cortical correspondence, and instead the values at the corresponding deflated surface location are used.

5.1.3 Optimal path methods

Various methods based on the Hamilton-Jacobi approach have been proposed to overcome some of the difficulties arising in tractography [58, 62, 63]. The main idea for these methods is to compute the shortest path where the cost associated with each path is an integral dependent on position and path orientation. The cost function also involves a penalty for paths that are unlikely given the tensors. These formulations result in first- or higher-order partial differential equations which model evolving fronts whose speeds are determined by information from the diffusion tensor [58]. These methods are inherently more robust to noise in diffusion data than the other tractography methods. However, these methods are generally not suitable for usage as features in cortical correspondence, since they typically do not generate connectivity probability maps for the entire cortex, with the notable exception of Lenglet et al. [64] who define a connectivity measure in an optimal control framework. Furthermore, these methods can not handle crossing fibers effectively, whereas this can be incorporated into explicit tractography models.

5.2 Mapping Connectivity to the Cortical Surface

Once connectivity maps for a given ROI is computed, it is necessary to map these values to the cortical surface in order to use them as features in the cortical correspondence framework (Fig. 5.1-C). Sec. 5.2.1 discusses the naive approach to illustrate its problems. Then, in Sec. 5.2.2, I describe a novel method of computing the connectivity probability at the cortical surface using a surface deflation algorithm (Fig. 5.1-B).

5.2.1 Kernel-based averaging

The intuitive way to assign connectivity values to the vertices of the cortical mesh is to look up the connectivity at the equivalent volume location via interpolation between the nearest voxels. However, given the isotropic nature of the diffusion data near the WM/GM boundary, this results in inconsistent features on the surface. A basic approach to solve this problem is by simply averaging using a small Gaussian kernel in an attempt to get rid of the excessively noisy values. This approach can be improved by averaging only the values of those voxels that reside within a white-matter mask, in order to avoid averaging values that lie in the gray matter where the connectivity values are 0.

However, a visual inspection of the connectivity values thus obtained on the surface, as shown in the right side of Figure 5.2, reveals a major problem. The stochastic tractography results are heavily dependent on the local sulcal depth, even after averaging. This observation is intuitively evident: the uncertainty in tracking increases near the WM/GM boundary and thus decreases the connectivity probability, as the DTI values become isotropic. Furthermore, the longer the fiber has to travel in such a high-uncertainty region, the lower the connectivity values will become, since the uncertainty accumulates. Therefore, in the valleys of sulci, high connectivity values




Figure 5.2: Impact of brain deflation algorithm on surface connectivity values. The stochastic tractography algorithm gives connectivity probabilities for the brainstem for this subject (*A*). The noisy tracking around temporal lobe is reflected on the connectivity map that uses simple averaging (*B*). The surface deflation method ignores the noisy signal and reflects a more accurate connectivity map (*C*). Note how strongly the averaging method depends on sulcal depth (*D*), illustrated in highlighted regions. The connectivity map obtained by the surface deflation method shows a pattern that closely agrees with anatomical knowledge: the brainstem is connected to the motor and premotor cortex, primary sensory cortex, parietal lobe, temporal lobe and occipital lobe, via the corticospinal tract, the medial lemniscus and the superior, medial and inferior cerebellar peduncles [65].

are observed, whereas at the ridges of the gyri, connectivity becomes close to 0. This is an artificial side-effect of the stochastic algorithm, and the connectivity values as they are form an undesirable and unstable feature for correspondence.

5.2.2 Brain deflation

In order to get accurate readings of fiber connectivity probability values, I propose using a deflated smooth white matter surface with one-to-one correspondence to the WM/GM boundary. This surface should not only be sufficiently far from the boundary, but also without the convolutions caused by sulci and gyri. Without such a

deflated surface, the probabilistic fiber connectivity values become heavily dependent on the local sulcal depth, yielding high connectivity values near sulci and low connectivity values near the gyri (see Fig. 5.2).

I propose a surface evolution method that evolves the WM surface by progressively smoothing out the gyri. To prevent the local sulcal depth from dominating the connectivity values, it's important to have a surface that is not only sufficiently far from the WM/GM boundary but also much smoother. This is accomplished by a mean-curvature-based smoothing algorithm, described in [19, 66]. This iterative method smoothes the surface mesh using a relaxation operator, such that the vertices are repositioned according to

$$V_i^{t+1} = (1 - \lambda)V_i^t + \lambda\bar{V}_i^t, \quad (5.1)$$

where V_i is the position of the i th vertex, t is the number of iterations, $\lambda \in [0, 1]$ is a smoothing parameter, and $\bar{V}_i$ is the average vertex position, which is the average position of neighboring triangle centers weighted by the triangle areas.

However, an unconstrained mean-curvature-based smoothing algorithm results in an inflated surface that evolves towards a sphere in the limit. This is an unsuitable approach, since the desired result is a deflated surface that closely follows the valleys of the sulci of the brain. To this end, I propose an adapted algorithm such that vertices located near the valleys of the sulci are fixed (by forcing the velocity λ to 0 at these vertices), which results in the smoothing of only the gyri, while keeping the rest of the mesh intact. However, keeping the valleys of the sulci indefinitely fixed would result in the creation of singular points due to the hard constraint posed on these vertices. To prevent this, the fixed locations are progressively released during the evolution based on thresholding of the vertex curvature, computed on the discrete

mesh as [66]

$$H(x_i) = \frac{1}{4A_i} \left| \sum_{j \in N_1(i)} (\cot \alpha_{ij} + \cot \beta_{ij})(x_i - x_j) \right| \quad (5.2)$$

where A_i is the area of the Voronoi cell associated with the vertex, $N_1(i)$ is the set of 1-neighbors of the i^{th} vertex and α_{ij} and β_{ij} are the two angles opposite to the edge in the two triangles sharing the edge (x_i, x_j) .

The vertices to be fixed initially are determined based on the sulcal depth. All positive local maxima of the sulcal depth are marked as fixed, and all vertices of the mesh that are located between already fixed locations are also fixed, in order to create merged surface patches rather than standalone points. It is preferable to start with too many fixed vertices rather than too few, because the progressive relaxation stage ensures no vertices remain fixed for longer than necessary. Constraining the fixed points to positive sulcal depth values ensures that vertices located on the gyri are free to move at all times, whereas the sulci can start moving only after the surrounding gyri have been smoothed out.

The progressive relaxation of the fixed locations is necessary to avoid any patches from remaining fixed indefinitely despite the fact that the rest of the mesh around it has sufficiently deflated. Once the fixed sulcal region becomes flat (detected by the mean curvature threshold), the zero-velocity constraint on the vertex is released, and the vertex is free to move.

As a final step in the deflation, I shrink the entire surface inwards by about one voxel, to ensure that the vertices at the sulci move away from the WM/GM boundary. Without this shrinking step, the probabilistic connectivity values at the sulci and gyri would be treated differently, which would introduce unwanted bias by only moving the gyri away from the WM/GM boundary. Note that the shrinking has to be done in small increments to avoid introducing topological changes to the surface. Fig. 5.3




Figure 5.3: Brain deflation progress for one subject. The surface outline is shown in contrasting colors overlaid on an axial slice of the brain. *A* shows the original WM surface, and *B-E* show the progress of the deflation at 1000 iteration intervals. The surface shown in *E* is used for retrieving probabilistic connectivity images after a final scaling step, shown in *F*. Note the progressive smoothing of the gyri as the surrounding regions become flat, which relaxes the velocity constraint on these vertices.

shows intermediate results of the surface deflation on an axial slice of the brain scan as well as the final scaling.

Normalizing connectivity maps

One of the problems with using probabilistic connectivity results for cortical correspondence is the issue of comparability given different ROI segmentation and different brain sizes. Given the same DWI data, additional voxels in the ROI tend to lower the connectivity probability of voxels, especially if the additional voxels are erroneous and do not truly belong to the same anatomical structure, because the probability is defined as the number of samples that visit the voxel divided by number of all samples. Even in the case of correct segmentations, different parts of the ROI might show different connectivity patterns, which can reduce the assigned probabilities significantly. Figure 5.4 illustrates this problem.

To normalize for such effects, I perform a histogram equalization on the connectivity feature values for each individual and for each ROI (Fig. 5.1-D). Each connectivity map is a map on the cortical surface with values defined on the n vertices. The connectivity values in this map can be discretized by simply multiplying them by the




Figure 5.4: Motivation for normalizing connectivity maps. ROI voxels are shown in yellow and non-ROI voxels are shown in blue. On the left, all fiber samples generated from the ROI (consisting of one voxel, B) lead to voxel A. Therefore, A has a connectivity of $3/3 = 1$. On the right, an additional voxel is included in the ROI for a different brain. All the fiber samples leaving D lead to C; similarly, all samples leaving ROI voxel E lead to F. Consequently, C and F both have connectivity $3/6 = 0.5$. However, the connection strength of these voxels is not any different than that of A. Normalization for ROI size effects is clearly necessary if these connectivity maps are to be effectively used for correspondence optimization.

total number of fiber samples, by definition of probabilistic connectivity. Then, if n_i is the number of occurrences of connectivity value i , the probability of an occurrence of a vertex with connectivity i in the map is $p_{\text{conn}}(i) = p(\text{conn} = i) = \frac{n_i}{n}$. The associated cumulative distribution function is given by $\text{cdf}_{\text{conn}}(i) = \sum_{j=0}^i p_{\text{conn}}(j)$. Then, a new map $\{\text{conn}_{\text{norm}}\}$ can be produced using $\text{conn}_{\text{norm}}(v) = \text{cdf}_{\text{conn}}(\text{conn}(v))$ at each vertex v . This normalized map is used for cortical correspondence optimization. It is important to note that this approach does not change the ordering of the connectivity strengths associated with the various parts of the cortex, but it merely stretches the histogram to ensure comparability across the population.

An alternative approach for normalization would be to divide the number of fiber samples that pass through the voxel by the maximum number of samples that passes any given boundary voxel rather than the total number of samples. In Fig. 5.4, this would assign a connectivity strength of $3/3 = 1$ to A, C, and F. It would indeed be generally the case that the connectivity values would not significantly change depending on ROI size, unless a voxel with a very strong connectivity to a different region is added to the ROI. However, this approach lacks the additional benefit of applying histogram equalization to the connectivity maps, namely, contrast enhancement.




Figure 5.5: Comparison of connectivity maps with linear normalization and histogram equalization. Color map denotes connectivity probabilities (blue represents a connectivity probability of 0.0, whereas red represents a connectivity probability of 1.0). With a linear normalization, the connectivity values on the surfaces are typically weak everywhere outside the initial ROI, and the connectivity map does not provide any information that can be useful for establishing correspondence between the two subjects with the exception of a small patch in the frontal area. The normalization via histogram equalization reveals several cortical regions that exhibit similar connectivity patterns (notably near the Sylvain fissure, the motor cortex and the somatosensory cortex, in addition to the frontal areas, all plausible connections based on anatomical knowledge). This information can be subsequently used for improving cortical correspondence.

Since ideal features for correspondence optimization should have strong contrast, I chose the histogram equalization approach rather than the division by the maximum connectivity. Figure 5.5 illustrates the advantage of the histogram equalization approach over simple linear normalization for creating suitable connectivity features for correspondence optimization.

It should be noted that the contrast enhancement is not a necessary step but merely is an additional performance booster for the correspondence algorithm. As a general principle, features with more distinct values add more useful information to the correspondence algorithm (as compared to the extreme case where $f(x) = c$

with a constant value c for all x locations). However, if the connectivity data is such that histogram equalization would reduce contrast rather than enhance it (such as might be the case with nonzero connectivity on the majority of the cortex), linear normalization is likely to be more suitable than a histogram equalization.

5.3 Results

I applied the methodology presented to a dataset of 9 healthy subjects with 1.5T DTI scans as well as structural MRI scans (this is the same dataset as the second population described in Section 4.5). The DTI scans had 60 gradient directions and 10 baselines, with $b = 700s/mm^2$ and $(2mm)^3$ voxel size. Cortical surfaces were reconstructed via FreeSurfer from T1 images that have been corrected for bias via the itkEMS tool using both T1 and T2 scans. No manual interventions were made to the FreeSurfer pipeline. Only left hemispheres were used.

The fiber tracking was performed using Ngo’s implementation of the stochastic tractography algorithm [11] laid out by Friman et al. [13], as described in Sec. 5.1.2. I use the output of the itkEMS algorithm described in Section 4.5, coregistered with the DWI data (by registering the T2-weighted image with the DWI baseline using an affine transformation with 12 degrees of freedom), as the necessary soft WM segmentation input. The ROI’s are obtained from the FreeSurfer segmentation.

I compare several methods of correspondence computation: location-based particle system (denoted *XYZ-entropy* throughout the following), the generalized entropy system based on sulcal depth only (*SD-entropy*) and the generalized entropy based on connectivity combined with sulcal depth and location (*XYZ-SD-PC-entropy*). For XYZ-SD-PC-entropy, I use probabilistic connectivity measurements to the corpus callosum, the brainstem and the left caudate, with the ROI segmentations provided

by FreeSurfer. Therefore, the local function used by the entropy system is

$$f(x, y, z) = \{x, y, z, sd(x, y, z), conn_{cc}(x, y, z), conn_{bs}(x, y, z), conn_{caud}(x, y, z)\},$$

where *conn* denotes the connectivity probability to each ROI as computed by the tractography component.

As described in Sec. 3.4 and 4.2, each feature channel is weighted such that the variance of the feature values across the population has a mean value of 1.0 across the surface, and thus weighted equally. This is useful to prevent features with large absolute values (such as spatial location, typically in the range $[-128..128]$) from dominating the features with small absolute values (such as connectivity probabilities, in the range $[0..1]$).

I compare these performance measures with results obtained from FreeSurfer, which is one of the most commonly used methods for cortical correspondence. As discussed in Sec. 2.1.1, the FreeSurfer correspondence method is essentially a two-step correspondence computation: it is initialized with spatial correspondence (while computing a spherical parameterization of the surfaces) and then it optimizes the sulcal depth correspondence. This method is also fundamentally different from the entropy-based techniques in that it focuses on a pairwise correspondence (subject to average), whereas the entropy-based method emphasizes a groupwise approach.

As in Chapter 4, I analyze two different measurements for correspondence evaluation purposes, namely, sulcal depth and cortical thickness. As already pointed out in Sec. 4.5, the sulcal depth evaluation is biased and is only used to demonstrate that the generalized entropy framework can be effectively used to achieve a better sulcal depth match better than FreeSurfer, which is also optimizing for sulcal depth similarity. However, in order to evaluate the quality of the correspondence results in an unbiased fashion, I analyze the cortical thickness as an unbiased measurement.

	Mean of Sulcal Depth Variance	Mean of Cortical Thickness (CT) Variance	Mean of CT Variance in Temporal Lobe
XYZ-entropy	0.109 (0.003)	0.262 (0.006)	0.275 (0.01)
SD-entropy	0.009 (8e-5)	0.253 (0.008)	0.233 (0.01)
XYZ-SD-PC-entropy	0.110 (0.003)	0.248 (0.006)	0.211 (0.01)
FreeSurfer	0.039 (0.003)	0.277 (0.010)	0.264 (0.04)

Figure 5.6: Average variances of cortical thickness and sulcal depth measurements across the whole cortical surface as well as across the temporal lobe, given different correspondence maps. The values in parentheses show the standard error associated with each mean.

For the cortical thickness, in addition to local variability reduction, the generalization and specificity measurements, discussed in detail in Sec. 2.3.3, are reported to further analyze the correspondence quality.

In general, XYZ-SD-PC-entropy method is expected to produce improved correspondence over certain regions (specifically, the ones that are strongly identifiable by fiber tract connections to subcortical regions chosen as ROI's) and smaller improvement in other regions where no relevant additional local information is provided. The goal of this approach is to improve local cortical correspondence in given regions by using relevant data. Note that it would be up to each individual application to define what regions are important for the given context, and what additional data can be used to improve the correspondence in these critical regions.

In particular, for this study, because fiber connections to the temporal lobe from both the corpus callosum and the left caudate are observed, significantly improved correspondence is expected in the temporal lobe. Therefore, in addition to the cortical thickness variance averaged across the entire surface, I also report the same values computed over the temporal lobe only.

The local variability reduction analysis is summarized in Fig. 5.6. FreeSurfer and SD-entropy both yield much tighter sulcal depth distributions than the other




Figure 5.7: Cortical thickness based generalization and specificity comparison. For both evaluation metrics, a lower value indicates a better correspondence. Therefore, it can be seen that the connectivity-based method outperforms the other two algorithms regarding these two metrics.

methods, which is to be expected since these methods optimize the sulcal depth similarity itself. Even so, it should be noted that SD-entropy generates a correspondence map that matches sulcal depth values much better. However, the cortical thickness measurements show, as seen in the second column, that the connectivity-based entropy system yields the tightest cortical thickness distribution overall among all the different algorithms tested. In particular, the correspondence quality was significantly enhanced in the temporal lobe, which appears to present a ‘problem area’ for most of the other algorithms (as evidenced by higher than average cortical thickness variance). The incorporation of additional connectivity information clearly improves correspondence.

Fig. 5.7 summarizes the generalization and the specificity analysis based on cortical thickness for the FreeSurfer, XYZ-entropy and XYZ-SD-PC-entropy methods. As pointed out earlier, a lower value of generalization or specificity is associated with a better correspondence map. Thus, the connectivity-based entropy system clearly has the best generalization ability, as well as slightly better specificity properties, independent of the number of shape eigenmodes used (M). Note that generalization ability is a more important measure for shape analysis purposes, whereas specificity is more oriented towards shape modeling.

In this chapter, I presented a novel method that allows using data from diffusion weighted images along with structural MRI scans in a cortical correspondence setting. This technique allows for the fiber connectivity information extracted from the DWI to be effectively projected on the cortical surface using a novel surface deflation technique. I use the entropy-based dynamic particle framework to seamlessly integrate this information with geometrical cues, such as spatial location and sulcal depth, in order to improve cortical correspondence.

My results illustrate the powerful generalizability of this technique: the user can improve the correspondence in all regions of the cortical surface, as long as strongly identifiable local features can be provided. Such local features can be extracted from structural images, DTI, or other imaging modalities such as magnetic resonance angiography (MRA).

Chapter 6

Discussion

6.1 Summary of contributions and thesis

At the end of this dissertation, I revisit the contributions and the thesis that were presented in Section 1.5 to summarize how each one was addressed. The contributions of this dissertation are as follows:

1. *I demonstrate that the use of an approach allowing for the incorporation of arbitrary local features into the similarity metric to be used for correspondence optimization enhances correspondence, as measured by objective evaluation criteria.*

This claim summarizes my thesis and has been dealt with throughout this dissertation. In particular, a novel parametric correspondence algorithm (presented in Chapter 3) and a novel nonparametric correspondence algorithm (presented in Chapter 4) were developed, both of which allow using arbitrary local features for defining the similarity metric. These algorithms were applied to a variety of anatomical surfaces (caudate, femur, striatum, lateral ventricle, cortex) using a variety of local features (curvature, sulcal depth, fiber connectivity). In all these experiments, objective evaluation criteria presented in Sec. 2.3.3 were used to demonstrate that correspondence quality significantly improves when

additional local information is used. These results were presented in Sec. 3.4, 4.5 and 5.3.

Thus, the usage of additional features improves anatomical correspondence. The key insight is that both frameworks presented here allow for the use of arbitrary local features as long as their absolute difference defines a metric in the feature space. Therefore, given a particular application, the user can incorporate all the information about the surfaces, whether these come from geometric computations, different image modalities or any other source, depending on what is available as well as what the user considers to be an anatomically meaningful similarity criterion.

2. *I present a novel parametric groupwise correspondence optimization method that allows using arbitrary local features for establishing correspondence.*

Chapter 3 presents an extension to the Minimum Description Length (MDL) algorithm that incorporates local features into the correspondence optimization. This is accomplished by extending MDL’s data matrix to represent not just spatial locations of sample points but the feature values at these sample locations.

3. *I demonstrate that using geometric information, such as local curvature measures, as additional local features improves correspondence quality when the objects in the population exhibit complex geometry.*

The original MDL algorithm that only considers spatial proximity of sample points typically performs poorly on objects of complex geometry, suggesting that additional geometric information can improve correspondence for such objects. In Sec. 3.4, I show that using curvature measures in addition to spatial location significantly improves correspondence quality for objects of complex geometry such as the femur and the striatum. For objects of simpler geometry,

such as the caudate and the lateral ventricle, no significant improvement was observed, since location-based MDL already performs quite well on these objects. However, the generalized MDL technique has limited applicability to large datasets with high numbers of vertices due to computational complexity associated with both the algorithm itself as well as with the necessary preprocessing stage to generate the surface parameterizations.

4. *I present a novel nonparametric groupwise correspondence optimization method that allows using arbitrary local features for establishing correspondence.*

The entropy-based particle framework, presented in Chapter 4, is a nonparametric groupwise correspondence optimization method. I extended it to incorporate arbitrary local features. This method is far more efficient compared to MDL and can process large datasets in relatively short time. This is an important consideration since the computational cost of MDL becomes prohibitive for large datasets as in the case of human cortical surfaces.

5. *I show that this nonparametric groupwise correspondence technique can be applied to the human cortex despite the geometric challenges presented by the convoluted surface; inflation of the surfaces as a preprocessing step solves this problem by producing surfaces smooth enough to avoid these challenges. Furthermore, I show that cortical correspondence significantly improves when sulcal depth is used as an additional local feature.*

Even though computational time is no longer an issue when the nonparametric correspondence approach is used, the human cortex still presents a challenging application target, because smoothly and slowly varying tangent planes is an underlying assumption for this algorithm. Circumventing this assumption by incorporating geodesic distances would lead to significant increase in computational time, making the method unsuitable for large datasets. In Sec. 4.3, I

showed that inflating the cortical surface sufficiently solves this problem, as reflected by the experimental results presented in Fig. 4.6. It can be seen that the particle framework performs better than FreeSurfer even without the additional local features (*XYZ-entropy*). Moreover, using sulcal depth to incorporate geometric information further improves cortical correspondence, as demonstrated on two different datasets in Sec. 4.5.

6. *I present a novel framework for integrating white matter (WM) fiber connectivity information into cortical correspondence, the first such method that uses fiber connectivity patterns to establish structural correspondence. To this end, I compute probabilistic connectivity maps from diffusion weighted images via a stochastic tractography algorithm. I project these connectivity values to the cortical surface by a new cortical deflation algorithm. I present empirical evidence showing that using connectivity features enhances cortical correspondence.*

This is the featured result of this dissertation and the main focus of Chapter 5. Structural MRI scans show WM as a mostly homogeneous region, making it impossible to infer information about WM fiber tracts. Such information on the WM fiber tracts can be determined from diffusion-weighted images. I presented a method for mapping the fiber tract structure to the cortical surface and used this information for correspondence optimization. This is achieved by executing a stochastic tractography algorithm to generate connectivity maps to subcortical ROI's. I showed that using the connectivity values at the WM/GM boundary is undesirable since the highly isotropic diffusion in this region makes the tractography results strongly dependent on the sulcal depth. Instead, I developed a mean curvature evolution algorithm constrained by sulcal depth and local curvature to deflate the cortical surface. Connectivity values at this deflated surface are used as correspondence features.

Sec. 5.3 presented empirical evidence that using these connectivity features

in addition to structural information (i.e., spatial location and sulcal depth) significantly improves correspondence quality. Correspondence is enhanced especially in the regions that show strong connectivity to the chosen set of ROI's. This finding supports the argument that given a particular application's context, the user can choose the features whose similarity is most anatomically meaningful to drive the correspondence optimization. These features can be chosen to cover the entire surface, such as the sulcal depth; they can also be only provided for the regions where an accurate correspondence is critical, such as the connectivity values that are non-zero only in limited surface patches.

7. *I develop open-source software that implements all the above techniques, as well as a visualization tool that allows qualitative examination of the surfaces, the local features associated with them and the surface samples used in the correspondence algorithm.*

Throughout this dissertation I presented results obtained using the parametric (MDL) and nonparametric (entropy-based) correspondence optimization software that I extended to allow the generalized similarity functions. Furthermore, I developed software that implements the analytical computation of various geometric features (curvature measures, first and second fundamental forms, etc.) based on SPHARM representation, in order to compute the various geometric feature candidates discussed in Sec. 3.3. For the entropy-based particle correspondence method, I implemented a front-to-end correspondence framework that takes surface meshes as input, computes the level-set representation used by the correspondence algorithm and reconstructs the surfaces after the correspondence optimization, using the technique discussed in Sec. 4.4. I also implemented the brain deflation algorithm discussed in Sec. 5.2.2. All the surface visualizations in this dissertation were created using KWMeshVisu, a tool that I developed for visualizing surfaces, the local features associated with these

surfaces whether in scalar, vector or tensor form, as well as the surface samples used for representing the correspondence.

All the above software is open-source and publicly available online through the UNC NeuroLib repository [67], both as stand-alone programs and as Slicer3 modules. Articles describing some of this software have been published [68, 45] to facilitate the dissemination of these tools.

Thesis: Statistical shape analysis of anatomical structures, which is essential to understanding the structural changes in anatomy in various stages of growth or disease, requires establishing accurate correspondence across object populations. However, anatomical correspondence is rarely a direct result of spatial proximity of sample points on the surface. A generalized correspondence framework that incorporates the similarity of non-spatial local features provides a more accurate correspondence of sample points across populations of surfaces. In particular, incorporating features based on cortical geometry as well as the fiber connectivity of the white matter significantly improves correspondence of the human cortical surfaces.

I showed throughout this dissertation that correspondence methods that are based only on spatial location produce suboptimal results. In particular, I showed that location-only MDL performs poorly on populations of surfaces of complex geometry. Several datasets were presented in Section 3.4 to support the argument that using local curvature information in addition to spatial location significantly improves correspondence in these populations. This is the result summarized in Claim 3. Furthermore, I showed that in the nonparametric setting presented in Chapter 4, incorporating geometric features improves correspondence compared to methods that only use spatial location. In particular, I showed that cortical correspondence can be improved using sulcal depth as a feature to drive the similarity measure (Section 4.5), as summarized in Claim 4 and 5. Finally, in Chapter 5 I presented a methodology that allows using DTI-based connectivity maps as a correspondence feature. As

discussed above regarding Claim 6, I showed that incorporating these connectivity maps as a correspondence feature in addition to spatial location and sulcal depth significantly improves cortical correspondence, as evidenced by significantly reduced variability in sulcal depth and cortical thickness in the population as well as improved generalization properties (Section 5.3).

I have thus shown that incorporating local features into the similarity metric significantly improves correspondence quality, as assessed by objective evaluation criteria. I presented a framework where arbitrary features can be used for correspondence. The only constraint is that the absolute difference of the feature should define a metric in the feature space. The choice of particular features to be used for correspondence is otherwise left to the user. The user is in the best position to decide what features constitute anatomically meaningful correspondences depending on the particular application context. This dissertation explored the effects of using certain geometric features as well as DTI-based connectivity features. Some further alternatives will be discussed in the next section; however, it is impossible to make an exhaustive list of features that might prove useful for correspondence optimization, and it therefore is the user's responsibility to decide what features provide relevant information on the population.

One of the striking results of my research is that when the connectivity features were used, correspondence was improved significantly even when the additional information was only available in limited patches of the surface. The connectivity data illustrates this finding: even though the connectivity probability is simply 0 for most of the cortical surface, the correspondence is nevertheless improved, particularly in those regions where the probability is nonzero. This finding further suggests that the choice of features should be indeed left to the user to allow the framework to be adapted to different application contexts. For instance, if the correspondence accuracy is more critical in the frontal lobe than elsewhere for a particular study, the user

might choose DTI features based on ROI's that are known to be strongly connected to the frontal lobe.

6.2 Future work and discussion

There are a number of questions and directions for future research related to this dissertation. Some of these are reviewed in this section.

There are several different kinds of local attributes that can be used for correspondence that were not discussed in this dissertation. Finding suitable mappings of these attributes to the surfaces and investigating whether they do indeed improve correspondence quality is a major future research direction. Three kinds of attributes that stand out among these are discussed in the following.

Magnetic resonance angiography (MRA) is a technique for imaging blood vessels. Since blood vessels carrying oxygen and other nutrients are closely coupled with the development and the function of the body, it is conceivable that incorporating knowledge about the vasculature structure into the similarity metric would improve correspondence. This information can be extracted from MRA images and mapped onto the surface in form of distance to the nearest vessel, distance to a particular set of major vessels or the nearest vessel of a given minimum size.

Image intensities from MRI scans is another candidate for useful correspondence features. These intensities can be themselves used directly, or quantities derived from them, such as image gradients or texture-based features that represent the local intensity patterns, can be used for correspondence optimization. Furthermore, the intensity-based entropy can be either computed on the surfaces, as the other features presented in this dissertation, or on the entire volumes, which would amount to an entropy-based groupwise registration approach. This last effect can be achieved by removing the constraint that the particles must live on the zero level set of the implicit

representation. In situations where the image contrast clearly differentiates between tissue types, using the image intensities as a correspondence feature seems sensible. On the other hand, ignoring the geometry of the surfaces and focusing solely on image registration might lead to insufficient correspondence results, such as misaligned sulci patterns. Using the intensities as an additional feature on the surface rather than the sole feature on the volume might provide a satisfactory middle ground.

Different ways of incorporating the WM fiber structure knowledge into cortical correspondence need to be explored. Using cortical regions rather than subcortical structures as the initial ROI's might produce better connectivity maps since the tracts might not have to pass through the densely populated areas where many fiber-crossings are likely to occur; however, these connections might be weaker and not produce suitable connectivity values. Better connectivity maps would reduce the amount of cortex deflation needed, leading to a more accurate representation of the connectivity of the cortical surface itself.

There are also alternative ways of projecting the connectivity probabilities to the cortical surface instead of using the values on the corresponding deflated surface location. One example is to stop the fiber tracking before the fiber samples reach the region of isotropic diffusion near the cortical boundary and to project the connectivity values to the surface by following the last known direction of the fiber. For this approach, it is necessary to detect which fiber samples will be projected to the surface, because voxels of isotropic diffusion that are located far from the WM/GM boundary are not suitable for this purpose. Furthermore, if only one direction per fiber is used for projection to the surface, it is not guaranteed that all the cortical voxels will be assigned connectivity values, even if they are very close to fiber path endpoints. Using several projection directions and introducing a soft penalty for the angle between the projection direction and the last known fiber direction might be a more appropriate approach.

Whether the cortical surface is deflated or the fiber tracts are followed outwards, it is clear that the projection of the connectivity values to the cortical surface requires some guesswork. This is a limitation of my work inherited from the current DTI and fiber tracking technology; more sophisticated tractography algorithms or better DTI scans can reduce the amount of noise in the connectivity probabilities obtained, leading to a more robust correspondence optimization.

One of the limitations of the generalized correspondence approach that I presented is that the computation time increases with the number of feature channels used. This might not be a problem in smaller datasets or with a relatively low number of features, but it could become prohibitive with larger datasets or when a quick, coarse correspondence computation is aimed for (perhaps as an initialization to an algorithm that requires roughly corresponding surface sample points). Therefore, the user must decide whether the correspondence accuracy gained from using a particular feature is sufficient to justify the increase in computational time.

Since the framework I presented in this dissertation gives the user complete freedom regarding the choice of features, selecting “good” features is very important. However, there is not an objective way of judging how useful a particular feature is for correspondence optimization without actually trying them. In my experience, features that represent different levels of detail complement each other well (i.e., a coarse-level feature and a fine-level feature). Furthermore, features that represent different aspects of the population are more valuable as opposed to several features that represent similar information (i.e., a geometric feature and a connectivity feature as opposed to 3 different curvature features). However, validating these intuitive ideas as well as formulating them as objective criteria for feature selection is beyond the scope of this dissertation and remains as a future research direction. It would be also interesting to investigate whether “default” features can be extracted given all the information that is available regarding the population (such as T1-, T2- and

diffusion-weighted images for each subject). This scenario might be useful in a situation where the user does not have sufficient insight into the application context or might want to perform a preliminary exploration of the data before committing to a final feature set.

The relative weighting of these features is a further area of investigation. In all the experiments presented in this dissertation, I have weighted the features such that each feature channel had the same variance (1.0) across the population. However, it is again up to the user to judge the relative importance of each feature given the application. The weights that would lead to the best correspondence results from a pure optimization point of view can be computed by increasing the dimensionality of the search space; however, not only this would lead to significant increase in computation time but also it would be unclear what such a weighting scheme means from an anatomical point of view. A more sensible approach might be to compute these weights based on the information content of each feature, using ideas from information theory. Such an approach might also help compensate for a poor choice of features, by assigning lower weights to features that contain similar information (such as several curvature measures) so that the other features can be allowed to have a more significant impact on the correspondence results.

In addition to different features that can be used for improving correspondence, other evaluation methods also need to be explored. For the cortical correspondence, functional MRI (fMRI) scans can provide a means of validating the cortical correspondence. Although there is no clear evidence that functional correspondence is equivalent to structural correspondence, it is nonetheless valuable to explore whether functional correspondence is improved when additional features are used, especially in the case of diffusion-based features that offer a significantly different source of information than merely the geometry of the cortex. Whether functional data or cortical thickness (as presented in this dissertation) is used for evaluation, hypothesis

tests can be applied to confirm that the correspondence improvement is statistically significant.

The methodology presented in this dissertation has been also applied to populations other than cortical surfaces, demonstrating a wide range of applications for this work. Lee et al. have applied it to populations of mouse brains in a cortical thickness study. Paniagua et al. have applied it to populations of human mandibles of healthy controls and osteoarthritis patients in a clinical study. Liu et al. have applied it to populations of lung surfaces reconstructed from CT scans. The methodology has successfully computed correspondences for these populations, and the results are in various stages of publication. Only local curvature has been used in addition to spatial position in these studies, and the development of application-specific correspondence features remains future work.

Finally, applying this framework to larger datasets, especially ones that include both healthy and diseased anatomies, is an important next step to validate the methodology presented in this dissertation.

Appendix A: Spherical Harmonics

This appendix presents a mathematical overview of the spherical harmonics basis functions and describes how to use them to represent surfaces. Given the SPHARM representation, it is possible to analytically compute differential surface features such as first and second fundamental form as well as curvatures.

A-1 Surface representation

Spherical harmonic basis functions Y_l^m , $-l \leq m \leq l$ of degree l and order m are defined on $\theta \in [0, \pi] \times \phi \in [0, 2\pi)$ based on the associated Legendre polynomials P_l^m via the following equations:

$$Y_l^m(\theta, \phi) = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} P_l^m(\cos(\theta)) e^{im\phi} \quad (\text{A-1})$$

$$Y_l^{-m}(\theta, \phi) = (-1)^m Y_l^{*m}(\theta, \phi) \quad (\text{A-2})$$

$$P_l^m(w) = \frac{(-1)^m}{2^l l!} (1-w^2)^{\frac{m}{2}} \frac{d^{m+l}}{dw^{m+l}} (w^2-1)^l \quad (\text{A-3})$$

where $*$ denotes complex conjugation.

To express a surface using spherical harmonics, the coordinate functions are decomposed and the surface $\underline{x}(\theta, \phi) = (x(\theta, \phi), y(\theta, \phi), z(\theta, \phi))^T$ takes the form:

$$\underline{x}(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l \underline{c}_l^m Y_l^m(\theta, \phi) \quad (\text{A-4})$$

where the coefficients $\underline{c}_l^m$ are three-dimensional vectors due to the three coordinate functions. These coefficients are obtained by solving a least-squares problem, as discussed in [14] and [15].

A-2 Partial derivatives

The first and second order partial derivatives of the Cartesian coordinates $\underline{x}(\theta, \phi)$ with respect to the spherical coordinates θ and ϕ are necessary for the computation of fundamental forms and therefore curvatures.

The first order partial derivatives $\partial \underline{x}(\theta, \phi)/\partial \theta$ and $\partial \underline{x}(\theta, \phi)/\partial \phi$ can be computed using the partial derivatives of the spherical harmonic basis functions by applying the chain rule.

$$\frac{\partial \underline{x}(\theta, \phi)}{\partial \phi} = \frac{\partial}{\partial \phi} \sum_{l=0}^{\infty} \sum_{m=-l}^l \underline{c}_l^m Y_l^m(\theta, \phi) \quad (\text{A-5})$$

$$= \sum_{l=0}^{\infty} \sum_{m=-l}^l \underline{c}_l^m \frac{\partial Y_l^m(\theta, \phi)}{\partial \phi} \quad (\text{A-6})$$

with

$$\frac{\partial Y_l^m(\theta, \phi)}{\partial \phi} = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} P_l^m(\cos(\theta)) \frac{\partial}{\partial \phi} e^{im\phi} \quad (\text{A-7})$$

$$= \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} P_l^m(\cos(\theta)) (-m \sin(m\phi) + im \cos(m\phi)) \quad (\text{A-8})$$

$$= \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} P_l^m(\cos(\theta)) e^{im\phi} im \quad (\text{A-9})$$

$$= im Y_l^m(\theta, \phi) \quad (\text{A-10})$$

Similarly,

$$\frac{\partial \underline{x}(\theta, \phi)}{\partial \theta} = \sum_{l=0}^{\infty} \sum_{m=-l}^l \underline{c}_l^m \frac{\partial Y_l^m(\theta, \phi)}{\partial \theta} \quad (\text{A-11})$$

with

$$\frac{\partial Y_l^m(\theta, \phi)}{\partial \theta} = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} \frac{dP_l^m(\cos(\theta))}{d\theta} (-\sin(\theta)) e^{im\phi} \quad (\text{A-12})$$

The Legendre polynomial's derivative can be computed as

$$\begin{aligned}\frac{dP_l^m(w)}{dw} &= \frac{(-1)^m}{2^l l!} [(1-w^2)^{\frac{m}{2}-1} \frac{m}{2} (-2w)] \frac{d^{m+l}}{dw^{m+l}} (w^2-1)^l \\ &\quad + \frac{(-1)^m}{2^l l!} (1-w^2)^{\frac{m}{2}} \frac{d^{m+l+1}}{dw^{m+l+1}} (w^2-1)^l\end{aligned}\tag{A-13}$$

$$= P_l^m(w) \frac{-mw}{1-w^2} + P_l^{m+1}(w) \frac{-1}{\sqrt{1-w^2}}\tag{A-14}$$

Substituting,

$$\begin{aligned}\frac{\partial Y_l^m(\theta, \phi)}{\partial \theta} &= -\sin\theta \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} e^{im\phi} \\ &\quad (P_l^m(\cos(\theta)) \frac{-m\cos(\theta)}{\sin^2(\theta)} + P_l^{m+1}(\cos(\theta)) \frac{-1}{\sin(\theta)})\end{aligned}\tag{A-15}$$

$$= m\cot(\theta) Y_l^m(\theta, \phi) + \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} e^{im\phi} P_l^{m+1}(\cos(\theta))\tag{A-16}$$

$$= m\cot(\theta) Y_l^m(\theta, \phi) + \sqrt{(l-m)(l+m+1)} e^{-i\phi} Y_l^{m+1}(\theta, \phi)\tag{A-17}$$

The second order partial derivatives can be then computed using the first order derivatives.

$$\frac{\partial^2 \underline{x}(\theta, \phi)}{\partial \phi^2} = \sum_{l=0}^{\infty} \sum_{m=-l}^l \underline{c}_l^m \frac{\partial^2 Y_l^m(\theta, \phi)}{\partial \phi^2}\tag{A-18}$$

with

$$\frac{\partial^2 Y_l^m(\theta, \phi)}{\partial \phi^2} = \frac{\partial}{\partial \phi} \left(\frac{\partial Y_l^m(\theta, \phi)}{\partial \phi} \right)\tag{A-19}$$

$$= \frac{\partial}{\partial \phi} (im Y_l^m(\theta, \phi))\tag{A-20}$$

$$= im \frac{\partial Y_l^m(\theta, \phi)}{\partial \phi}\tag{A-21}$$

$$= -m^2 Y_l^m(\theta, \phi)\tag{A-22}$$

Similarly,

$$\frac{\partial^2 \underline{x}(\theta, \phi)}{\partial \phi \partial \theta} = \sum_{l=0}^{\infty} \sum_{m=-l}^l \underline{c}_l^m \frac{\partial^2 Y_l^m(\theta, \phi)}{\partial \phi \partial \theta} \quad (\text{A-23})$$

with

$$\frac{\partial^2 Y_l^m(\theta, \phi)}{\partial \theta \partial \phi} = \frac{\partial}{\partial \theta} \left(\frac{\partial Y_l^m(\theta, \phi)}{\partial \phi} \right) \quad (\text{A-24})$$

$$= \frac{\partial}{\partial \theta} (im Y_l^m(\theta, \phi)) \quad (\text{A-25})$$

$$= im(m \cot(\theta) Y_l^m(\theta, \phi) + \sqrt{(l-m)(l+m+1)} e^{-i\phi} Y_l^{m+1}(\theta, \phi)) \quad (\text{A-26})$$

Finally,

$$\frac{\partial^2 \underline{x}(\theta, \phi)}{\partial \theta^2} = \sum_{l=0}^{\infty} \sum_{m=-l}^l \underline{c}_l^m \frac{\partial^2 Y_l^m(\theta, \phi)}{\partial \theta^2} \quad (\text{A-27})$$

with

$$\frac{\partial^2 Y_l^m(\theta, \phi)}{\partial \theta^2} = \frac{\partial}{\partial \theta} \left(\frac{\partial Y_l^m(\theta, \phi)}{\partial \theta} \right) \quad (\text{A-28})$$

$$= \frac{\partial}{\partial \theta} (m \cot(\theta) Y_l^m(\theta, \phi) + \sqrt{(l-m)(l+m+1)} e^{-i\phi} Y_l^{m+1}(\theta, \phi)) \quad (\text{A-29})$$

$$\begin{aligned} &= m \cot(\theta) \frac{\partial}{\partial \theta} Y_l^m(\theta, \phi) - m \csc^2(\theta) Y_l^m(\theta, \phi) + \sqrt{(l-m)(l+m+1)} e^{-i\phi} \\ &\quad ((m+1) \cot(\theta) Y_l^{m+1}(\theta, \phi) + \sqrt{(l-m-1)(l+m+2)} e^{-i\phi} Y_l^{m+2}(\theta, \phi)) \end{aligned} \quad (\text{A-30})$$

Let $a = \sqrt{(l-m)(l+m+1)}$ and $b = \sqrt{(l-m-1)(l+m+2)}$ for brevity. Then,

$$\begin{aligned} \frac{\partial^2 Y_l^m(\theta, \phi)}{\partial \theta^2} &= m \cot(\theta) (m \cot(\theta) Y_l^m(\theta, \phi) + a e^{-i\phi} Y_l^{m+1}(\theta, \phi)) - m \csc^2(\theta) Y_l^m(\theta, \phi) \\ &\quad + a e^{-i\phi} (m+1) \cot(\theta) Y_l^{m+1}(\theta, \phi) + a b e^{-2i\phi} Y_l^{m+2}(\theta, \phi) \end{aligned} \quad (\text{A-31})$$

$$\begin{aligned} &= Y_l^m(\theta, \phi) m (m \cot^2(\theta) - \csc^2(\theta)) + Y_l^{m+1}(\theta, \phi) a e^{-i\phi} \cot(\theta) (1+2m) \\ &\quad + Y_l^{m+2}(\theta, \phi) a b e^{-2i\phi} \end{aligned} \quad (\text{A-32})$$

A-3 Surface normal

The unit surface normal is given by the cross product of the two tangent vectors

$$\underline{n} = \frac{\frac{\partial \underline{x}}{\partial \theta} \times \frac{\partial \underline{x}}{\partial \phi}}{\left| \frac{\partial \underline{x}}{\partial \theta} \times \frac{\partial \underline{x}}{\partial \phi} \right|}. \quad (\text{A-33})$$

A-4 First and second fundamental forms

The scalar coefficients of the first and second fundamental form, (E, F, G) and (e, f, g) respectively, can be calculated from the following dot-products [69]:

$$E = \frac{\partial \underline{x}}{\partial \theta} \cdot \frac{\partial \underline{x}}{\partial \theta} \quad (\text{A-34})$$

$$F = \frac{\partial \underline{x}}{\partial \theta} \cdot \frac{\partial \underline{x}}{\partial \phi} \quad (\text{A-35})$$

$$G = \frac{\partial \underline{x}}{\partial \phi} \cdot \frac{\partial \underline{x}}{\partial \phi} \quad (\text{A-36})$$

$$e = \frac{\partial^2 \underline{x}}{\partial \theta^2} \cdot \underline{n} \quad (\text{A-37})$$

$$f = \frac{\partial^2 \underline{x}}{\partial \theta \partial \phi} \cdot \underline{n} \quad (\text{A-38})$$

$$g = \frac{\partial^2 \underline{x}}{\partial \phi^2} \cdot \underline{n} \quad (\text{A-39})$$

A-5 Curvature metrics

The principal curvatures κ_1 and κ_2 can be computed from [69]

$$\kappa_i = \frac{e + 2f\lambda_i + g\lambda_i^2}{E + 2F\lambda_i + G\lambda_i^2}, i = 1..2 \quad (\text{A-40})$$

where λ_1 and λ_2 are the solutions to the quadratic equation

$$(Fg - Gf)\lambda^2 + (Eg - Ge)\lambda + Ef - Fe = 0 \quad (\text{A-41})$$

Once the principal curvatures are obtained, mean curvature H , Gaussian curvature K , shape index S and curvedness C are

$$H = \frac{1}{2}(\kappa_1 + \kappa_2) \quad (\text{A-42})$$

$$K = \kappa_1 \kappa_2 \quad (\text{A-43})$$

$$S = -\frac{2}{\pi} \tan^{-1} \left(\frac{\kappa_1 + \kappa_2}{\kappa_1 - \kappa_2} \right) \quad (\text{A-44})$$

$$C = \frac{2}{\pi} \ln \left(\sqrt{\frac{\kappa_1^2 + \kappa_2^2}{2}} \right) \quad (\text{A-45})$$

Bibliography

- [1] Davies, R., Twining, C., Cootes, T., Waterton, J.: 3D statistical shape models using direct optimisation of description length. European Conference on Computer Vision, Proceedings (Part III) (Jan 2002) 1 – 17 2, 32
- [2] Davies, R., Twining, C., Cootes, T., Waterton, J., Taylor, C.: A minimum description length approach to statistical shape modeling. Medical Imaging, IEEE Transactions on **21**(5) (May 2002) 525 – 537 2, 24, 31, 32
- [3] Cates, J., Meyer, M., Fletcher, P.T., Whitaker, R.: Entropy-based particle systems for shape correspondence. MFCA Mathematical Foundations of Computational Anatomy Workshop at MICCAI (Jul 2006) 90 – 99 4, 25, 51, 55
- [4] Cates, J., Fletcher, P.T., Styner, M., Shenton, M.E., Whitaker, R.T.: Shape modeling and analysis with entropy-based particle systems. Information Processing in Medical Imaging, Proceedings **4584** (Feb 2007) 333 – 345 4, 25, 51, 52, 55
- [5] Segonne, F., Dale, A., Busa, E., Glessner, M., Salat, D., Hahn, H., Fischl, B.: A hybrid approach to the skull stripping problem in MRI. NeuroImage **22**(3) (Jul 2004) 1060 – 1075 5, 18
- [6] Dale, A.M., Fischl, B., Sereno, M.I.: Cortical surface-based analysis, I: Segmentation and surface reconstruction. NeuroImage **9**(2) (Jan 1999) 179 – 194 5, 18, 59
- [7] Fischl, B., Sereno, M., Tootell, R., Dale, A.: High-resolution intersubject averaging and a coordinate system for the cortical surface. Human Brain Mapping **8**(4) (Nov 1999) 272 – 284 5, 18
- [8] Fischl, B., Sereno, M.I., Dale, A.M.: Cortical surface-based analysis, II: Inflation, flattening, and a surface-based coordinate system. NeuroImage **9**(2) (Jan 1999) 195 – 207 5, 18
- [9] Fischl, B., Liu, A., Dale, A.M.: Automated manifold surgery: Constructing geometrically accurate and topologically correct models of the human cerebral cortex. Medical Imaging, IEEE Transactions on **20**(1) (Mar 2001) 70 – 80 5, 18
- [10] Fischl, B., Dale, A.M.: Measuring the thickness of the human cerebral cortex from magnetic resonance images. PNAS **97**(20) (Sep 2000) 11050 – 11055 5, 18
- [11] Ngo, T.: A stochastic tractography system and applications. Master’s thesis, Massachusetts Institute of Technology (May 2007) 7, 73, 82

- [12] Friman, O., Westin, C.F.: Uncertainty in white matter fiber tractography. International Conference on Medical Image Computing and Computer-Assisted Intervention, Proceedings **8**(Pt 1) (Jan 2005) 107 – 114 7, 73
- [13] Friman, O., Farneback, G., Westin, C.: A bayesian approach for stochastic white matter tractography. Medical Imaging, IEEE Transactions on **25**(8) (Aug 2006) 965 – 978 7, 73, 82
- [14] Brechbühler, C., Gerig, G., Kubler, O.: Parametrization of closed surfaces for 3-D shape description. Computer Vision and Image Understanding **61**(2) (Feb 1995) 154 – 170 15, 32, 99
- [15] Styner, M., Oguz, I., Xu, S., Brechbuhler, C., Pantazis, D., Levitt, J.J., Shenton, M.E., Gerig, G.: Framework for the statistical shape analysis of brain structures using SPHARM-PDM. Insight Journal (Jul 2006) 15, 99
- [16] Tagare, H.D.: Shape-based nonrigid correspondence with application to heart motion analysis. IEEE Transactions on Medical Imaging **18**(7) (July 1999) 570 – 579 16
- [17] Tosun, D., Prince, J.: Cortical surface alignment using geometry driven multi-spectral optical flow. Information Processing in Medical Imaging, Proceedings **3565** (Jul 2005) 480 – 492 17, 40
- [18] Koenderink, J.J.: Solid Shape. Artificial Intelligence Series. MIT Press, Cambridge, MA, USA (1990) 17, 18, 37, 38, 39, 40
- [19] Tosun, D., Rettmann, M., Prince, J.: Mapping techniques for aligning sulci across multiple brains. Medical Image Analysis **8**(3) (Sep 2004) 295 – 309 17, 77
- [20] Meier, D., Fisher, E.: Parameter space warping: Shape-based correspondence between morphologically different objects. IEEE Transactions on Medical Imaging **21**(1) (Sep 2002) 31 – 47 18
- [21] Goebel, R., Esposito, F., Formisano, E.: Analysis of Functional Image Analysis Contest (FIAC) data with brainVoyager QX: From single-subject to cortically aligned group general linear model analysis and self-organizing group independent component analysis. Human Brain Mapping **27**(5) (May 2006) 392 – 401 19
- [22] Cachier, P., Mangin, J.F., Pennec, X., Riviere, D., Papadopoulos-Orfanos, D., Regis, J., Ayache, N.: Multisubject non-rigid registration of brain MRI using intensity and geometric features. MICCAI LNCS **2208** (Sep 2001) 734 – 742 20
- [23] Wang, Y., Peterson, B.S., Staib, L.H.: Shape-based 3D surface correspondence using geodesics and local geometry. CVPR **2** (Apr 2000) 644 – 651 20, 21

- [24] Talairach, J., Tournoux, P.: Co-planar stereotaxic atlas of the human brain. 3-Dimensional proportional system: an approach to cerebral imaging. Thieme Medical Publishers, New York (1988) 21
- [25] Steinmetz, H., Furst, G., Freund, H.: Cerebral cortical localization: application and validation of the proportional grid system in MR imaging. *Journal of computer assisted tomography* **13** (Jan 1989) 10 – 19 22
- [26] Hunton, D.L., Miezin, F.M., Buckner, R.L., van Mier, H.I., Raichle, M.E., Petersen, S.E.: An assessment of functional-anatomical variability in neuroimaging studies. *Human Brain Mapping* **4** (Mar 1996) 122 – 139 22
- [27] Ashburner, J., Friston, K.J.: Nonlinear spatial normalization using basis functions. *Human Brain Mapping* **7** (Jan 1999) 254 – 266 22
- [28] Jenkinson, M., Smith, S.: A global optimisation method for robust affine registration of brain images. *Medical Image Analysis* **5** (Jun 2001) 143 – 156 22
- [29] Andersson, J., Smith, S., Jenkinson, M.: FNIRT - FMRIB's non-linear image registration tool. *Human Brain Mapping Poster* 496 (2008) 22
- [30] Christensen, G., Rabbitt, R., Miller, M.: 3D brain mapping using a deformable neuroanatomy. *Physics in Medicine and Biology* **39** (Mar 1994) 609 – 619 23
- [31] Shen, D., Davatzikos, C.: Hammer: Hierarchical attribute matching mechanism for elastic registration. *Medical Imaging, IEEE Transactions on* **21**(11) (Nov 2002) 1421 – 1439 23
- [32] Cootes, T., Taylor, C.J., Cooper, D., Graham, J.: Active shape models - their training and application. *Computer Vision and Image Understanding* **61** (Jan 1995) 38 – 59 23
- [33] Bookstein, F.L.: Landmark methods for forms without landmarks: morphometrics of group differences in outline shape. *Medical Image Analysis* **1** (Jul 1996) 225 – 243 24
- [34] Kotcheff, A.C., Taylor, C.J.: Automatic construction of eigenshape models by direct optimization. *Medical Image Analysis* **2**(4) (Jul 1998) 303 – 314 24
- [35] Davies, R.H.: Learning Shape: Optimal Models for Analysing Shape Variability. PhD thesis, University of Manchester, Manchester, UK (2002) 24, 27
- [36] Ward, A.D., Hamarneh, G.: The groupwise medial axis transform for fuzzy skeletonization and pruning. *Transactions on Pattern Analysis and Machine Intelligence* (2009) 25
- [37] Styner, M., Rajamani, K., Nolte, L., Zsemlye, G., Szekely, G., Taylor, C.J., Davies, R.H.: Evaluation of 3D correspondence methods for model building. *Information Processing in Medical Imaging, Proceedings* **2732** (Aug 2003) 63 – 75 25, 41, 51

- [38] Munsell, B.C., Dalal, P., Wang, S.: Evaluating shape correspondence for statistical shape analysis: A benchmark study. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **30**(11) (Nov 2008) 2023 – 2039 26
- [39] Jeong, J.Y.: Estimation of probability distribution on multiple anatomical object complex. PhD thesis, University of North Carolina, Chapel Hill, NC, USA (2008) 29
- [40] Heimann, T., Wolf, I., Williams, T., Meinzer, H.P.: 3D active shape models using gradient descent optimization of description length. *Information Processing in Medical Imaging, Proceedings* **3565** (Aug 2005) 566 – 577 33
- [41] Thodberg, H.H.: Minimum description length shape and appearance models. *Information Processing in Medical Imaging, Proceedings* **2732** (Aug 2003) 51 – 62 33
- [42] Twining, C., Davies, R.H., Taylor, C.J.: Non-parametric surface-based regularisation for building statistical shape models. *Information Processing in Medical Imaging, Proceedings* (Jul 2007) 738 – 750 35
- [43] Praun, E., Hoppe, H.: Spherical parametrization and remeshing. *SIGGRAPH '03: ACM SIGGRAPH 2003 Papers* (Jul 2003) 340 – 349 35
- [44] Styner, M., Oguz, I., Heimann, T., Gerig, G.: Minimum description length with local geometry. *Proc IEEE Symposium on Biomedical Imaging* (May 2008) 1283 – 1286 36
- [45] Heimann, T., Oguz, I., Wolf, I., Styner, M., Meinzer, H.: Implementing the automatic generation of 3D statistical shape models with ITK. *Open Science Workshop at MICCAI 2006, Insight Journal* (2006) 36, 92
- [46] Styner, M., Lieberman, J.A., McClure, R.K., Weinberger, D.R., Jones, D.W., Gerig, G.: Morphometric analysis of lateral ventricles in schizophrenia and healthy controls regarding genetic and disease-specific factors. *PNAS* **102**(13) (Mar 2005) 4872 – 4877 43
- [47] Huang, X., Lee, Y.Z., Mckeown, M., Gerig, G., Gu, H., Lin, W., Lewis, M.M., Ford, S., Tröster, A.I., Weinberger, D.R., Styner, M.: Asymmetrical ventricular enlargement in Parkinson's disease. *Movement Disorders* **22**(11) (Aug 2007) 1657 – 1660 43
- [48] Aguirre, M.R., Linguraru, M.G., Marias, K., Ayache, N., Nolte, L.P., Ballester, M.A.G.: Statistical shape analysis via principal factor analysis. *IEEE International Symposium on Biomedical Imaging, Proceedings* (Apr 2007) 1216 – 1219 44
- [49] Cover, T.M., Thomas, J.A.: *Elements of Information Theory*. Wiley-Interscience (1991) 51

- [50] Meyer, M.D., George, I.P., Whitaker, R.T.: Robust particle systems for curvature dependent sampling of implicit surfaces. *Proceedings of the International Conference on Shape Modeling and Applications* (2005) 124 – 133 54
- [51] Oguz, I., Cates, J., Fletcher, T., Whitaker, R., Cool, D., Aylward, S., Styner, M.: Entropy-based particle systems and local features for cortical correspondence optimization. *Proc IEEE Symposium on Biomedical Imaging* (May 2008) 1637 – 1640 57
- [52] Thompson, P.M., Toga, A.W.: A surface-based technique for warping 3-Dimensional images of the brain. *Medical Imaging, IEEE Transactions on* **15** (Aug 1996) 402 – 417 59
- [53] Angenent, S., Haker, S., Tannenbaum, A., Kikinis, R.: On the Laplace-Beltrami operator and brain surface flattening. *IEEE TMI* **18**(8) (Oct 1999) 700 – 711 59
- [54] Gu, X., Wang, Y., Chan, T., Thompson, P., Yau, S.T.: Genus zero surface conformal mapping and its application to brain surface mapping. *Medical Imaging, IEEE Transactions on* **23**(8) (Aug 2004) 949 – 958 59
- [55] Prastawa, M., Gilmore, J., W.L., Gerig, G.: Automatic segmentation of MR images of the developing newborn brain. In: *MICCAI*. Volume 9. (Oct 2005) 457 – 466 66
- [56] Oguz, I., Niethammer, M., Cates, J., Whitaker, R., Fletcher, T., Vachet, C., Styner, M.: Cortical correspondence with probabilistic fiber connectivity. *Information Processing in Medical Imaging, Proceedings* (Jul 2009) 651 – 663 69
- [57] Mukherjee, P., Berman, J., Chung, S., Hess, C., Henry, R.: Diffusion tensor MR imaging and fiber tractography: Theoretic underpinnings. *American Journal of Neuroradiology* **29**(4) (Apr 2008) 632 – 641 71, 72
- [58] P.T.Fletcher, Tao, R., Jeong, W.K., Whitaker, R.T.: A volumetric approach to quantifying region-to-region white matter connectivity in diffusion tensor MRI. *Information Processing in Medical Imaging, Proceedings* (Jul 2007) 346 – 358 72, 74
- [59] Lazar, M., Alexander, A.L.: Bootstrap white matter tractography. *NeuroImage* **24** (Jan 2005) 524 – 532 72
- [60] Jones, D., Pierpaoli, C.: Confidence mapping in DT-MRI tractography using a bootstrap approach. *Mag.Res.Med.* **53** (2005) 1143 – 1149 72
- [61] Behrens, T., Woolrich, M., Jenkinson, M., Johansen-Berg, H., Nunes, R., Clare, S., Matthews, P., Brady, J., Smith, S.: Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Mag.Res.Med.* **50** (2003) 1077 – 1088 73

- [62] O'Donnell, L., Haker, S., Westin, C.: New approaches to estimation of white matter connectivity in diffusion tensor MRI: elliptic PDE's and geodesics in a tensor-warped space. MICCAI (2002) 459 – 466 74
- [63] Pichon, E., Westin, C., Tannenbaum, A.: A Hamilton-Jacobi-Bellman approach to high angular resolution diffusion tractography. MICCAI (2005) 180 – 187 74
- [64] Lenglet, C., Prados, E., Pons, J.P., Deriche, R., Faugeras, O.: Brain connectivity mapping using riemannian geometry, control theory and PDE's. SIAM Journal of Imaging Sciences **2**(2) (Apr 2009) 285 – 322 74
- [65] Stieltjes, B., Kaufmann, W., van Zijl, P., Fredericksen, K., Pearlson, G., Solaivappan, M., Mori, S.: Diffusion tensor imaging and axonal tracking in the human brainstem. NeuroImage **14** (Sep 2001) 732 – 735 76
- [66] Meyer, M., Desbrun, M., Schröder, P., Barr, A.: Discrete differential-geometry operators for triangulated 2-manifolds. VisMath (2003) 35 – 57 77, 78
- [67] UNC: NeuroLib. <http://www.ia.unc.edu/dev/> 92
- [68] Oguz, I., Gerig, G., Barre, S., Styner, M.: KWMeshVisu: A mesh visualization tool for shape analysis. Open Science Workshop at MICCAI 2006, Insight Journal (2006) 92
- [69] Graustein, W.C.: Differential Geometry. Dover Publications Inc., New York (1966) 103