

GEOMETRIC METHODS FOR IMAGE REGISTRATION AND ANALYSIS

by

Anand Arvind Joshi

A Dissertation Presented to the
FACULTY OF THE GRADUATE SCHOOL
UNIVERSITY OF SOUTHERN CALIFORNIA
In Partial Fulfillment of the
Requirements for the Degree
DOCTOR OF PHILOSOPHY
(ELECTRICAL ENGINEERING)

August 2008

Copyright 2008

Anand Arvind Joshi

Dedication

I dedicate this thesis to my parents Arvind and Suvarna Joshi.

Acknowledgments

I would like to express my deepest respect and gratitude to my advisor, Prof. Richard Leahy. His contribution has been invaluable; he supported me financially throughout my graduate studies in USC, guided me in selecting and solving research problems, introduced me to distinguished researchers and scientists, provided me a friendly and welcome environment, gave me freedom to select my own research directions and let me enjoy long holidays. He is a great mentor and a source of inspiration for me, and it has really been an honor working with him. I also wish to thank the members of my guidance committee: Dr. Krishna Nayak and Dr. Francis Bonahon for their suggestions and valuable feedback. My appreciations go to Dr. David Shattuck and Dr. Paul Thompson at the University of California, Los Angeles for many fruitful collaborations and discussions. I could not have selected better colleagues than Dr. Abhijit Chaudhari, Sangeetha Somayajula, Sanghee Cho, Joyita Dutta, Dr. Sangtae Ahn, Dr. Quanzheng Li and Dr. Dimitrios Pantazis. We have spend quality time together, discussing research problems as well as life experiences. I must also mention that Abhijit, Quanzheng and Dimitrios motivated me from time to time to look for research problems. I also would like to express my gratitude to Dr. Ilya Eckstein for fruitful collaborations.

In the University of Southern California I have enjoyed the cozy and warm environment of an excellent research lab. I want to thank its members: Sangtae Ahn, Abhijit Chaudhari, Sanghee Cho, Belma Dogdas, Hua Hui, Zheng Li, Sangeetha Somayajula,

Juan Luis Poletti Soto, Evren Asma, YuTeng Chang, David Wheland and Syed Ashrafulla for adding positively to my research and academic experience.

Most importantly, I want to thank my family for providing unlimited support and helping me realize my dreams.

Table of Contents

Dedication	ii
Acknowledgments	iii
Abstract	xii
Chapter 1: Introduction	1
Chapter 2: Cortical Surface Parameterization	5
2.1 Parameterization and the Coordinate System	6
2.1.1 Validation of p -harmonic mappings	9
Chapter 3: Cortical Surface Registration	13
3.1 Thin Plate Splines Registration in the Intrinsic Geometry of the Cortical Surface	16
3.1.1 Mathematical Formulation	16
3.1.2 Discretization Algorithm	20
3.1.3 Bending Energy Minimization	21
3.1.4 Validation TPS surface registration	24
3.2 A Finite Element Method for Simultaneous Registration and Parameter- ization	26
3.2.1 Surface Registration	27
3.2.2 Mathematical Formulation	28
3.2.3 Finite Element Formulation	29
3.2.4 Results and Validation	33
3.3 Optimum Choice of Sulcal Subset for Registration	35
3.3.1 Registration Error	36
3.3.2 Probabilistic Model of the Sulcal Errors	39
3.3.3 Results	41
3.3.4 Discussion	45
Chapter 4: Processing of Data in the Surface Geometry	49
4.1 Image Filtering on Surfaces	50

4.2	Mathematical Formulation	51
4.2.1	Isotropic filtering	52
4.2.2	Anisotropic filtering	53
4.3	Discretization and Numerical Method	55
4.3.1	Discretization Algorithm	55
4.3.2	The Heat Equation in the Intrinsic Geometry	61
4.4	The Heat Kernel as a PDF	62
Chapter 5: Volumetric Registration using Harmonic Maps		66
5.1	Introduction	66
5.2	Problem Statement and Formulation	68
5.3	Indirect Mapping Approach	70
5.3.1	Mathematical Formulation	73
5.3.2	Initialization Procedure	75
5.3.3	Mapping to the Unit Ball $B(0, 1)$	75
5.3.4	Harmonic Mapping Between the Two Brains	77
5.3.5	Implementation	78
5.4	Direct Mapping Approach	79
5.4.1	Mathematical Formulation	80
5.4.2	Harmonic Mapping	80
5.5	Volumetric Intensity Registration	81
5.5.1	Formulation	82
5.5.2	Implementation	84
5.6	Results and Validation	86
5.7	Conclusion	91
Chapter 6: Conclusions and Future Work		94
6.1	Geometric Features and Manual Landmarks based Surface Registration	96
6.2	Registration of DTI images	97
Bibliography		98

List of Figures

1.1	The cortical surface of the human brain depicted on a MR data (top row) and rendered as a surface (bottom row).	2
2.1	Sulcal Tracing Tools	6
2.2	The figure shows the cortical surface and its map to a square. The corpus callosum is constrained to lie on the boundary of the square.	9
2.3	The p -harmonic maps of the left hemisphere of an individual cortex. . .	10
2.4	The figure shows smoothed histograms for angle distortion and area distortion respectively. In the angle distortion plot, angle distortion increases with the value of p . In the area distortion plot, the distortions for $p=4,6,8$ are less than that for $p=2$ and most of the points have small angle distortion only. However there is no observable trend for the value of p in either case.	12
3.1	(a) A cortical surface with hand labeled sulci; (b) A flat map of the two cortical surface. The arrows show connectivity at points along the boundary of the square. Due to the spherical topology of the cortical surface, we can assign to it a coordinate system that allows us to compute partial derivatives across the interhemispheric fissure. (c) Chessboard texture mapped to the surface using the square maps.	15
3.2	(upper) The figure shows the warping field computed on the surface. The deformation field is smoothly varying. This is achieved because the bending energy regularization was performed in the intrinsic geometry of the surface. The color indicates the magnitude of the deformation. (lower) The thin-plate spline deformation field applied to a regular grid representing left and right hemispheres.	22

3.3	Alignment of the sulcal landmarks: 6 brains are registered to a common cortical surface using their p -harmonic maps in the plane. They are approximately aligned by the p -harmonic maps justifying our small deformation linear model (thin plate bending energy model) which is used for landmark alignment. After applying the covariant TPS deformation field to the surface parameterization, we can see that the sulci show better alignment.	26
3.4	(a),(b) The two cortical surfaces with hand labeled sulci as colored curves; (c),(d) flat maps of a single hemisphere for the two brains without the sulcal alignment constraint; (e) overlay of sulcal curves on the flat maps without alignment; (f),(g) flat maps with sulcal alignment; (h) overlay of sulcal curves on the flat maps with alignment.	30
3.5	RMS error and percentage overlap in the flattened map as a function of σ .	33
3.6	Mapping of sulcal landmarks from 5 subjects to the atlas brain (left) without and (right) with the sulcal alignment constraint.	34
3.7	The complete set of candidate sulcal curves from which we select an optimal subset for constrained cortical registration	37
3.8	(a) Registration of two cortical surfaces based on the flat mapping method; (b) Parcellation of the cortex into regions surrounding the traced sulci; (c) Registration error for two corresponding sulci where $e_n(s)$ are samples of the registration error.	38
3.9	Sample covariance matrices for the x, y, and z components of the registration error, represented as color coded images.	42
3.10	Optimal subsets of sulci for cortical registration. Each row gives the indices of the optimal subset of sulci that minimize the registration error against all other combinations with equal number of constrained curves (also see Fig. 3.7). The three right columns show that the estimated (est.) error is close to the calculated actual (act.) error when actual registrations with the same constrained curves are performed. Our method predicts the registration error both for the training (trn) and the testing (tst) set of brains.	43
3.11	Optimal sulcal sets for 5, 10, and 15 curves.	45

3.12	Top row: subjective selection of 6 curves, with preference on long sulci distant from each other that are expected to minimize cortical registration error; bottom row: optimal sulcal set with the 6 curves selected by our method.	46
4.1	The impulse response of the isotropic smoothing filters are displayed in the parameter space and on the surface [JSTL05]. It can be seen that when the surface metric is used to compute the Laplace-Beltrami, the impulse response kernel is not isotropic in the parameter space, however it is isotropic on the surface.	58
4.2	left: The mean curvature of the cortical surface plotted on a smoothed representation (for improved visualization of curvature in sulcal folds; right: The mean curvature plotted in 2D parameter space for a single cortical hemisphere. Isotropic diffusion blurs the regions as well as the edges separating them while while anisotropic diffusion reduces noise while preserving edges.	59
4.3	The figures shows the heat kernels estimated to fit the two datasets for MEG somatosensory data. For each of the datasets the estimated pdf is displayed in the parameter space and on the cortical surface.	64
4.4	The classifier: Red and Blue regions shows the two decision regions . . .	64
5.1	Cortical surface alignment after using AIR software for intensity based volumetric alignment using a 168 parameter 5^{th} order polynomial. Note that although the overall morphology is similar between the brains, the two cortical surfaces do not align well.	70
5.2	Illustration of our general framework for surface-constrained volume registration. We first compute the map v from brain manifold (N, I) to the unit ball to form manifold (N, h) . We then compute a map \tilde{u} from brain (M, I) to (N, h) . The final harmonic map from (M, I) to (N, I) is then given by $u = v^{-1} \circ \tilde{u}$	74

5.3	Initialization for harmonic mapping from M to N . First we generate flat square maps of the two brains, one for each hemisphere, with pre-aligned sulci. The squares corresponding to each hemispheres are mapped to a disk and the disks are projected onto the unit sphere. We then generate a volumetric maps from each of the brains to the unit ball. Since all these maps are bijective, the resulting map results in a bijective point correspondence between the two brains. However, this correspondence is not optimal with respect to the harmonic energy of maps from the first brain to the second, but is used as an initialization for minimization of (5.6).	76
5.4	Illustration of the deformation induced with respect to the Euclidean coordinates by mapping to the unit ball. Shown are iso-surfaces corresponding to the Euclidean coordinates for different radii in the unit ball. Distortions become increasingly pronounced towards the outer edge of the sphere where the entire convoluted cortical surface is mapped to the surface of the ball.	77
5.5	Schematic of the intensity alignment procedure. Once harmonic maps u^M and u^N are computed, we refine these with intensity driven warps w^M and w^N while imposing constraints so that the final deformations are inverse consistent.	82
5.6	Illustration of the effects of the two stages of volumetric matching is shown by applying the deformations to a regular mesh representing one slice. Since the deformation is in 3D, the third in-paper value is represented by color. (a) Regular mesh representing one slice in the subject; (b) the regular mesh warped by the harmonic mapping which matches the subject cortical surface to the template cortical surface. Note that deformation is largest near the surface since the harmonic map is constrained only by the cortical surface; (c) Regular mesh representing one slice in the harmonically warped subject; (d) the intensity-based refinement now refines the deformation of the template to improve the match between subcortical structures. In this case the deformation is constrained to zero at the boundary and are confined to the interior of the volume.	85
5.7	Examples of direct mapping approach. (a) Original subject volume; (b) original template; (c) registration of subject to template using surface constrained harmonic mapping, note that the surface matches that of the template; (d) intensity-based refinement of the harmonic map of subject to template to complete registration procedure	88

5.8	<p>Volumetric registration using direct mapping approach: (a) Illustration of the extrapolation of the surface mapping to the 3D volume by harmonic mapping. The pairs of surfaces are shown in red and green. The deformation field is represented by placing a regular grid in the central coronal slice of the brain and deforming it according to the harmonic map. The projection of this deformation onto a 2D plane is shown with the in-plane value encoded according to the adjacent color bar. (b) The result of harmonic mapping and linear elastic refinement of the subject brain to the template brain. Note that the inner and outer cortical surfaces, by constraint, are exactly matched. The linear elastic refinement produces an approximate match between subcortical structures. The deformation field here shows the result of cortically constrained intensity-driven refinement. Note that the deformations are zero at the boundary and nonzero in the vicinity of the ventricles, thalamus and other subcortical structures.</p>	90
5.9	<p>Examples of surface constrained volumetric registration. (a) Original subject volume; (b) template; (c) registration of subject to template using surface constrained harmonic mapping, note that the cortical surface matches that of the template; (d) intensity-based refinement of the harmonic map of subject to template</p>	91
6.1	<p>Geometric framework for registration and analysis</p>	95

Abstract

Registration and analysis of neuro-imaging data presents a challenging problem due to the complex folding patterns in the human brain. Specifically, the cortical surface of the human brain can be modeled as a highly convoluted 2D surface. Since it is non-flat, the non-Euclidean geometry of the cortex needs to be accounted for while performing registration and subsequent signal processing of anatomical and functional signals on the cortex. Techniques from differential geometry offer a powerful set of tools to deal with the convoluted nature of the cortex. We present a method based on p-harmonic mapping for performing cortical surface parameterization. A 2D coordinate system induced by the flat mapping is then used to compute the surface metric and discretize derivatives in the surface geometry. For performing inter-subject cortical registration based on sulcal landmarks, we generalize thin-plate splines to non-flat surfaces by using covariant derivatives. We also present an FEM based method for simultaneous parameterization and registration of sulcal landmarks based on elastic energy minimization. The manual effort required for selecting the sulcal landmarks can be minimized if we choose an optimal set of such landmarks. We present a method for optimally selecting a subset of any size from a set of candidate sulcal landmarks and also predict the associated registration error for that subset using conditional distributions. Surface signals from individual brains can be brought to a common atlas surface by using these surface based registration techniques.

Isotropic and anisotropic diffusion filtering methods are formulated for processing of the cortical data. This is performed by using parameterization-based methods which use covariant diffusion operators in the flat space. When the surface data is a point-set on the cortex, we propose a method to quantify its mean and variance with respect to the surface geometry.

The registration techniques presented for surface alignment are extended to volumes to perform full surface and volume registration. This is done by using volumetric harmonic mappings that extend the surface point correspondence to the cortical brain volume. Finally, the volumetric registration is refined by using inverse-consistent linear elastic intensity registration. This set of methods presents a unified framework for registration and analysis of brain signals for inter-subject neuroanatomical studies.

Chapter 1

Introduction

Brain is home to our mind and personality. It houses our cherished memories and future hopes. It orchestrates the symphony of consciousness that gives us purpose and passion, motion and emotion. Understanding the workings of human mind can only be achieved if we understand the structure and function of the human brain.

The outer part of the brain comprises of grey matter which is internally supported by white matter. The two hemispheres of the brain are separated by the central fissure and connect to each other at the corpus callosum. Cerebellum is found at the posterior-inferior part of the brain. Cerebral cortex or simply cortex is the outermost layer of the cerebrum and is the place where most of the neuronal activations take place. The human cerebral cortex is 2-4 mm thick and plays a central role in many complex brain functions including memory, attention, perceptual awareness and language. Due to the relatively small thickness of the cortex, it can be modeled as a 2D highly convoluted surface with more than two third buried in the grooves, called *sulci*. The sulcal folding pattern varies across individuals, however, some of the major sulci are seen across individuals [OKA90]. It is known that the sulci are related to the function of the brain and therefore inter-subject alignment of the cortex should be carried out with the constraint that these sulci are aligned.

Medical imaging modalities acquire various anatomical (CT, MR, etc.) and functional (PET, EEG, MEG, etc.) neuro-imaging data. Intersubject analysis of this data allows us to study group differences and similarities. The anatomical variability across individuals needs to be normalized before such a study can be carried out. Medical




Figure 1.1: The cortical surface of the human brain depicted on a MR data (top row) and rendered as a surface (bottom row).

image registration performs this normalization by aligning the coordinate systems of the various medical images to register them to a common template or atlas. One of the most challenging problems in image registration is the alignment of human brains.

Registration of surface models of the cerebral cortex has important applications in inter-subject studies of neuroanatomical data for mapping and analyzing progression of disorders such as Alzheimer’s disease [TMV⁺01] and studying growth patterns in developing human brains [TMT00, GGL⁺04]. Investigators have studied several anatomical and functional aspects of the human brain such as genetic influences [THdZ⁺02] and the influence of medication and drugs abuse on the structure and function of the brain [NB97, Cha01]. Inter-subject analysis, or intra-subject analysis over a period of time of such data, present difficult problems due to the inter-subject variability and convoluted geometry of the cortical surface. Since most neural and metabolic activity takes place in the cortex, and because the thickness of the cortex is small relative to the resolution

of most functional imaging techniques, it is plausible to model the cortex as a surface rather than as a volume.

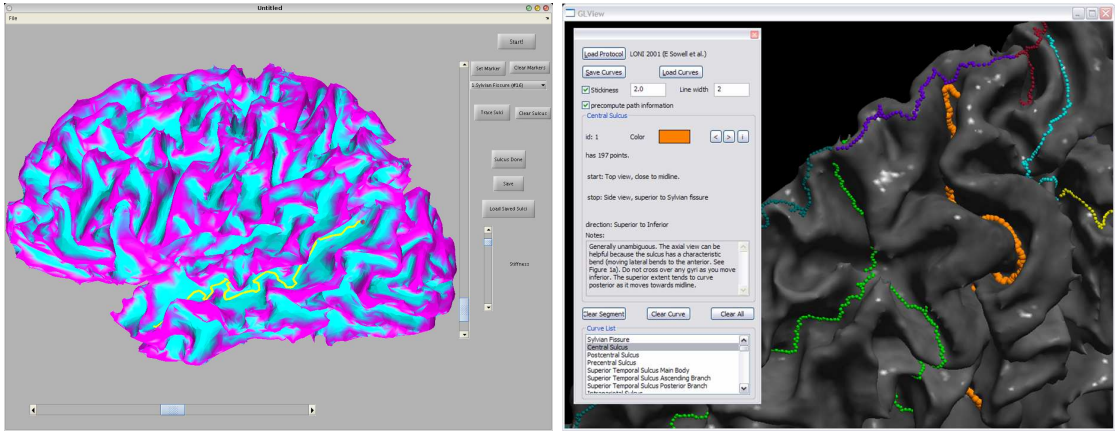
Since the cortex is non-flat, the non-euclidean geometry of the cortex needs to be accounted for while doing registration and subsequent signal processing of anatomical and functional signals on the cortex. In this report we propose a surface parameterization method which computes a 2D coordinate system on the cortex. We use p -harmonic mapping of the cortex to R^2 to assign 2D coordinates to the surface points. This coordinate system is then used to compute the associated surface metric for the assigned coordinates which are then used to discretize derivatives in the surface geometry. In order to bring several brain surfaces in a common template space, we present two surface registration techniques to find a point to point correspondence between the two surfaces. The first technique involves p -harmonic mapping of the cortex to a plane and then reparameterization with thin-plate bending energy as a regularizing function. Alternatively, the second technique incorporates sulcal landmark matching in the parameterization method itself. This is done by using a more general elastic model for parameterization. The point correspondence set by surface registration can be used to bring surface functional signals such as MEG dipoles or neuronal activations, fMRI and anatomical signals such as cortical thickness from individual brains to a common atlas surface. Isotropic and anisotropic diffusion filtering methods are formulated for different kinds of smoothing of such cortical data. When the surface data is a point-set on the cortex, we propose a method to quantify its mean and variance with respect to the surface geometry. The registration technique presented for surface alignment is then extended to volumes to perform full surface and volume registration using harmonic mappings. Inverse-consistent elastic intensity registration is then used to further improve the volumetric alignment. Various validation techniques were used to assess the performance of

the above tools and to compare them with existing methods as presented in subsequent chapters.

Chapter 2

Cortical Surface Parameterization

The surface area of the cerebral cortex is approximately 1570 cm^2 [HSB⁺00]. 60-70% of the surface area is buried in the folds and creases (sulci). There is considerable variability and individual differences in the size, location and extent of the sulci and gyri across human subjects. Bringing multiple brain surfaces into a common coordinate system is helpful in studying variability of these sulcal patterns across subjects, for integrating and averaging functional data across subjects, and in studying patterns in cortical development over time. Since the cortex can be modeled as a convoluted sheet with the topology of a sphere, it is natural to parameterize it using spherical coordinates[FSTD98]. Eck et al.[EDD⁺95] and Kanai et al.[KSK98b] model a triangulated surface as a configuration of springs with one spring placed along each edge of each triangle. The resulting energy functional, the *harmonic energy*, is shown to be a quadratic form and is minimized using gradient descent to transform the surface into a planar disk. Desbrun et al.[DMA02] propose a parameterization technique which uses the *cot* of angles in the given triangulation. The resulting energy functional (Tutte energy) is argued to be a measure of angle distortion and a new parameterization is obtained by minimizing it. Haker et al.[AHTK99] presented a method for conformally mapping the cortical surface to a sphere. Their method uses the Laplace-Beltrami operator and the fact that for a conformal map, the Laplace-Beltrami of the parameterization function is zero everywhere on the surface. Though these methods ensure a perfectly conformal map, the stereographic projections involved can introduce a large amount of length and area distortion.



(a) Matlab tool for marking sulcal anatomy

(b) BrainSuite's sulcal tracing module

Figure 2.1: Sulcal Tracing Tools

Circle packing is introduced as a parameterization method in [HSB⁺00]. Analytic surfaces can be approximated by circle packing, but for general surfaces, the surface packing method considers only the connectivity and not geometry [WGH⁺05]. Fischl et al. used mechanical models to simulate an inflation of the cortical surface to produce an inflated surface and a spherical map [FSTD98].

We proposed a parameterization technique for the cortical surface based on p -harmonic energy minimization [JLTS04]. Angle and area distortion metrics were computed to evaluate the performance of this flattening procedure.

2.1 Parameterization and the Coordinate System

In this section, we describe our method to parameterize a triangulated surface mesh. In the context of our work, this mesh will typically represent the surface of the cerebral cortex; thus we will refer to this mesh model as the cortical surface. We use our p -harmonic mapping technique [JLTS04] for parameterization. The parameterization can be viewed as an assignment of complex numbers or vectors in \mathbb{R}^2 to each vertex in the

triangulated surface and the assignment is performed in such a way that the resulting p -harmonic energy is minimized. Let S be a surface with boundary. We define $\phi : S \rightarrow \mathbb{R}^2$ to be a function such that the p -harmonic energy given by $E_s = \int \|\nabla\phi\|^p dS$ is minimized. We impose constraints on this minimization by fixing the location of the inter-hemispheric fissure so that it is mapped to a unit square. We rewrite the energy functional as the sum of two energy functionals $\phi = [\alpha, \beta]'$, one for each coordinate, such that the corresponding arguments are scalars,

$$E_s = \int \|\nabla\phi\|^p dS, \quad p \in (1, \infty)$$

This minimization can be performed by minimizations over two real-valued functions. Discretization is done using finite elements. We make the assumption that both of them are piecewise linear. Let α be a piecewise linear real-valued scalar function defined over the surface, and α^i is α restricted to triangle i . Since α^i is linear on the i^{th} triangle we can write,

$$\alpha^i(x, y) = a_0^i + a_1^i x + a_2^i y \quad (2.1)$$

The three coefficients can be determined if values of the function α are known at the three vertices of the triangle. These equations can be written in matrix form as

$$\underbrace{\begin{bmatrix} 1 & x_1^i & y_1^i \\ 1 & x_2^i & y_2^i \\ 1 & x_3^i & y_3^i \end{bmatrix}}_{D^i} \begin{bmatrix} a_0^i \\ a_1^i \\ a_2^i \end{bmatrix} = \begin{bmatrix} \alpha^i(x_1, y_1) \\ \alpha^i(x_2, y_2) \\ \alpha^i(x_3, y_3) \end{bmatrix} \quad (2.2)$$

The coefficients a_0^i , a_1^i and a_2^i can be obtained by inverting the 3×3 matrix. From (2.1) and by inverting the matrix in (2.2), we obtain

$$\begin{aligned}\overrightarrow{\nabla\alpha^i} &= \begin{bmatrix} a_1^i \\ a_2^i \end{bmatrix} \\ &= \frac{1}{|D^i|} \underbrace{\begin{bmatrix} y_2^i - y_1^i & y_3^i - y_1^i & y_1^i - y_2^i \\ x_1^i - x_2^i & x_1^i - x_3^i & x_2^i - x_1^i \end{bmatrix}}_{B^i} \underbrace{\begin{bmatrix} \alpha^i(x_1, y_1) \\ \alpha^i(x_2, y_2) \\ \alpha^i(x_3, y_3) \end{bmatrix}}_{\Gamma^i} \\ \overrightarrow{\nabla\alpha^i} &= \frac{1}{2A^i} B^i \Gamma^i\end{aligned}$$

We use the fact that for any triangle, $|D^i| = 2A^i$ where A^i is the area of the triangle. Since α^i is piecewise linear, its gradient is constant over each triangle i , so that:

$$\int \|\nabla\alpha\|^p dS = \sum_i \|\nabla\alpha\|^p A^i$$

where the sum is over all triangles. Therefore,

$$\begin{aligned}\arg \min_{\alpha} \int \|\nabla\alpha\|^p dS &= \arg \min_{\Gamma} \sum_i \|M^i \Gamma^i\|^p \\ &= \arg \min_{\Gamma} \|M\Gamma\|^p\end{aligned}$$

where $M^i = \frac{1}{(A^i)^{(p-1)/p}} B^i$, M is composed using coefficients of B^i and Γ is a vector with coefficients α for each vertex. The vector $M\Gamma$ can be split into two parts: free vertices and constrained vertices. Values of α at constrained vertices are known.

$$\arg \min_{\alpha} \int \|\nabla\alpha\|^p dS = \arg \min_{\Gamma_f} \|M_f \Gamma_f + M_c \Gamma_c\|^p$$




Figure 2.2: The figure shows the cortical surface and its map to a square. The corpus callosum is constrained to lie on the boundary of the square.

where M_f , Γ_f and M_c, Γ_c are free and constrained parts of the M and Γ matrices.

This results in an unconstrained minimization problem. The fact that matrix M_f is sparse allows us to use the computationally efficient conjugate gradient method for obtaining the solution. The Jacobi preconditioner reduces the execution time considerably. The resulting maps are known to be bijections because the target domain is convex and flat [ES64, Ham75, FR02]. Using this scheme, we map each cortical hemisphere onto a unit square by constraining the inter-hemispheric fissure to lie on the boundary of the square.

2.1.1 Validation of p -harmonic mappings

In this section we present our method of evaluating the performance of the p -harmonic mappings described in section 2.1. We start by extracting a high-resolution triangulated cerebral cortical surface model for each brain. Each brain surface is represented by approximately 200,000 triangles. The BrainSuite software we use for

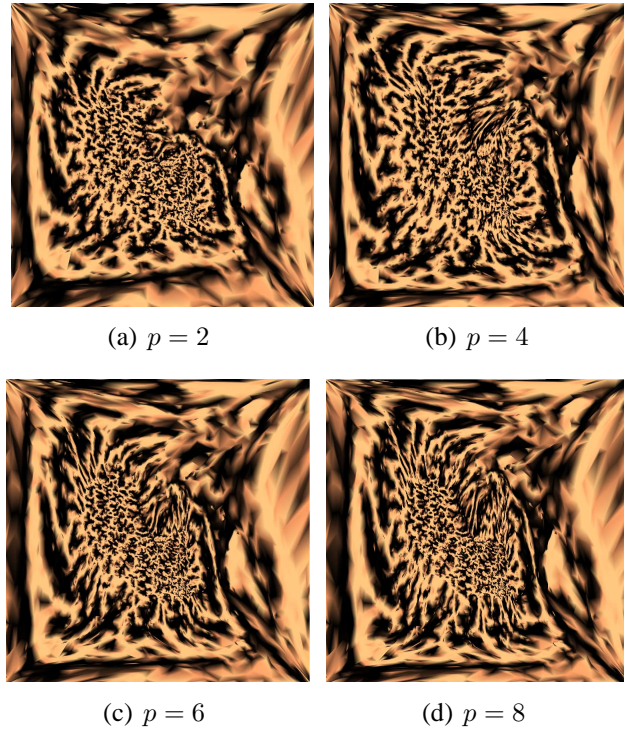


Figure 2.3: The p -harmonic maps of the left hemisphere of an individual cortex.


extraction also labels and separates the two cortical hemispheres and delineates the closed contour representing the inter-hemispheric fissure that separates the two hemispheres. We then parameterize the contour according its length and constrain it to lie on the boundary of the unit square. The mapping described in section 2.1 is then computed by minimizing the p -harmonic functional by conjugate gradient with Jacobi preconditioner [Smi85]. The minimization is very fast compared to other methods [AHTK99, FSTD98, HSB⁺00] and takes on the order of 20 seconds on 3GHz Intel Pentium 4 processor. Performing this operation for both hemispheres produces a bijective mapping of the cortical surface to a pair of unit squares.

In order to explore and evaluate the performance of such mappings and their dependence on the value of p chosen, we computed these mappings for $p = 2, 4, 6, 8$ and used the following metrics N_{angle} and N_{area} as measures of angle and area distortions.


$$N_{angle} = \sqrt{|g_{11}g_{22} + g_{12}g_{21}|} / g \quad (2.3)$$

$$N_{area} = \|g - g_{avg}\| \quad (2.4)$$

N_{angle} can be interpreted as a normalized inner product of the two columns of the metric tensor. It is zero if the mapping preserves angles (conformal). N_{area} is deviation of the differential area dS from its mean value. We evaluate these metrics at all the vertices and then plot their histograms for different values of p . The Fig. 2.4 shows p -harmonic maps of the cortical surfaces for the values of $p = 2, 4, 6$ and 8 . It can be seen from Fig. 2.4(a) that maps for $p = 4, 6, 8$ have much less area distortion than the map for $p = 2$. However there was no consistent trend for the values of p . Also it can be seen from Fig. 2.4(b) that the angle distortions are comparable for all values of p . From a computational point of view, though the use of covariant derivatives can make the subsequent processing independent of the value of p , we choose $p = 4$ because it has less area distortion than $p = 2$ case and hence the numerical error introduced during the resampling of the cortical surface on a regular grid is less.



(a) Angle Distortion



(b) Area Distortion

Figure 2.4: The figure shows smoothed histograms for angle distortion and area distortion respectively. In the angle distortion plot, angle distortion increases with the value of p . In the area distortion plot, the distortions for $p=4,6,8$ are less than that for $p=2$ and most of the points have small angle distortion only. However there is no observable trend for the value of p in either case.

Chapter 3

Cortical Surface Registration

Various surface-based techniques have been developed for inter-subject registration of two cortical models. These techniques can be used to register subject surfaces to a common atlas which in turn registers cortical data representing structure and function of the human brain to the atlas. There are two main categories of methods that align the cortex from a subject to an atlas: manual landmark based methods [JSTL07c] and automatic methods based on alignment of geometric features [WCT05] or surface indices [TRP05]. The main advantage of automatic methods is that there is no manual input required for performing the alignment. However they may be less reliable in the sense that they do not incorporate higher level knowledge of sulcal anatomy. While they have been successfully applied in several settings, their accuracy may not be satisfactory for expert neuroanatomists, particularly in the presence of the wide variation that may be present in neuroanatomy or the image acquisition quality. Data from subjects exhibiting abnormal cortical shape, such as individuals with Alzheimers disease, may be handled better by manual delineation. It is likely that landmarks defined by experts, who have been trained to make consistent decisions when faced with ambiguities that frequently arise in the analysis of cortical geometry, will produce improved registration results. In some cases a particular area, such as the visual cortex, may be of interest and constraints specific to that area may provide more appropriate registration.

One class of techniques involves flattening the two cortical surfaces to a plane [HSB⁺00] or to a sphere [FSTD98] using mechanical models or variational methods and then analyzing the data in the common flattened space. Other surface based techniques

work in the surface geometry itself rather than a plane or a sphere and choose to account for the surface metric in the inter subject registration [TWMT00, THS⁺04, LTPH04, MST04, WGH⁺05]. The advantage of such techniques is that they make the registration results independent of the intermediate flat space resulting in a more consistently accurate registration throughout the cortex. In this chapter we present a technique that is a generalization of the popular thin-plate spline methods from R^n to a non-Euclidean surface, as well as a Finite Element-based technique.

We presented our p -harmonic mapping method in Chapter 2, which maps each individual cortical hemisphere to the unit square. Our p -harmonic method results in a very fast parameterization of high-resolution cortical surfaces and always results in a bijective map. We use the resulting square maps of the cortical hemispheres to assign a coordinate system to the cortex. We then use these coordinates to compute the metric tensor and Christoffel symbols of the mapping. In order to register one brain to another, we warp coordinates of one brain with respect to another using sulcal landmarks such that the bending energy is minimized within the true geometry of the surface. This is achieved by solving the resulting variational problem using covariant derivatives and thus making the warping results independent of the coordinate system. Our warping approach is derived from the one presented in [TWMT00]. However, we use thin plate splines as a regularizing function. This is because the availability of p -harmonic maps allows us to have an approximate orthogonal coordinate system on the cortical surface and therefore we are able to decompose the deformation into two orthogonal components. Also availability of a smooth parameterization from 3D space to unit square means that the deformations are low dimensional in the parameter space too. Therefore we use DCT basis functions to represent the warping field. These techniques result in a considerable speed up and stability in the registrations. As an improvement over this method, we also

present a simultaneous parameterization and alignment technique as discussed further in Sec. 3.2. We also present evaluations of these registration techniques.




Figure 3.1: (a) A cortical surface with hand labeled sulci; (b) A flat map of the two cortical surface. The arrows show connectivity at points along the boundary of the square. Due to the spherical topology of the cortical surface, we can assign to it a coordinate system that allows us to compute partial derivatives across the interhemispheric fissure. (c) Chessboard texture mapped to the surface using the square maps.

3.1 Thin Plate Splines Registration in the Intrinsic Geometry of the Cortical Surface

3.1.1 Mathematical Formulation

The parameterization method presented in Chapter 2 gives us an initial approximate alignment of the labeled sulcal landmarks as shown in Fig. 3.3. It can be seen that the alignment is not perfect, however the deformation required to align the sulci perfectly is relatively small compared to the brain size. Therefore we use linear models from continuum mechanics which approximate small deformations to regularize the required deformation field. Here we discuss the widely used thin-plate splines, but we generalize them to the non-Euclidean geometry of the cortical surface. Having parameterized each of the cortical surfaces, we now align coordinate systems between one surface, which we denote the “atlas”, and another which we call the “subject”. The alignment uses a set of manually labeled sulci, sampled uniformly along their lengths, as a set of point constraints [TT96b]. To compute a smooth warping field ϕ from one coordinate system to the other we use the thin plate spline bending energy on the subject surface as a regularizing function. The p -harmonic maps serve two purposes. First, they set up an initial alignment of the features across multiple subjects. Second, they are used as our computational space to align the cortices. However, the thin-plate spline based alignment uses covariant derivatives, and is therefore invariant with respect to the specific parameterization [TMV⁺01].

Thin plate biharmonic splines [Boo89] are a very popular method for landmark-based registration of 2D or 3D images. These splines are solutions of the biharmonic equation

$$\frac{\partial^4 \phi}{\partial u^4} + 2 \frac{\partial^4 \phi}{\partial u^2 \partial v^2} + \frac{\partial^4 \phi}{\partial v^4} = 0 \quad (3.1)$$

or equivalently, they solve a variational problem that minimizes the bending energy E_b of a thin metal plate:

$$E_b = \int \left(\frac{\partial^2 \phi}{\partial x^2} \right)^2 + 2 \left(\frac{\partial^2 \phi}{\partial x \partial y} \right)^2 + \left(\frac{\partial^2 \phi}{\partial y^2} \right)^2 dx dy \quad (3.2)$$

We minimize this bending energy subject to the point landmark constraints, implemented here using a quadratic penalty function approach. Since we wish to minimize the bending energy in the surface, we must account for the intrinsic geometry of the surface when computing the integral. While we use the parameter space for doing the calculations required for evaluation of the bending energy, we account for the effect of the parameterization while calculating the integral. This is achieved using covariant derivatives which results in the property that given a set of homologous landmarks in some initial alignment, the deformation is independent of the parameterization used for the computation of the TPS deformation field. The use of covariant derivatives eliminates the effect of the initial parameterization on the resulting warping field.

Let \mathbf{x} denote the 3-D position vector of a point on the cortical surface. Let u^1, u^2 denote the coordinates in the parameter space. The metric tensor coefficients required in the computation are given by:

$$g_{11} = \left\| \frac{\partial \mathbf{x}}{\partial u^1} \right\|^2, \quad (3.3)$$

$$g_{22} = \left\| \frac{\partial \mathbf{x}}{\partial u^2} \right\|^2, \quad (3.4)$$

$$g_{12} = g_{21} = \left\langle \frac{\partial \mathbf{x}}{\partial u^1}, \frac{\partial \mathbf{x}}{\partial u^2} \right\rangle, \quad (3.5)$$

$$g = \sqrt{g_{11}g_{22} - (g_{12})^2} \quad (3.6)$$

We note that the eigenfunctions of the biharmonic operator on the surfaces are dependent on the surface itself. Therefore we do not expand the deformations in terms of a common

eigenfunction basis as in [Boo89]. Instead we take a more direct approach and minimize the integral numerically. The bending energy is minimized in the intrinsic geometry after replacing the first and second partial derivatives in (3.2) by the corresponding covariant derivatives. Integration over the surface can be carried out by integration in the parameter space while compensating with the surface metric g . The differential form ds^2 for the integration in the surface S is related to its counterpart in the parameter space (u, v) by $ds^2 = gdudv$. Let S be the set of all vertices, and let S_c denote the set of constrained vertices (landmarks). Let d_j^1 and d_j^2 denote the u and v displacements required at the j^{th} landmark, $1 \leq j \leq N$, to take it to its location in the atlas space. Cartesian tensors suffice for flows in 2D or 3D Euclidean spaces. However the cortical surface is a two dimensional non-Euclidean space and from the outset demands a full tensorial treatment. We do this by replacing the usual partial derivatives by covariant derivatives as done in continuum mechanics on manifolds [Kre99]. Although we want the deformation field with respect to the cortical surface to be independent of the specific choice of parameterization, the deformation field expressed in the 2D parameter space invariably does depend on the initial parameterization. This property is desirable since it ensures the covariance properties of the deformation vector field. Small deformations expressed in the parameter space can be modeled as contravariant vectors [TMT00, Kre99] since, with respect to two different parameterizations u and \bar{u} , the respective values of the deformations ϕ and $\bar{\phi}$ are related by $\bar{\phi}^\beta = \phi^j \frac{\partial \bar{u}^\beta}{\partial u^\alpha}$. In order to preserve their tensorial nature, we need to use covariant derivatives instead of the usual partial derivatives. The covariant derivative $\phi_{,\sigma}^\beta$ of a contravariant tensor ϕ^β is given by:

$$\phi_{,\sigma}^\beta = \frac{\partial \phi^\beta}{\partial u^\sigma} + \phi^\kappa \Gamma_{\kappa\sigma}^\beta \text{ where } \alpha, \beta, \kappa \in \{1, 2\} \quad (3.7)$$

where $\Gamma_{\kappa\sigma}^{\beta}$ denote the Christoffel symbols of the second kind [Kre99] given by:

$$\Gamma_{\alpha\alpha}^{\alpha} = \frac{1}{2g} \left[g_{\beta\beta} \frac{\partial g_{\alpha\alpha}}{\partial u^{\alpha}} + g_{12} \left(\frac{\partial g_{\alpha\alpha}}{\partial u^{\beta}} - 2 \frac{\partial g_{12}}{\partial u^{\alpha}} \right) \right] \quad (3.8)$$

$$\Gamma_{\alpha\alpha}^{\beta} = \frac{1}{2g} \left[g_{\alpha\alpha} \left(2 \frac{\partial g_{12}}{\partial u^{\alpha}} - \frac{\partial g_{\alpha\alpha}}{\partial u^{\beta}} \right) - g_{12} \frac{\partial g_{\alpha\alpha}}{\partial u^{\alpha}} \right] \quad (3.9)$$

$$\Gamma_{\alpha\beta}^{\beta} = \Gamma_{\beta\alpha}^{\beta} = \frac{1}{2g} \left[g_{\alpha\alpha} \frac{\partial g_{\beta\beta}}{\partial u^{\alpha}} - g_{12} \frac{\partial g_{\alpha\alpha}}{\partial u^{\beta}} \right] \quad (3.10)$$

where $\alpha, \beta \in \{1, 2\}$. The first covariant derivative of a contravariant tensor ϕ^{ζ} is a mixed tensor $\phi^{\zeta}_{,\beta}$. Covariant derivatives $\phi^{\zeta}_{,\beta\sigma}$ of such a tensor are given by:

$$\phi^{\zeta}_{,\beta\sigma} = \frac{\partial \phi^{\zeta}_{,\beta}}{\partial u^{\sigma}} - \phi^{\zeta}_{,\mu} \Gamma_{\beta\sigma}^{\mu} + \phi_{\beta}^{\nu} \Gamma_{\nu\sigma}^{\zeta}$$

where $\sigma, \beta, \zeta, \mu, \nu \in \{1, 2\}$ (3.11)

The warping field (ϕ^1, ϕ^2) with respect to the parameter space (u, v) that minimizes bending energy in the surface while matching the constraints is then given by:

$$\phi^1 = \arg \min_{\psi^1} \int_P \left((\psi^1_{,11})^2 + (\sqrt{2}\psi^1_{,12})^2 + (\psi^1_{,22})^2 \right) g du dv,$$

with $\phi^1(u_j, v_j) = d_j^1, \forall j \in S_c$ (3.12)

$$\phi^2 = \arg \min_{\psi^2} \int_P \left((\psi^2_{,11})^2 + (\sqrt{2}\psi^2_{,12})^2 + (\psi^2_{,22})^2 \right) g du dv,$$

with $\phi^2(u_j, v_j) = d_j^2, \forall j \in S_c$ (3.13)

The warping field (ϕ^1, ϕ^2) at the interhemispheric fissure is not forced to be zero as described in the next section.

3.1.2 Discretization Algorithm

In order to solve (3.12) and (3.13) for the thin-plate spline registration, we need to discretize the integral in that equation. We use the p -harmonic square maps of the triangulated tessellation of the cortical surface for defining a coordinate system. The square maps for each hemisphere are then resampled on a regular 256x256 grid. Because the interhemispheric fissure is fixed on the boundary of the square for each hemisphere, one can visualize the (u, v) parameter space as two squares placed on each other and connected at the boundaries of the squares. The main advantage of this space is the ease of composing and solving various partial differential equations in discrete form since this allows us to calculate partial derivatives across the two hemispheres and to include explicitly the connectivity of the two cortical hemispheres in subsequent analysis. This boundless space is then used for discretizing the partial derivatives with respect to u and v spatial coordinates in the solution of the differential equations. For instance, assume that $f : M \rightarrow \mathbb{R}$ is a scalar-valued function defined on the cortical surface M . We arrange its discretized representation at each vertex in the triangulation of the surface in a vector $\vec{f} = f_i$. In order to discretize $\frac{\partial f}{\partial u}$ by central difference, we calculate the usual central difference at the interior points in the squares. On the boundary of the squares, we consider the connectivity relationship shown in Fig. 3.1 for the neighborhood in the central difference approximation. Using these relations, we compose a central difference matrix D_u^c and obtain discretization of $\frac{\partial f}{\partial u}$ as $D_u^c \vec{f}$. Similarly we compose matrices $D_u^f, D_u^b, D_v^c, D_v^f$ and D_v^b , — the central, forward and backward difference operators — for the u coordinate, and D_v^c, D_v^f and D_v^b , — the central, forward and backward difference operators — for the v coordinate.

We carry out the discretization of the linear operator corresponding to the bending energy in (3.12) and (3.13) in the following steps.

1. Parameterize the cortical surface to map it into two squares and assign to it the coordinate system described above.

2. Form the forward, backward and central difference matrices $D_u^f, D_v^f, D_u^b, D_v^b$ and D_u^c, D_v^c for u and v coordinates respectively.
3. Compute the surface metric coefficients g_{11}, g_{12}, g_{21} and g_{22} . This is accomplished by replacing partial derivatives in (3.3), (3.4), (3.5) and (3.6) by their discrete versions from step 1.
4. Compute the Christoffel symbols according to (3.8), (3.9) and (3.10) by replacing partial derivatives in that equation by finite difference matrices from step 1.
5. Compute the first and second covariant derivative operators using (3.7) and (3.11). This can be done by first computing the operator corresponding to (3.7) and then using it to compose the operator corresponding to (3.11). Replace the partial derivatives in their expressions by finite difference matrices from step 1 and concatenate them to form a covariant bending energy functional matrix which is used to minimize the covariant bending energy.

3.1.3 Bending Energy Minimization

We discretized the bending energy integral in (3.12) and (3.13) in the parameter space over a 256x256 regular grid for each hemisphere. We denote the covariant differential operator in these equations by L . As described previously, our parameter space takes into account the neighborhood relationships between the two hemispheres and thus the covariant operator L is discretized in such a way that derivatives at the interhemispheric fissure are calculated correctly. In our current implementation our constraints are enforced by adding a quadratic penalty term rather than the exact matching constraints

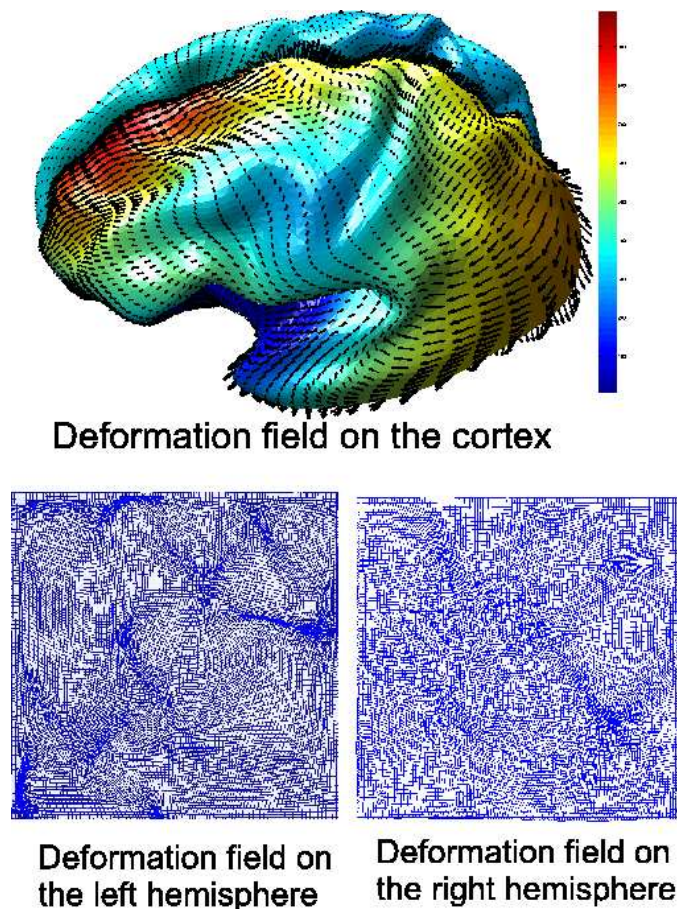


Figure 3.2: (upper) The figure shows the warping field computed on the surface. The deformation field is smoothly varying. This is achieved because the bending energy regularization was performed in the intrinsic geometry of the surface. The color indicates the magnitude of the deformation. (lower) The thin-plate spline deformation field applied to a regular grid representing left and right hemispheres.

in (3.12) and (3.13). Let $\Phi = (\phi^1, \phi^2)$ denote the deformation field. The discretized cost function then takes the form

$$\Phi = \arg \min \sum_{i \in S} \|\sqrt{g} L_i \Phi_i\|^2 + \sigma^2 \sum_{j \in S_c} \|\sqrt{g}(L_j \Phi_j - d_j)\|^2 \quad (3.14)$$

The resulting least squares problem is very high-dimensional (256x256x2x2 parameters), but it could be solved directly since the matrix L is sparse. However, we reduce the dimensionality of the problem by projecting onto a subset of the discrete cosine transform (DCT) basis functions. Provided the constraints can be satisfied with a relatively smooth deformation, this approach will work with fewer basis functions than the original 256x256 samples in (u, v) space. Let B denotes the DCT basis matrix, $T = LB$, $\Psi = B^T\Phi$ and $T_i = L_iB$. The optimization problem

$$\begin{aligned} \Phi = \arg \min & \sum_{i \in S} \|\sqrt{g}L_iBB^T\Phi_i\|^2 \\ & + \sigma^2 \sum_{i \in S_c} \|\sqrt{g}L_iBB^T\Phi_i - d_i\|^2 \end{aligned} \quad (3.15)$$

reduces to:

$$\Psi = \arg \min \sum_{i \in S} \|\sqrt{g}T\Psi\|^2 + \sigma^2 \sum_{i \in S_c} \|\sqrt{g}T_i\Psi - d_i\|^2 \quad (3.16)$$

In this way, we calculate the deformations in the DCT transform space. Use of this basis leads to a significant increase in speed. We observe that choosing a higher value of the parameter σ will lead to more accurate alignment of the sulcal landmark points, but in practice a very high value leads to non-bijective deformation of the coordinate space. Due to this trade-off, we pick a value of σ by trial and error. For certain individual subjects σ is decreased if the deformation field is non-bijective. The warps thus obtained are then applied to the (u, v) coordinates of each cortical surface to coregister them to the template. This process is illustrated in Fig. 3.2 where we show the sulci traced on the original cortical surface and their corresponding locations in flat space. We then show the relative locations of these sulcal features in flat space for the subject and atlas before and after matching. Note that we use a quadratic penalty function to match the

landmarks so that they do not exactly align after registration. Cortical regions near the boundary of the unit square exhibit larger metric distortion relative to the cortical surface than do regions near the center. Since the warp bending energy is computed with respect to the intrinsic geometry of the surface rather than flat space, we see that the warp in flat space exhibits larger deformations near the boundaries than at the center, following the pattern of metric distortion.

3.1.4 Validation TPS surface registration

Alignment of two cortical surfaces was performed using the intrinsic TPS method presented above. For our purpose, we used 16x16 DCT basis functions in each of the u and v directions. We found that the resulting warps closely resembled the warping field computed without using basis functions. The use of basis functions resulted in a runtime of 2 min. as compared to the runtime of 2 to 3 hours in the case of computation without using basis functions. Fig. 3.3 shows alignment of sulcal maps before and after registration. The warping field is smooth on the cortex since the surface geometry was considered during the regularization.

There is no gold standard for evaluating the performance of registration algorithms such as the one presented here. However, there are several properties that are desirable for any surface registration algorithm. Our method for evaluating the quality of our registration results is based on the following properties:

1. Insensitivity to the anatomical variability between multiple subjects. Though it is difficult to expect any automatic registration algorithm to align the anatomical features accurately, we expect a sulcus to be aligned more accurately if the remaining sulci are used as landmarks and are forced to align.

2. Insensitivity to a small amount of noise in the extracted surface coordinates. The process of extracting the cortical surface involves several stages and the results of each stage are sensitive to various parameters. We model this variability introduced during the extraction process as additive Gaussian noise in the x, y, z coordinates. The warping process should be relatively insensitive to this noise and should depend only on the global structure of the brain.
3. Insensitivity to small linear (affine) scaling of the surface coordinates. These kind of volumetric warps can get introduced in the imaging process. Also brains from different age groups have different sizes and registration should not depend on factors such as the overall size of the brain.

The error results presented here are in terms of the root mean squared error. In order to evaluate performance with respect to (1), we carried out a leave-one-out validation scheme. We aligned cortices of 6 subjects with one another using 22 out of 23 labeled sulci leaving one sulcus out of the registration each time. For each of the ${}^6C_2 * 23 = 345$ registrations, we measured how well the sulcus that was left out of the registration process aligns across the subjects before and after registration. Before carrying out the warping, there was mean squared error of 28.6 mm in the free sulcus. After aligning all but the free sulcus, the remaining root mean squared error was 2.81 mm. For (2), we added Gaussian noise in each of the x, y, z coordinates and register each of the brains with the noiseless brains. In this case, since we know the correct point correspondence between the noiseless and the noisy brains, we calculated the alignment error for the entire surface rather than just the sulci. Before applying TPS warping, there was 40.9 mm mean squared error. After warping there was 3.58 mm alignment error. For (3) we applied affine warps to the cortical surfaces and aligned the affine warped surface with the original surfaces. In this case also we calculated error for the entire surface as in (2). Before warping there was 35.8 mm error which reduced to 3.18 mm error after warping.

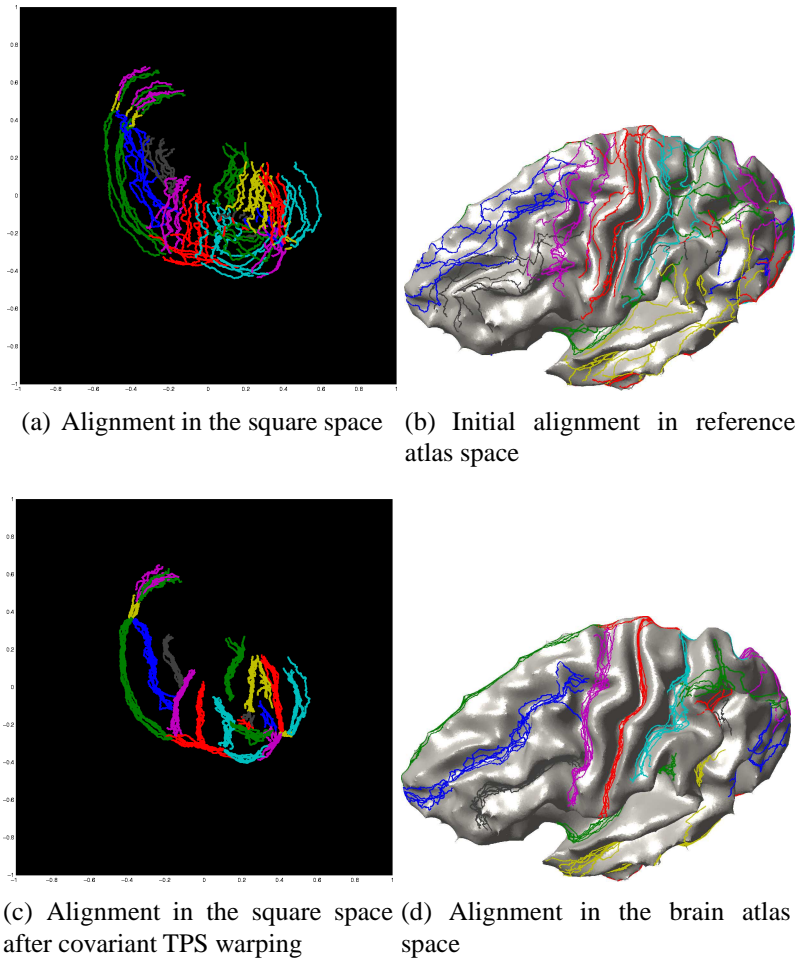


Figure 3.3: Alignment of the sulcal landmarks: 6 brains are registered to a common cortical surface using their p -harmonic maps in the plane. They are approximately aligned by the p -harmonic maps justifying our small deformation linear model (thin plate bending energy model) which is used for landmark alignment. After applying the covariant TPS deformation field to the surface parameterization, we can see that the sulci show better alignment.

3.2 A Finite Element Method for Simultaneous Registration and Parameterization

The method presented in Sec. 3.1 is a two step procedure which first maps the two cortical surfaces to a plane and then computes a deformation vector field that

aligns sulcal landmarks with respect to their planar coordinates. Similar methods were presented by various researchers which use plane, sphere or some intermediate representation of the cortex as a common space for performing the alignment [HSB⁺00, BGKM98, FSTD98, TRP05, THS⁺04, WGH⁺05, JSTL05]. In our two step procedure, in order to solve the resulting variational minimization problem, numerical derivatives were computed by resampling the brain on a uniform grid with respect to the parameterization. In addition to the computational cost of resampling and interpolation, this step results in a loss of resolution since the regular or semi-regular grid in flat space is not necessarily optimal for representing the brain in 3D space. In our new approach, we incorporate sulcal landmark alignment directly in our parameterization method and thus avoid the resampling and reparameterization step completely.

We propose an FEM based elastic mapping method that avoids the use of an intermediate surface flattening step for landmark matching. It incorporates the landmark registration into the parameterization method itself. We use the Cauchy-Navier elastic equilibrium equation for performing this matching as explained in the next section. This approach also has the advantage that the computation cost is relatively small and that the resulting alignment is inverse consistent [JC02] as will become clear from the symmetry of the cost function defined below.

3.2.1 Surface Registration

To perform cortical surface registration and parameterization with labeled sulcal curves as constraints, we model the cortical surface as an elastic sheet and solve the associated elastic equilibrium equation using an FEM. We choose the more general elastic model over a surface based harmonic mapping method [AHTK99, TSC00, JLTS04, WLCT05] because we found that the surface based harmonic mappings do not remain bijective

when multiple sulcal landmark constraints are imposed on the interior of the flat parameter space. However, for the elastic model we have so far always obtained a near bijective map by adjusting the model parameters λ and μ appropriately. The reason for this situation, intuitively, is that relative to the power of the Laplacian alone, the Cauchy-Navier elasticity operator provides additional control over the gradient of the divergence of the surface vector field, and this indirectly controls the Jacobian of the mapping, constraining it from taking on extreme values and thereby violating the smoothness assumption.

3.2.2 Mathematical Formulation

We assume as input a pair of genus-zero, tessellated cortical surfaces extracted from a volumetric MR image [SL00]. Our goal is to map the surfaces of each cortical hemisphere in the two brains to the unit square such that in the flat map a set of manually delineated sulcal landmarks are aligned with respect to the flat space coordinates. Point landmarks are generated by sampling uniformly along each sulcal curve. Let $\phi = [\phi_1, \phi_2]^T$ be the 2D coordinates assigned to every point on a given cortical surface such that the coordinates ϕ satisfy the Cauchy-Navier elastic equilibrium equation with Dirichlet boundary conditions on the boundary of each cortical hemisphere, represented by the corpus callosum. We constrain the corpus callosum to lie on the boundary of the unit square mapped as a uniform speed curve. We solve the equilibrium equation in the geometry of the cortical surface using the form:

$$\mu\Delta\phi + (\mu + \lambda)\nabla(\nabla \cdot \phi) = 0. \quad (3.17)$$

where μ and λ are Lamé's coefficients. The operators Δ and ∇ represent the Laplace-Beltrami and covariant gradient operators, respectively, with respect to the surface

geometry. The solution of this equation can be obtained variationally by minimizing the following integral on the cortical surface [HCF02]:

$$E(\phi) = \int_S \frac{\lambda}{4} (\text{trace}((D\phi)^T + D\phi))^2 + \frac{\mu}{2} \text{trace}(((D\phi)^T + D\phi)^2) dS. \quad (3.18)$$

where $D\phi$ is the covariant derivative of the coordinate vector field ϕ . The integral $E(\phi)$ is the total *strain energy*. Although the elastic equilibrium equation models only small deformations, we have found that in practice we can always compute a flat map of the cortex by setting the parameters $\mu = 1$ and $\lambda = 10$.

Minimizing (3.18) produces a flat map of each hemisphere but will not constrain the locations of the sulcal landmarks. To do this, we introduce the following constraints. Let ϕ_S and ϕ_A denote the 2D coordinates to be assigned to the subject and atlas brain hemispheres respectively. Then we define the Lagrangian cost function $C(\phi_S, \phi_A)$ as

$$C(\phi_S, \phi_A) = E(\phi_S) + E(\phi_A) + \sigma \sum_{k \in M} (\phi_S(k) - \phi_A(k))^2 \quad (3.19)$$

where $\phi_S(k)$ and $\phi_A(k)$ denote the coordinates assigned to the set of sulcal landmarks M , and σ is a Lagrange multiplier. Note that we do not constrain the locations of the sulci in the flat map but simply constrain homologous landmarks in the two maps to lie at the same coordinates.

3.2.3 Finite Element Formulation

To minimize (3.19) on a tessellated surface we use an FEM to discretize the strain energy $E(\phi)$. Since the integrand in (3.19) is a tensor, it is justifiable to compute it locally at each vertex point by assigning a local coordinate system (x, y) to its neighborhood.

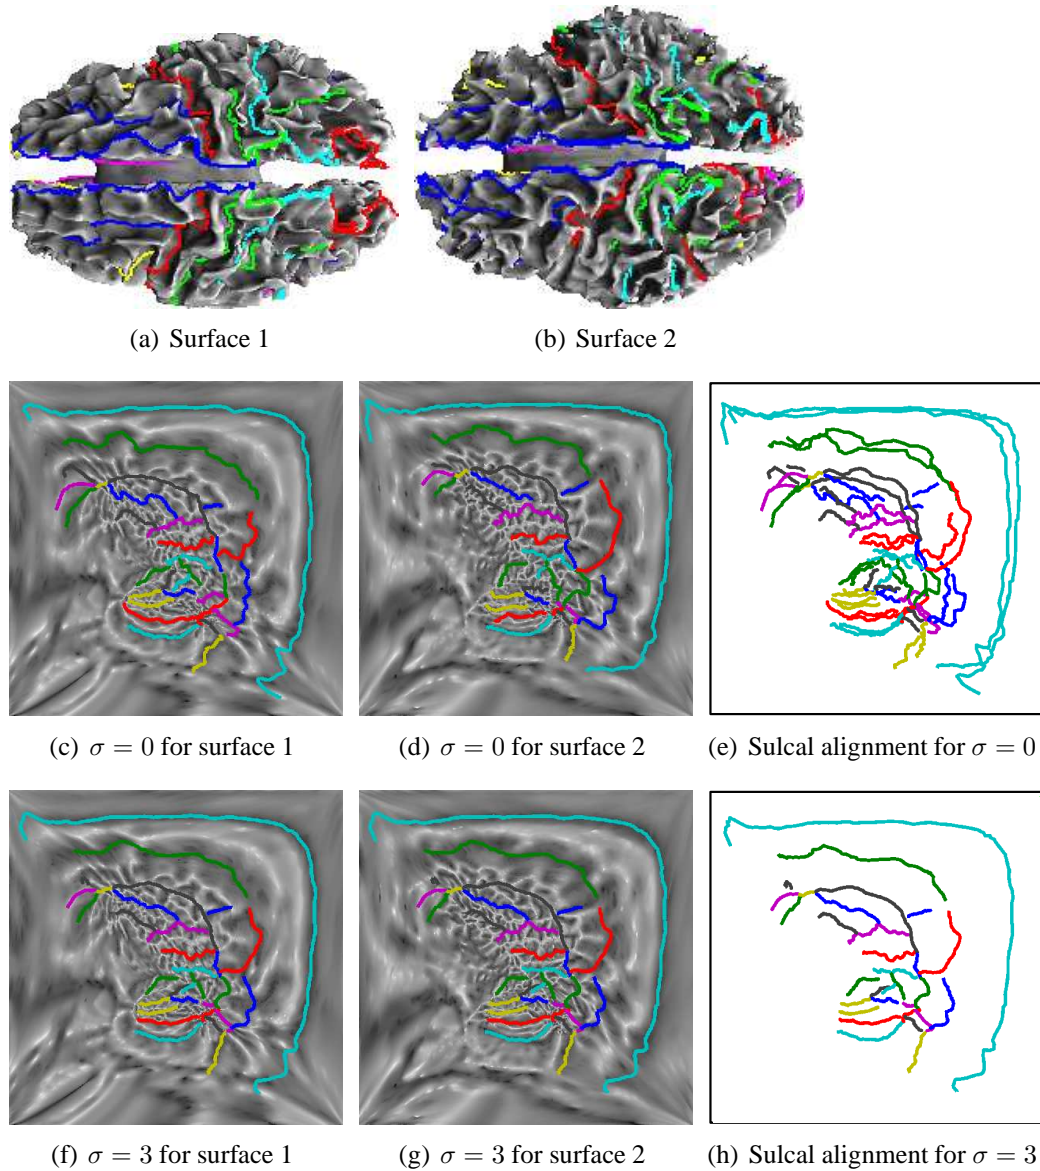


Figure 3.4: (a),(b) The two cortical surfaces with hand labeled sulci as colored curves; (c),(d) flat maps of a single hemisphere for the two brains without the sulcal alignment constraint; (e) overlay of sulcal curves on the flat maps without alignment; (f),(g) flat maps with sulcal alignment; (h) overlay of sulcal curves on the flat maps with alignment.

For each triangle the covariant derivative $D\phi$ in the local coordinates x, y becomes the Jacobian matrix:

$$D\phi = \begin{pmatrix} \frac{\partial\phi_1}{\partial x} & \frac{\partial\phi_1}{\partial y} \\ \frac{\partial\phi_2}{\partial x} & \frac{\partial\phi_2}{\partial y} \end{pmatrix} \quad (3.20)$$

From (3.18), the strain energy $E_i(\phi)$ for the i^{th} triangle Δ_i is given by:

$$E_i(\phi) = \int_{\Delta_i} (2\mu + \lambda) \left(\left(\frac{\partial \phi_1}{\partial x} \right)^2 + \left(\frac{\partial \phi_2}{\partial y} \right)^2 \right) + 2(\mu + \lambda) \left(\frac{\partial \phi_1}{\partial y} \right) \left(\frac{\partial \phi_2}{\partial x} \right) + \mu \left(\left(\frac{\partial \phi_1}{\partial y} \right)^2 + \left(\frac{\partial \phi_2}{\partial x} \right)^2 \right) dS. \quad (3.21)$$

We now describe the FEM discretization of the partial derivatives with respect to the local coordinates. Let α be any piecewise linear real-valued scalar function defined over the surface, and α_i the function restricted to triangle i with local coordinates x, y . Also denote the local coordinates of the three vertices as $(x_1, y_1), (x_2, y_2)$ and (x_3, y_3) respectively. Since α_i is linear on the i^{th} triangle, we can write,

$$\alpha_i(x, y) = a_0^i + a_1^i x + a_2^i y \quad (3.22)$$

Writing this expression at three vertices of the triangle i in matrix form,

$$\underbrace{\begin{pmatrix} 1 & x_1^i & y_1^i \\ 1 & x_2^i & y_2^i \\ 1 & x_3^i & y_3^i \end{pmatrix}}_{D^i} \begin{pmatrix} a_0^i \\ a_1^i \\ a_2^i \end{pmatrix} = \begin{pmatrix} \alpha^i(x_1, y_1) \\ \alpha^i(x_2, y_2) \\ \alpha^i(x_3, y_3) \end{pmatrix} \quad (3.23)$$

The coefficients a_0^i, a_1^i and a_2^i can be obtained by inverting the matrix D^i . From (3.22) and by inverting the matrix in (3.23), we obtain

$$\begin{pmatrix} \frac{\partial \alpha^i}{\partial x} \\ \frac{\partial \alpha^i}{\partial y} \end{pmatrix} = \begin{pmatrix} a_1^i \\ a_2^i \end{pmatrix} \quad (3.24)$$

$$= \frac{1}{|D^i|} \begin{pmatrix} y_2^i - y_1^i & y_3^i - y_1^i & y_1^i - y_2^i \\ x_1^i - x_2^i & x_1^i - x_3^i & x_2^i - x_1^i \end{pmatrix} \begin{pmatrix} \alpha^i(x_1, y_1) \\ \alpha^i(x_2, y_2) \\ \alpha^i(x_3, y_3) \end{pmatrix} \quad (3.25)$$

Denote the discretization of $\frac{\partial}{\partial x}$ and $\frac{\partial}{\partial y}$ at triangle i by D_x^i and D_y^i respectively. Also note that $|D^i| = 2A_i$ where A_i is the area of the i^{th} triangle. Then we have:

$$D_x^i = \frac{1}{2A_i} \begin{pmatrix} y_2^i - y_1^i & y_3^i - y_1^i & y_1^i - y_2^i \end{pmatrix} \quad (3.26)$$

$$D_y^i = \frac{1}{2A_i} \begin{pmatrix} x_1^i - x_2^i & x_1^i - x_3^i & x_2^i - x_1^i \end{pmatrix}. \quad (3.27)$$

Substituting these in (3.21) and (3.19), we have

$$E(\phi) = \sum_i \frac{1}{4A_i} \begin{pmatrix} \phi_1^i & \phi_2^i \end{pmatrix} K \begin{pmatrix} \phi_1^i \\ \phi_2^i \end{pmatrix} \quad (3.28)$$

$$= \sum_i \left\| \frac{1}{2\sqrt{A_i}} \begin{pmatrix} \sqrt{\lambda}D_x^i & \sqrt{\lambda}D_y^i \\ \sqrt{\mu}D_y^i & \sqrt{\mu}D_x^i \\ \sqrt{2\mu}D_x^i & 0 \\ 0 & \sqrt{2\mu}D_y^i \end{pmatrix} \phi^i \right\|^2. \quad (3.29)$$

where K is given by

$$K = \begin{pmatrix} (\lambda + 2\mu)D_x^{it}D_x^i + \mu D_y^{it}D_y^i & \lambda D_x^{it}D_y^i + \mu D_y^{it}D_x^i \\ \lambda D_y^{it}D_x^i + \mu D_x^{it}D_y^i & (\lambda + 2\mu)D_y^{it}D_y^i + \mu D_x^{it}D_x^i \end{pmatrix} \quad (3.30)$$

This method is used to discretize both $E(\phi_S)$ and $E(\phi_A)$. It can be seen from (3.29) and (3.19) that the cost function is quadratic. We minimize (3.19) with respect to both ϕ_S and ϕ_A , with the corpus callosum fixed at the boundary of the unit square, to compute the sulcally-coregistered flat maps for both brains simultaneously. The minimization is performed by using a preconditioned conjugate gradient method with Jacobi preconditioner. In practice the minimization algorithm converges in approximately 500 iterations, requiring 3-4 mins on a desktop computer for surfaces with approximately 200,000 vertices.




Figure 3.5: RMS error and percentage overlap in the flattened map as a function of σ .

3.2.4 Results and Validation

We first extract cortical surfaces from MRI for each subject using the BrainSuite software [SL02] to produce a genus-zero tessellated representation of the inner gray/white cortical boundary. We then manually delineate 23 major sulci on each of these extracted cortical hemisphere meshes. Delineation is performed in accordance with a sulcal labeling protocol with established intra- and inter-rater reliability [THdZ⁺02]. This protocol specifies that sulci do not intersect and that individual sulci are continuous curves that are not interrupted. If interruptions are present the curves are simply interpolated across any interrupting gyri. In cases where a full set cannot be defined, a subset can be used without any change in the algorithm defined here. Uniform samples along the sulcal curves serve as landmarks in our registration.

Fig. 3.4 illustrates the alignment process. Fig. 3.5 shows the RMS error in matching of sulcal landmarks and the percentage area of overlap or folding in the flat maps as a function of the Lagrange multiplier σ . Enforcing a more accurate sulcal alignment by increasing σ results in an increase in overlap in the mappings. We choose $\sigma = 3$ for further analysis. Although the elastic mappings are not formally guaranteed to produce a bijective registration, we found that by setting $\mu = 1$, $\lambda = 10$ and $\sigma = 3$, we can




Figure 3.6: Mapping of sulcal landmarks from 5 subjects to the atlas brain (left) without and (right) with the sulcal alignment constraint.

achieve a nearly bijective map with an average overlap of approximately 0.4% of the surface area. By inspection we see that the overlap occurs in the vicinity of pairs of landmarks that are closely spaced in one brain and distant in the other. One solution to this problem is to locally reparameterize in the neighborhood of the overlap once the flat maps are computed.

We performed a leave-one-out validation for examining the performance of our method. We choose one brain as a representative ‘atlas’ and align cortices of 5 subjects with the atlas using 22 of the 23 labeled sulci leaving one sulcus out of the registration each time. For each of the registrations, we measured how well the sulcus that was left out of the registration process aligned across the subjects with ($\sigma = 0$) and without ($\sigma = 3$) sulcal alignment. Without alignment, there was an RMS error of 33.1 mm in the free sulcus. With alignment using all but the free sulcus, the remaining rms error was 3.2 mm for the free sulcus.

Incorporating sulcal landmark alignment directly in our parameterization method not only avoids the resampling and reparameterization steps and reduces computational cost while maintaining high resolution in the surface tessellations, but also makes the registration inverse consistent. The improved speed and resolution of the registration may help in large scale and detailed comparisons of cortical data.

3.3 Optimum Choice of Sulcal Subset for Registration

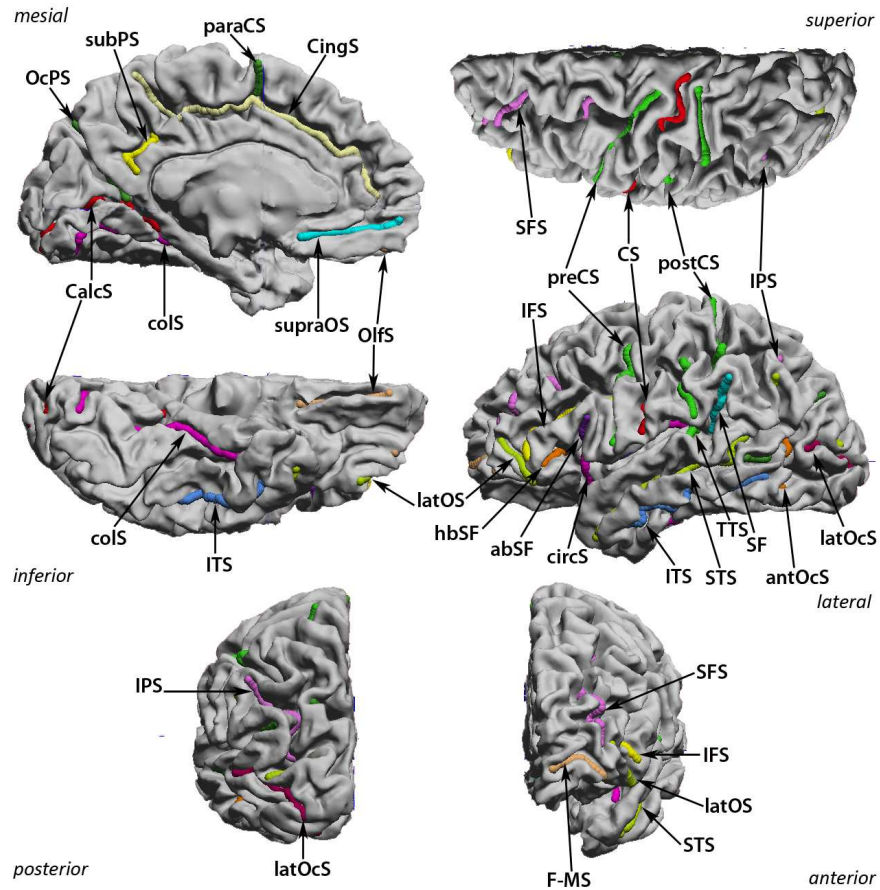
The objective of landmark based manual registration methods presented in Sec. 3.1 and Sec. 3.2 is to minimize the alignment error in sulcal curves. Their disadvantage is that the individual tracers need to be trained, and even then inter-rater variability introduces some uncertainty into the procedure. In registration applications, errors in automatic sulcal identification may propagate into errors in the registration accuracy. There is an inherent tradeoff between manual effort for tracing sulcal landmarks and registration accuracy. Increasing the number of sulcal landmarks achieves more accurate registration, but it also increases the required manual effort. Due to this, for large scale studies, manual procedures may be infeasible unless we minimize the number of sulcal curves required in the manual tracing protocol. Here, we address this issue.

In this section, we present an algorithm that finds an optimal subset of sulcal landmarks with a given number of sulci, which leads to minimum error in registration. We begin with a large set of sulcal curves that have been identified by the neuroanatomist on our team as candidate landmarks for cortical registration. Our objective is to select an optimal subset from this set such that, for a given number of curves, the sulcal registration error is minimized when computed over all sulci. One straightforward approach is to actually perform registration of the sulcal curves for a set of training images using all possible subsets and then measure the error in the remaining unconstrained sulcal curves. The difficulty with this approach is that there are a huge number of combinations possible. In our case we have 26 candidate curves. Suppose we want to define a protocol that uses 10 curves, the number of combinations to be tested is $\binom{26}{10} \approx 5.3$ million. If the error is to be estimated by performing pairwise registrations of 20 brains, i.e. $\binom{20}{2}$ registrations, then the total number of registrations required is $\binom{20}{2} \binom{26}{10} \approx 1$ billion. This is a prohibitively large number considering the fact that surface registrations are computationally expensive.

Instead of performing actual brain registrations with multiple subsets of constrained sulci, we perform only pair-wise unconstrained registrations using the elastic energy minimization procedure described in Sec. 3.2. The resulting maps produce reasonable correspondences so that we can model the measured sulcal registration errors using a multivariate Gaussian distribution. Using conditional probabilities, we then analytically predict the registration error that would result if we constrained a subset of the curves to match using hand labeled sulci. These errors can be rapidly computed using conditional covariances, and as we show below, lead to reasonably accurate estimates of the true errors that result when constraining the curves. For a fixed number of constrained curves, we estimate the error for all possible subsets of that size and select the one with the smallest predicted error. We investigate the prediction accuracy of our model by doing actual registrations using the optimal sulcal constraint set. Our algorithm reveals the trade-off between the number of curves and registration accuracy. An appropriate optimal subset of sulci can be chosen for a particular study based on manual effort and desired registration accuracy. Once such a subset is chosen, only the sulci from that subset need to be manually labeled on the brains used for a neuroanatomical study.

3.3.1 Registration Error

The point correspondence defined by registration allows us to map a point set from one brain to another brain. For every pair of registered hemispheres, we map the traced curves of one brain to the other, which is arbitrarily defined as a target. The registration is either unconstrained for error prediction, or constrained for validation. We parameterize each sulcal curve n over $[0, 1]$ and then compute S equidistant points on each sulcus corresponding to $s = \{0.1/S, 0.2/S, \dots, 1\}$. The point to point errors $e_n(s)$ are treated as S different samples of the error e_n , as illustrated in Fig. 3.8, where $e_n(s)$ is



1) central sulcus (CS)	14) sup. temporal with upper branch (STS)
2) precentral sulcus (preCS)	15) inferior temporal sulcus (ITS)
3) superior frontal sulcus (SFS)	16) occipeto temporal sulcus (OTS, not shown)
4) inferior frontal sulcus (IFS)	17) collateral sulcus (colS)
5) ascending branch of sylvian fissure (abSF)	18) transverse temporal sulcus (TTS)
6) horizontal branch of sylvian fissure (hbSF)	19) circular sulcus (circS)
7) lateral orbital sulcus (latOcs)	20) postcentral sulcus (postCS)
8) frontomarginal sulcus (F-MS)	21) intraparietal sulcus (IPS)
9) Cingulate sulcus (CingS)	22) parieto occipital sulcus (OcPS)
10) paracentral sulcus (paraCS)	23) subparietal sulcus (subPS)
11) supra orbital sulcus (supraOS)	24) calcarine sulcus (CalcS)
12) olfactory or medial orbital sulcus (OlfS)	25) transverse occipital sulcus (TOS, not shown)
13) sylvian fissure terminal split (SF)	26) lateral occipital sulcus (latOcs)

Figure 3.7: The complete set of candidate sulcal curves from which we select an optimal subset for constrained cortical registration

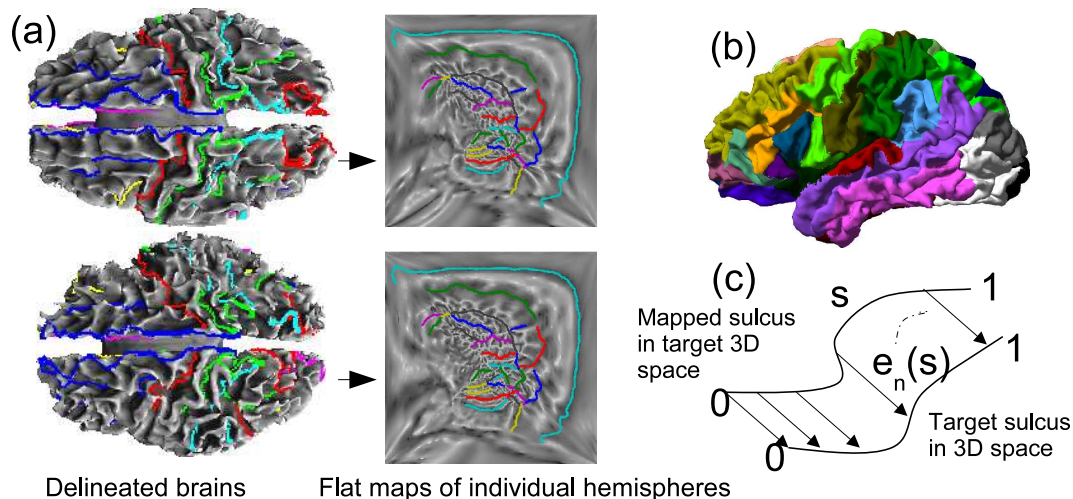


Figure 3.8: (a) Registration of two cortical surfaces based on the flat mapping method; (b) Parcellation of the cortex into regions surrounding the traced sulci; (c) Registration error for two corresponding sulci where $e_n(s)$ are samples of the registration error.

the registration error in 3D coordinates for location s on the n^{th} curve. For symmetry, we repeat the procedure by interchanging subject and target brains.

The alignment error in a sulcus causes a registration error in the surrounding cortical area. Therefore, isolated sulci will have more impact on registration, since their misregistration will affect large cortical regions. To compensate for this effect, we parcellate the cortex into $N = 26$ regions by assigning each cortical point to the nearest sulcal curve (Fig 3.8b). The parcellation was performed for all $M = 24$ available brain hemispheres. We then defined a weight function for the n^{th} sulcus to be $w_n = \frac{1}{M} \sum_{i=1}^M A_n^i / A^i$, where A_n^i is the surface area of the n^{th} parcellated region in the i^{th} brain, and A^i is the total surface area of the i^{th} hemisphere.

Finally, the surface registration error metric was defined as

$$E_R = \mathbb{E} \left(\sum_n w_n (e_n^x)^2 + w_n (e_n^y)^2 + w_n (e_n^z)^2 \right), \quad (3.31)$$

where e_n^x, e_n^y and e_n^z represents x, y , and z components of e_n , and $\mathbb{E}(\cdot \cdot \cdot)$ is the expectation operator. Below, we substitute $E_n^x = \sqrt{w_n}e_n^x$ in order to simplify subsequent analysis. The objective of the surface registration procedure is to minimize this registration error E_R .

3.3.2 Probabilistic Model of the Sulcal Errors

We model the sulcal errors E_1^x, \dots, E_N^x as jointly Gaussian random variables, since these errors are drawn from a large population of brain pairs. We describe computations for the x component of the error; similar computations are performed for y and z . The distribution model of E_j^x for $j \in \{1, \dots, N\}$ is:

$$f_{E^x}(E_1^x, \dots, E_N^x) = \frac{1}{(2\pi)^{N/2} |\Sigma^x|^{1/2}} \exp\left(-\frac{1}{2} E^{xT} (\Sigma^x)^{-1} E^x\right) \quad (3.32)$$

where Σ^x denotes the covariance matrix of E^x . Therefore, the registration error can be expressed as:

$$E_R^x = \mathbb{E}\left\{\sum_{i=1}^N (E_i^x)^2\right\} = \text{trace}(\Sigma^x) \quad (3.33)$$

We now want to predict the registration error when some of the sulci are explicitly constrained to register. We partition the curves into two sets: sulci F which are free and sulci C which are constrained so that $\{1 \dots N\} = F \cup C$. We assume that the registration algorithm is well behaved in a sense that it does not create unnatural deformations on the unconstrained sulci when a subset of them are constrained. In other words, if we constrain some sulci to register, the distribution of the remaining ones would be the same as if the constrained ones matched simply by chance, conditioned on the constrained sulci having zero error. Therefore, we model the registration errors in unconstrained sulci as

the conditional distribution of the original joint Gaussian density. The probability density of a jointly Gaussian vector, conditioned on some of its elements being zero, is also jointly Gaussian. Therefore, the registration error E_R^{xc} after matching the sulci from C can be obtained using the conditional expectation:

$$E_R^{xc} = \mathbb{E} \left(\sum_{i \in F} E_i^{x2} | E_j^x = 0 \forall j \in C \right) = \text{trace}(\Sigma_C^x) \quad (3.34)$$

where Σ_C^x is the conditional covariance matrix of the error terms corresponding to free sulci. By rearranging sulci so that free sulci precede the constrained ones, we can partition the covariance matrix as:

$$\Sigma^x = \begin{pmatrix} \Sigma_{ff}^x & \Sigma_{fc}^x \\ \Sigma_{cf}^x & \Sigma_{cc}^x \end{pmatrix}. \quad (3.35)$$

where Σ_{ff}^x and Σ_{cc}^x are the error covariances for free sulci and constrained sulci respectively, and Σ_{fc}^x and Σ_{cf}^x are the cross-covariances.

The conditional covariance is given by:

$$\Sigma_C^x = \Sigma_{ff}^x - \Sigma_{fc}^x (\Sigma_{cc}^x)^{-1} \Sigma_{cf}^x. \quad (3.36)$$

which is the Schur complement of Σ_{cc}^x in Σ^x [MKB79]. The estimated registration error E_R^{xc} after constraining a subset of sulci is then:

$$E_R^{xc} = \text{trace}(\Sigma_C^x). \quad (3.37)$$

This formula allows us to estimate the x component of the registration error for a particular combination of constrained sulci and free sulci. The total registration error is evaluated by adding the x,y, and z components.

$$E_R^c = \text{trace}(\Sigma_C^x) + \text{trace}(\Sigma_C^y) + \text{trace}(\Sigma_C^z). \quad (3.38)$$

We use this formula to estimate the total registration errors for all $\binom{N}{N_c}$ combinations of sulcal subsets, where N_c is the number of constrained sulci, and choose the subset that minimizes this error.

3.3.3 Results

A total of 24 brains, or equivalently 48 hemispheres, were delineated. Our tracings, consisting of 26 candidate sulci per hemisphere (Fig. 3.7), were verified and corrected whenever necessary by a neuroanatomist. We assigned the hemispheres into two subsets, a training set of 24 hemispheres and a testing set of 24 hemispheres, in order to check:

- Accuracy of the estimator: if the errors predicted by the our method are close to the actual errors after registration.
- Generalizability of the results to other datasets: if we chose a different dataset (testing set) of brains and sulci, whether the registration errors remain similar to the ones from the training dataset.

We performed unconstrained mappings for all the 24 training hemispheres by directly minimizing Eq. 3.18 for each hemisphere separately, instead of doing pairwise registrations using Eq. 3.19 with $\sigma = 0$, since the optimization in Eq. 3.19 becomes separable in the unconstrained case $\sigma = 0$. Using the flat maps of the 24 hemispheres




Figure 3.9: Sample covariance matrices for the x, y, and z components of the registration error, represented as color coded images.

we computed samples of sulcal errors E_n^x , E_n^y , and E_n^z , with $S = 10$ samples for each sulcus, for all possible pairwise combinations of hemispheres as described in Sec. 3.3.1. Whenever a sulcus were missing from either cortical surface, we assumed abnormal anatomy in that region and assigned zero registration error for those sulci. The resulting sample covariance matrices of the errors are shown in Fig. 3.10 using color code. The non-zero off-diagonal elements indicate that the errors are correlated among sulci, and thus constraining some of them would affect the registration error of the others. The correlation structure of the sulcal errors depends on the sulcal locations in the flat maps. Here we used square maps as discussed in Sec. 3.3.1, but we expect that our results are robust with respect to the mapping method.

By applying Eq. 3.38 to all subsets of a given number of constrained curves, we identified the subset that minimizes the registration error. The optimal subsets of curves are given in Fig. 3.10 for all numbers of constrained sulci from 1 to 26. We also calculated the sulcal registration errors for each of these optimal subsets by doing actual registrations. In order to perform actual registration, we chose $\sigma = 3$ as discussed in




Figure 3.10: Optimal subsets of sulci for cortical registration. Each row gives the indices of the optimal subset of sulci that minimize the registration error against all other combinations with equal number of constrained curves (also see Fig. 3.7). The three right columns show that the estimated (est.) error is close to the calculated actual (act.) error when actual registrations with the same constrained curves are performed. Our method predicts the registration error both for the training (trn) and the testing (tst) set of brains.

[JSTL07a] for the constrained subset of sulci. Comparing estimated and actual registration errors, also in Fig. 3.10, we see that the predicted values are close to those obtained when actually constraining the curves.

A Lilliefors test rejected the null hypothesis of normality for the errors E_n^x , E_n^y , and E_n^z for many sulci, and therefore our Gaussianity assumption is not fully satisfied. This is not surprising, since for example errors are naturally bounded by the size of the brains

and therefore some deviation from normality are anticipated. However, the distributions were unimodal and the predicted errors of our model are in accordance with the actual ones, indicating that our distributional assumptions are reasonable for this application.

In order to check the generalizability of the results, we used the 24 cortical hemispheres from the testing set, which are different from the original 24 hemispheres of the training set. We performed pairwise registrations of the testing brains constraining the optimal subset of sulci, as shown in Fig. 3.7. The average registration errors in this case were again close to the predicted errors as shown in the same figure. Therefore, our results are valid to other brain datasets.

To further test our method, we subjectively selected a set of 6 curves to be constrained, namely CS, SFS, CingS, STS, IPS, and CalcS, which seemed a reasonable a priori selection based on sulcal extent and spatial distribution around the cortex. The algorithm predicted an error of $194.36mm^2$ and the actual error was $200.03mm^2$. The optimal set (SFS, STS, OTS, postCS, IPS, and CalcS) found by our method had predicted an error of $167.19mm^2$ and the actual error was $164.77mm^2$, which is better than our subjectively selected subset. We anticipate that in general the curves selected by our method should be superior to those selected on an intuitive basis, since various confounding effects due to elastic flat mapping as well as correlations in errors are accounted for in the algorithm. In this specific case, we notice that our algorithm preferred more sulcal curves on and around temporal lobe. This can be explained by the fact that temporal lobe maps to a very small area in the flat square space. Therefore any alignment error made in that region in the flat space gets amplified in the 3D space.




Figure 3.11: Optimal sulcal sets for 5, 10, and 15 curves.

3.3.4 Discussion

We have described a general procedure for selecting subsets of sulcal landmarks for use in constrained cortical registration. The procedure can be used to reduce the time required for manual labeling of sulci in group studies of cortical anatomy and function.

The optimal subsets of curves, shown in Fig. 3.10, provide an intuition on the criteria our method uses to select curves. First notice that the central sulcus is not selected for protocols with a small number of curves (less than 16). This is probably because sulci

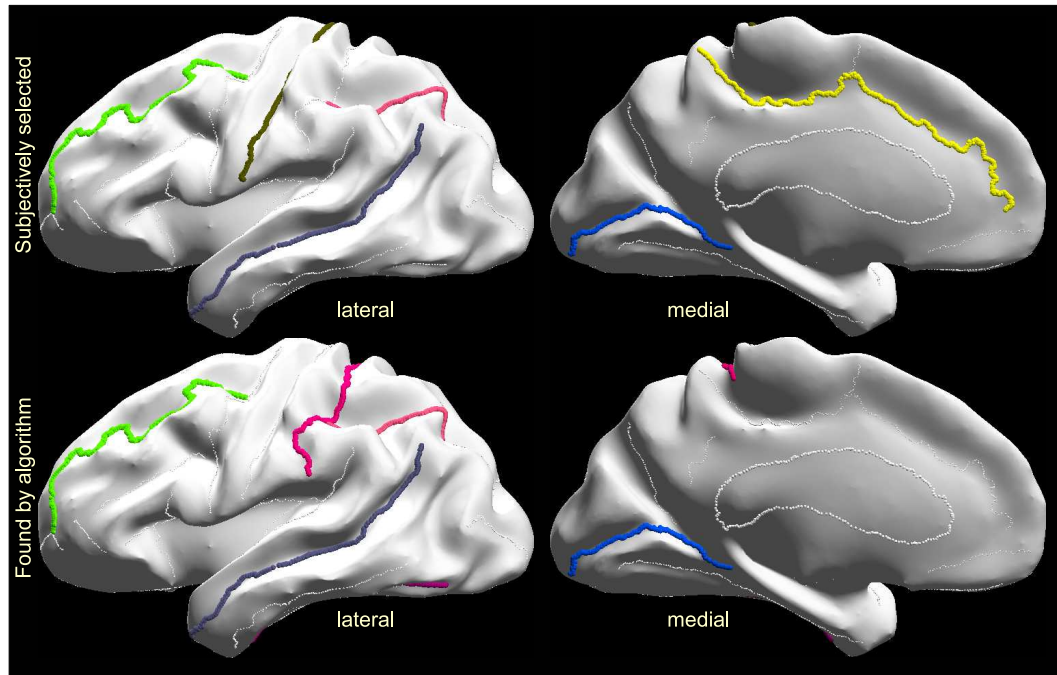


Figure 3.12: Top row: subjective selection of 6 curves, with preference on long sulci distant from each other that are expected to minimize cortical registration error; bottom row: optimal sulcal set with the 6 curves selected by our method.

that are most stable and consistent among brains, such as the central sulcus, may tend to align well even without explicitly tracing them. Therefore, they may not improve the registration error sufficiently to justify their inclusion in the tracing protocol. Furthermore, short sulci neighboring other candidate curves are excluded from the optimal protocol, such as the paracentral sulcus which is close to the cingulate sulcus, and the subparietal sulcus which is close to the cingulate and the parieto-occipital sulcus.

On the other hand, the most important sulcus for surface based registration seems to be the superior temporal sulcus with its upper branch. This is possibly for two reasons: (1) it is one of the longest sulcus and hence aligning it will register a large region of the brain; (2) in cortical flat maps where the corpus callosum is mapped on the edges of a unit square, such as in our method in Fig. 3.8a, the temporal lobe is mapped to a small

area near the center of the unit square. Since it is away from the corpus callosum, there is a significant misregistration error if it is unconstrained. Therefore it is important to align it accurately, and so it is selected by our method.

A registration error always remains when only a subset of sulci is used for registration. Whether this is acceptable or not depends on the particular neuroscience study. For example, anatomical studies [TMV⁺01, STR⁺02, Cha01, GGL⁺04] require high accuracy and might need more sulci, whereas functional studies, such as low resolution magnetoencephalography data [PNBL05], can tolerate higher registration error. Fig. 3.10 can be used as a guideline: based on the degree of registration accuracy required, a different number of curves may be used.

Our method provides the subset of sulci to be delineated in a registration study based solely on the registration error. However, some changes in the selected subset can be made based on other practical considerations, such as convenience in tracing. For instance, identification and tracing of the central sulcus is always easy and it could be helpful in identifying the surrounding sulci. Therefore we expect that it would be typically included in a tracing protocol. Furthermore, for neuroscience studies focusing on particular cortical regions, for example language studies interested in activation in the temporal lobe and Broca's area, the registration error metric defined in Sec. 3.3.1 can be modified by assigning more weight to the regions of interest; thus custom optimal curve protocols can be defined, tailored to the needs of individual neuroscience studies.

Errors and variability in identifying cortical landmarks are a common problem concerning all landmark based techniques and can affect the registration error. However, they are beyond the scope of this mathematical formulation. For this particular study the curves were carefully identified and cross-checked by a neuroanatomy expert. Inter and intra-rater variability is typically minimized by appropriate training and cross checking

of traces. A possible extension of our method could allow modeling of intra/inter-rater variability in identifying sulci, so that unreliable ones are excluded from the protocol.

Our methodology readily extends to other landmark based registration methods in which the goal is to select an optimal subset of landmarks for large scale studies, from a set of candidate landmarks. Finally, it can possibly be applied to other areas of computer vision [MA98, Ols00, GI94] for aiding optimal landmark selection.

The surface based surface registration techniques presented in this chapter set up point to point correspondence between two surfaces based on manually traced sulcal landmarks. We also presented a method to optimally select the set the sulcal landmarks in order to minimize the manual effort. These methods can be used to register neuroanatomical data from individual surfaces to a common atlas. This data can then be analyzed in the surface geometry by using the techniques presented in the next chapter.

Chapter 4

Processing of Data in the Surface

Geometry

Neuroimaging data, such as cortical thickness or neural activation, can often be analyzed more informatively with respect to the cortical surface rather than the entire volume of the brain. This analysis should be carried out in the intrinsic geometry of the surface rather than in the ambient space. We present methods for generalizations of two commonly used image filtering methods to non-Euclidean surface geometries. First we describe a method for isotropic diffusion filtering, which is equivalent to Gaussian filtering in Euclidean space. We then describe its extension to anisotropic filtering. In order to discretize and numerically compute the isotropic and anisotropic geometric operators, we first parameterize the surface using a p -harmonic mapping. We then use this parameterization as our computational space and account for the surface metric while carrying out isotropic and anisotropic filtering. We illustrate these methods in an application to smoothing of mean curvature maps on the cortical surface. For the cases when the cortical data is a point set on the surface, we present a method to quantify its mean and variance. This is illustrated in the analysis of MEG dipole locations corresponding to finger tapping.

4.1 Image Filtering on Surfaces

Gaussian kernel smoothing has been widely used in 3D medical imaging as a tool to increase signal-to-noise ratio. However, in many medical imaging applications neuro-anatomical [TT96a][THdZ⁺03], functional [JSTL05] and statistical [CRD⁺05] data are defined with respect to the non-Euclidean cortical surface and ideally should be processed with respect to the non-Euclidean geometry of the surface. The Gaussian kernel is isotropic in Euclidean space, but on curved surfaces the notion of a Gaussian function needs to be generalized. One existing approach, called diffusion smoothing [TSC00], replaces the Gaussian filter by the heat equation which is then solved on the surface. Thus filtering is formulated as the process of heat diffusion by explicitly solving an isotropic diffusion equation with the given data as an initial condition [CWT⁺01]. The drawback of this approach is the complexity of setting up a finite element method formulations or implicit PDEs and difficulty in making the numerical scheme stable [Chu05]. Here we describe an alternative approach to smoothing using the heat equation, which is based on a parameterized representation of the surface.

Anisotropic filtering or Perona-Malik flow [PM90] has been widely used in region selective smoothing and edge preserving filtering of 2D and 3D images. Anisotropic diffusion filtering on non-Euclidean surfaces has been applied to processing and modification of surface geometries [HP04][CDR04]. In contrast, here we focus on anisotropic filtering of anatomical or functional images which are scalar functions defined on these surfaces [TSC00]. In order to solve the isotropic and anisotropic diffusion equations on non-flat surfaces, the associated Laplace-Beltrami operators needs to be discretized. The existing approaches to this discretization use FEM formulations [LPDS04][BX03][CWT⁺01]. We present an alternative to the FEM approach. We first generate a global parameterization of the surface, compute the metric tensor for

the parameterization, and use this to compute the isotropic and anisotropic Laplace-Beltrami operators. We first parameterize the cortical surface using a p -harmonic mapping technique . We then resample the surface on a regular lattice with respect to the 2D parameterization and solve the associated PDEs using this discretization while accounting for the p -harmonic mapping transformation. Our method explicitly accounts for the metric of the surface and does not need the local flatness assumption made in FEM methods [CWT⁺01][BX03]. In Euclidean case, discretization of the time derivative in the diffusion equations can be carried out using the Crank-Nicolson method [Smi85] due to its numerical accuracy and stability. Our approach allows us to generalize this method to non-Euclidean cortical surface thus making our method both stable and accurate.

4.2 Mathematical Formulation

We assume a genus zero cortical surface on which we define a scalar valued field representing the anatomical or functional image of interest. We also assume that a 2D coordinate system is assigned to this surface through a parameterization process. We summarize our approach to generating this parameterization in Section 4.3. Our goal is to define filtering operations on this image which are computed with respect to the intrinsic geometry of the surface.

Throughout this chapter we use Einstein's summation convention [Do 76][Kre99] in order to simplify the notation. Let $I(s, t)$ be a scalar function which denotes the image given on the cortical surface S and t denotes time. $I(s, 0)$ represents the original unsmoothed image. Let $g, g_{ij} : i, j \in 1, 2$ denote the metric tensor associated with S for a given coordinate system and $g^{ij} : i, j \in 1, 2$ denote inverse of the metric g_{ij} .

4.2.1 Isotropic filtering

The isotropic diffusion equation on surface S , with the original image $I(s, 0)$, $s \in S$ as the initial condition, is given by

$$\frac{\partial}{\partial t} I(s, t) = \Delta I(s, t), \quad (4.1)$$

where Δ is the Laplace-Beltrami operator that generalizes the Laplacian in Euclidean space to Riemannian space S :

$$\Delta I(s, t) = \frac{1}{\sqrt{g}} \frac{\partial}{\partial u^\nu} \left(\sqrt{g} g^{\mu\nu} \frac{\partial I(s, t)}{\partial u^\mu} \right). \quad (4.2)$$

We discretize this operator using the discretizations of the metric tensor and thus explicitly model the geometry of the surface in our method.

The discretization of the time derivative on the left hand side of (4.1) can be carried out by explicit discretization methods for hyperbolic PDEs. In the explicit scheme for solving (4.1), time t is discretized using a forward difference and $I(s, n)$ is used for calculation of the left-hand side of (4.1), where $I(s, n)$ denotes the image value at iteration n . Let L denote the discretization of Δ and δ the time step; the resulting discretized equation is given by

$$\frac{I(s, n + 1) - I(s, n)}{\delta} = LI(s, n). \quad (4.3)$$

Rearranging terms we have the update equation:

$$I(s, n + 1) = I(s, n) + \delta LI(s, n). \quad (4.4)$$

This is an explicit method for solving the heat equation. While it has the advantage of being fast to compute, the choice of δ is critical in the implementation, with large values

of δ resulting in numerical instability producing oscillatory solutions. A theoretical upper limit on the size of δ depends on grid size and the metric tensor coefficients $\{g_{i,j}\}$ and is hard to determine. Violating the upper limit on the value of δ causes amplification of numerical errors which in turn results in divergence of the solution. [Smi85]

In order to overcome this difficulty, here we adapt the Crank-Nicolson scheme to suite our particular equation. While it is slower than the explicit method (4.4), it has the advantage of being stable as well as accurate [Smi85]. In this case, (4.1) is discretized as

$$\frac{I(s, n + 1) - I(s, n)}{\delta} = \frac{1}{2}L(I(s, n) + I(s, n + 1)). \quad (4.5)$$

Rearranging terms gives:

$$\begin{aligned} I(s, n + 1) - \frac{\delta}{2}LI(s, n + 1) &= I(s, n) + \frac{\delta}{2}LI(s, n) \\ L_1I(s, n + 1) &= b, \end{aligned} \quad (4.6)$$

where $L_1 = I - \frac{\delta}{2}L$ and $b = I(s, n) + \frac{\delta}{2}LI(s, n)$ This linear system of equations is then solved at each iteration using conjugate gradient to compute $I(s, n + 1)$ from $I(s, n)$.

4.2.2 Anisotropic filtering

Anisotropic diffusion filtering of planar images was first described by Perona and Malik [PM90]. Here we generalize this idea to non-Euclidean surfaces, which allows us to perform spatially variant and image dependent nonlinear filtering of surface constrained image data within the geometry of the surface itself. The anisotropic diffusion filter is formulated as a diffusion process that encourages smoothing within regions of slowly

varying intensity while inhibiting smoothing across boundaries characterized by large image gradients. The anisotropic diffusion equation has the form:

$$\frac{\partial I(s, t)}{\partial t} = \nabla \cdot (D(s, t) \nabla I(s, t)), \quad (4.7)$$

where the diffusion coefficient $D(s, t)$ is a monotonically decreasing function of image gradient magnitude:

$$D(s, t) = f(\|\nabla I(s, t)\|). \quad (4.8)$$

Varying the diffusion coefficient with image gradient allows for locally adaptive edge preserving smoothing. Two choices for f were suggested [PM90]:

$$f_1(s, t) = \exp\left(-\left(\frac{\|\nabla I(s, t)\|}{\chi}\right)^2\right) \quad (4.9)$$

$$f_2(s, t) = \frac{1}{\exp\left(-\left(\frac{\|\nabla I(s, t)\|}{\chi}\right)^{\alpha}\right)} \quad \alpha > 0, \quad (4.10)$$

Where χ is referred to as the flow constant. Since these filters are expressed using PDEs, they generalize to non-Euclidean spaces. For the cortical surface, we replace (4.7) by

$$\frac{\partial I(s, t)}{\partial t} = \frac{1}{\sqrt{g}} \frac{\partial}{\partial u^\nu} \left(\sqrt{g} D(s, t) g^{\mu\nu} \frac{\partial I(s, t)}{\partial u^\mu} \right). \quad (4.11)$$

To compute the diffusion constants we also need an estimator of the gradient. We replace $\|\nabla I(s, t)\|^2$ by the *differentiator of the first order* [Kre99] given by

$$\nabla(I(s, t), I(s, t)) = g^{\alpha\beta} \frac{\partial I(s, t)}{\partial u^\alpha} \frac{\partial I(s, t)}{\partial u^\beta}. \quad (4.12)$$

With these substitutions, the anisotropic heat equation is well-defined on the cortical surface independently of the particular parameterization used for its computation.

4.3 Discretization and Numerical Method

We use a p -harmonic map presented in Chapter 2 for parameterization in which we minimize a p -harmonic energy function while constraining a closed curve in the interhemispheric fissure, which divides the brain into two hemispheres, to map to the boundary of a unit square. This procedure produces a one-to-one mapping from each hemisphere to a unit square.

Let \mathbf{x} denote the 3-D position vector of a point on the cortical surface. Let u^1, u^2 denote the coordinates in the parameter space. The metric tensor coefficients required in the computation are given by:

$$g_{11} = \left\| \frac{\partial \mathbf{x}}{\partial u^1} \right\|^2, \quad g_{22} = \left\| \frac{\partial \mathbf{x}}{\partial u^2} \right\|^2, \quad (4.13)$$

$$g_{12} = g_{21} = \left\langle \frac{\partial \mathbf{x}}{\partial u^1}, \frac{\partial \mathbf{x}}{\partial u^2} \right\rangle, \quad (4.14)$$

$$g = g_{11}g_{22} - (g_{12})^2, \quad (4.15)$$

The inverse metric coefficients g^{ij} are given by:

$$g^{11} = \frac{g_{22}}{g}, \quad g^{12} = g^{21} = -\frac{g_{12}}{g}, \quad g^{22} = \frac{g_{11}}{g}, \quad (4.16)$$

4.3.1 Discretization Algorithm

In order to solve the diffusion equations numerically, we need to discretize the isotropic and anisotropic Laplace-Beltrami operators in (4.2) and (4.11). We use the unit-square p -harmonic maps of the triangulated tessellation of the cortical surface to define a 2D

coordinate system. The square maps for each hemisphere are resampled on a regular 256x256 grid. The co-ordinate system we assign to the cortical surface is depicted in Fig. 3.1. The two squares in the (u^1, u^2) parameter space represent the two hemispheres in the x^1, x^2, x^3 space. The boundaries of the squares correspond to the common interhemispheric fissure between the two cortical hemispheres. The neighborhood relations between the edges of the two squares is depicted by different arrows in the figure. Because the interhemispheric fissure is fixed on the boundary of the squares representing the two hemispheres, one can visualize the $u^1 - u^2$ parameter space as two squares placed on each other and connected at the boundaries of the squares. We follow these neighborhood relations when discretizing the partial derivatives at the boundary of the two squares. This allows us to compute partial derivatives across the two cortical hemispheres making the boundary separating them completely transparent to the numerical discretizations. This boundaryless parameter space is then used for discretizing the partial derivatives with respect to the u^1 and u^2 spatial coordinates in the solution of the differential equations. For instance, assume that $f : M \rightarrow \mathbb{R}$ is a scalar valued function defined on the cortical surface M . We arrange its discretized representation at each vertex in the regular grid of the surface in a vector \vec{f} . In order to discretize $\frac{\partial f}{\partial u^1}$ by central differences on the entire surface, we calculate usual central differences at the points which are not on the boundary of the squares. On the boundary points of the squares, we use the connectivity relationship shown in Fig. 3.1 for the neighborhood in the central difference approximation. Using these relations, we compose a central difference matrix $D_{u^1}^c$ and obtain discretization of $\frac{\partial f}{\partial u^1}$ as $D_{u^1}^c \vec{f}$. Similarly we compose matrices $D_{u^1}^f, D_{u^1}^b$, the forward and backward difference operators for the u^1 coordinate, and $D_{u^2}^c, D_{u^2}^f$ and $D_{u^2}^b$, the central, forward and backward difference operators for the u^2 coordinate.

We carry out the discretization of the isotropic and anisotropic operators as described in the following steps:

1. Parameterize the cortical surface to map it in two squares and assign it the coordinate system described above.
2. Form the forward, backward and central difference matrices $D_{u^1}^f, D_{u^2}^f, D_{u^1}^b, D_{u^2}^b$ and $D_{u^1}^c, D_{u^2}^c$ for u^1 and u^2 coordinates respectively.
3. Compute the surface metric coefficients $g, g_{11}, g_{12}, g_{21}$ and g_{22} and also the inverse metric coefficients g^{11}, g^{12}, g^{22} . This is done by replacing partial derivatives in (4.13), (4.14), (4.15), (4.16) by their discrete versions from step 2.
4. Compute the isotropic or anisotropic Laplace-Beltrami operators using (4.2) or (4.11).


In the case of isotropic diffusion discretization of the diffusion operator needs to be carried out only once before starting the time iterations. On the other hand, for anisotropic diffusion the diffusion operator depends on $I(s, t)$ and hence needs to be updated by carrying out the last step 4 repeatedly after every time step. In order to decrease this numerical cost, we update the operator after every 100 iterations assuming that the incremental change in $I(s, n)$ is small.

The impulse response of the isotropic diffusion filter is shown in Fig. 4.1 both on the cortical surface and in the parameter space of one hemisphere. It can be seen that use of the surface metric results in a more circularly symmetric impulse response on the cortical surface. Note that because of the non-linear nature of the anisotropic diffusion filter, its behavior cannot be characterized by its impulse response.

We performed numerical simulations on an Intel Pentium 4 3.2 GHz computer with 2GB of RAM using MATLAB. The cortical surface was extracted from a 256 x 256




(a) The heat kernel computed using the Laplacian in the (u, v) parameter space





(b) The heat kernel computed using the Laplace-Beltrami operator on the cortical surface

Figure 4.1: The impulse response of the isotropic smoothing filters are displayed in the parameter space and on the surface [JSTL05]. It can be seen that when the surface metric is used to compute the Laplace-Beltrami, the impulse response kernel is not isotropic in the parameter space, however it is isotropic on the surface.



x 170 voxel T1-weighted MR image of a volunteer subject. Processing time from the raw MR volume to extraction of the topologically corrected and tessellated cortex using BrainSuite took 7 mins. The tessellated cortex had a total of 1.4 million nodes. The p harmonic parameterization of the 1.4 million node cortical surface took 37 secs. Note that adding the parameterization step does not add significantly to the total computational cost compared to a direct FEM method [BX03][LPDS04]. The number of iterations n along with the size of time step δ decide the amount of smoothing applied. Smaller values of δ result in more numerically accurate solutions while the execution time is directly proportional n . We chose $\delta = 1 \times 10^{-5}$ and $n = 40000$. Isotropic diffusion on the resampled surface took 20 mins with this choice of n and δ while



(a) Noisy mean curvature



(b) Isotropic diffusion



(c) Anisotropic diffusion




Figure 4.2: left: The mean curvature of the cortical surface plotted on a smoothed representation (for improved visualization of curvature in sulcal folds; right: The mean curvature plotted in 2D parameter space for a single cortical hemisphere. Isotropic diffusion blurs the regions as well as the edges separating them while while anisotropic diffusion reduces noise while preserving edges.

anisotropic diffusion took 1.5 hours. The difference in execution times is mainly due to the non-linear nature of the anisotropic diffusion which requires re-computation of the diffusion operator L repeatedly during the iterations. The code through parameterization was implemented in C/C++ with substantial effort directed at optimizing run-times while the diffusions were computed in MATLAB and, based on our earlier experience, we can expect a several-fold speed up when these are reprogrammed directly in C/C++.

We illustrate the diffusion operations by running them on mean curvature maps computed on the cortical surface. We compute the mean curvature using the method described in [CMR⁺03] and resample it on the regular grid. However we note that as an added advantage of our approach is that we can also compute the mean curvature by using our discretization of the metric tensor. In particular, the mean curvature H can be computed as:

$$H = \frac{1}{2} b_{\alpha\beta} g^{\alpha\beta}, \quad (4.17)$$

where the *second fundamental form* $b_{\alpha\beta}$ is given by

$$b_{\alpha\beta} = \frac{1}{\sqrt{g}} \left| \frac{\partial \mathbf{x}}{\partial u^1} + \frac{\partial \mathbf{x}}{\partial u^2} + \frac{\partial^2 \mathbf{x}}{\partial u^\alpha \partial u^\beta} \right|. \quad (4.18)$$

The minima and saddle points of the mean curvature of the cortical surface are known to follow the sulcal patterns [CMR⁺03] and therefore are vital features for automatic labeling of the sulci [RHXP02]. However, as can be seen in Fig. 4.2(a), there is a considerable amount of noise in the mean curvature computed on the cortical surface. This is primarily due to the fact that the mean curvature is a local feature and is therefore prone to errors in extraction and discretization of the cortical surface. We see that the isotropic diffusion filtering smooths out this noise, but since this filtering is not region selective, it also blurs the regions between sulci (positive mean curvature) and

gyri (negative mean curvature) as seen in Fig. 4.2(b). On the other hand, anisotropic filtering removes noise while carrying out the smoothing only within regions and thus respecting boundaries between sulci and gyri as seen in Fig. 4.2(c), thus illustrating the advantage of anisotropic filtering. Isotropic filtering can be used where such selective smoothing is not required such as in smoothing of functional data when smoothness is required for application of parametric random-field methods for control of false positives in multiple hypothesis testing [WMN⁺96]. These techniques can also be used for multiscale representations of functional activation [JSTL05], statistical data [CRD⁺05] and neuro-anatomical variability [THdZ⁺03].

4.3.2 The Heat Equation in the Intrinsic Geometry

The heat equation in the intrinsic geometry of the surface is given by:

$$\left(\Delta - \frac{\partial}{\partial t}\right)\zeta = 0 \quad \text{where } \Delta = \frac{1}{\sqrt{g}} \frac{\partial}{\partial u^i} \sqrt{g} g^{ij} \frac{\partial}{\partial u^j}$$

where Δ denotes the Laplace-Beltrami operator and ζ is the scalar field being diffused. We discretized the operator using the metric tensor calculations described in the Appendix. Using this discretized operator, we set up the Crank-Nicolson scheme [Smi85] for solving the heat equation since it is known to be stable. We illustrate the differences between using the usual Laplacian and the Laplace-Beltrami operator in Fig. 4.1. In the former, diffusion is computed with respect to the 2D Euclidean space and produces a 2D Gaussian distribution in the flat parameter space which maps to a clearly anisotropic distribution on the surface. Conversely, the Laplace-Beltrami form computes the diffusion directly on the surface, on which it produces an isotropic distribution while exhibiting anisotropic behavior with respect to the parameter space. Solutions of linear partial differential equations, such as the heat equation, can be characterized by

Green's functions. The Green's function of the heat equation, also known as the *heat kernel*, has been a topic of extensive research in spectral theory [Ros97]. Though the heat kernel can only be implicitly defined in arbitrary surfaces, several of its properties in Euclidean spaces extend to Riemannian spaces and, in particular, to surfaces.

Here we list a few properties we will use later in this chapter. Proofs can be found in [Ros97]. Let M be a geodesically complete Riemannian manifold. Then the heat kernel $K_t(x, y)$ exists and satisfies

1. $K_t(x, y) = K_t(y, x)$
2. $\lim_{t \rightarrow 0} K_t(x, y) = \delta_x(y)$
3. $(\Delta - \frac{\partial}{\partial t})K = 0$
4. $K_t(x, y) = \int_M K_{t-s}(x, z)K_s(z, y)dz$
5. $K_t(x, y) = \sum_{i=0}^{\infty} e^{-\lambda_i t} \phi_i(x)\phi_i(y)$

4.4 The Heat Kernel as a PDF

We know that the heat kernel is positive everywhere. It integrates to one on the manifold [Dav89] and therefore it is a suitable candidate for modeling the probability density function of sample points lying in the manifold. Moreover, in Euclidean space, the heat kernel is identical to the Gaussian pdf. Therefore we propose replacing the Gaussian density with the covariant heat kernel in our surface-based analysis [Hsu02].

Just as we can characterize an isotropic Gaussian distribution in the Euclidean plane through its mean and standard deviation, so we can characterize distributions on the surface through mean and variance-like parameters that characterize the location of the heat kernel and the 'time' at which it is observed. Estimation of these parameters is in

turn analogous to maximum likelihood parameter estimation, i.e. parameter estimation for a set of sample points on the surface can be viewed as the problem of finding the kernel of a covariant differential operator that best fits these points.

For isotropic distributions the corresponding heat kernel $K(m, t)$ on a Riemannian manifold can be completely specified by two parameters: m , the location of the initial impulse, and the time t . Parameters m and t play the role of the mean and variance in the Gaussian case. Thus the probability of finding a sample at x is modeled as $p(x|m, t) = K_t(m, x)$. So the problem of fitting the heat kernel in the given sample points can be reduced to the problem of estimating these two parameters of the heat kernel.

If the sample points are x_i , we define the likelihood function for m and t as:

$$L(m, t) = \prod_{n=1}^N K_t(m, x_i)$$

Because of property 2 above, $K_t(m, x)$ can be calculated explicitly by placing a delta function at point m and solving the heat equation up to time t . The problem with this approach is that the parameter m (the location of the mean) is unknown. However, since the heat kernel is symmetric (property 1), we can instead place the delta function at the sample points x_i whose locations are known, rather than at the unknown mean location m , and running the heat equation up to time t . This allows us to explicitly compute the likelihood function (4.19) for a set of sample points x_i for any time point t . The values of m and t for which the likelihood function $L(m, t)$ attains its maximum are then our estimates of the mean and variance.

To use this scheme for supervised classification of two clusters of points, we first compute ML estimates of the parameters (m_1, t_1) and (m_2, t_2) for the two clusters. We then define a likelihood ratio as the ratio of the two heat kernels: $R =$



(a) pdf estimated for digit 1



(b) pdf estimated for digit 5

Figure 4.3: The figures shows the heat kernels estimated to fit the two datasets for MEG somatosensory data. For each of the datasets the estimated pdf is displayed in the parameter space and on the cortical surface.




Figure 4.4: The classifier: Red and Blue regions shows the two decision regions

$K_1(m_1, t_1)/K_2(m_2, t_2)$ and compute this ratio at each point on the surface. The surface is then partitioned into two regions $R > 1$ and $R \leq 1$.

We illustrate the technique presented above for classification of point localizations of S1 somatosensory regions. For each of 5 subjects we simulated 6 points each representing locations of thumb and index figure on the postcentral gyrus; the 6 points could, for example, represent localizations from 6 separate studies on a single subject. We brought

the cortical surfaces for all subjects into register, using one of the subjects as the atlas, as described above. We then used the pooled data from all subjects in the atlas-coordinates to compute the mean and standard deviation for the thumb and index finger respectively as illustrated in Fig. 4.3. We then applied the likelihood ratio statistic to partition the cortex as illustrated in Fig. 4.4. Note that this two-class problem classifies the entire brain, including both hemispheres, into two regions. With more somatosensory areas involved we could perform a finer partitioning of somatosensory cortex producing maps of the most probable areas to which each sensory unit would map. While this is a somewhat artificial problem, it is clear that an extension of this analysis would allow us to produce probabilistic maps of functional localization in the atlas space.

Chapter 5

Volumetric Registration using Harmonic Maps

5.1 Introduction

In this chapter we describe an approach to brain image registration based on harmonic maps that combines surface based and volume based approaches producing a volumetric alignment in which there is also a one-to-one correspondence between points on the two cortical surfaces.

Talairach normalization based on a piecewise affine transformation [TT88] was the first commonly used volumetric alignment technique. Because it uses a restricted set of anatomical landmarks and is piecewise affine, it results in relatively poor alignment and has been largely replaced by automated intensity-based alignment methods that also allow non-rigid deformations [AF99, WGH⁺98]. There are a vast array of such methods, differing in how they measure the fit between the two images (e.g., squared error, correlation, mutual information), the parameterization of the transformation (e.g., polynomial, splines, discrete cosine transform or other eigenfunction bases), and the procedure used to regularize the transformation (e.g., elastic, biharmonic, or viscous fluid models) [HBHH01]. Polynomial warps and linear elastic deformations implicitly assume that deformations are small and do not guarantee preservation of topology for larger deformations [CRM⁺95]. The viscous fluid approach [CRM96] and more recent approaches using large-deformation diffeomorphic metric mapping [GVM04, AG04]

were developed to address the problem of ensuring diffeomorphic maps and are better able to register objects whose alignment requires large deformations while conserving their topology.

Since these intensity-based methods do not explicitly model the cortical surface, alignment can be rather poor. An illustration of this is shown in Fig. 5.1, where we have used the Automated Image Registration (AIR) software [WGH⁺98, WGW⁺98] to align two brain volumes using a 5th order polynomial (168 parameters). While the regions of cortical grey matter exhibit reasonably good correspondence between the two images, the cortical surfaces themselves do not align well. Since cytoarchitectural and functional parcellation of the cortex is intimately related to the folding of the cortex, it is important when comparing cortical anatomy and function in two or more subjects that the surfaces are aligned. For this reason, there has been an increasing interest in analyzing the cerebral cortex based on alignment of surfaces rather than volumes.

Various surface-based techniques have been developed for inter-subject registration of two cortical models. One class of techniques involves flattening the two cortical surfaces to a plane [HSB⁺00] or to a sphere [FSTD98] using mechanical models or variational methods and then analyzing the data in the common flattened space [BGKM98]. Other surface based techniques work in the surface geometry itself rather than a plane or a sphere and choose to account for the surface metric in the inter-subject registration [TWMT00, JSTL05]. The advantage of such techniques is that they produce registration results that are independent of the intermediate flat space (or, equivalently, the specific parameterization of the cortex) resulting in a more consistently accurate registration throughout the cortex. These approaches involve manually delineated sulcal landmark matching [JSTL05] in the intrinsic surface geometry. While some progress has been made recently towards automating the matching process using mutual information [WCT05] or optical flows of mean-curvature images in the surface parameter

space [TP05, TRP05], fully automatic alignment of high resolution cortical surfaces remains a challenging problem.

While the volume registration methods described above do not provide suitable cortical alignment, the cortical registration methods do not define any volumetric correspondence. One approach is to combine landmark points, curves and surfaces as additional constraints in an intensity-based warping method [PCS⁺89, TT96a, KL99, DLF99, HHCS⁺02, DPB96, DP94]. For example, landmarks, curves [DP94] and image matching [DPB96] are applied in a hierarchical manner in a large deformations framework ensuring generation of diffeomorphisms [JM00, GJF⁺06]. Registration methods such as the Hierarchical Attribute Matching Mechanism for Image Registration (HAMMER) algorithm [LSD04] incorporate surface as well as volume information for the alignment. For brain images, the desired deformation fields need to be obtained incrementally by using large deformation or fluid models [CYV⁺00, JC02] and hence are computationally expensive. Additionally, accurate alignment of the cortical surface as well as the cortical volume remains a challenging task mainly due to the complex folding pattern variability of the cortex.

In this chapter, we propose a new method based on harmonic mappings for extending the surface matching to the entire cortical volume, and present a modified intensity alignment based on [Chr99] to compute the final map. The resulting method, comprising the three steps outlined above, gives an inverse consistent map which is capable of aligning both subcortical and sulcal features.

5.2 Problem Statement and Formulation

Here we address the following problem: produce a one-to-one mapping between two brain volumes such that subcortical structures and sulcal landmarks are aligned and that

there is also a one-to-one correspondence between the cortical surfaces of the two volumes. Equivalently, given 3D manifolds M and N representing the two brain volumes, with boundaries ∂M and ∂N representing their respective cortical surfaces, we want to find a map from M to N such that ∂M , the surface of M , maps to ∂N , the surface of N , and the intensities of the images in the interior of M and N are matched. In addition the map must satisfy a sulcal matching constraint so that labelled sulci on the surface ∂M map onto homologous sulci on ∂N . The boundaries, ∂M and ∂N , are assumed to have a spherical topology.

We solve the mapping problem in three steps:

1. Surface matching, which computes a map between ∂M and ∂N , the cortical surfaces of the two brains. The mapping is based on minimization of an elastic strain energy subject to the constraint that a set of interactively labelled sulci are aligned, as described in Chapter 3.
2. Extrapolation of the surface map to the entire cortical volume such that the cortical surfaces remain aligned. This is done by computing a harmonic map between M and N subject to a surface matching constraint. As we describe in Section 5.3, an intermediate spherical representation is used to facilitate enforcement of this constraint. We note also that while the sulci are constrained to remain in correspondence, the cortical surfaces can flow with respect to each other when computing the volume harmonic map provided we retain the one-to-one mapping between ∂M and ∂N .
3. Refinement of the harmonic map on the interiors of M and N to improve intensity alignment of subcortical structures. For this step we use an inverse consistent linear elastic registration method as described in Section 5.5.




Figure 5.1: Cortical surface alignment after using AIR software for intensity based volumetric alignment using a 168 parameter 5^{th} order polynomial. Note that although the overall morphology is similar between the brains, the two cortical surfaces do not align well.

The mapping in Step 2 requires large scale deformation to ensure that ∂M and ∂N are aligned. Linear elastic or thin-plate spline registration based on landmarks cannot be used for this purpose [EÅ05]. Harmonic maps on the other hand are suitable since they are bijective provided that the boundary (the cortical surface in this case) is mapped bijectively. Conversely, the final step requires a more local refinement of the mapping to align subcortical structures so that use of linear elastic methods is appropriate.

5.3 Indirect Mapping Approach

The surface registration procedure described in Chapter 3 sets up a point to point correspondence between the two cortical surfaces, which represent the boundary of the two cerebral volumes. Extrapolating this correspondence from the boundary surface to the entire cerebral volume in a one-to-one manner is challenging due to the convoluted nature of the cortex. In fact, most of the linear models such as linear elastic or thin-plate splines become non-bijective under relatively mild landmark matching constraints [EÅ05]. 3D harmonic maps are attractive for this purpose due to their tendency to be

bijjective if the boundary (cortical surface) is mapped bijectively, which is the case here. In this section we describe a framework for computing a harmonic mapping between two 3D volumes as well as the computational approach used for implementation. Details of harmonic maps and their properties can be found in [Jos02].

Let $u : M \rightarrow N$ be a C^∞ map from a 3 dimensional Riemannian manifold (M, g) to an 3 dimensional Riemannian manifold (N, h) where g and h are Riemannian metrics for M and N respectively. A Riemannian metric defines an inner product at every point in the manifold and thus helps in defining the notion of distance on the manifold [Jos02]. Let $\{g_{ij}; i, j \in \{1, 2, 3\}\}$ denote components of the Riemannian metric tensor g and $\{h_{\alpha\beta}; \alpha, \beta \in \{1, 2, 3\}\}$ denote the components of the Riemannian metric tensor h . The inverse of the metric $g = \{g_{ij}\}$ is denoted by $\{g^{ij}\}$. Let (x^1, x^2, x^3) and (u^1, u^2, u^3) be local coordinates for x and $u(x)$ respectively. Let Du denote the derivative (generalized Jacobian) of the map. The *energy density function* $e(u)$ of map u is defined to be norm of Du [Nis01] and is given by

$$e(u)(x) = \frac{1}{2} |Du_x|^2 \quad (5.1)$$

$$= \frac{1}{2} \sum_{i,j=1}^3 \sum_{\alpha,\beta=1}^3 g^{ij}(x) h_{\alpha\beta}(u(x)) \frac{\partial u^\alpha(x)}{\partial x^i} \frac{\partial u^\beta(x)}{\partial x^j}, \quad (5.2)$$

which can be thought of as the rate of expansion of the map u in orthogonal directions, at point $x \in M$ [Nis01]. The *mapping energy* is defined as:

$$E(u) = \int_M e(u)(x) d\mu_g. \quad (5.3)$$

Therefore, in coordinate form [Nis01], it is given by

$$E(u) = \frac{1}{2} \int_M \sum_{i,j=1}^3 \sum_{\alpha,\beta=1}^3 g^{ij}(x) h_{\alpha\beta}(u(x)) \frac{\partial u^\alpha(x)}{\partial x^i} \frac{\partial u^\beta(x)}{\partial x^j} d\mu_g \quad (5.4)$$

where the integration is over the manifold M with respect to the intrinsic measure $d\mu_g$ induced by its Riemannian metric g .

A *harmonic map* from (M, g) to (N, h) is defined to be a critical point of the mapping energy $E(u)$. In this sense harmonic maps are the least expanding maps in $C^\infty(M, N)$, the space of all smooth maps from M to N . Therefore, among all possible smooth maps between two manifolds, the harmonic maps have the tendency to avoid overlaps and folds in the map, resulting in a bijective map.

A number of existence, uniqueness, and regularity results have been proven for harmonic maps [Xin96]. Eells and Sampson [ES64] proved the existence of a harmonic map from any compact Riemannian manifold to a compact Riemannian manifold of non-positive sectional curvature. Hamilton [Ham75] generalized this result to manifolds with boundaries. In medical imaging, harmonic mappings and p -harmonic mappings, their generalized counterparts [FR02], have been used for various applications such as surface parameterization and registration [AHTK99, KSK98a, JLTS04] and image smoothing [TSC00]. Wang et al. [WGY04] describe a method for volumetric mapping of the brain to the unit ball $B(0, 1)$. Here we use harmonic maps to align two brain volumes so that both the brain volumes and cortical surfaces are aligned.

When computing the harmonic maps we could fix the correspondence between the two surfaces using the method from Chapter 3 and map only the interior of the two volumes. This would result in a suboptimal mapping with respect to the 3D mapping energy. To overcome this limitation, we instead allow the surface M to flow within the surface of N when computing the map. The only constraints placed on the surfaces are that the maps are aligned at the set of user defined sulcal landmarks and that the boundary ∂M maps onto ∂N . This less restrictive surface mapping constraint cannot be formulated directly in the ambient Euclidean 3D space since there is no analytical expression for the surfaces. It could be accomplished without parameterizing the surface

using a level set approach [TSC00, MSO04]. Here we use an intermediate representation for the manifolds which allows us to enforce the boundary matching constraint while allowing one boundary to flow within the other. We achieve this by first mapping to the unit ball as described below. This mapping to the unit ball results in a non-Euclidean representation of N thus requiring the use of the Riemannian metric in computing the harmonic map.

5.3.1 Mathematical Formulation

We find the map v of the 3D brain manifold N to the 3D unit ball $B(0, 1)$ [WGY04] using the method described in Sec. 5.3.3. Let $v = (v^1, v^2, v^3)$ denote three components of the map v . This map is bijective and therefore we can treat the unit ball $B(0, 1)$ as an alternative representation (N, h) of the manifold N , with associated metric h , that has the advantage over the Euclidean space (N, I) that the cortical surface lies on the surface of the sphere (here I represents the identity metric for the Euclidean space); h is the metric induced by the map v . With this alternative representation of N , the components of its metric $h_{\alpha\beta}$ at point $x = (x^1, x^2, x^3)$ are given by:

$$h_{\alpha\beta} = \sum_{i=1}^3 \frac{\partial x^i}{\partial v^\alpha} \frac{\partial x^i}{\partial v^\beta}. \quad (5.5)$$

Now instead of needing to directly compute the harmonic map $u : (M, I) \rightarrow (N, I)$, we instead find the harmonic map $\tilde{u} : (M, I) \rightarrow (N, h) \approx B(0, 1)$ subject to the constraint that the cortical surface ∂M maps to the spherical boundary of the unit ball, as illustrated in Fig. 5.2.

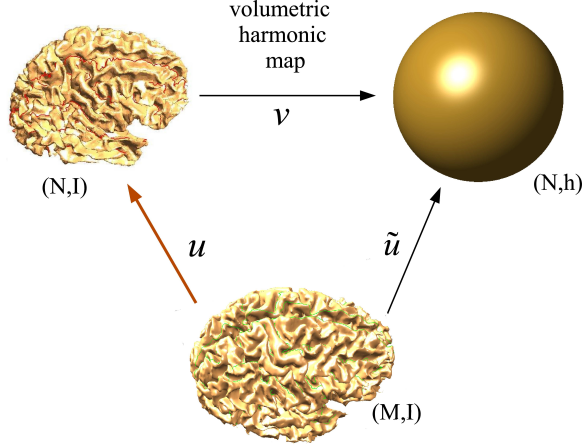


Figure 5.2: Illustration of our general framework for surface-constrained volume registration. We first compute the map v from brain manifold (N, I) to the unit ball to form manifold (N, h) . We then compute a map \tilde{u} from brain (M, I) to (N, h) . The final harmonic map from (M, I) to (N, I) is then given by $u = v^{-1} \circ \tilde{u}$.

Since M remains in the Euclidean space, its metric is I , so $g^{ij}(x)$ is the identity operator and the harmonic mapping problem (5.4) becomes:

$$\tilde{u} = \arg \min_{\gamma} \int_M \sum_{i=1}^3 \sum_{\alpha, \beta=1}^3 h_{\alpha\beta}(\gamma(x)) \left(\frac{\partial \gamma^\alpha(x)}{\partial x^i} \right) \left(\frac{\partial \gamma^\beta(x)}{\partial x^i} \right) d\mu_g \quad (5.6)$$

subject to $\|u(\tilde{x})\|^2 = 1$ for $x \in \partial M$, the surface of M . Note that this constraint allows the surface map to flow within the spherical boundary. We also want to constrain the maps so that predefined sulcal landmarks are aligned. To achieve this we impose the additional constraints that $\tilde{u}(c) = u_c$ for $c \in M_c$ where M_c are the set of sulcal landmark points in M and u_c are the locations of the homologous landmarks in (N, h) . Having obtained \tilde{u} by minimizing the integral in (5.6), the final harmonic mapping from $u : (M, I) \rightarrow (N, I)$ can then be computed as $u = v^{-1} \circ \tilde{u}$ as illustrated in Fig. 5.2.

5.3.2 Initialization Procedure

Because the minimization problem (5.6) is nonlinear, it is important to have a good initial estimate of the map \tilde{u} in order to achieve convergence in reasonable time. We therefore generate an initial estimate \tilde{u}_0 of \tilde{u} by computing a map of the second manifold (M, I) to the unit ball, just as we do for the first manifold (N, I) (Fig. 5.2). Thus our initialization generates a bijective initial map, which is not necessarily harmonic. The procedure is illustrated in Fig. 5.3.

The initialization consists of the following steps. We first compute flat maps to the unit square for each hemisphere of the two brains with aligned sulci as described in Chapter 3. A stereographic projection then maps the two hemispheres of each brain to the unit sphere so that the corpus callosum that forms the boundary of the unit squares maps to the equator. Using these surface maps as constraints, we then map N and M to the unit ball to provide, respectively, the unit ball manifold (N, h) and an initial estimate \tilde{u}_0 of the desired map \tilde{u} from (M, I) to (N, h) . The initial map obtained in this manner is smooth and bijective. With this initialization, the 3D harmonic map is computed by minimizing (5.6) to obtain the final harmonic mapping from M to N .

5.3.3 Mapping to the Unit Ball $B(0, 1)$

In the special case when (M, g) and (N, h) are 3D Euclidean manifolds, then $h_{\alpha\beta} = \delta_{\alpha}^{\beta}$, $g_{ij} = g^{ij} = \delta_i^j$, the Kronecker delta, or identity tensor, for $\alpha, \beta, i, j \in 1, 2, 3$, and the mapping energy simplifies to

$$E(u) = \int_M |\nabla u|^2 dV \quad (5.7)$$

where ∇ is the usual gradient operator in 3D Euclidean space and dV is the volume integral [WGY04]. In order to map the given cortical brain volume M to the unit ball,

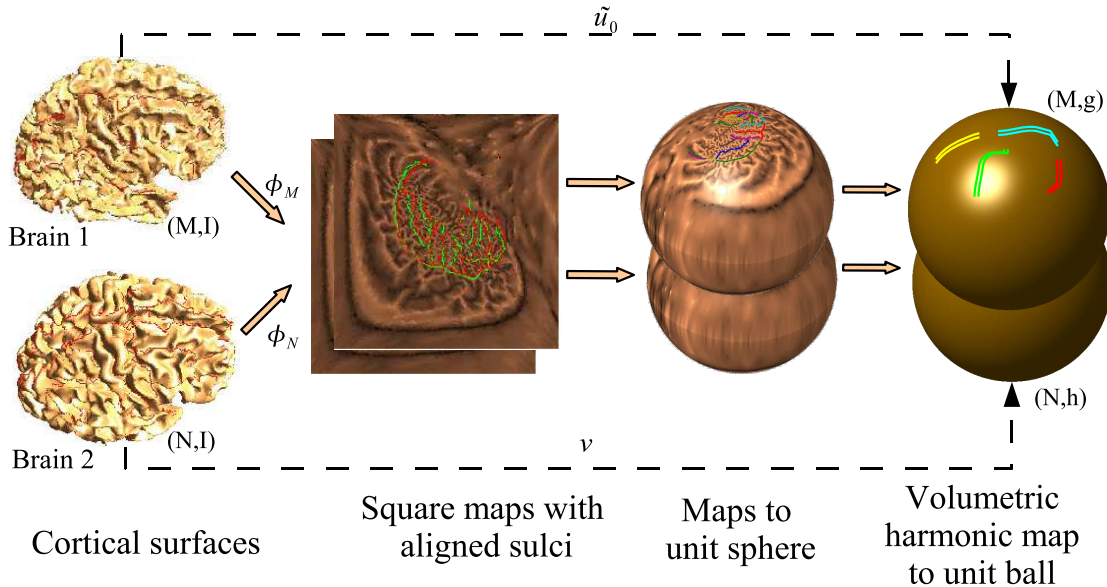


Figure 5.3: Initialization for harmonic mapping from M to N . First we generate flat square maps of the two brains, one for each hemisphere, with pre-aligned sulci. The squares corresponding to each hemisphere are mapped to a disk and the disks are projected onto the unit sphere. We then generate a volumetric map from each of the brains to the unit ball. Since all these maps are bijective, the resulting map results in a bijective point correspondence between the two brains. However, this correspondence is not optimal with respect to the harmonic energy of maps from the first brain to the second, but is used as an initialization for minimization of (5.6).

this energy is minimized subject to the constraint that the surface of M maps to the surface of the unit ball using the point-to-point correspondence defined by the flat mapping obtained as described in Chapter 3. This is computed by numerical integration over the voxel lattice using finite differences to approximate the gradients in (5.7). The resulting function is minimized using a preconditioned conjugate gradient method. The process of mapping to the unit ball is illustrated in Fig. 5.4 where we show iso-surfaces in brain coordinates corresponding to different radii, r , within the unit ball. At $r = 1$ we are at the outer surface of the brain and see the full cortical surface. As r is reduced we




Figure 5.4: Illustration of the deformation induced with respect to the Euclidean coordinates by mapping to the unit ball. Shown are iso-surfaces corresponding to the Euclidean coordinates for different radii in the unit ball. Distortions become increasingly pronounced towards the outer edge of the sphere where the entire convoluted cortical surface is mapped to the surface of the ball.

see successively less distortion since the harmonic map is driven entirely by the surface constraint.

5.3.4 Harmonic Mapping Between the Two Brains

The mapping to the unit ball is applied to both brain volumes M and N . The mapping of the Euclidean coordinates in M to the unit ball provides the initial estimate \tilde{u}_0 of the harmonic map \tilde{u} . We then refine this map by minimizing the harmonic energy in (5.6) from (M, I) to (N, h) , the unit ball representation of N . Again, the problem is solved using numerical integration and finite difference operators, in this case accounting for the metric h according to (5.6) when computing these derivatives. In this mapping, the locations of the sulci in M are constrained using their initial mappings \tilde{u}_0 computed when flattening and matching the cortical surfaces. Other points on the surface are allowed to move freely to minimize the harmonic energy, subject to the constraint that all points on the surface map to $\|\tilde{u}\|^2 = 1$, which is achieved by adding a penalty function to the discretized form of (5.6).

5.3.5 Implementation

We first describe a numerical method for computation of the metric $h_{ij}(x)$ and then outline the harmonic mapping method.

Computation of Metric

The metric $h_{ij}(x)$, $x \in N$ is associated with the unit ball coordinates $B(0, 1)$ given to N by the map $v = (v^1, v^2, v^3)$ (Fig. 5.2). It is given by $h_{\alpha\beta}(p) = \sum_{i=1}^3 \frac{\partial x^i}{\partial v^\alpha} \frac{\partial x^i}{\partial v^\beta}$ with $\alpha, \beta \in \{1, 2, 3\}$ at $x = (x^1, x^2, x^3)$. Note that although $x \in N$ is in the regular grid, $v(x) \in B(0, 1)$ is not necessarily so, and hence computation of partial derivatives with respect to v directly is difficult. In order to compute $\frac{\partial v^\alpha}{\partial x^i}$, first compute $\frac{\partial v^\gamma}{\partial x^j}$ using finite differences and then use the chain rule identity

$$\sum_{\gamma=1}^3 \frac{\partial x^i}{\partial v^\gamma} \frac{\partial v^\gamma}{\partial x^j} = \frac{\partial x^i}{\partial x^j} = \delta_j^i \quad (5.8)$$

to solve for $\frac{\partial v^\gamma}{\partial x^j}$. The metric h_{ij} is computed by substituting these partial derivatives in the above equation.

Harmonic Mapping

The harmonic mapping procedure can now be summarized as follows:

1. Align the surfaces of both the brains M and N using the procedure described in Chapter 3.
2. Map the unit squares to unit disks by the transformation $(x, y) \rightarrow \left(\frac{x}{\sqrt{x^2+y^2}}, \frac{y}{\sqrt{x^2+y^2}} \right)$ and then project them onto two hemispheres using $(x, y) \rightarrow (x, y, \pm\sqrt{x^2+y^2})$.

3. Using this mapping of the cortical surface to the unit sphere as the boundary condition, generate volumetric harmonic maps of M and N to the unit ball $B(0, 1)$ as described in Sec. 5.3.3.
4. Compute the metric h associated with the unit ball $B(0, 1)$ coordinates of N as described above.
5. Minimize (5.6) holding the matched sulci fixed, and letting the cortical surface ∂M slide along boundary of the unit ball. This is done by minimizing (5.6) with the constraint that $\|\tilde{u}(x)\|^2 = 1$ for $x \in \partial M$ and $\tilde{u}(c) = \tilde{u}_0(c)$ for $c \in M_c$ where $M_c \subset M$ denotes the set of sulcal points on M . The partial derivatives in (5.6) are discretized by finite differences and the minimization is done by gradient descent.
6. Compute the deformation vector field $u(x) - x$ where $u = v^{-1} \circ \tilde{u}$ and apply this to map brain volume M to N . Trilinear interpolation is used for this deformation.

5.4 Direct Mapping Approach

The limitation of the approach presented in the previous sections is that by using the map to the unit ball, the method is restricted to mapping only the cerebral volume contained within the cortical surface. Here we avoid this restriction by computing the harmonic map directly in Euclidean space so that the entire brain volume can be registered. However, this approach keeps the surface points fixed during volumetric harmonic maps and hence the surface registration is suboptimal with respect to the volumetric energy [JSTL07b][JSTLed]. Since the map between the cortical surfaces is fixed, there is no longer a need for the intermediate spherical representation. While this approach places a more restrictive constraint on the mapping of the surface, in practice we see only a small difference between the two methods in the mapping of the interior of the cerebrum.

5.4.1 Mathematical Formulation

The registration problem is formulated in a similar manner to the approach used in Sec. 5.3. We start by aligning the cortical surfaces, semi-automatically, using sulcal landmarks.

Given two 3D manifolds M and N representing brain volumes, with ∂M_1 , ∂M_2 and ∂N_1 , ∂N_2 representing surfaces corresponding to cortical grey/white matter and grey/CSF boundaries, we want to find a map from M to N such that (i) ∂M_1 , the grey/white matter surface of M , maps to ∂N_1 , the grey/white matter surface of N ; (ii) ∂M_2 , the grey/CSF surface of M , maps to ∂N_2 , the grey matter/CSF surface of N ; and (iii) the intensities of the images in the interior of M and N are matched. The surfaces, ∂M_1 , ∂M_2 and ∂N_1 , ∂N_2 , are assumed to have a spherical topology. We solve the mapping problem in three steps:

1. Surface matching which computes maps between surface pairs - the cortical surfaces and the grey matter/csf surfaces of the two brains, with sulcal alignment constraints (Chapter 3);
2. extrapolation of the surface map to the entire cortical volume. This is done by computing a harmonic map between M and N subject to a surface matching constraint (Section 5.4.2), and
3. Refinement of the harmonic map on the interiors of M and N to improve intensity alignment of subcortical structures (Section 5.5).

5.4.2 Harmonic Mapping

The surface registration procedure described in Chapter 3 sets up a point to point correspondence between the pairs of surfaces ∂M_1 , ∂M_2 and ∂N_1 , ∂N_2 . As noted earlier,

treating these surfaces as landmarks is not helpful since they are highly convoluted and finding a volumetric diffeomorphism consistent with the surface map is non-trivial. One approach that can achieve such a diffeomorphism is to compute a harmonic map. A harmonic map $u = (u^1, u^2, u^3)$ from 3D manifold M to 3D manifold N is defined as the minimizer of the harmonic energy [Jos02],

$$E_h(u) = \frac{1}{2} \int_M \sum_{i=1}^3 \sum_{\alpha=1}^3 \left(\frac{\partial u^\alpha(x)}{\partial x^i} \right)^2 dV. \quad (5.9)$$

Note that (5.9) is quadratic in u^α and that the summands are decoupled with respect to α . Consequently the harmonic energy $E_h(u)$ can be separately minimized with respect to each component u^α , $\alpha \in \{1, 2, 3\}$.

We compute the minimizer of $E_h(u)$ using a conjugate gradient method with Jacobi preconditioner. The mapping of the two surfaces computed in the previous sections act as constraints such that ∂M_1 maps to ∂N_1 and ∂M_2 maps to ∂N_2 . This harmonic mapping extrapolates the surface mappings to the entire volume such that the surface alignments are retained.

5.5 Volumetric Intensity Registration

The surface constrained harmonic mapping procedure of the direct mapping approach in Sec. 5.3 or the direct mapping approach in Sec. 5.4 described above produces a bijective mapping between the two brain volumes. However, it uses only surface shape and sulcal labels and does not use the MRI intensity values to compute the map. The result is a large scale deformation that aligns surface features but will benefit from an intensity-based refinement aimed at aligning subcortical features. In order to do this refinement and also make the final map inverse consistent, we use linear elastic inverse consistent registration based on Christensen’s approach [Chr99] with the modifications




Figure 5.5: Schematic of the intensity alignment procedure. Once harmonic maps u^M and u^N are computed, we refine these with intensity driven warps w^M and w^N while imposing constraints so that the final deformations are inverse consistent.

described below to ensure that the entire mapping process, rather than just this last step, is inverse consistent.

5.5.1 Formulation

The surface constrained volumetric harmonic mapping procedure described above can be used to generate two maps $u^M : M \rightarrow N$ and $u^N : N \rightarrow M$, each harmonic, but not necessarily inverses of each other. The corresponding deformation fields for these maps can be expressed as $d_u^M(x) = u^M(x) - x, x \in M$ and $d_u^N(x) = u^N(x) - x, x \in N$. Note that both of these deformation fields accurately align the two surfaces and

corresponding sulci, and are one-to-one. These deformations are used to initialize the volumetric inverse consistent intensity registration procedure that we now describe.

Let $f_M(x), x \in M$ denote intensity at point $x \in M$ and $f_N(x), x \in N$ denote intensity at point $x \in N$. The situation can be summarized as follows and is illustrated in Fig. 5.5): We have harmonic maps $u^M : M \rightarrow N$ $u^N : N \rightarrow M$ that change the shapes of domains M and N to match their respective targets N and M . In order to align the intensities, we seek refinement maps $w^M : M \rightarrow M$ and $w^N : N \rightarrow N$ such that the mapped intensity value $f_M \circ w^M \circ u^N$ matches f_N (or equivalently $f_M \circ w^M$ matches $f_N \circ (u^N)^{-1}$), and $f_N \circ w^N \circ u^M$ matches f_M (or $f_N \circ w^N$ matches $f_M \circ (u^M)^{-1}$). For inverse consistency, we need $w^N \approx (u^M \circ w^M \circ u^N)^{-1}$ and $w^M \approx (u^N \circ w^N \circ u^M)^{-1}$. Let d_w^M, d_w^N denote the deformation fields corresponding to w^M, w^N and let $\tilde{d}_w^M, \tilde{d}_w^N$ denote the deformation fields for $(u^N \circ w^N \circ u^M)^{-1}, (u^M \circ w^M \circ u^N)^{-1}$.

The inverse consistency similarity cost function $C(d_w^M, d_w^N)$, can now be defined as the sum of three terms:

$$\begin{aligned}
C(d_w^M, d_w^N) &= C_{REG}(d_w^M, d_w^N) + \alpha C_{SIM}(d_w^M, d_w^N) \\
&+ \beta C_{ICC}(d_w^M, d_w^N) \text{ subject to } d_w^N(u^M(x)) = 0, x \in \partial M \text{ and} \\
& d_w^M(u^N(x)) = 0, x \in \partial N \quad (5.10)
\end{aligned}$$

where the boundary constraints ensure that the cortices remain aligned after registration and the three constituent terms are defined as follows:

$$\begin{aligned}
C_{REG}(d_w^M, d_w^N) &= \|L_M d_w^M\|^2 + \|L_N d_w^N\|^2 \\
C_{SIM}(d_w^M, d_w^N) &= \|f_M(x + d_w^M(x)) - f_N(u^{N^{-1}}(x))\|^2 + \\
&\quad \|f_N(x + d_w^N(x)) - f_M(u^{M^{-1}}(x))\|^2 \\
&\approx \|f_M(x) + \nabla_M f_M(x) \cdot d_w^M(x) - f_N(u^{N^{-1}}(x))\|^2 + \\
&\quad \|f_N(x) + \nabla_N f_N(x) \cdot d_w^N(x) - f_M(u^{M^{-1}}(x))\|^2 + \\
C_{ICC}(d_w^M, d_w^N) &= \|d_w^M(x) - \tilde{d}_w^M(x)\|^2 + \\
&\quad \|d_w^N(x) - \tilde{d}_w^N(x)\|^2 \\
C(d_w^M, d_w^N) &= C_{REG}(d_w^M, d_w^N) + \alpha C_{SIM}(d_w^M, d_w^N) \\
&+ \beta C_{ICC}(d_w^M, d_w^N) \text{ subject to } d_w^N(u^M(x)) = 0, x \in \partial M \text{ and} \\
&\quad d_w^M(u^N(x)) = 0, x \in \partial N \quad (5.11)
\end{aligned}$$

The first term is the regularizer where $L_M = \alpha \nabla_M^2 + \beta \nabla_M(\nabla_M \cdot) + \gamma$ and $L_N = \alpha \nabla_N^2 + \beta \nabla_N(\nabla_N \cdot) + \gamma$ denote the Cauchy Navier elasticity operators in M and N respectively. The second term measures the intensity match between the transformations in both directions and the third term is measure of deviation from the inverse consistent condition. This is a quadratic cost function and can be minimized by the conjugate gradient method. We use a preconditioned conjugate gradient method with Jacobi preconditioner for this purpose.

5.5.2 Implementation

1. First, the harmonic maps $u^M : M \rightarrow N$ and $u^N : N \rightarrow M$ are computed using the procedures described in Sec. 5.3.4 or Sec. 5.4.

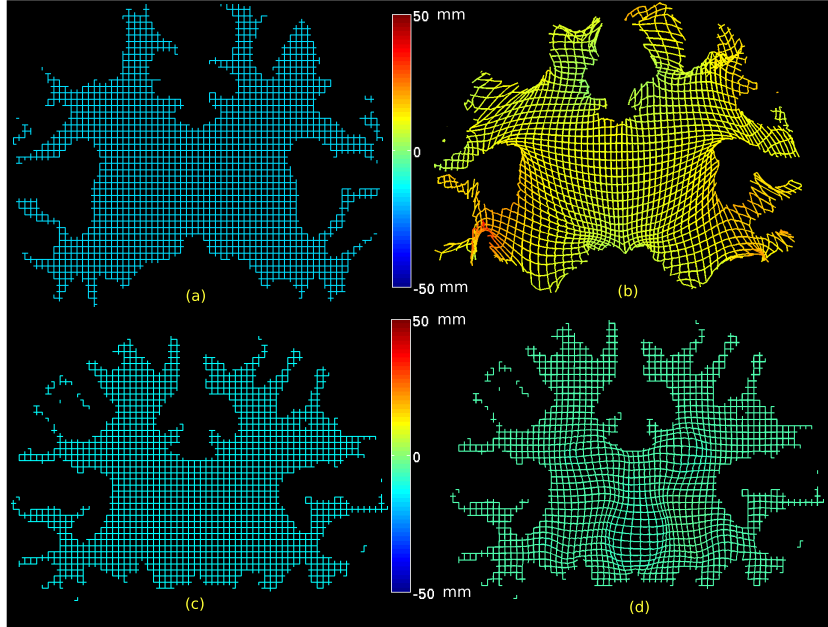


Figure 5.6: Illustration of the effects of the two stages of volumetric matching is shown by applying the deformations to a regular mesh representing one slice. Since the deformation is in 3D, the third in-paper value is represented by color. (a) Regular mesh representing one slice in the subject; (b) the regular mesh warped by the harmonic mapping which matches the subject cortical surface to the template cortical surface. Note that deformation is largest near the surface since the harmonic map is constrained only by the cortical surface; (c) Regular mesh representing one slice in the harmonically warped subject; (d) the intensity-based refinement now refines the deformation of the template to improve the match between subcortical structures. In this case the deformation is constrained to zero at the boundary and are confined to the interior of the volume.

2. The inverses of the map $u_M^{-1} : N \rightarrow M$ is computed. This is done by interpolating the correspondence $u_M^{-1} : u_M(x) \mapsto x$ from points to the regular voxel grid of N using Matlab's *griddata3* function with linear interpolation. This function implements the method based on Delauney triangulation as described in [BDH96] although it can also be computed using the method described in [Chr99]. $u_N^{-1} : M \rightarrow N$ is computed similarly.

3. Set $d_w^M = 0$ and $d_w^N = 0$.

4. Compute the maps $w^N(y) = y + d_w^N(y)$, $y \in N$, $\tilde{w}^M = (u^N \circ w^N \circ u^M)^{-1}$ and $\tilde{d}_w^M(x) = \tilde{w}^M(x) - x$.
5. Compute the difference term $f_N(x) - f_M(u^{N-1}(x))$.
6. Compute an updated estimate of the deformation field \hat{d}_w^M from (5.10) using a preconditioned conjugate gradient method.
7. Repeat steps 4-6 with M and N interchanged.
8. Test inverse consistency error C_{ICC} for convergence, otherwise go to Step 4.

This final refinement completes the surface-constrained registration procedure. While there are several steps required to complete the registration, each step can be reduced to either a surface or a volume mapping cast as an energy minimization problem, possibly with constraints, and can be effectively computed using a preconditioned conjugate gradient method. The different effects of the harmonic mapping, producing large scale deformations, and the linear elastic intensity-driven refinement, producing small scale deformations, are illustrated in Fig. 5.6

5.6 Results and Validation

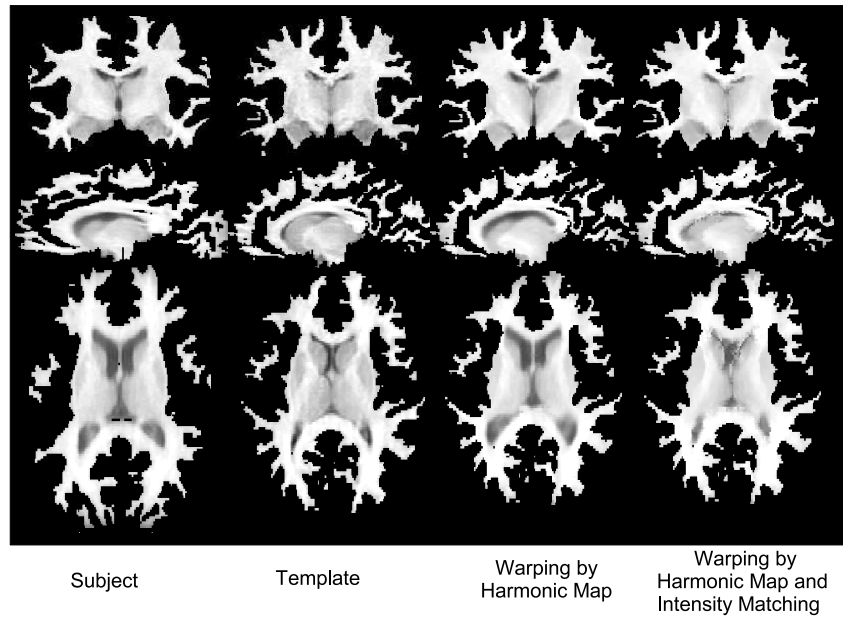
In order to illustrate the application of our surface constrained registration procedure to T1-weighted MR brain images and validate its performance, we obtained labeled brain data from the Internet Brain Segmentation Repository (IBSR) dataset at the Center for Morphometric Analysis at Massachusetts General Hospital. This consists of T1-weighted MR images with $1.5mm$ slice thickness as well as expert segmentations of 43 individual structures. The cortical masks were obtained and their topology corrected using the BrainSuite software as described in Chapter 3. The cortical surfaces were then interactively labelled with 23 sulcal curves on each hemisphere using a standard labeling protocol [THdZ⁺02]. Our registration algorithm was applied by performing

surface matching, harmonic mapping and volumetric intensity registration as described above. Shown in Fig. 5.7 and Fig. 5.9 are three orthogonal views of a subject before and after alignment to the template image. Note that before alignment the surfaces of the subject and template are clearly different, while after the harmonic mapping the deformed subject surface almost exactly matches the morphology of that of the template. However, since at this point we do not take the image intensities into account, the interior structures do not align well. Following the final intensity-based alignment procedure the subcortical structures of the warped subject show improved agreement with those in the template. Also shown in Fig. 5.9 and Fig. 5.7 are the labels provided by the IBSR data set before and after mapping.

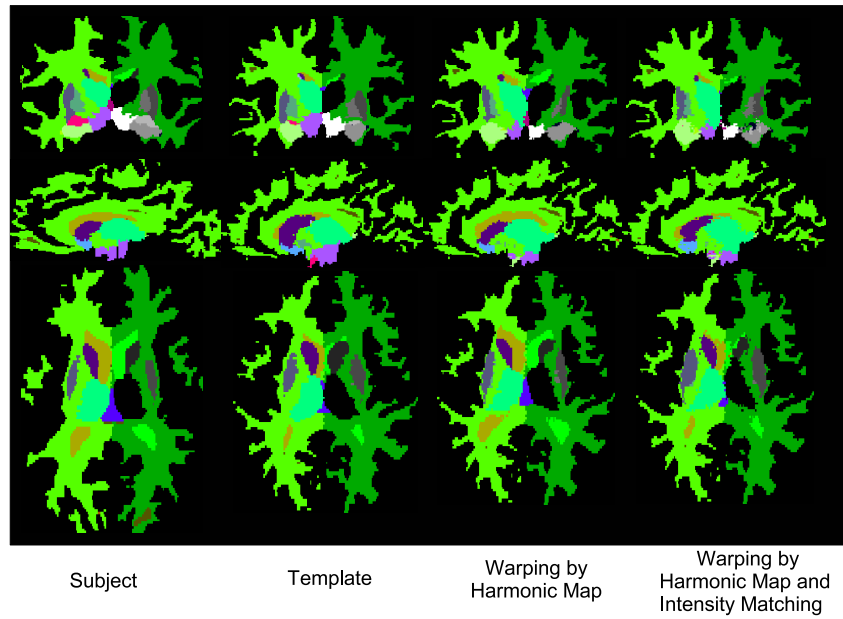
Our method for evaluating the quality of our registration results is based on the following two desirable features:

1. Alignment of the cortical surface and sulcal landmarks. We expect the sulcal landmarks to be accurately aligned after registration and for the two surfaces to coincide.
2. Alignment of subcortical structures. We also expect the boundaries of subcortical structures (thalamus, lateral ventricles, corpus callosum) to be well aligned after registration.

To evaluate performance with respect to 1 and 2 we used a set of 6 MR volumes on which we labeled 23 sulci in each hemisphere. For comparison we use a 5th order polynomial intensity-driven warp computed using the AIR software [WGH⁺98, WGW⁺98]. We also compare performance with the HAMMER [LSD04, SD02] algorithm. HAMMER is an automated method for volume registration which is able to achieve improved alignment of geometric features by basing the alignment on an attribute vector that includes a set of geometric moment invariants rather than simply the voxel intensities.



(a)



(b)

Figure 5.7: Examples of direct mapping approach. (a) Original subject volume; (b) original template; (c) registration of subject to template using surface constrained harmonic mapping, note that the surface matches that of the template; (d) intensity-based refinement of the harmonic map of subject to template to complete registration procedure

We note that since our approach uses explicitly labelled sulci we can expect better performance than either AIR or HAMMER in terms of the alignment of these features. However, AIR and HAMMER provide a basis for comparison from some of the most widely used and best performing algorithms for volumetric registration.

We measured the mean squared distance between pairs of homologous landmarks corresponding to uniform samples along each of the 23 labeled sulci. We repeated this procedure for each of the 30 possible pairwise registrations of two from six brains and computed the average mean squared distance over all registrations. We found that the mean squared misalignment between sulcal landmarks was $11.5mm$ for HAMMER, $11mm$ for AIR and $2.4mm$ for our cortically constrained method. The significantly lower error for our approach is unsurprising since matching of sulci is imposed as a constraint. The reason that the error is not zero is that the constraint is imposed using a penalty function rather than strictly using Lagrange multipliers.

To evaluate performance in terms of subcortical structures we used the manually labeled regions in the IBSR data set. To evaluate accuracy, we computed the Dice coefficients between the template and warped subject for each subcortical structure, where the structure names and boundaries were taken from the IBSR database. The Dice coefficient measures overlap between two sets representing regions S_1 and S_2 , and is defined as $\frac{2|S_1 \cap S_2|}{|S_1| + |S_2|}$ where $|\cdot|$ denotes size of the set [ZDMP94]. Values range from zero for disjoint sets to unity for identical sets. A comparison of the Dice coefficients for some major subcortical organs is shown in Fig. 5.7, where we show Dice coefficients for our method before and after application of the intensity-based alignment step. This comparison shows similar results for all three methods, with each producing superior results in some subcortical structures. For example, HAMMER produced superior results in thalamus, while our proposed method produced superior results in hippocampus. Thus the geometric invariants in HAMMER seem to improve performance relative to our

intensity based alignment of deeper subcortical structures, while our use of a cortical constraint leads to superior performance with respect to sulcal alignment and structures that are more superficial with respect to the cerebral cortex, such as the hippocampus. This is a preliminary validation and larger scale validation is needed on a larger population with a larger range of brain structures.

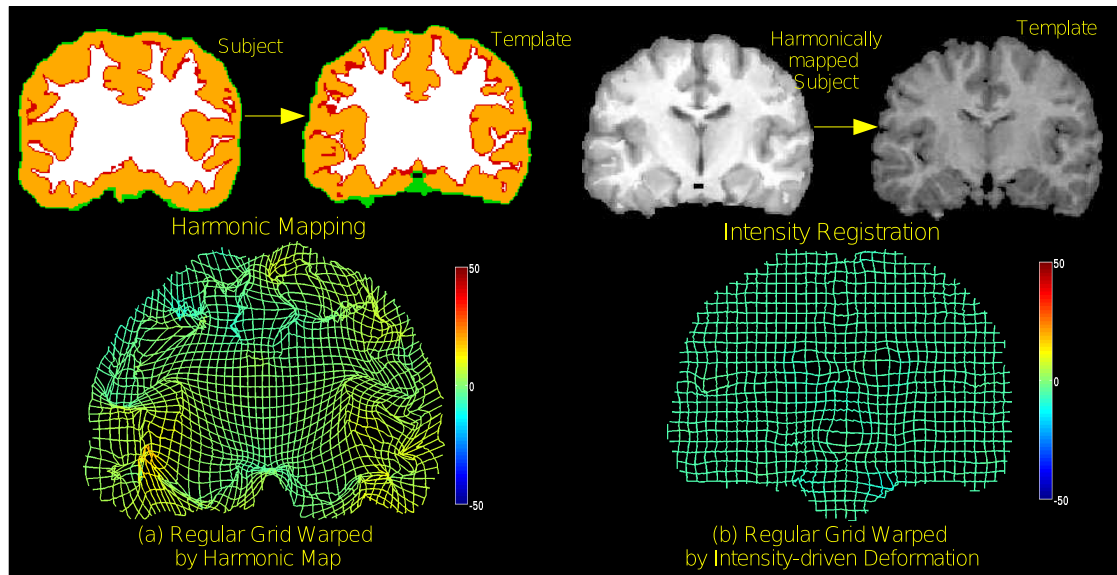


Figure 5.8: Volumetric registration using direct mapping approach: (a) Illustration of the extrapolation of the surface mapping to the 3D volume by harmonic mapping. The pairs of surfaces are shown in red and green. The deformation field is represented by placing a regular grid in the central coronal slice of the brain and deforming it according to the harmonic map. The projection of this deformation onto a 2D plane is shown with the in-plane value encoded according to the adjacent color bar. (b) The result of harmonic mapping and linear elastic refinement of the subject brain to the template brain. Note that the inner and outer cortical surfaces, by constraint, are exactly matched. The linear elastic refinement produces an approximate match between subcortical structures. The deformation field here shows the result of cortically constrained intensity-driven refinement. Note that the deformations are zero at the boundary and nonzero in the vicinity of the ventricles, thalamus and other subcortical structures.

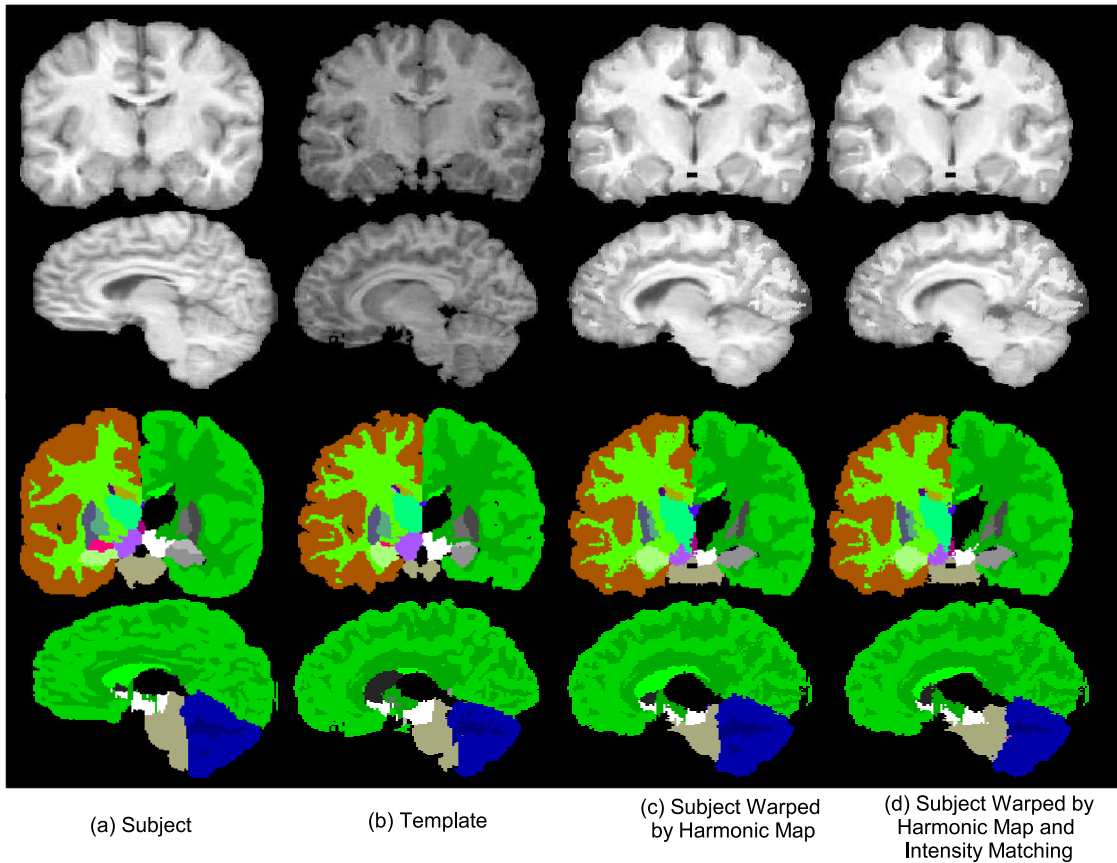


Figure 5.9: Examples of surface constrained volumetric registration. (a) Original subject volume; (b) template; (c) registration of subject to template using surface constrained harmonic mapping, note that the cortical surface matches that of the template; (d) intensity-based refinement of the harmonic map of subject to template

5.7 Conclusion

We have presented a framework for coregistration of brain volume data using harmonic maps. Through the use of an intermediate spherical map, we are able to constrain the surfaces of the two brain volumes to align while enforcing point matching only at a set of hand labeled sulcal curves. Using harmonic maps we are able to compute large scale deformations between brain volumes.

We have also described, as an initialization procedure, a new method for cortical surface parameterization and sulcal alignment in which the two problems are solved in a single step using a finite element method. This method has the properties that it is inverse consistent between the two brains and can be computed directly from a tessellated representation of the surface, rather than requiring resampling using a regular grid with respect to the induced parameterization.

The examples shown here demonstrate the cortical matching properties and the ability to also align subcortical structures. One of the limitations of this evaluation was that cortical grey matter was not included in the registration since the cortical surfaces were generated by BrainSuite [SL00], which selects the inner grey/white boundary as the cortical surface. However, this is a limitation of the preprocessing step rather than the method itself, and the process can be applied to the full cerebral volume provided that a genus-zero brain volume and sulcal labels are supplied. A second limitation is that the cerebellum and brainstem are not included in the analysis since the volume of interest that is mapped is restricted to the cerebrum, bounded by the outer cortical sheet. We can address this issue in practice by modifying the final intensity-based matching step by first adding the brainstem and cerebellum back to the cerebrum. This would also require extrapolation of the deformation field from the harmonic map outwards to these structures as an initialization of the intensity based warp. Alternatively, the cerebellum could also be explicitly modelled using a surface based approach (see, e.g. Hurdal et al. [HSB⁺00]), and its surface and enclosed volume could be treated in a similar fashion to the cerebrum.

Table 5.1: Comparison of Dice coefficients

Subcortical Structure	AIR	Harmonic	HAMMER	Harmonic with intensity
Left Thalamus	0.6588	0.5294	0.7303	0.5856
Left Caudate	0.4426	0.4336	0.5688	0.5716
Left Putamen	0.4079	0.3497	0.4905	0.5092
Left Hippocampus	0.4676	0.3069	0.3916	0.3930
Right Thalamus	0.6326	0.5018	0.7495	0.6230
Right Caudate	0.3671	0.3572	0.5098	0.5116
Right Putamen	0.3096	0.2358	0.4111	0.4679
Right Hippocampus	0.5391	0.3455	0.1989	0.4342
Avg. Dice coeff. for all structures	0.3021	0.3821	0.3621	0.4019
Std. Dev. of Dice coeff.	0.1937	0.2547	0.2390	0.2671

Acknowledgment

The authors would like to thank the Center for Morphometric Analysis at Massachusetts General Hospital for providing the MR brain data sets and their manual segmentations. The MR and segmentation data sets are available at <http://www.cma.mgh.harvard.edu/ibsr/>.

Chapter 6

Conclusions and Future Work

We have presented a set of geometric methods constituting a framework for registration and analysis of brain images. Various tools and their interconnections are depicted in Fig. 6. The sulcal tracing tool can be used to delineate a set of sulci on the cortex. These sulcal sets are then used as input to surface-based registrations presented in Chapter 3. We also presented a method to optimize the set of sulcal landmarks in order to minimize the manual effort (Sec. 3.3). We presented two surface registration techniques in that chapter: (1) covariant thin-plate splines (Sec. 3.1) (2) FEM-based Surface registration (Sec. 3.2). We concluded that the second approach leads to faster computation and accurate registration. This surface registration can be used for integrating surface-based functional or anatomical data from individual subjects to a common atlas. Intersubject analysis of such data can be carried out in the geometry of the atlas surface. We presented parameterization-based numerical methods for isotropic and anisotropic smoothing filters in the surface geometry (Chapter 4). Smoothing can be performed on the atlas surface, or on the original subject surfaces. When the data is a point-set, we presented a method for quantifying its mean and variance with respect to surface geometry. Again, this analysis can be carried out for a single subject or for a coregistered dataset to the atlas surface. The surface registration method was extended to volumetric registration using harmonic maps (Chapter 5). We presented two approaches: (1) Indirect approach using intermediate representation (Sec. 5.3) and (2) Direct mapping approach (Sec. 5.4). While the direct mapping approach is faster, it does not guarantee diffeomorphic mappings. On the other hand the indirect approach is significantly slower, but results

in diffeomorphic mappings. As a result of this method, we get a full 3D volumetric registration of the brain in which cortical surfaces as well as the subcortical structures are aligned.

There are a number of directions in which this work can be extended. In the following sections, we present a few possible extensions as well as applications.

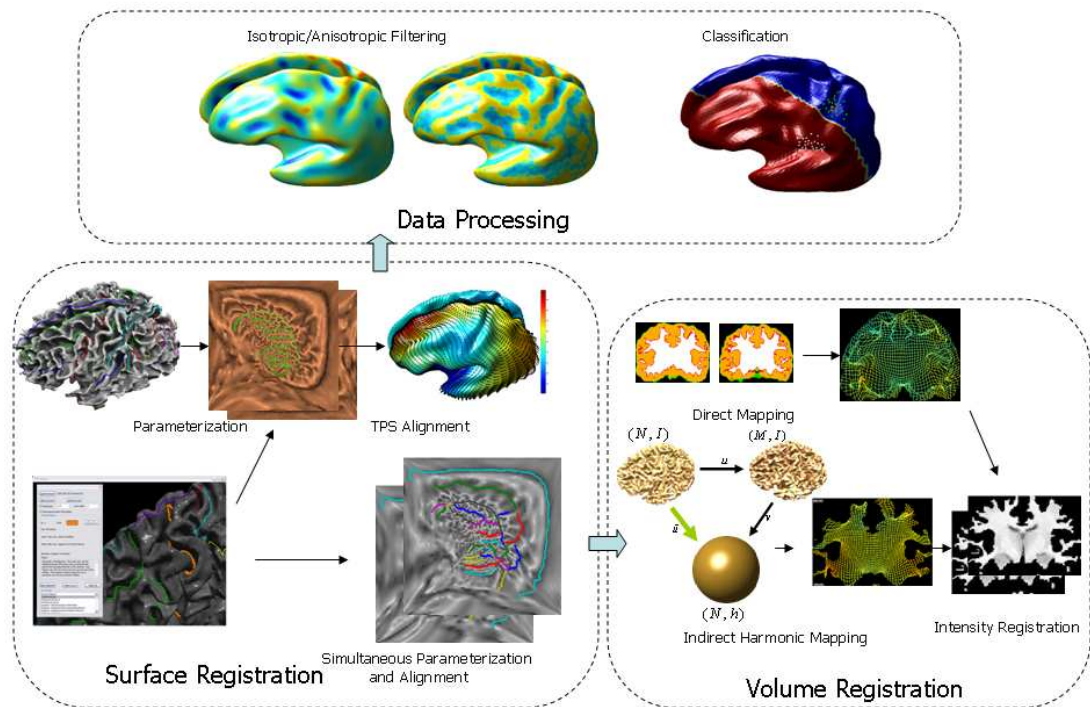


Figure 6.1: Geometric framework for registration and analysis

6.1 Geometric Features and Manual Landmarks based Surface Registration

Our method for cortical surface registration is a manual landmark based method. Alternatively, there are a number of automatic surface registration methods which perform such an alignment based on geometric features such as curvature, shape indices, etc. [TP05, FSTD98, HSB⁺00]. The advantage of automatic methods is that they do not involve manual input and therefore they are ideally suited for large-scale studies. However, accurate alignment of the brain anatomy involves higher level knowledge which is difficult to incorporate in such methods. These methods show consistent misalignment of certain areas, such as Broch's area. Sulcal folds are sometimes misregistered when there are branches. Manual landmark based methods overcome these difficulties by using user input in the form of expert labeled sulci. Moreover, these methods are ideally suited for abnormal anatomy. Also they could be useful for performing more accurate registration in a region of interest by marking more landmarks in this region. The disadvantage of manual techniques is that a considerable amount of training is required. Also manual effort is needed in order to identify the sulcal curves. In order to address these issues, and take advantage of both types of methods, we would like to formulate a semi-automatic method where only some of the sulcal curves are labeled manually when automatic geometric feature based methods do not give a correct registration. This can lead to minimization of the manual tracing effort without sacrificing accuracy and control in the surface registration process.

6.2 Registration of DTI images

Diffusion Tensor Imaging (DTI) produces in vivo images weighted with characteristics of water molecule diffusion inside a tissue [LAS⁺02]. In each voxel, it produces a 3×3 diffusion tensor which indicates the principle directions of water diffusion. This imaging modality is particularly useful to infer the white-matter connectivity of the brain [BJW⁺03]. The tensor data produced by the DTI images is used to reconstruct fiber tracts in the white matter (tractography). Recently, more advanced models of the diffusion process have been proposed that aim to overcome the weaknesses of the diffusion tensor model. Amongst others, these include q-space imaging [HYN⁺08] and generalized diffusion tensor imaging [OM].

In order to perform intersubject comparison and analysis of DTI data, accurate alignment of white matter is important. Particularly, since the sulcal curves are closely related to the function of the brain, any such comparison needs accurate alignment of the sulci. The volumetric registration technique presented in Chapter 5 makes such an alignment possible. We plan to use our volumetric registration techniques for intersubject comparisons of DTI data and fiber tracks. We will perform intersubject alignment of brains using T1 weighted MR data. The deformation field obtained this way can then be applied to the diffusion tensors to reorient them appropriately [APBG01]. Their variance can be quantified across subjects using the Lie group structure of the diffusion tensors [LRD06]. This kind of analysis can help identify similarities and differences in white matter connectivity across a population of subjects.

Bibliography

- [AF99] J. Ashburner and K.J. Friston. Spatial normalization. In A.W. Toga, editor, *Brain Warping*, pages 27–44. Academic Press, 1999.
- [AG04] Brian B. Avants and James C. Gee. Shape averaging with diffeomorphic flows for atlas creation. In *ISBI*, 2004.
- [AHTK99] S. Angenent, S. Haker, A. Tannenbaum, and R. Kikinis. Laplace-Beltrami operator and brain surface flattening. *IEEE Transactions on Medical Imaging*, 18:700–711, 1999.
- [APBG01] D. C. Alexander, C. Pierpaoli, P. J. Basser, and J. C. Gee. Spatial transformations of diffusion tensor magnetic resonance images. *IEEE Trans. Med. Imaging*, 20(11):1131–1139, Nov 2001.
- [BDH96] C. B. Barber, D. P. Dobkin, and H.T. Huhdanpaa. The quickhull algorithm for convex hulls. *ACM Transactions on Mathematical Software*, 1996.
- [BGKM98] M. Bakircioglu, U. Grenander, N. Khaneja, and M. I. Miller. Curve matching on brain surfaces using frenet distances. *Human Brain Mapping*, 6:329–333, 1998.
- [BJW⁺03] T. E. Behrens, H. Johansen-Berg, M. W. Woolrich, S. M. Smith, C. A. Wheeler-Kingshott, P. A. Boulby, G. J. Barker, E. L. Sillery, K. Sheehan, O. Ciccarelli, A. J. Thompson, J. M. Brady, and P. M. Matthews. Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging. *Nat Neurosci*, 6(7):750–757, 2003.
- [Boo89] F L Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11:567–585, June 1989.
- [BX03] Chandrajit L. Bajaj and Guoliang Xu. Anisotropic diffusion of surfaces and functions on surfaces. *ACM Trans. Graph.*, 22(1):4–32, 2003.

- [CDR04] U. Clarenz, U. Diewald, and M. Rumpf. Processing textured surfaces via anisotropic geometric diffusion. *IEEE Trans. on Image Processing*, 13(2):248–261, Feb. 2004.
- [Cha01] Jean-Pierre Changeux. Drug use and abuse. In G. M. Edelman and J. Changeux, editors, *The Brain*. Transaction Publishers, 2001.
- [Chr99] Gary E. Christensen. Consistent linear-elastic transformations for image matching. *Lecture Notes in Computer Science*, 1613:224–237, 1999.
- [Chu05] M. K. Chung. Introducing heat and geodesic kernel smoothing on cortical manifolds. *11th Annual Meeting of the OAHM, Toronto*, 2005.
- [CJM97] Gary E. Christensen, Sarang C. Joshi, and Michael I. Miller. Volumetric transformation of brain anatomy. *IEEE TMI*, 16(6), December 1997.
- [CMR⁺03] A. Cachia, J.-F. Mangin, D. Riviere, F. Kherif, N. Boddaert, A. Andrade, D. Papadopoulos-Orfanos, J.-B. Poline, I. Bloch, M. Zilbovicius, P. Sonigo, F. Brunelle, , and J. Regis. A primal sketch of the cortex mean curvature: a morphogenesis based approach to study the variability of the folding patterns. *IEEE Trans. Med. Imaging*, 22(6):754–765, Jun 2003.
- [CRD⁺05] M. K. Chung, S. M. Robbins, K. M. Dalton, R. J. Davidson, A. L. Alexander, and A.C. Evans. Cortical thickness analysis in autism with heat kernel smoothing. *NeuroImage*, 25(4):1256–1265, May 2005.
- [CRM94] Gary E. Christensen, Richard Rabbitt, and Michael I. Miller. 3D brain mapping using a deformable neuroanatomy. *Physics in Medicine and Biology*, 39:609–618, March 1994.
- [CRM⁺95] G. E. Christensen, R. D. Rabbitt, M. I. Miller, S. C. Joshi, U. Grenander, T. A. Coogan, and D. C. V. Essen. Topological properties of smooth anatomic maps. In *In 14 Conference on Information Processing in Medical Imaging, France*, pages 101–112. Kluwer Academic Publishers, 1995.
- [CRM96] G. E. Christensen, R. D. Rabbit, and M. I. Miller. Deformable templates using large deformation kinematics. *IEEE Transactions on Image Processing*, 5(10):1435–1447, 1996.
- [CWT⁺01] M. K. Chung, K.J. Worsley, J. Taylor, J.O. Ramsay, S. Robbins, and A.C. Evans. Diffusion smoothing on the cortical surface. *NeuroImage*, 13(6S1):95, Jun 2001.
- [CY01] Vincent Camion and Laurent Younes. Geodesic interpolating splines. *Lecture Notes in Computer Science*, pages 513–527, 2001.

- [CYV⁺00] G. E. Christensen, P. Yin, M. W. Vannier, K. S. C. Chao, J. L. Dempsey, and J. F. Williamson. Large-deformation image registration using fluid landmarks. In *Image Analysis and Interpretation, 2000. Proceedings. 4th IEEE Southwest Symposium*, pages 269–273, 2000.
- [Dav89] E B Davies. *Heat kernels and spectral theory*. Cambridge University Press, 1989.
- [DLF99] J. H. Downs, J. L. Lancaster, and P. T. Fox. Surface based spatial normalization using convex hulls. In *Brain Warping*, San Diego, CA, 1999. Academic.
- [DMA02] M. Desbrun, M. Meyer, and P. Alliez. Intrinsic parameterizations of surface meshes. In *Proceedings of Eurographics, 2002*.
- [Do 76] M. Do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, 1976.
- [DP94] C. Davatzikos and J. Prince. Brain image registration based on curve mapping. In *IEEE Workshop Biomedical Image Anal.*, pages 245–254, 1994.
- [DPB96] C. Davatzikos, J. Prince, and R. Bryan. Image registration based on boundary mapping. *IEEE Transactions on Medical Imaging*, 1996.
- [EÅ05] A. P. Eriksson and K. Åström. On the bijectivity of thin plate transforms. In *Swedish Symposium on Image Analysis*, pages 53–56, 2005.
- [EDD⁺95] M. Eck, T. DeRose, T. Duchamp, H. Hoppe, M. Lounsbery, and W. Stuetzel. Multiresolution analysis of arbitrary meshes. In *Computer Graphics (Proceedings of SIGGRAPH 95)*, August 1995.
- [ES64] J. Eells and J. H. Sampson. Harmonic mappings of Riemannian manifolds. *Ann. J. Math.*, pages 109–160, 1964.
- [FR02] Ali Fardoun and Rachid Regbaoui. Heat flow for p-harmonic maps between compact Riemannian manifolds. *Indiana Univ. Math. J.*, 51:1305–1320, 2002.
- [FSTD98] B. Fischl, M. I. Sereno, R. B. H. Tootell, and A. M. Dale. High-resolution inter-subject averaging and a coordinate system for the cortical surface. *Human Brain Mapping*, 8:272–284, 1998.
- [GAM⁺95] J. S. George, C. J. Aine, J. C. Mosher, D. M. Schmidt, D. M. Ranken, H. A. Schlitt, C. C. Wood, J. D. Lewine, J. A. Sanders, and J W. Belliveau. Mapping function in the human brain with magnetoencephalography, anatomical magnetic resonance imaging, and functional magnetic resonance imaging. *J Clin Neurophysiol*, 12(5):406–431, Sep. 1995.

- [GFK⁺95] Y. Ge, J. M. Fitzpatrick, R. M. Kessler, M. Jeske-Janicka, and R. A. Margolin. Intersubject brain image registration using both cortical and sub-cortical landmarks. In *Proc. SPIE Vol. 2434, p. 81-95, Medical Imaging 1995: Image Processing, Murray H. Loew; Ed.*, pages 81–95, May 1995.
- [GGL⁺04] Nitin Gogtay, Jay N. Giedd, Leslie Lusk, Kiralee M. Hayashi, Deanna Greenstein, A. Catherine Vaituzis, III Tom F. Nugent, David H. Herman, Liv S. Clasen, Arthur W. Toga, Judith L. Rapoport, and Paul M. Thompson. Dynamic mapping of human cortical development during childhood through early adulthood. *PNAS*, 101:8174–8179, 2004.
- [GI94] Russell Greiner and Ramana Isukapalli. Learning to select useful landmarks. In *National Conference on Artificial Intelligence*, pages 1251–1256, 1994.
- [GJF⁺06] G. Gerig, S. Joshi, T. Fletcher, K. Gorcowski, S. Xu, S. M. Pizer, and M. Styner. Statistics of population of images and its embedded objects: Driving applications in neuroimaging. In *ISBI*, pages 1120–1123, April 2006.
- [GM98] Ulf Grenander and Michael I. Miller. Computational anatomy: an emerging discipline. *Q. Appl. Math.*, LVI(4):617–694, 1998.
- [GVM04] Joan Glaunés, Marc Vaillant, and Michael I. Miller. Landmark matching via large deformation diffeomorphisms on the sphere. *J. Math. Imaging Vis.*, 20(1-2):179–200, 2004.
- [Ham75] R. Hamilton. Harmonic maps of manifolds with boundary. In *Lecture Notes in Mathematics*, 471. Springer, 1975.
- [HBHH01] Derek L. G. Hill, Philipp G. Batchelor, Mark Holden, and David J. Hawkes. Medical image registration. *Phys. Med. Biol.*, 46(4):R1–R45, March 2001.
- [HCF02] G. Hermosillo, C. Chéfd’Hotel, and Olivier Faugeras. Variational methods for multimodal image matching. *International Journal of Computer Vision*, 50(3):329–343, December 2002.
- [HHCS⁺02] T. Hartkens, D.L.G. Hill, Andy D. Castellano-Smith, D.J. Hawkes, C.R. Maurer, A.J. Martin, W.A. Hall, and C.L. Truwit H. Liu. Using points and surfaces to improve voxel-based non-rigid registration. In *MICCAI*, pages 565–572, 2002.
- [HP04] K. Hildebrandt and K. Polthier. Anisotropic filtering of non-linear surface features. *Computer Graphics Forum*, 23(3):391–400, Sept. 2004.

- [HSB⁺00] M. K. Hurdal, K. Stephenson, P. L. Bowers, D. W. L. Sumners, and D. A. Rottenberg. Coordinate system for conformal cerebellar flat maps. *NeuroImage*, 11:S467, 2000.
- [Hsu02] E. P. Hsu. *Stochastic Analysis on Manifolds*. American Mathematical Society, Providence, RI, 2002.
- [HYN⁺08] Keigo Hikishima, Kazuo Yagi, Tomokazu Numano, Kazuhiro Homma, Naotaka Nitta, Tetsu Nakatani, and Koji Hyodo. Volumetric q-space imaging by 3d diffusion-weighted mri. *Magnetic Resonance Imaging*, 2008.
- [JC02] H. J. Johnson and G. E. Christensen. Consistent landmark and intensity-based image registration. *IEEE Transactions on Medical Imaging*, 21(5):450–461, 2002.
- [JLTS04] Anand A. Joshi, Richard M. Leahy, Paul M. Thompson, and David W. Shattuck. Cortical surface parameterization by p -harmonic energy minimization. In *ISBI*, pages 428–431, 2004.
- [JM00] S. C. Joshi and M. I. Miller. Landmark matching via large deformation diffeomorphisms. *IEEE Transactions on Image Processing*, 9(8), August 2000.
- [Jos02] J. Jost. *Riemannian geometry and geometric analysis*. Springer Verlag, 2002.
- [JSTL05] Anand A. Joshi, David W. Shattuck, Paul M. Thompson, and Richard M. Leahy. A framework for registration, statistical characterization and classification of cortically constrained functional imaging data. In *Lecture Notes in Computer Science*, volume 3565, pages 186–196, July 2005.
- [JSTL07a] A. A. Joshi, D. W. Shattuck, P. M. Thompson, and R. M. Leahy. A finite element method for elastic parameterization and alignment of cortical surfaces using sulcal constraints. In *Proc. of ISBI*, 2007.
- [JSTL07b] A. A. Joshi, D. W. Shattuck, P. M. Thompson, and R. M. Leahy. Simultaneous surface and volumetric registration using harmonic maps. In *Proceedings of SPIE*, Jan 2007.
- [JSTL07c] Anand A. Joshi, David W. Shattuck, Paul M. Thompson, and Richard M. Leahy. Surface-constrained volumetric brain registration using harmonic mappings. *IEEE Trans. Med. Imaging*, 26(12):1657–1669, 2007.
- [JSTLed] A. Joshi, D. Shattuck, P. Thompson, and R. Leahy. Simultaneous surface and volumetric brain registration using harmonic mappings. *IEEE TMI*, accepted.

- [KL99] N. Krahnstover and C. Lorenz. Development of point-based shape representation of arbitrary three-dimensional medical objects suitable for statistical shape modeling. In *Proc. SPIE-Medical Imaging 1999:Image Processing*, volume 3661, pages 620–631, 1999.
- [Kre99] I. Kreyzig. *Differential Geometry*. Dover, 1999.
- [KSK98a] T. Kanai, H. Suzuki, and F. Kimura. Three-dimensional geometric metamorphosis based on harmonic maps. *The Visual Computer*, 14(4):166–176, 1998.
- [KSK98b] T. Kannai, H. Suzuki, and F. Kimura. Three-dimensional geometric metamorphosis based on harmonic maps. In *The Visual Computer*, volume 14, pages 166–176, 1998.
- [LAS⁺02] N. F. Lori, E. Akbudak, J. S. Shimony, T. S. Cull, A. Z. Snyder, R. K. Guillery, and T. E. Conturo. Diffusion tensor fibre tracking of human brain connectivity: aquisition methods, reliability analysis and biological result. *NMR Biomed*, 15(7-8):494–515, 2002.
- [LPDS04] L. Lopez-Perez, R. Deriche, and N. Sochen. The beltrami flow over triangulated manifolds. In *ECCV Workshops CVAMIA and MMBIA*, pages 135–144, 2004.
- [LRD06] Christophe Lenglet, Mika el Rousson, and Rachid Deriche. Dti segmentation by statistical surface evolution. *IEEE Trans. Med. Imaging*, 25(6), 2006.
- [LSD04] Tianming Liu, Dinggang Shen, and Christos Davatzikos. Deformable registration of cortical structures via hybrid volumetric and surface warping. *NeuroImage*, 22(4):1790–1801, 2004.
- [LTPH04] Alexia Leow, Paul M. Thompson, Hillary Protas, and Sung-Cheng Huang. Brain warping with implicit representations. In *ISBI*, pages 603–606. IEEE, 2004.
- [MA98] C. B. Madsen and C. S. Andersen. *Journal of Robotics and Autonomous Systems*, 23:277–292, 1998.
- [MKB79] KV Mardia, JT Kent, and JM Bibby. *Multivariate Analysis*. Academic Press, 1979.
- [MSO04] Facundo Mémoli, Guillermo Sapiro, and Stanley Osher. Solving variational problems and partial differential equations mapping into general target manifolds. *J. Comput. Phys.*, 195(1):263–292, 2004.

- [MST04] Facundo Mémoli, Guillermo Sapiro, and Paul Thompson. Implicit brain imaging. *NeuroImage*, 23(1):S179–S188, 2004.
- [NB97] G. G. Nahas and T. F. Burks, editors. *Drug Abuse in the Decade of the Brain*. IOS Press, 1997.
- [Nis01] Seiki Nishikawa. *Variational Problems in Geometry*, volume 205 of *Translations of Mathematical Monographs*. AMS, 2001.
- [OKA90] M Ono, S Kubik, and CD Abernathey. *Atlas of the Cerebral Sulci*. Stuttgart, New York:Thieme, 1990.
- [Ols00] C. Olson. Landmark selection for terrain matching, 2000.
- [OM] E. Ozarslan and T. H. Mareci. Generalized diffusion tensor imaging and analytical relationships between diffusion tensor imaging and high angular resolution diffusion imaging. *Magnetic Resonance in Medicine*, 50(5):955–965.
- [PCS⁺89] C. A. Pelizzari, G. T. Y. Chen, D. R. Spelbring, R. R. Weichselbaum, and C. T. Chen. Accurate three-dimensional registration of CT, PET and/or MR images of the brain. *J. Comput. Assist. Tomogr.*, 13(1):22–26, 1989.
- [PM90] Pietro Perona and Jitendra Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-12(7):629–639, July 1990.
- [PNBL05] D. Pantazis, Thomas E Nichols, Sylvain Baillet, and R.M. Leahy. A comparison of random field theory and permutation methods for the statistical analysis of meg data. *Neuroimage*, 25(2):383–394, Apr 2005.
- [RHXP02] M. Rettmann, X. Han, C. Xu, and J. Prince. Automated sulcal segmentation using watersheds on the cortical surface. *NeuroImage*, 15(2):329–344, 2002.
- [RL03] Nicolas Ray and Bruno Levy. Hierarchical least squares conformal map. In *PG '03: Proceedings of the 11th Pacific Conference on Computer Graphics and Applications*, page 263, Washington, DC, USA, 2003. IEEE Computer Society.
- [Ros97] S Rosenberg. *The Laplacian on a Riemannian manifold*. Cambridge University Press, 1997.
- [SD02] Dinggang Shen and Christos Davatzikos. Hammer: Hierarchical attribute matching mechanism for elastic registration. *IEEE TRANSACTIONS ON MEDICAL IMAGING*, 21(11), 2002.

- [SL00] David W. Shattuck and Richard M. Leahy. BrainSuite: An automated cortical surface identification tool. In Scott L. Delp, Anthony M. DiGioia, and Branislav Jaramaz, editors, *MICCAI*, volume 1935 of *Lecture Notes in Computer Science*, pages 50–61. Springer, 2000.
- [SL02] D. W. Shattuck and R. M. Leahy. Brainsuite: An automated cortical surface identification tool. *Medical Image Analysis*, 8(2):129–142, 2002.
- [SMGT04] D. W. Shattuck, A. MacKenzie-Graham, and A. W. Toga. Duff: software tools for visualization and processing of neuroimaging data. *IEEE International Symposium on Biomedical Imaging: Macro to Nano, 2004*, 1:644–647, April 2004.
- [Smi85] G. D. Smith. *Numerical solution of partial differential equations: Finite difference methods*. Oxford Applied Mathematics and Computing Science Series. Oxford : Clarendon Press, 3 edition, 1985.
- [SPT⁺03] Elizabeth R. Sowell, Bradley S. Peterson, Paul M. Thompson, Suzanne E. Welcome, Amy L. Henkenius, and Arthur W. Toga. Mapping cortical change across the human life span. *Nature Neuroscience*, 6:309–315, 2003.
- [STR⁺02] E. R. Sowell, P. M. Thompson, D. Rex, D. Kornsand, K. D. Tessner, T. L. Jernigan, and A. W. Toga. Mapping sulcal pattern asymmetry and local cortical surface gray matter distribution in vivo: maturation in perisylvian cortices. *Cereb. Cortex*, 12(1):17–26, Jan 2002.
- [THdZ⁺02] Paul M. Thompson, Kiralee M. Hayashi, Greig. de Zubicaray, Andrew L. Janke, Stephen E. Rose, James Semple, David M. Doddrell, Tyrone D. Cannon, and Arthur W. Toga. Detecting dynamic and genetic effects on brain structure using high dimensional cortical pattern matching. In *Proceedings of ISBI, 2002*.
- [THdZ⁺03] Paul M. Thompson, Kiralee M. Hayashi, Greig de Zubicaray, Andrew L. Janke, Stephen E. Rose, James Semple, David Herman, Michael S. Hong, Stephanie S. Dittmer, David M. Doddrell, and Arthur W. Toga. Dynamics of gray matter loss in Alzheimer’s disease. *The Journal of Neuroscience*, 23(3):994–1005, 2003.
- [THS⁺04] P. M. Thompson, K. M. Hayashi, E. R. Sowell, N. Gogtay, J. N. Giedd, J. L. Rapoport, G. I. de Zubicaray, A. L. Janke, S. E. Rose, J. Semple, D. M. Doddrell, Y. L. Wang, TGM van Erp, T. D. Cannon, and A. W. Toga. Mapping cortical change in Alzheimer’s disease, brain development and schizophrenia. *NeuroImage*, 23(1):S2–S18, Sept. 2004.

- [TMM⁺97] Paul M. Thompson, David MacDonald, Michael S. Mega, Colin J. Holmes, Alan C. Evans, and Arthur W. Toga. Detection and mapping of abnormal brain structure with a probabilistic atlas of cortical surfaces. *Journal of Computer Assisted Tomography*, 21(4):567–581, Jul.-Aug. 1997.
- [TMT00] Paul M Thompson, Michael S Mega, and Arthur W Toga. Disease-specific probabilistic brain atlases. In *Proceedings of IEEE International Conference on Computer Vision and Pattern Recognition*, pages 227–234, 2000.
- [TMV⁺01] Paul M. Thompson, Michael S. Mega, Christine Vidal, Judith Rapoport, and Arthur W. Toga. Detecting disease-specific patterns of brain structure using cortical pattern matching and a population-based probabilistic brain atlas. In *Proc. 17th IPMI2001, Davis, CA, USA*, pages 488–501, 2001.
- [TP05] Duygu Tosun and Jerry L. Prince. Cortical surface alignment using geometry driven multispectral optical flow. In *Information Processing in Medical Imaging*, volume 3565 of *LNCS*, pages 480–492, 2005.
- [TRP05] Duygu Tosun, Maryam E. Rettmann, and Jerry L. Prince. Mapping techniques for aligning sulci across multiple brains. *Medical Image Analysis*, 8(3):295–309, 2005.
- [TSC00] Bei Tang, Guillermo Sapiro, and Vicent Caselles. Diffusion of general data on non-flat manifolds via harmonic maps theory: The direction diffusion case. *International Journal of Computer Vision*, 36(2):149–161, 2000.
- [TT88] J. Talairach and P. Tournoux. *Co-planar Stereotaxic Atlas of the Human Brain: 3-Dimensional Proportional System - an Approach to Cerebral Imaging*. Thieme Medical Publishers, New York, NY, 1988.
- [TT96a] P. M. Thompson and A. W. Toga. A surface-based technique for warping 3-dimensional brain. *IEEE Transactions on Medical Imaging*, 15(4):1–16, 1996.
- [TT96b] Paul M. Thompson and Arthur W. Toga. A surface based technique for warping 3-dimensional images of the brain. *IEEE Transactions on Medical Imaging*, 15(4):1–16, 1996.
- [TVG⁺01] Paul M. Thompson, Christine Vidal, Jay N. Giedd, Peter Gochman, Jonathan Blumenthal, Robert Nicolson, Arthur W. Toga, and Judith L. Rapoport. Mapping adolescent brain change reveals dynamic wave of accelerated gray matter loss in very early-onset schizophrenia. *PNAS*, 98(20):11650–11655, 2001.

- [TWMT00] P. M. Thompson, R. P. Wood, M. S. Mega, and A. W. Toga. Mathematical/computational challenges in creating deformable and probabilistic atlases of the human brain (invited paper). *Human Brain Mapping*, 9(2):81–92, Feb. 2000.
- [WCT05] Y. Wang, M. C. Chiang, and P. M. Thompson. Automated surface matching using mutual information applied to Riemann surface structures. In J. Duncan and G. Gerig, editors, *MICCAI 2005, LNCS 3750*, pages 666–674. Springer-Verlag Berlin Heidelberg, 2005.
- [WGH⁺98] R. P. Woods, S. T. Grafton, C. J. Holmes, S. R. Cherry, and J. C. Mazziotta. Automated image registration: I. General methods and intrasubject, intramodality validation. *Journal of Computer Assisted Tomography*, 22:139–152, 1998.
- [WGH⁺05] Y. Wang, X. Gu, K.M. Hayashi, T.F. Chan, P.M. Thompson, and S.T. Yau. Brain surface parameterization using Riemann surface structure. In *MICCAI 2005*, pages 657–665, 2005.
- [WGW⁺98] R. P. Woods, S. T. Grafton, J. D. G. Watson, N. L. Sicotte, and J. C. Mazziotta. Automated image registration: II. Intersubject validation of linear and nonlinear models. *Journal of Computer Assisted Tomography*, 22:153–165, 1998.
- [WGY04] Y. Wang, X. Gu, and S. T. Yau. Volumetric harmonic map. *Communications in Information and Systems*, 3(3):191–202, 2004.
- [WLCT05] Y. Wang, L. M. Lui, T. F. Chan, and P. M. Thompson. Optimization of brain conformal mapping with landmarks. In *MICCAI 2005: Proceedings, Part II*, pages 675–683, 2005.
- [WMN⁺96] K. Worsley, S. Marrett, P. Neelin, A. Vandal, K. Friston, and A. Evans. A unified statistical approach for determining significant signals in images of cerebral activation. *HBM*, 4:58–73, 1996.
- [Xin96] Yuanlong Xin. *Geometry of harmonic maps*. Birkhäuser, 1996.
- [ZDMP94] A. P. Zijdenbos, B. M. Dawant, R. A. Margolin, and A. Palmer. Morphometric analysis of white matter lesions in mr images. *IEEE Transactions on Medical Imaging*, 13:716–724, Dec. 1994.