

Medical Image Analysis via Fréchet Means of Diffeomorphisms

Bradley C. Davis

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

Chapel Hill
2008

Approved by:

Sarang C. Joshi, Advisor

Stephen M. Pizer, Co-principal Reader

Elizabeth Bullitt, Reader

J. S. Marron, Reader

Jack Snoeyink, Reader

© 2008
Bradley C. Davis
ALL RIGHTS RESERVED

Abstract

**Bradley C. Davis: Medical Image Analysis via Fréchet Means of
Diffeomorphisms.
(Under the direction of Sarang C. Joshi.)**

The construction of average models of anatomy, as well as regression analysis of anatomical structures, are key issues in medical research, e.g., in the study of brain development and disease progression. When the underlying anatomical process can be modeled by parameters in a Euclidean space, classical statistical techniques are applicable. However, recent work suggests that attempts to describe anatomical differences using *flat Euclidean spaces* undermine our ability to represent natural biological variability. In response, this dissertation contributes to the development of a particular *nonlinear* shape analysis methodology.

This dissertation uses a nonlinear deformable model to measure anatomical change and define geometry-based averaging and regression for anatomical structures represented within medical images. Geometric differences are modeled by coordinate transformations, i.e., deformations, of underlying image coordinates. In order to represent local geometric changes and accommodate large deformations, these transformations are taken to be the group of diffeomorphisms with an associated metric.

A mean anatomical image is defined using this deformation-based metric via the Fréchet mean—the minimizer of the sum of squared distances. Similarly, a new method called *manifold kernel regression* is presented for estimating systematic changes—as a function of a predictor variable, such as age—from data in nonlinear spaces. It is defined by recasting kernel regression in terms of a kernel-weighted Fréchet mean. This method is applied to determine systematic geometric changes in the brain from a random design dataset of medical images. Finally, diffeomorphic image mapping is extended to accommodate *extraneous structures*—objects that are present in one image and absent in another and thus change image topology—by deflating them prior to the estimation of geometric change. The method is applied to quantify the motion of the prostate in the presence of transient bowel gas.

In honor of Lenore, Raycliffe, Sarah Kathryn, and Robert.

Acknowledgments

It has been a privilege and honor for me to be associated with my friends, colleagues, and teachers at the University of North Carolina at Chapel Hill and beyond. Writing a dissertation is a long process; I found a lot of help, luck, and encouragement along the way.

In particular I thank my advisor Sarang Joshi for his integrity, support, confidence, friendship, and unrelenting, infectious enthusiasm. It has been a privilege to learn from him; the work in these pages builds upon and is heavily influenced by his ideas.

My committee played a vital role in my education and the development of this dissertation. I thank Stephen Pizer for his leadership in the multidisciplinary Medical Image Display and Analysis Group, from which I learned so much. He is a first rate teacher, friend, and mentor. Elizabeth Bullitt supported much of my work through her grants and offered an important clinical prospective to my research. J. Steve Marron's interest and expertise in object oriented data analysis contributed heavily to the development of the ideas in this dissertation. Jack Snoeyink offered a number of important suggestions for the improvement of this document.

To the faculty, staff, and students of the Medical Image Display and Analysis Group and the Departments of Computer Science, Mathematics, Statistics, and Radiation Oncology at UNC who have created an incredible educational environment: it has been a pleasure to learn from all of you through my courses, group meetings, team projects and one-on-one interactions. You are great friends, colleagues, and role models. A special thanks to P. Thomas Fletcher, Peter Lorenzen, Guido Gerig, Martin Styner, and Clement Vachet for their large contributions—both explicit and implicit—to this work.

I especially thank the staff members of the Department of Computer Science who were always there to offer a helping hand. Special thanks go to Delphine Bull, Kim Jones, Janet Jones, Michael Stone, Tammy Pike, Murray Anderegg, and Sandra Neely.

I thank Will Schroeder, Stephen Aylward, and the gang at Kitware for their patience and encouragement as I finished this dissertation.

A special thanks goes to Katie Belle, my steadfast companion during many long hours in front of the computer.

I thank my parents, Edward and Kathy, and my sister, Audrey, for their love and invaluable support throughout the entirety of my life and especially during the last few months.

Finally, I thank my wife, Alexandra, for her boundless patience, love, and unrelenting support (not to mention her excellent copy editing skills). You are my inspiration.

Table of Contents

List of Tables	xii
List of Figures	xiii
List of Abbreviations	xv
List of Symbols	xvi
1 Introduction	1
1.1 Motivation	1
1.1.1 Accommodating topological change in large-deformation diffeomorphic image mapping	4
1.1.2 Inferring mean anatomical configurations from medical imagery	4
1.1.3 Regression analysis for medical imagery	6
1.2 Thesis and contributions	9
1.3 Overview of chapters	10
2 Large-deformation diffeomorphic image matching	12
2.1 Introduction	12
2.2 The large-deformation diffeomorphic framework	13
2.3 Image matching and an image-to-image distance using large-deformation diffeomorphisms	16
2.4 Euler-Lagrange equations for the image matching problem	18
3 Accommodating topological change for practical large-deformation diffeomorphic image mapping: Application to the measurement of prostate motion	20

3.1	Introduction	20
3.1.1	The need for automatic segmentation in adaptive radiation therapy of the prostate	21
3.1.2	Estimating prostate motion via image mapping	24
3.2	Methodology	28
3.2.1	Image mapping and topology	28
3.2.2	Accommodating topological change	30
3.2.3	Registration pipeline for intra-patient registration	34
3.3	Results: Automatic segmentation of the prostate from in-treatment-room CT scans	37
3.3.1	Implementation: Beamlock	40
3.3.2	Centroid analysis	40
3.3.3	Volume overlap analysis	44
3.4	Conclusions	46
4	Variational optimization for large-deformation diffeomorphic image averaging	48
4.1	Introduction	48
4.2	Methods	50
4.2.1	Arithmetic mean in Hilbert spaces	51
4.2.2	Fréchet mean image	52
4.2.3	Computing Fréchet mean images with LDDMM and greedy solutions	55
4.2.4	Fréchet mean image algorithm properties	58
4.2.5	Implementation notes	60
4.3	Algorithm demonstration with synthetic data	61
4.3.1	Ovals	62
4.3.2	Bullseyes	64

4.4	Related work	66
4.5	Conclusion	68
5	Manifold kernel regression	70
5.1	Introduction	70
5.2	Methods	73
5.2.1	Review of univariate kernel regression	73
5.2.2	Kernel regression on Riemannian manifolds	75
5.2.3	Bandwidth selection	76
5.2.4	Kernel regression for populations of anatomical images	78
5.2.5	Diffeomorphic population growth model	79
5.3	Synthetic data experiment	82
5.4	Conclusion	86
6	The aging brain: Measuring change using regression	88
6.1	Introduction	88
6.2	Brain image database	88
6.3	Manifold kernel regression for MR brain images	89
6.4	Local, instantaneous volumetric change	101
6.5	Investigation of kernel widths	101
6.6	Computational details	106
6.7	Conclusion	110
7	Discussion	111
7.1	Summary of contributions and thesis	111
7.2	Benefits and limitations of diffeomorphisms for the study of shape	116
7.3	Hypothesis testing with manifold kernel regression	117
7.3.1	Hypothesis testing for independence of predictor and response . .	118

7.3.2	Hypothesis testing to determine if population trends coincide . . .	120
7.4	Future work	120
7.4.1	Properties of means and kernel regression estimators on manifolds	121
7.4.2	Parametric regression	121
7.4.3	Application to other shape representations and metrics	122
7.4.4	The inverse problem	123
7.4.5	Managing outliers and topological change	124
7.5	Other application areas	126
Appendix A: Mathematical Preliminaries		127
A-1	Introduction	127
A-2	Topological spaces	127
A-3	Metric spaces	127
A-4	Vector spaces	128
A-4.1	Hilbert space	129
A-4.2	Images as L^2 functions	129
A-5	Properties of mappings	130
A-6	Group theory	130
A-7	Differential geometry	131
A-7.1	Measuring distances on manifolds	133
Appendix B: Euler-Lagrange equation for LDDMM image mapping .		136
Appendix C: Derivation of the differential operator for “fluid” image registration		140
C-1	Introduction	140
C-2	The derivation of L	140
Appendix D: Numerical solution for velocity fields in fluid registration		143

D-1	Introduction	143
D-2	Operators	143
D-3	Finite difference approximations	146
D-4	Fourier domain solution for L	149
D-5	Fourier domain solution for L_B	152
	D-5.1 Note on interpretation of α and γ	152
	Bibliography	154

List of Tables

3.1	Segmentation results: prostate centroid difference summary statistics . .	42
3.2	Segmentation results: prostate centroid distance summary statistics . . .	43
3.3	Segmentation results: prostate Dice similarity summary statistics	46
3.4	Automatic and manual segmentation comparison	46

List of Figures

1.1	Fréchet mean and manifold kernel regression illustrations	6
2.1	Velocity integration	15
3.1	Prostate motion over the course of radiation therapy treatment	22
3.2	Patient setup error	23
3.3	In-treatment-room CT scanner	25
3.4	Quantifying tissue motion via deformable image registration	26
3.5	Automatic prostate segmentation via deformable image registration . . .	27
3.6	The geometric effect of bowel gas	27
3.7	Improved segmentation via gas deflation	28
3.8	Gas deflation schematic	33
3.9	Visualization of organ segmentations	37
3.10	Automatic prostate segmentation results	38
3.11	Segmentation accuracy results: prostate centroid differences	41
3.12	Segmentation accuracy results: prostate Dice similarity coefficients	44
4.1	Fréchet mean image schematic	54
4.2	Synthetic data experiment: Fréchet mean ellipse images	63
4.3	Synthetic data experiment: Fréchet means of bullseye images	65
5.1	Illustration of univariate kernel regression	72
5.2	Brain image database	73
5.3	Manifold kernel regression schematic	77
5.4	Population growth model schematic	80

5.5	Synthetic bullseye data set construction	83
5.6	Random design database of 2D bullseye images	84
5.7	Regression results for synthetic bullseye database	85
6.1	Healthy subject MRI brain image database	90
6.2	Subject age distribution and kernel weights	91
6.3	Regressed anatomical images: female cohort	92
6.4	Regressed anatomical images: male cohort	93
6.5	Regressed anatomical images: combined cohort	94
6.6	Average, aging brain: female cohort, axial view	95
6.7	Average, aging brain: female cohort, sagittal view	96
6.8	Average, aging brain: female cohort, coronal view	96
6.9	Average, aging brain: male cohort, axial view	97
6.10	Average, aging brain: male cohort, sagittal view	98
6.11	Average, aging brain: male cohort, coronal view	98
6.12	Average, aging brain: combined cohort, axial view	99
6.13	Average, aging brain: combined cohort, sagittal view	100
6.14	Average, aging brain: combined cohort, coronal view	100
6.15	Aging brain shape change: female cohort	102
6.16	Aging brain shape change: male cohort	103
6.17	Aging brain shape change: combined cohort	104
6.18	Bandwidth selection example	105
6.19	Regressed brains with different kernel widths	107
6.20	Fréchet mean image experimental timing and memory results	108
6.21	Fréchet mean image experimental threading results	109

List of Abbreviations

ART.....	Adaptive Radiation Therapy
CT.....	Computed Tomography
DSC.....	Dice Similarity Coefficient
EBRT.....	External Beam Radiation Therapy
FFT.....	Fast Fourier Transform
IQR.....	Interquartile Range
LDDMM.....	Large-Deformations Diffeomorphic Metric Mapping (also LDMM)
LDSC.....	Logit transform of Dice Similarity Coefficient
LSCV.....	Least-squares cross-validation
MKR.....	Manifold Kernel Regression
MKRE.....	Manifold Kernel Regression Estimator
MRA.....	Magnetic Resonance Angiography
MRI.....	Magnetic Resonance Image
PCA.....	Principal Component Analysis
RSS.....	Residual Sum of Squares
STD.....	Standard Deviation

List of Symbols

\mathbb{R}	the real numbers
Ω	subset of \mathbb{R}^3 where images and deformations are defined ($[0,1] \times [0,1] \times [0,1]$)
$L^2(\Omega)$	Hilbert space of square-integrable functions on Ω
$\ \cdot\ _{L^2}$	L_2 (root-integrated-squared) norm
\mathcal{I}	set of images (same as $L^2(\Omega)$)
I, J	images (elements of \mathcal{I})
I_ϕ	deformed image ($I \circ \phi^{-1}$)
I^F, I^M	fixed (will not deform) and moving (will deform) images
V	Hilbert space of velocity fields $\Omega \rightarrow \mathbb{R}^3$
$\ \cdot\ _V$	Sobolev norm associated with V
L, K	differential and related Green's function operators
v_t	time-varying flow of velocity fields
$\text{Diff}_V(\Omega)$	set of diffeomorphisms produced by flows v_t
ϕ	diffeomorphism (elements of $\text{Diff}_V(\Omega)$)
ϕ_t	diffeomorphism at time t (of flow v_t)
ϕ_{t_1, t_2}	diffeomorphism mapping coordinates at time t_1 to time t_2 ($\phi_{t_2} \circ \phi_{t_1}^{-1}$)
$\dot{\phi}_t$	time derivitive of diffeomorphism ϕ (i.e., $\frac{d}{dt}\phi_t$)
$D\phi$	the 3×3 Jacobian matrix of ϕ
u, v	velocity fields (elements of V)
$\nabla^2, \nabla, \nabla \cdot$	Laplacian, gradient, and divergence operators
α, β, γ	parameters of differential operator L
\dagger	the adjoint operator
Id_X	the identity element in the space X
σ	a weighting factor between terms in image matching functional

Chapter 1

Introduction

1.1 Motivation

During the last 40 years, astonishing breakthroughs in imaging technology have provided scientists with unprecedented access to anatomical structures through three-dimensional, *in vivo*, non-invasive imaging modalities. Magnetic resonance (MR) imaging, computed tomography, and ultrasound provide structural diagrams of the human body at sub-millimeter resolution. Diffusion-weighted and diffusion-tensor MR imaging modalities provide insight into local anatomical structure by measuring the local characteristics of water diffusion within tissue. Functional imaging modalities such as functional MR, contrast MR, and arterial spin labeling imaging enable the study of blood volume and flow within the brain. Tagged MR imaging techniques allow for precise recording of how tissue, such as the heart wall, moves over time.

The effective utilization of medical image data impacts modern health care by providing improved therapeutic methods, advances in disease detection and understanding, and quantitative assessment of therapeutic protocols. For example, high-resolution structural computed tomography (CT) and MR images allow physicians to precisely define and target cancerous tissue during radiation therapy treatments (Halperin et al., 2007). Diffusion weighted and diffusion tensor images are used to diagnose vascular

strokes and to study the white matter connectivity in the brain (Warach et al., 1995; Mori et al., 2005). Contrast enhanced MR images are used to monitor the effect—and thus judge the efficacy of—drugs and therapies for diseases ranging from cancer to degenerative arthritis (Beckmann et al., 2004).

Given the abundance of information that medical images provide and their expanding applicability in clinical and research medicine, there is a growing need for the development of mathematical and statistical methods for the analysis of medical images and the anatomical structures that are represented within them. The fields of medical image analysis and statistical shape analysis aim to address this need through the combination of analytic methods, computer processing, and computer-based robotics and visualization systems. For example, clinical neurosurgery systems combine preoperative and intraoperative patient images in order to guide surgical procedures (e.g., BrainLab). Modern radiation therapy treatment methods rely on detailed computer models of patient anatomy and treatment beam geometry in order to deliver precise radiation doses to cancerous tumors while sparing neighboring radiosensitive structures (Tewell and Adams, 2004). Statistical models representing the variability of structures within the brain are being used to study the progression of diseases such as schizophrenia, Fragile X, and Alzheimer’s disease (Styner et al., 2005).

However, the diverse and intricate geometrical structures within anatomy, as well as the nonlinear nature of anatomical shape change, pose a significant challenge for the accurate representation and analysis of these structures. For example, classical statistical techniques such as principal components analysis (PCA) and regression rely on the vector-space structure of observations and are not appropriate for nonlinear models of anatomical shape or shape change (Grenander and Miller, 1998; Miller, 2004; Pizer et al., 2003; Fletcher et al., 2004). Therefore, a currently active area of research seeks to develop shape descriptions and statistical methods that are applicable for non-Euclidean data.

There are two common frameworks for image-based analysis of anatomical structures; both are capable of capturing nonlinear variability among shapes. The first uses either parametric or non-parametric shape models to explicitly represent the geometry of anatomical structures observed within images. Examples of these models include collections of point landmarks (Cootes et al., 1995), medial-based representations (Pizer et al., 2003), and level sets (Malladi et al., 1995). Within this framework, anatomical differences and motion are represented via manipulation of the underlying shape model. The statistics of anatomical shape can be studied through the statistical analysis of these shape models.

The second framework, used in this dissertation, is motivated by the observation that the comparison of anatomical structures is inherently related to the construction of spatial transformations that map one anatomy to another. For example, large-deformation mappings of the underlying coordinate system of image volumes are capable of representing intricate anatomical changes within the brain (Miller and Younes, 2001). In Grenander’s pattern theory (Grenander, 1996), and in particular in *computational anatomy* (Grenander and Miller, 1998), it is the analysis of these mappings themselves, rather than the analysis of explicit shape representations, which leads to insight into geometric change and variability of anatomical structure.

This dissertation is focused on utilizing this image deformation, or image mapping, framework to measure anatomical motion and apply statistical methods to anatomical structures represented by medical images. There are three primary contributions. The first contribution is a method for the quantification of anatomical motion in the presence of *extraneous structures*—objects that are present in one image and absent in another. The second contribution is a novel technique for inferring a mean anatomical configuration from a collection of medical images, as well as a mapping from each image in the collection to the mean. The third contribution of this thesis is a new technique called *manifold kernel regression* for estimating systematic changes from data in nonlinear

spaces, and its application to medical image data.

1.1.1 Accommodating topological change in large-deformation diffeomorphic image mapping

Within medical image analysis, it is often necessary to quantify organ motion between anatomical observations. For example, in order to accurately record the radiation dose delivered to the prostate during successive radiation treatments, the motion of the prostate with respect to the planned treatment position must be measured. One common approach for quantifying organ motion uses spatial mappings to establish a correspondence between the underlying coordinate systems of successive anatomical images.

Large-deformation diffeomorphic transformations are often used to capture local, intricate changes while preventing folding or tearing of space. However, by definition, diffeomorphic transformations cannot be applied to match structures with different topologies. For image mapping, this implies that no correspondence will exist for structures that are present in one image and absent in the next. I present a novel method that extends large-deformation diffeomorphic image mapping to accommodate such topological changes by deflating these structures before computing diffeomorphic transformations. In this way correspondence is determined where it can exist. The method is applied to quantify the motion of the prostate in the presence of transient bowel gas.

1.1.2 Inferring mean anatomical configurations from medical imagery

For anatomy represented within medical images, a natural problem is the construction of a statistical representative anatomical configuration for a population. Such representatives are important, for example, when investigating clinical hypotheses related to the

shape of structures indicated in brain development or mental disorders. Another use for such a representative is to define a common spatial coordinate system so that spatial information from across a population can be accumulated, statistically analyzed, and presented in a single frame of reference.

One common method for generating representative anatomical images is to simply select an individual image from the population. However, unless there is an *a priori* reason to choose one individual over the rest, this choice will lead to a bias in the analysis toward that particular anatomical configuration presented by the individual. This method is also commonly used to establish spatial correspondence between image coordinate systems: one image is chosen as a reference and all other images are spatially aligned with this reference image. Unfortunately, this is biased as well since the image-to-image correspondences rely on the choice of the reference image.

A sensible approach for generating an anatomical representative image that is not biased by any particular individual is to generate an anatomical configuration that is in some sense centrally located—in terms of the configuration of anatomy within an image—with respect to the population under study. In statistics this notion is captured by the mean, and the arithmetic mean is easily computed from data that are elements of a vector space. However, the naive approach of averaging an image voxel-wise clearly neither produces a realistic anatomical image nor captures the notion of mean anatomical configuration.

Recently, the notion of Fréchet mean has been used to define mean shapes in non-linear shape spaces that have a metric space structure (Fletcher et al., 2004; Pennec, 2006). For example, Fletcher et al. (2004) extend concepts such as averaging and principal components analysis to manifolds representing anatomical shape variability. This dissertation presents a novel method for generating a Fréchet mean anatomical image from a collection of images representing a population. The mean image is defined as the image that requires least amount of squared deformation energy in order to match each




Figure 1.1: **Fréchet mean and manifold kernel regression illustrations.** (a) Illustration of a Fréchet mean on a curved manifold. The mean μ minimizes the sum of squared distances, along the manifold, to the input points p_i . (b) Diffeomorphic changes of coordinate allow spatial information to be mapped between the mean and each input image. (c) Illustration of manifold kernel regression. For any value of a predictor variable t , such as age, the manifold-valued observations p_i are summarized by the weighted Fréchet mean point $\hat{n}_h(t)$.

image in the collection (Figure 1.1 (a)). The resulting image serves a natural average representation for the anatomy contained within the images and is not biased by the choice of any particular image or ordering of the population images. The Fréchet mean image also provides a natural coordinate system for the collection of images. Since the method uses large deformation diffeomorphic registration to establish correspondence between the image coordinate systems during the process of forming the mean, spatial data is easily mapped between images within the collection and is easily accumulated within the mean coordinate system on the mean image (see Figure 1.1 (b)).

1.1.3 Regression analysis for medical imagery

Another important area of medical image analysis is the development of methods for automated and computer-assisted determination of systematic anatomical change with respect to a predictor variable such as time. Shape-change and growth models are used

to analyze and better understand healthy anatomical structure, function, and change such as the beating heart or aging brain. Similarly, these methods are used to analyze the onset and progression of diseases, enabling researchers to better understand disease processes and uncover disease predictors. Finally, the analysis of systematic anatomical change provides a method for judging and comparing the efficacy of therapeutic protocols. For example, the measurement of tumor size and shape over time is critical in judging severity of cancer and the efficacy of various cancer treatments.

A number of longitudinal growth models have been developed to provide this type of analysis for time series imagery of a single subject (for example Beg, 2003; Clatz et al., 2005; Miller, 2004; Thompson et al., 2000). These models infer a description of the geometrical change in anatomy from the images, which are indexed by a time measurement such as age, elapsed time, or duration of treatment. While longitudinal growth models provide detailed information about anatomical change in an *individual case*, they cannot be applied directly to a *population* in order to study systematic, time-related trends that occur on average within a population.

In order to determine systematic, population-wide trends from anatomical imagery, it is necessary to incorporate data from a population of individuals. This entails studying a collection of medical images—an image database—where each image represents the anatomy of a particular individual and is associated with a particular value of the predictor variable. There are several classical frameworks for the analysis of population changes over time within statistics. For these frameworks there is a tradeoff between the amount of information that can be inferred from the data and the practicality of obtaining the data.

Longitudinal studies track a specific subset of a population over time. However, their practical use for investigating long-term trends in anatomical data is limited. Longitudinal image data is costly and impractical to obtain because of the need to coordinate patients and staff and control the imaging protocol over an extended period of time.

Also, longitudinal datasets—by definition—may take years or decades to acquire while studies must often be carried out more quickly.

On the other hand, random design studies, in which individual patients are not tracked over time, are more practical because data can be collected within a relatively short amount of time. Furthermore, unlike cross-sectional studies, there are no fixed and uniform time points at which the observations must be collected. For example, a database of brain images from a population of healthy adults, with the age of the individual associated with each image, forms a random design database of images.

Therefore, random design medical image databases provide a rich and practical environment for the study of anatomical change. However, in order to determine systematic trends from random design data, it is necessary to separate two distinct aspects of anatomical variation: individual variation and systematic effect. For example, a study of brain atrophy as a function of age, for healthy adults, must factor out individual brain size.

For data that can be described in terms of vector-valued measurements a variety of parametric and non-parametric regression techniques can be applied to describe systematic change over time for a population. However, vector spaces do not adequately represent the highly variable and nonlinear geometric changes that characterize anatomical motion and anatomical differences. While linear regression techniques can be applied to nonlinear shape data by embedding it in vector spaces, the results can be erroneous and misleading.

This dissertation presents a new technique, called *manifold kernel regression* for estimating systematic changes from data in nonlinear spaces. Manifold kernel regression is a generalization of a standard technique in statistics called kernel regression, which is a method used to estimate the relationship, on average, between two random variables. Manifold kernel regression is based on the notion of Fréchet expectation, which can be used to define averages on manifolds (Figure 1.1 (c)). When applied to anatomical

images, manifold kernel regression can be used to regress a population representative image as a function of time from an image database. The analysis of this regressed image gives insight into the systematic anatomical changes, as a function of age, within the population.

The driving anatomical problem for mean and regression work in this dissertation is the analysis of shape change in the brain, as a function of age, from a database of three-dimensional MR volumes. While this problem is in the field of medical image analysis, manifold kernel regression can be applied more generally to any manifold-valued data. For example, it is applied to rotational pose in Davis et al. (2007).

1.2 Thesis and contributions

Thesis: Manifold kernel regression is a natural generalization of kernel regression that enables regression analysis for points on a manifold. It extends classical kernel regression in order to estimate, from a collection of observations, the relationship—on average—between an independent predictor variable, such as time, and a dependent variable represented by points on a manifold. In particular, this method is useful for determining population average anatomical shape change over time from a random design database of medical images. Because it provides a quantitative link between a predictor variable and anatomical structure, manifold kernel regression is an effective tool for improving our understanding of anatomical changes within populations.

The contributions of this dissertation are as follows:

1. I present a novel method, called manifold kernel regression, that enables regression analysis of manifold-valued data.
2. I apply manifold kernel regression to the study of anatomical change from a random design database of medical images. In particular, it is defined for images using the framework of large-deformation diffeomorphic image mapping.

3. I demonstrate manifold kernel regression by measuring average geometric change in the aging brain from a random design dataset of 3D MR images. The effect of regression kernel width on the regressed shape is explored.
4. I present a novel method for computing a Fréchet mean image from a collection of images.
5. I describe a program for efficiently computing Fréchet mean images and applying the manifold kernel regression analysis method to 2D and 3D images on shared-memory, multi-processor machines. Performance measurements are included.
6. I present a novel method for extending diffeomorphic image mapping to accommodate certain topological changes. The method is applied to track the changing position of the prostate relative to the pelvis in the context of transient bowel gas. The effectiveness of this method is tested in a retrospective study involving 40 3D computed tomography images from 3 patients undergoing adaptive radiation therapy.

1.3 Overview of chapters

The remainder of this dissertation is organized as follows:

Chapter 2 presents an overview of the mathematical topics that are used within this dissertation. These topics include metric spaces, differentiable manifolds, and diffeomorphisms. The Riemannian metric space structure of diffeomorphisms is reviewed.

Chapter 3 applies large-deformation diffeomorphic image registration to the problem of automatic segmentation of the prostate for radiation therapy. A novel method for extending image registration to accommodate topological changes is described. A study comparing the segmentation results given by this method to manual segmentations is presented.

Chapter 4 presents a variational optimization method for computing a Fréchet mean image from a collection of input images using the large-deformation diffeomorphic image registration framework. A summary of recent work in computing Fréchet means in nonlinear spaces and various methods for computing mean images is included.

Chapter 5 presents manifold kernel regression, which is a method for applying kernel regression to manifold-valued data. An algorithm is presented for applying manifold kernel regression to collections of medical images within the large-deformation diffeomorphic framework.

Chapter 6 applies the manifold kernel regression algorithm from Chapter 5 to study changes in brain structure from a random design database of 3D MR brain images of healthy adult subjects. Regressed population average brain images, as a function of subject age, are generated for male-only, female-only, and combined cohorts. A large deformation diffeomorphic growth model for longitudinal data is applied to the regressed images in order to measure local, population average geometrical changes as a function of age. An exploration of regression kernel width selection in the diffeomorphic setting is also presented.

Chapter 7 contains a discussion of the contributions of this thesis and an outline of future work.

Appendix A reviews the basic mathematical structures used within this dissertation, including vectors spaces, function spaces, groups, and differential manifolds.

Appendix B reviews the Euler-Lagrange equation for the large-deformations diffeomorphic metric mapping (LDDMM) solution to the diffeomorphic image matching problem.

Appendix C reviews the derivation of the differential operator for “fluid” image registration.

Appendix D describes a method for numerically inverting this differential operator within the fluid registration algorithm.

Chapter 2

Large-deformation diffeomorphic image matching

2.1 Introduction

The goal of this chapter is to summarize the large-deformation diffeomorphic image matching framework that is used throughout this dissertation. In this framework, diffeomorphic changes of coordinates are used to describe geometric change for objects represented within images. For medical images, this geometric change may be due to change over time for an individual, or it may represent geometric differences between two different individuals. The analysis of these deformations provides insight into shape changes or geometric differences in the underlying geometric structures.

Using the algebraic and differential geometric structure of diffeomorphisms, it is possible to define a metric that provides a well-defined notion of “amount of geometric change.” This metric is used in chapters 4 and 5 to develop methods for computing mean images and for applying regression to collections of images.

No novel contributions are presented in this chapter. The large-deformation diffeomorphic image matching framework has its roots in pattern theory (Grenander, 1996) and has been the focus of active development by a number of authors. For more details see, for example, Christensen et al. (1996); Dupuis et al. (1998); Grenander and Miller

(1998); Trouvé (1998); Miller and Younes (2001); Miller et al. (2002); Younes (2005); Beg et al. (2005); Miller et al. (2006).

Appendix A reviews the mathematical concepts used in this chapter.

2.2 The large-deformation diffeomorphic framework

In this dissertation, geometric structures, such as anatomical tissue and organs, are represented by 2-dimensional and 3-dimensional images. These images are modeled as real-valued L^2 functions on the domain $\Omega \subset \mathbb{R}^3$.¹ The space of images is denoted by

$$\mathcal{I} \equiv L^2(\Omega). \quad (2.1)$$

Spatial transformations are used to deform images by deforming the underlying coordinate space of Ω .

These transformations $\phi \in \text{Diff}_V(\Omega)$ are elements of a subgroup of diffeomorphisms $\text{Diff}(\Omega)$, $\phi : \Omega \rightarrow \Omega$ that are generated by flows of smooth, time-varying velocity fields with support on Ω for a simulated time parameter $t \in [0, 1]$. The introduction of the velocity fields enables large-deformation transformations to be produced while maintaining the diffeomorphic property (Dupuis et al., 1998).

These flows $v_t : [0, 1] \rightarrow V$ are generated from velocity fields that are elements a Hilbert space V with associated inner product $\langle \cdot, \cdot \rangle_V$. For $u, v \in V$, this inner product is defined using a linear differential operator L (with associated adjoint L^\dagger)

$$\langle u, v \rangle_V \equiv \langle Lu, Lv \rangle_{L^2} = \langle L^\dagger Lu, v \rangle_{L^2} = \int_{\Omega} \langle L^\dagger Lu(x), v(x) \rangle_{\mathbb{R}^3} dx, \quad (2.2)$$

¹For the sake of clarity, the notation in this dissertation is restricted to 3-dimensional space. This is convenient since most medical applications use 3-dimensional images. However, this work is also applicable to 2-dimensional images.

where $\langle \cdot, \cdot \rangle_{\mathbb{E}^3}$ is the Euclidean inner product. This inner product on velocity fields induces the norm

$$\|v\|_V \equiv \sqrt{\langle v, v \rangle_V}. \quad (2.3)$$

The form of the differential operator L is taken from fluid mechanics (Christensen et al., 1996; Dupuis et al., 1998) to be

$$L = \alpha \nabla^2 + \beta (\nabla \cdot) \nabla + \gamma \quad (2.4)$$

where α and β govern the viscous properties of the deforming medium and γ ensures that L is invertible. Appendix C describes this operator in more detail.

The operator L is associated with the compact self-adjoint operator K by

$$\langle u, v \rangle_{L^2} = \langle Ku, v \rangle_V \quad (2.5)$$

which implies that

$$u = KL^\dagger Lu. \quad (2.6)$$

Practically, one can think of K as a smoothing operator.

The flow v_t is related to the diffeomorphism ϕ via the Lagrangian ODE

$$\frac{d}{dt}\phi_t(x) \equiv \dot{\phi}_t = v_t(\phi_t(x)). \quad (2.7)$$

In particular, ϕ is generated from v_t according to

$$\phi_t(x) = x + \int_0^t v_t \circ \phi_t(x) dt \quad (2.8)$$



Figure 2.1: **Velocity integration.** The diffeomorphism $\phi \in \text{Diff}_V(\Omega)$ evolves in time as a flow ϕ_t . This flow is governed by the time-varying velocity field v_t . This process is demonstrated for a single point $x \in \Omega$.

subject to

$$\phi_0(x) = x \quad \phi(x) = \phi_1(x) \quad \dot{\phi}_t(x) = v_t \circ \phi_t(x) \quad \text{for all } x \in \Omega. \quad (2.9)$$

Figure 2.1 demonstrates this process for a single point $x \in \Omega$. Dupuis et al. (1998) establish sufficient conditions, in the form of smoothness constraints, on $\langle \cdot, \cdot \rangle_V$ for these flows to generate diffeomorphisms.

A differentiable manifold structure is defined for $\text{Diff}_V(\Omega)$, where V is the tangent space at the identity. The combination of group structure and differentiable structure allow $\text{Diff}_V(\Omega)$ to behave very much like a Lie group. In particular, a right-invariant Riemannian distance is defined on $\text{Diff}_V(\Omega)$ based on $\langle \cdot, \cdot \rangle_V$ at the identity by

$$d_{\text{Diff}_V(\Omega)}(\text{Id}_{\text{Diff}_V(\Omega)}, \phi) = \inf_{v: \dot{\phi}_t = v_t(\phi_t)} \int_0^1 \sqrt{\langle v_t, v_t \rangle_V} dt = \inf_{v: \dot{\phi}_t = v_t(\phi_t)} \int_0^1 \|v_t\|_V dt \quad (2.10)$$

$$(2.11)$$

subject to

$$\phi(x) = x + \int_0^1 v_t \circ \phi_t(x) dt \quad \text{for all } x \in \Omega. \quad (2.12)$$

The distance between any two elements of $\text{Diff}_V(\Omega)$ is defined by

$$d_{\text{Diff}_V(\Omega)}(\phi_1, \phi_2) = d_{\text{Diff}_V(\Omega)}(\text{Id}_{\text{Diff}_V(\Omega)}, \phi_2 \circ \phi_1^{-1}). \quad (2.13)$$

With this structure the length of curves can be measured along the manifold $\text{Diff}_V(\Omega)$. This distance provides a metric space structure for $\text{Diff}_V(\Omega)$ (see Appendix A, Proposition A-7.1). When related back to the underlying geometric structures represented in images, this distance provides a well defined notion of “amount of geometric change.”

2.3 Image matching and an image-to-image distance using large-deformation diffeomorphisms

In this section the large-deformation diffeomorphism framework described above is applied to computing transformations that deform one image to match another. The resulting deformations are used to define an image-to-image metric that takes geometric change into account. Intuitively, the distance between two images is given by the “amount of deformation” required for one image to match another.

Consider a diffeomorphism $\phi \in \text{Diff}_V(\Omega)$. The action of ϕ on an image $I \in \mathcal{I}$ is defined by

$$I_\phi \equiv I \circ \phi^{-1}. \quad (2.14)$$

Given a fixed and moving (to-be-deformed) image, I^F and I^M in \mathcal{I} , the goal is to generate a deformation ϕ that best aligns I_ϕ^M with I^F .

For *exact matching*, where a deformation is all that is needed to explain the difference

between I^F and I^M , ϕ is defined by the optimization problem

$$\hat{\phi} = \underset{v: \dot{\phi}_t = v_t(\phi_t)}{\operatorname{arginf}} \int_0^1 \|v_t\|_V^2 dt \quad \text{such that} \quad I^F = I_{\hat{\phi}}^M. \quad (2.15)$$

In this case ϕ is the element of $\operatorname{Diff}_V(\Omega)$ that deforms I^M to match I^F with the smallest squared distance according to $d_{\operatorname{Diff}_V(\Omega)}$. The resulting squared distance value is used to define a squared image-to-image metric

$$d_{\mathcal{I}, \text{Exact}}^2(I^F, I^M) = \inf_{v: \dot{\phi}_t = v_t(\phi_t)} \int_0^1 \|v_t\|_V^2 dt \quad \text{such that} \quad I^F = I_{\phi}^M. \quad (2.16)$$

For *inexact matching*, a mechanism for penalizing residual image mismatch is required. This *image dissimilarity metric* is determined by the image modalities and image noise models. In this work the L^2 norm is used²

$$d_{\mathcal{I}, \text{Inexact}}^2(I^F, I^M) = \inf_{v: \dot{\phi}_t = v_t(\phi_t)} \int_0^1 \|v_t\|_v^2 dt + \frac{1}{\sigma^2} \|I_{\phi}^M - I^F\|_{L^2}^2. \quad (2.17)$$

This equation introduces the free parameter σ that governs the relative weight of the two terms. Small values of σ increase the importance of the image dissimilarity metric, forcing the images to match as well as possible; large values of σ produce deformations that require less “energy” according the metric on $\operatorname{Diff}_V(\Omega)$. Dupuis et al. (1998) show the existence of a minimizer for this equation.

Although this construction is motivated by the metric on $\operatorname{Diff}_V(\Omega)$, it does not strictly define a metric on $\operatorname{Diff}_V(\Omega)$ because of the second term in Equation (2.17).

A Bayesian interpretation (although not rigorous) for Equation (2.17) is that the first term acts like a prior on the distribution possible deformations while the second term describes how well the deformed moving image matches the fixed image given the

²Of course other metrics are possible; in Lorenzen et al. (2006) the Kullback-Leibler divergence is used as an image dissimilarity metric to align tissue probability maps.

current deformation (Dupuis et al., 1998).

2.4 Euler-Lagrange equations for the image matching problem

Beg et al. (2005) show that the Euler-Lagrange equations for the energy functional in Equation (2.17) are

$$2v_t - K \left(\frac{2}{\sigma^2} |D\phi_{t,1}| \nabla I_{\phi_{0,t}}^M (I_{\phi_{0,t}}^M - I_{\phi_{1,t}}^F) \right) = 0. \quad (2.18)$$

where $D\phi_{t,1}$ is the 3×3 Jacobian matrix of the transformation $\phi_{t,1} \equiv \phi_1 \circ \phi_t^{-1}$. This equation is also derived in Appendix B of this document. Beg calls these solutions the Large-Deformations Diffeomorphic Metric Mapping (LDDMM) solution for Equation (2.17). For any particular time point $t \in [0, 1]$ the gradient of the energy functional (Equation (2.17)) is

$$\nabla_{v_t} E_t = 2v_t - K \left(\frac{2}{\sigma^2} |D\phi_{t,1}| \nabla I_{\phi_{0,t}}^M (I_{\phi_{0,t}}^M - I_{\phi_{1,t}}^F) \right). \quad (2.19)$$

Greedy solution

Christensen et al. (1996) proposed a greedy solution to Equation (2.17). This solution separates the time dimension of the problem from the space dimensions. At each iteration, a new velocity is computed that optimizes the functional (2.17) given that the current deformation is fixed (i.e., the past velocity fields are fixed). Unlike the LDDMM approach, this optimization does not update velocity fields once they are first estimated or take future velocity fields into account. Using a step-size ϵ , these velocity fields are

integrated to produce the final deformation. In this case the gradient is

$$v_t = K \left(\frac{2}{\sigma^2} \nabla I_{\phi_0,t}^M (I_{\phi_0,t}^M - I_{\phi_1,t}^F) \right). \quad (2.20)$$

Chapter 3

Accommodating topological change for practical large-deformation diffeomorphic image mapping: Application to the measurement of prostate motion

3.1 Introduction

As described in the introduction, diffeomorphic image registration is commonly used to define a spatial correspondence between geometrical structures that are represented within images. However, these diffeomorphic mappings are not suitable for matching objects with different topological properties, such as when a new object is present in one image and absent in another. In this chapter a novel method is presented for extending large-deformation diffeomorphic image mapping to accommodate this type of topological change. The method is applied in the context of external beam radiation therapy to quantify the motion of the prostate in the presence of transient bowel gas.

This work was carried out in conjunction with Drs. Sarang Joshi, Mark Foskey, Lav Goyal, Julian Rosenman, Edward Chaney, and Sha Chang of the department of Radiation Oncology at the University of North Carolina at Chapel Hill. Variations of

this work have been published in (Davis et al., 2005) and (Foskey et al., 2005).

3.1.1 The need for automatic segmentation in adaptive radiation therapy of the prostate

External beam radiation therapy (EBRT) is one major treatment method for prostate cancer. In EBRT, cancerous tissue is destroyed through the delivery of high energy x-rays in a series of 40 or more daily treatments (DeVita et al., 2004; Halperin et al., 2007). To be safe and effective, the radiation dose to the cancer-containing prostate should be as high as possible while the dose to surrounding organs, such as the rectum and bladder, must be limited. This effect is achieved by using multiple radiation beams that overlap on the tumor and are shaped to exclude normal tissue as much as possible. However, internal organ motion and patient setup errors present a serious challenge to EBRT (see figures 3.1 and 3.2). The prostate, rectum, bladder and other organs move with respect to the fixed patient position, and even small changes in their position can result in either tumor under-dosing, normal tissue over-dosing, or both.

In order to meet this challenge, adaptive radiation therapy (ART), which uses periodic intra-treatment CT images for localization of the tumor and radiosensitive normal structures, is being investigated. In ART a feedback control strategy (Yan et al., 2000) is used to correct for differences in the planned and delivered dose distributions due to spatial changes in the treatment volume early in the treatment period.

Although in-treatment-room CT scanners provide the enabling imaging *hardware* to implement ART, no *software* methods or tools for automatic image processing exist to enable the incorporation of these images in the adaptive treatment of prostate or other cancer. As a result, manual intervention is required to *segment*—i.e., define the location of—the tumor and other structures within each image. However, manual segmentation places an impractical burden on highly skilled and already overburdened




Figure 3.1: Prostate motion over the course of radiation therapy treatment. Visual depiction of prostate motion over the course of treatment for nine patients in our study. The actual location of the prostate on each treatment day is indicated by the superimposed white contours. These contours are taken from manual segmentations of treatment images. The discrepancies between the contours exhibit the effect of setup error and organ motion on the prostate position. Note that different patients exhibit different amounts of prostate motion; compare the close contour agreement for patient 3101 with the wide contour variability for patient 3109. For some patients (3102, 3109) motion is primarily noticeable in the anterior-posterior direction; for other patients (3106, 3107) motion is primarily noticeable in the lateral direction.




Figure 3.2: **Patient setup error.** First column: axial and sagittal slices from the planning image of patient 3102. Second column: the same slices (with respect to the planning image coordinate system) taken from a treatment image. Third column: the voxel-wise absolute difference between the planning and treatment images. Black represents perfect intensity agreement, which is noticeable in the interior of the bones and outside the patient. Brighter regions, indicating intensity disagreement, are especially apparent: (1) in regions where gas is present in one image and absent in the other, (2) around the bladder, which is large on the treatment day compared to the planning day, (3) uniformly along boundaries with high intensity gradient, indicating a global setup error such as a translation.

personnel. Moreover, clinically significant inter- and intra-rater¹ variability of manual segmentations introduces a source of treatment uncertainty that current adaptive radiation therapy techniques do not address (van Herk et al., 1995; Ketting et al., 1997).

¹‘Rater’ is used in the following sense: a *rater* is a person who segments an image. *Intra-rater* variability refers to differences in segmentations made by the same rater. *Inter-rater* variability refers to differences in segmentations made by two different raters. These differences can be measured for repeated segmentations of a single image by a single or multiple raters. These differences are caused by the fact that segmenting complicated structures with ill-defined boundaries in images is a difficult task. Different raters, or even the same rater over time, will inevitably segment the same structure differently because they guess at ill-defined object boundaries differently. They also use different segmentation strategies, display settings, and rules of thumb to segment images.

3.1.2 Estimating prostate motion via image mapping

Deformable image registration is one method for automatically quantifying organ motion over the course of EBRT. The CT image taken at planning time, the *planning image*, is used as a reference. A physician manually segments the prostate and nearby radiosensitive structures in this image. On each treatment day, the patient is positioned and then, prior to treatment, a new CT image is acquired using an in-treatment-room CT scanner that shares a table with the linear accelerator (Figure 3.3). Each *treatment image* characterizes the patient configuration at that treatment time. After establishing the spatial correspondence between the planning image and each treatment image using deformable image registration (Figure 3.4), manually segmented structures from the planning image can be mapped into these treatment images.

In this way, automatic segmentations of the treatment images are provided by the combination of a manually segmented planning image and automatic, deformable image registration (Figure 3.5). This registration-based (or atlas-based) segmentation procedure is *not* automatic in the sense that no manual work is required: a segmentation of the planning image is essential. However, no additional manual effort is required to segment treatment images once the planning image segmentation is available. Furthermore, planning image segmentations are always available since they are generated as part of routine clinical practice.

In CT images of the abdomen, however, the presence of bowel gas complicates the registration process since no correspondence exists for pockets of gas across different days. Figure 3.6 shows two rigidly aligned axial images of a patient taken on two different days. Bowel gas is present in one but absent in the other. Figure 3.7 shows a failed automatic segmentation of the rectum. Panel (b) shows the result of automatic segmentation using large-deformation image registration. Manually drawn contours of the prostate and rectum are mapped, using this correspondence, from the planning image (a) onto the daily image. Manual contours are drawn in red and mapped contours are



Figure 3.3: **In-treatment-room CT scanner.** A CT-on-rails (left) shares a treatment table with the linear accelerator (top right) that is used for external beam radiation therapy. This setup allows the patient to be imaged and treated in the same position.

drawn in yellow. The manual and automatically generated contours in the daily image are misaligned; the presence of bowel gas has caused correspondence errors around the rectum.

In order to apply large-deformation image registration in this setting, a novel method was developed that combines image registration with a bowel gas segmentation and deflation algorithm. Final image-to-image correspondence is defined by the composition of the deflation and registration transformations. A demonstration of the results of this algorithm is shown in Figure 3.7 panel (c). Notice the close alignment between the manual contours and the contours generated by this method. Section 3.3 presents a study of this method’s accuracy for determining prostate motion during the course of radiation therapy of the prostate.

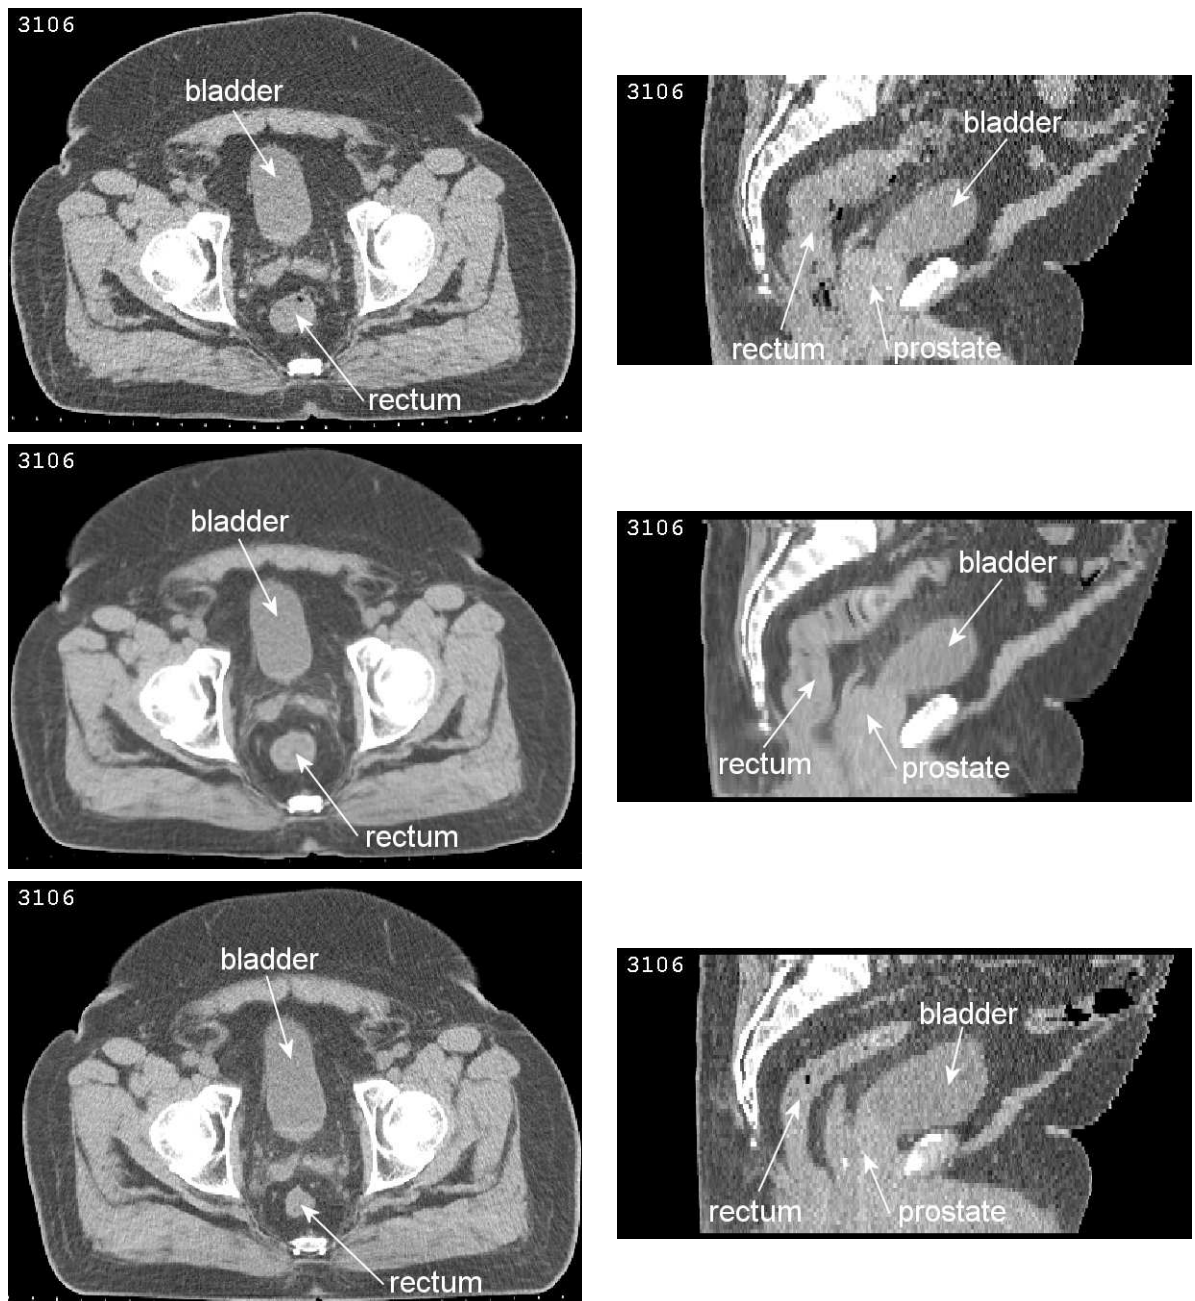


Figure 3.4: **Quantifying tissue motion via deformable image registration.** Example of deformable image registration. The first and last rows show axial and sagittal slices of the planning and treatment images. The second row shows the treatment image after deformable image registration, which brings the treatment image into alignment with the planning image. Note how the changes in size and shape of the bladder and rectum are accounted for.




Figure 3.5: **Automatic prostate segmentation via deformable image registration.** Example of automatic segmentation using deformable image registration. (a) Axial slice of a planning image with the prostate labeled by a white contour. (b) The same axial slice (in terms of planning coordinates) from a treatment image. The planned prostate position is shown in white, the actual prostate in black (both contours manual). (c) The same treatment image and manual (black) contour. The white contour is automatically generated by performing deformable image registration and applying the resulting deformation to the planning segmentation. The close agreement of the contours indicates that image registration accurately captures the prostate motion.




Figure 3.6: **The geometric effect of bowel gas.** Axial CT slice of the same patient acquired on different days, showing the effect of bowel gas.




Figure 3.7: **Improved segmentation via gas deflation.** Automatic segmentation of the prostate and rectum. Manually segmented structures in the planning image (a) are mapped to the daily image (b) *before* accounting for bowel gas, and (c) *after* accounting for bowel gas with the gas deflation algorithm. Manually drawn contours are shown in red and mapped contours are shown in yellow.

3.2 Methodology

This section describes a methodology for deformable image registration that accommodates the presence of transient bowel gas. First, a general notion of topological equivalence for images is defined. Next, a deflation algorithm for removing extraneous structures—such as bowel gas—from images is described. This algorithm can be combined with deformable registration in order to register images with extraneous structures. Finally, the details of applying this algorithm to CT image registration for automatic prostate segmentation are presented.

3.2.1 Image mapping and topology

Topology, often described intuitively as *rubber sheet geometry*, is the study of those properties of an object that remain unchanged under continuous transformations of the object. Two geometrical objects² are *topologically equivalent* if each can be continuously deformed to match the other. The archetypical example in topology is that—because they both have a single hole—a donut is topologically equivalent to a coffee cup.³

²More precisely, topological spaces using the standard Euclidean topology.

³And so, as the joke goes, a mathematician cannot tell one from the other.

The notion of topological equivalence is fundamentally related to image mapping because both are concerned with transformations between geometrical objects. On the one hand, one image can be deformed so that its intensity values exactly match another image. However, the alignment of image intensities does not imply—it merely suggests—the proper alignment of structures represented within the images. Instead, a notion of topological equivalence can be applied in a straightforward manner to images by relying on topological equivalence of geometrical objects represented within images. Intuitively, two images are topologically equivalent if each image can be continuously deformed—without tearing space or gluing distinct objects together—so that the *contents* of one image match the *contents* of the other. In the rest of this section this definition is made precise.

Let $\Omega \equiv [0, 1] \times [0, 1] \times [0, 1] \subset \mathbb{R}^n$ be a common ambient space. Images $I, J \in \mathcal{I}$ are defined as L^2 functions from Ω to the reals. Every geometrical object represented in an image can be described by a non-empty open subset $U^k \subset \Omega$, using the standard Euclidean topology, that encodes the location of the object labeled k . These subsets may overlap. All points $x \in \Omega$ should be labeled as part of at least one object.⁴ A family \mathcal{U} of such geometrical objects is associated with each image and represents an interpretation of the contents of that image.

Within this setting, the notion of topological equivalence of images can be built from the following definitions:

Definition 3.2.1 (Homeomorphism, adapted from Lee (2003), p. 541). Let X and Y be topological spaces. A continuous bijective map $F : X \rightarrow Y$ with continuous inverse is called a *homeomorphism*. If there exists a homeomorphism from X to Y , we say that X and Y are *homeomorphic*. A *diffeomorphism* $G : X \rightarrow Y$ is a homeomorphism with the additional requirement that F and its inverse are smooth (C^∞).

Definition 3.2.2 (Identifiability). Two families of geometrical objects \mathcal{U}_I and \mathcal{U}_J are

⁴Points without a natural label could be labeled as “background,” for example.

identifiable if there exists a natural bijection between elements of \mathcal{U}_I and \mathcal{U}_J that associates the same geometrical objects across the two images. That is, for all object labels, $U_I^k \in \mathcal{U}_I$ if and only if $U_J^k \in \mathcal{U}_J$. Clearly, if an object is present in one image and absent in another, no such bijection is possible.

Definition 3.2.3 (Homeomorphism of objects *over* Ω). Two families of geometrical objects \mathcal{U}_I and \mathcal{U}_J are *homeomorphic over* Ω if they are identifiable and there exists a homeomorphism $F : \Omega \rightarrow \Omega$ that is also a homeomorphism for every pair of sets $U_I^k \in \mathcal{U}_I, U_J^k \in \mathcal{U}_J$.

Finally, because the terminology of geometrical objects represented within images is cumbersome, I will abuse notation and give the following definition of topological equivalence for images.

Definition 3.2.4 (Topological Equivalence of Images). Two images $I, J \in \mathcal{I}$ are topologically equivalent if their associated geometrical objects, represented by \mathcal{U}_I and \mathcal{U}_J , are homeomorphic over Ω .

Figure 3.6 presents an example of images that are not topologically equivalent. It should be emphasized that this definition is not tied to the intensity values of the images, but instead relies on some interpretation of the objects represented within the images. If, for example, objects are interpreted differently within an image, the topology of that image will change.

3.2.2 Accommodating topological change

For images that are topologically equivalent, in the sense defined above, diffeomorphisms are suitable for representing *geometrical changes*. When *topological changes* are present, such as the addition of an object into an image, diffeomorphisms cannot, by definition, describe the underlying change. However, even in the presence of such topological

changes, it may be practically meaningful and helpful to use diffeomorphic image mapping to estimate correspondence in regions of the image *where correspondence does exist*, while working around the topological changes imposed by any extra objects.

Let I and J be a pair of identical images except that J contains exactly one extra geometrical object that is not represented in I . We call this extra object a *extraneous object* and denote it by $N \in U_J$. Suppose that the other objects, which are common to I and J , have undergone some geometrical change, described by the diffeomorphism $h : \Omega \rightarrow \Omega$ that we wish to recover.

In the intuitive language of topology the strategy is similar to performing surgery: cut the extraneous structure out of image J and glue the resulting hole closed before deforming J to match image I . However, rather than cutting, we *shrink* N , using a transformation $s : \Omega \rightarrow \Omega$, until it is negligible in size compared to the image resolution.⁵ During this process, the space around N expands smoothly to fill in the space created by the deflating structure.

More precisely, a diffeomorphic deflation transformation $s : \Omega \rightarrow \Omega$ is determined such that $J \circ s$ is the image J after a deformation that deflates N . This transformation is constructed using an algorithm that models the motion of a viscous fluid under the influence of external forces. This methodology is motivated by the fluid algorithm used for image registration (Christensen et al., 1996; Dupuis et al., 1998). An iterative solution for s is summarized in Algorithm 3.1.

The transformation s is defined by integrating velocity fields $v(x, t)$ forward in time:

$$s(x) = x + \int_0^1 v(s(x, t), t) dt. \quad (3.1)$$

⁵Practically, this approximates topological equivalence at the resolution that image mapping techniques work.

These velocity fields are induced by a force function

$$F(x, t) = \nabla(J_N \circ s_t)(x) \quad (3.2)$$

that is the gradient of the binary image that labels the structure N within the image J . The force function and velocity fields are related by the linear operator L (and it's adjoint L^\dagger)⁶ that defines the mechanical properties of the deforming continuum:

$$L^\dagger L v(x, t) = F(x, t). \quad (3.3)$$

Convergence is achieved when the incremental change in s falls below a predetermined threshold.⁷

Algorithm 3.1 Iterative greedy algorithm for extraneous structure deflation

Input: Image J containing extraneous structure N

Output: Deformation s such that N is deflated in $J \circ s$

- 1: $s^0 \leftarrow Id$ // start with identity mapping
 - 2: $i \leftarrow 1$
 - 3: **repeat**
 - 4: $F^i \leftarrow \nabla(J_N \circ s^{i-1})$ // induce inward force at boundary of extraneous structure
 - 5: $L^\dagger L v^i = F^i$ // solve for instantaneous velocity
 - 6: $s^i \leftarrow s^{i-1}(\epsilon v^i)$ // update mapping
 - 7: $i \leftarrow i + 1$
 - 8: **until** convergence
-

Figure 3.8 contains snapshots of this deflation algorithm applied to a CT image of the pelvic region. The large gas pocket present in the image has been deflated, resulting in an image that can be registered using large-deformation diffeomorphic image registration.

In the more general case multiple extraneous structures may appear, and extraneous structures may appear in both images. The same methodology can be applied by combining all extraneous structures in each image into a single structure. After shrinking

⁶See Chapter 2 for more details.

⁷More precisely, when $\operatorname{argmax}_{x \in \Omega} \|s^i(x) - s^{i-1}(x)\|^2$ falls below a specified threshold.




Figure 3.8: **Gas deflation schematic.** (a) Axial slice CT image with large pocket of bowel gas. (b) Zoomed in on the gas pocket. The gas is segmented using simple thresholding. Gas is deflated by a flow induced by the gradient of the binary image. (c) The image after application of the deflation transformation.

transformations s_I and s_J have been computed, correspondence $h_{I \circ s_I, J \circ s_J}$ is estimated between the ‘deflated’ images and final correspondence, $h_{I,J}$, is determined by composing these transformations:

$$h_{I,J} = s_I^{-1} \circ h_{I \circ s_I, J \circ s_J} \circ s_J. \quad (3.4)$$

Algorithm 3.2 summarizes this process.

Algorithm 3.2 Image matching pipeline

Input: Images I and J that each contain extraneous structures

Output: Transformation $h_{I,J}$ that defines a spatial correspondence between I and J in regions without extraneous structures

- 1: Segment extraneous structures within I and J
 - 2: Shrink extraneous structures via transformations s_I and s_J
 - 3: Estimate mapping $h_{I \circ s_I, J \circ s_J}$ between images $I \circ s_I$ and $J \circ s_J$
 - 4: Compose transformations to obtain $h_{I,J} \equiv s_J \circ h \circ s_I^{-1}$
-

Because the I and J are not topologically equivalent images, there is no diffeomorphism that can truly define a structural correspondence between them. However, in the regions that do not contain extraneous structures—e.g., $\Omega \setminus N$ —a diffeomorphic mapping is achieved. This mapping can be used to analyze the deforming structures represented within both images.

3.2.3 Registration pipeline for intra-patient registration

This section describes how extraneous structure deflation is applied in the context of the intra-patient CT image registration problem for ART. The goal is to estimate a transformation $h_{P \rightarrow T}$ that maps points in the planning image, I_P , to corresponding points in a treatment image, I_T . The transformation $h_{P \rightarrow T}$ is estimated using a three step process that combines rigid image registration, bowel gas deflation, and deformable image registration. Each of these stages produces a spatial transformation; $h_{P \rightarrow T}$ is defined as the composition of these transformations.

Rigid registration

First, a rigid transformation that aligns the pelvis in the planning and treatment images is computed. This quantifies the rigid patient setup error. The planning and treatment images are thresholded so that only bone is visible. The region of interest is restricted to the pelvis as it remains fixed while the femurs and spine can rotate or bend. The rigid transformation, r , is estimated using an intensity based gradient descent algorithm (Joshi et al., 2003).

Accommodating bowel gas

In order to remove bowel gas from I_P and I_T , the deflation algorithm is applied to produce deflated images I_{P_d} and I_{T_d} . The following process is applied separately to both the planning and treatment images. A binary segmentation of the gas is created using a simple thresholding operation. This is possible because the contrast between gas and surrounding tissue is high in CT images. This binary segmentation is refined using a morphological close operation, which eliminates small pockets of gas. Next, a deflation transformation s is estimated. As described in Section 3.2.2, s is based on a flow induced by the gradient of the binary image. Points along the gas-tissue border, where the gradient is non-zero, flow in the direction of the gradient. As a result, gas

filled regions collapse—deflating like a balloon.

Deformable image registration

Finally, non-rigid motion is quantified using large-deformation diffeomorphic image registration. A diffeomorphism $h_{\text{def}}: \Omega_{P_d} \rightarrow \Omega_{T_d}$ that defines a voxel-to-voxel correspondence between the two gas-deflated images I_{P_d} and I_{T_d} is generated. The registration is determined by finding the deformation field h_{def} that minimizes the mean squared error between I_{P_d} and the deformed image $I_{T_d} \circ h_{\text{def}}$,

$$D(h) = \int_{x \in \Omega} |I_{P_d}(x) - I_{T_d}(h_{\text{def}}(x))|^2 dx. \quad (3.5)$$

The transformation is constrained to be diffeomorphic by enforcing that it satisfy laws of continuum mechanics derived from visco-elastic fluid modeling (Christensen et al., 1996; Miller et al., 1999).

Composite transformation

Correspondence between the original images I_P and I_T is defined by concatenating the rigid, deflation, and deformable registration transformations, i.e.,

$$h_{P \rightarrow T} = r \circ s_T \circ h_{\text{def}} \circ s_P^{-1}. \quad (3.6)$$

That is, the point x in the planning image corresponds to the point $h_{P \rightarrow T}(x)$ in the treatment image.

Mapping segmentations via spatial transformations

Using the deformation $h_{P \rightarrow T}$, it is possible to carry manual organ segmentations from the planning image to a treatment image, deforming them to match the new image.⁸ This provides an automatic segmentation of the treatment image that is based on the manual segmentation of the planning image. In practice, the automatic segmentations must still be reviewed by a physician, but they need not be edited unless errors are found.

In this work a manual segmentation is represented by a collection of planar contours that each define the location of an anatomical structure within a particular 2D slice—or plane—of a 3D image. These contours are generated by a physician who uses a computer program to draw (using mouse-clicks) these contours, one slice at a time, over a displayed image. Each contour is defined by a collection of ordered, coplanar vertices v^i and the line segments l^i that connect them. That is, the line segment l^i connects vertex v^i with vertex v^{i+1} . The last line segment is a special case that closes the contour; it connects the last vertex with the first vertex.

The transformation $h_{P \rightarrow T}$ is used to move the vertices of the contours from their locations in the planning image to the corresponding points in the treatment image (Figure 3.5). This process does not result in a set of planar contours since vertices will typically be moved out of plane to varying degrees. Therefore, the vertices that define the contours are first converted to a surface model made up of triangles (Figure 3.9) using the Power Crust algorithm (Amenta et al., 2001). The surface is deformed by replacing each vertex v^i with $h_{P \rightarrow T}(v^i)$. Contours are regenerated by slicing the deformed surface model with planes parallel to the xy axis.

⁸Segmentations can also be carried from the daily image to the planning image since $h_{P \rightarrow T}$ is invertible.

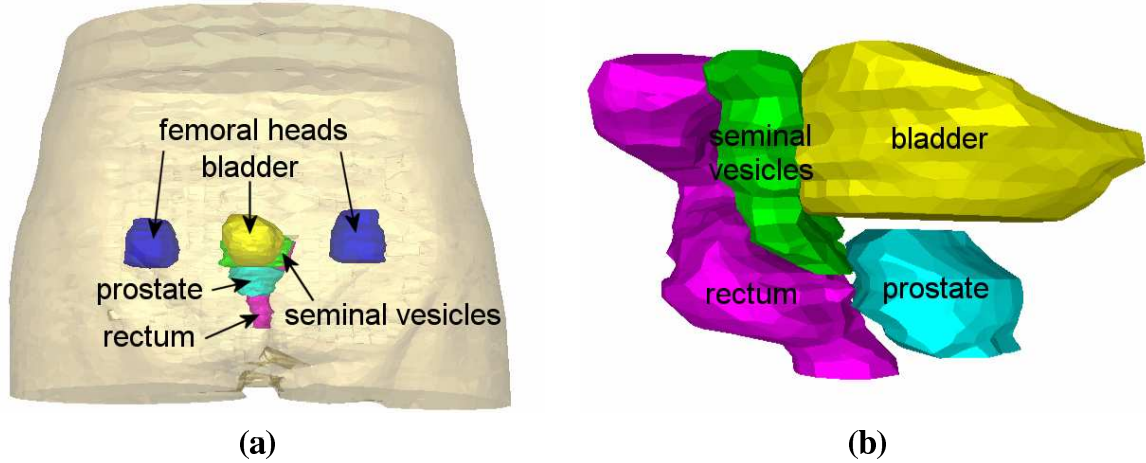


Figure 3.9: **Visualization of organ segmentations.** Panel (a) is an anterior view of a 3D rendering displaying segmentations of the skin, prostate, rectum, bladder, seminal vesicles, and femoral heads. Panel (b) shows a lateral view of the prostate, rectum, and bladder of the same patient. The surfaces are constructed by tiling manually drawn contours.

3.3 Results: Automatic segmentation of the prostate from in-treatment-room CT scans

Figure 3.10 permits a visual assessment of the accuracy of this method. This figure is similar to Figure 3.1 except that the white contours indicate the daily treatment prostate positions *deformed into the space of the planning image* (rather than the actual daily prostate positions). Discrepancies between the deformed segmentations measure not only image registration uncertainty, but also intra-rater variability of the manual, treatment-day segmentations. In this section the accuracy of this segmentation method is quantified, with attention to human variability, by measuring segmentation volume overlap and centroid differences.

The statistical analysis is based on comparing automatically generated segmentations to manual, hand-drawn segmentations. However, there is appreciable variation in manual segmentation, making it unreasonable to choose a particular manual segmentation as definitive. Groups have reported segmentation variation in a number of contexts,




Figure 3.10: **Automatic prostate segmentation results.** The images show manual segmentations of each daily image deformed into the space of the planning image. The close agreement (compare to Figure 3.1) of the deformed treatment-day segmentations with the position of the prostate in the planning images provides evidence for the accuracy of the image registration algorithm along the prostate boundary.

including brain tumors (Leunens et al., 1993), lung cancer (Valley and Mirimanoff, 1993; Ketting et al., 1997), and prostate MR (Zou et al., 2004). Rasch et al. (1999) reported inter-rater variabilities in the segmentation of the prostate in CT and MR images. For CT images, they found overall observer variation of 3.5mm (1 standard deviation) at the apex of the prostate and an overall volume variation of up to 5%.

Given this inter-rater variability, the automatically generated segmentations are assessed by comparing them with segmentations from manual raters. The segmentations from different manual raters are compared in the same way. The accuracy and reliability of the automatic segmentations is judged according to the standard of the measured inter-rater variability.

As of the date of this study (2005), the Department of Radiation Oncology at the University of North Carolina at Chapel Hill had acquired CT scans for a total of 138 treatment days from nine patients. All of these images were manually segmented by at least one rater. However, due to the time-consuming nature of manual segmentation, images from only five of these patients were manually segmented by a second rater. The 65 images from these five patients are used for the analysis in this section.

The experimental setup was as follows. Each daily CT image was collected prior to treatment using the Siemens Primatom scanner pictured in Figure 3.3, with a resolution of $0.098 \times 0.098 \times 0.3$ cm. Each planning image, as well as every treatment image, was manually segmented twice, once by rater *A* and once by rater *B*. For each patient the transformations h_i that deformably align the planning image with the treatment image for each treatment day i were computed using the method described above. Automatic segmentations were generated for each treatment image by applying h_i to segmentations in the planning image. The automatic segmentations are labeled as rater C (for “computer”). C_A and C_B represent treatment image segmentations that were automatically generated by deforming the manual planning image segmentations drawn by raters *A* and *B*, respectively. Therefore, there are a total of four segmentations for each treat-

ment image: two manual segmentations (A and B) and two automatic segmentations (C_A and C_B).

For each patient and for each treatment day, there are six pairwise comparisons that can be made from the set of four segmentations. We report data on five of these comparisons: AB , comparing manual segmentations by rater A against those by rater B; C_AA and C_BB , comparing automatic segmentations with manual segmentations produced by the same rater; and C_AB and C_BA , comparing automatic segmentations with manual segmentations produced by a different rater.

The automatic segmentations are produced by transforming manual planning segmentations produced by either rater A or rater B. Therefore, due to the influence of inter-rater variability, it is expected that—on average—segmentations generated by the same rater (e.g., A and C_A) will be closer in terms of centroid distance and volume overlap than segmentations generated by different raters (e.g., A and C_B).

3.3.1 Implementation: Beamlock

In order to efficiently apply this methodology to patient images, a program, called Beamlock, was developed. Beamlock automatically applies the gas shrinking and image mapping procedure to planning and treatment images. The program is implemented in C++ and requires approximately one hour of processing time per treatment image on a dual processor (Intel Xeon CPU 3.06GHz) computer with approximately 4 gigabytes of RAM.

3.3.2 Centroid analysis

The centroid of the prostate is especially important for radiation treatment planning and therapy because it is the origin, or *isocenter*, for the treatment plan. To measure the accuracy of the automatic segmentations, the centroid of each automatic segmentation is compared with the centroid of the corresponding manual segmentation.




Figure 3.11: **Segmentation accuracy results: prostate centroid differences.** (a) Centroid differences measured in the lateral (X), anterior-posterior (Y), and superior-inferior (Z) directions. The horizontal lines on the box plots represent the lower quartile, median, and upper quartile values. The whiskers indicate the extent of the rest of the data, except that outliers, which fall more than 1.5 times the interquartile range past the ends of the box, are denoted with the ‘+’ symbol. (b) Euclidean distance between segmentation centroids.

First consider the question: Are the centroids of the automatic segmentations systematically shifted with respect to the manual segmentations? Let S_A^i , S_B^i , $S_{C_A}^i$, and $S_{C_B}^i$ denote the prostate segmentations from raters A , B , C_A , and C_B , respectively, for image i . Let $\mathcal{C}(\cdot)$ be a function that returns the centroid (in \mathbb{R}^3) of a segmentation. The distribution of the centroid differences $\mathcal{C}(S_{C_A}^i) - \mathcal{C}(S_A^i)$, $i \in 1, 2, \dots, N$ (and similarly for C_B) can be examined in order to determine whether the centroids of the automatic segmentations are systematically shifted in any particular direction. Likewise, to test for systematic shifts between manual raters A and B , the distribution $\mathcal{C}(S_B^i) - \mathcal{C}(S_A^i)$ is examined. Figure 3.11 (a) shows box-and-whisker plots of these differences for the BA , $C_A A$, and $C_B B$ comparisons. The differences in the lateral (X), anterior-posterior (Y), and superior-inferior (Z) directions are measured separately. Summary statistics are provided in Table 3.1. According to these data there is no significant shift between centroids of the computer generated segmentations and rater A’s manual segmentations

Table 3.1: Summary statistics for centroid difference distributions. The mean, standard deviation, and 95% confidence interval for the mean are reported.

	Lateral (X)			A-P (Y)			Sup-Inf (Z)		
	BA	C_{AA}	C_{BB}	BA	C_{AA}	C_{BB}	BA	C_{AA}	C_{BB}
Mean	0.00	-0.01	0.03	-0.01	0.02	0.12	0.07	0.10	-0.07
STD	0.07	0.07	0.06	0.15	0.13	0.18	0.28	0.18	0.28
95% CI	-0.02	-0.03	0.01	-0.05	-0.02	0.07	0.00	0.05	-0.14
	0.02	0.00	0.04	0.02	0.05	0.16	0.14	0.14	0.00

in the lateral and A-P directions. There *is* a significant shift ($p < 0.001$ for two-tailed t-test) in the Sup-Inf direction of approximately 0.09 cm, which is less than one third of the Sup-Inf image resolution (0.3 cm). For the C_{BB} comparisons there are significant shifts in the lateral and A-P directions of approximately 0.03 cm and 0.12 cm, respectively, which are at or less than the voxel resolution in these dimensions. The comparison between manual raters shows that there is a significant shift in the Sup-Inf direction of approximately 0.07 cm.

In the lateral and Sup-Inf directions, the standard deviation of the manual AB comparisons is as large or larger than the standard deviation of the C_{AA} and C_{BB} comparisons. In the A-P direction, the standard deviation of the C_{BB} comparisons is slightly higher than the manual comparison.

In Figure 3.11 (a) there is a striking difference in the variance of the lateral (X) distributions and the variance for the A-P (Y) and Sup-Inf (Z) distributions. The large difference in physical voxel spacing in the (X) and (Z) directions—0.093 cm for (X) versus 0.3 cm for (Z)—is one important factor. However, the resolution is the same in the (X) and (Y) directions. One explanation is that the motion of the prostate is highly influenced by the expansion and contraction of the bladder and rectum. Because of their relative locations—the bladder is superior to, and the rectum is posterior to the prostate (see Figure 3.9)—these organs are likely to cause more motion of the prostate

Table 3.2: Summary statistics for centroid distance distributions.

	Euclidean Distance				
	AB	C_AA	C_BB	C_AB	C_BA
Mean	0.29	0.21	0.32	0.37	0.35
Median	0.25	0.20	0.32	0.31	0.35
Max	0.72	0.67	1.08	1.08	0.70
STD	0.16	0.13	0.17	0.22	0.15
IQR	0.23	0.21	0.19	0.26	0.24

in the A-P (Y) and Sup-Inf (Z) directions than in the lateral (X) direction.

Next the Euclidean distance measured between segmentation centroids is analyzed. Figure 3.11 (b) shows box-and-whisker plots of these distances. Summary statistics for these data are presented in Table 3.2. As the distributions of these distances are not approximately normal, medians and interquartile ranges are reported along with means and standard deviations.

All of the mean distances are within image resolution. Paired t-tests were used to test for equality of the means of these distributions. The C_AA mean distance is significantly less than the AB mean distance ($p < 0.001$) while there is no significant difference between the C_BB and AB mean distances. As expected, the centroids of the automatically generated segmentations are consistently closer to same-rater manual segmentations than cross-rater manual segmentations.

Based on these data, the automatic segmentation method is comparable to human raters in accuracy for estimating centroids and, as judged by the error bars and standard deviations, at least as reliable. However, there are outliers, with the maximum centroid distance being over 1 cm. For this reason, segmentations should be reviewed by a physician before use in radiation therapy treatment planning.

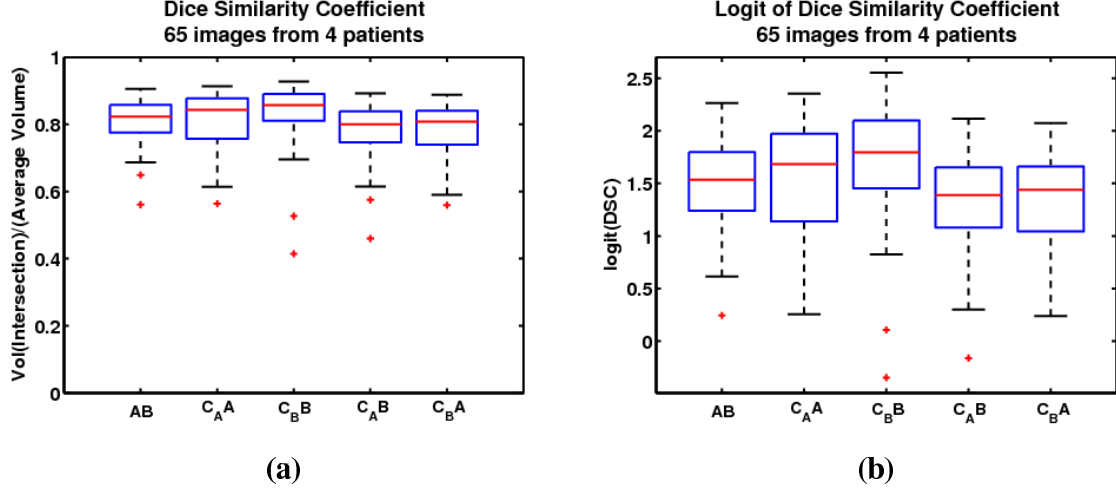


Figure 3.12: **Segmentation accuracy results: prostate Dice similarity coefficients.** Dice Similarity Coefficient (DSC) and logit DSC.

3.3.3 Volume overlap analysis

While the difference in centroids is an important estimate of segmentation alignment, it does not take into account shape or pose changes that may cause the segmentation boundaries to differ. The relative volume overlap of the segmentations is one method for measuring such differences. To measure the degree of volume overlap between the automatic and manual segmentations in this study, we use the Dice Similarity Coefficient (DSC) (Dice, 1945). For two segmentations, S_1 and S_2 , the DSC is defined as the ratio of the volume of their intersection to their average volume:

$$\text{DSC}(S_1, S_2) = \frac{\text{Volume}(S_1 \cap S_2)}{\frac{1}{2} (\text{Volume}(S_1) + \text{Volume}(S_2))} \quad (3.7)$$

The DSC has a value of 1 for perfect agreement and 0 when there is no overlap. The DSC can be derived from the κ statistic for measuring chance-corrected agreement between independent raters (Zijdenbos et al., 1994).

Figure 3.12 (a) shows a box-and-whisker plot of the Dice similarity coefficient for each prostate segmentation comparison. The mean DSCs for the $C_A A$ and $C_B B$ com-

parisons are 0.82 (STD=0.08) and 0.84 (STD=0.08), respectively. The mean DSC for the two manual raters was similar (mean=0.81, STD=0.06). This indicates the overlap of the automatic segmentations is on par with the overlap for the inter-rater manual segmentations.

A similar study, described in Zou et al. (2004), assessed the reliability of manual prostate segmentations in interoperative MR images. They report a mean DSC for manual raters of 0.838. The greater degree of agreement (DSC) with MR images is reasonable since prostate boundaries are more evident in MR images than in CT images, and thus manual raters are likely to segment MR images more reliably than CT images.

Better results for automatic segmentation of the prostate in CT images have also been achieved. A recent study by Pizer et al. (2008) applied principal geodesic analysis to train m-rep shape models of the prostate for prostate segmentation in CT images. In a leave-one-out study involving 80 total images from 5 patients, they achieved a median volume overlap score of 0.94 when m-rep segmentations were compared to manual segmentations. These superb results can at least partially be explained by the strong shape prior that is learned from the principal geodesic analysis. This prior informs the segmentation procedure when boundaries are not evident in image data.

To evaluate the DSC distributions, the logit of the DSC (LDSC) is used

$$\text{LDSC}(S_1, S_2) = \ln \left(\frac{\text{DSC}(S_1, S_2)}{1 - \text{DSC}(S_1, S_2)} \right). \quad (3.8)$$

Agresti (1990) showed that for large sample sizes (in the case of our prostate segmentations, the number of voxels is approximately 20,000), LDSC has a Gaussian distribution. Figure 3.12 (b) shows a box-and-whisker plot of the LDSC values for each comparison.

A paired t-test on the LDSC values was performed in order to test for significant differences between the AB and $C_A A$ or $C_B B$ comparisons. A one-tailed test shows that the DSCs for the $C_B B$ comparisons are significantly ($p < 0.001$) greater than the DSCs

Table 3.3: Summary statistics for the DSC measures.

	AB	C_AA	C_BB	C_AB	C_BA
Mean	0.81	0.82	0.84	0.78	0.78
Median	0.82	0.84	0.86	0.80	0.81
STD	0.06	0.08	0.08	0.08	0.08
IQR	0.08	0.12	0.08	0.09	0.10

Table 3.4: Comparison of automatic segmentation to manual segmenter A via the DSC and LDSC.

	DSC(S_1, S_2)		LDSC(S_1, S_2)	
	Prostate	Bladder	Prostate	Bladder
n	76	20	76	20
Mean	0.801	0.816	1.466	1.576
Median	0.825	0.826	1.554	1.557
STD	0.081	0.078	0.494	0.539
IQR	0.121	0.133	0.804	0.034

for the AB comparisons. No significant difference between the C_AA and AB comparisons were found ($p = 0.12$ for a two-tailed test). Therefore, the automatic segmentations coincide with the manual segmentations at least as well as a second manual rater.

Table 3.4 summarizes the manual-versus-automatic comparison for segmenter A only, for all patients that have been processed. There are fewer bladder segmentations because the bladder was typically not segmented after the first five treatment days.

3.4 Conclusions

This chapter describes a method for estimating mappings between images when extraneous structures are present. The method identifies extraneous structures within images and shrinks them out of an image before image mapping is applied. This method was applied to the clinical problem of estimating prostate motion over the course of treatment in the presence of transient bowel gas.

In the future, it would be interesting to apply this methodology to new clinical

problems. For example, this approach can be applied in a straightforward manner to the registration of brain images in the presence of pathology such as tumors.

Chapter 4

Variational optimization for large-deformation diffeomorphic image averaging

4.1 Introduction

For anatomy represented within medical images, a natural problem is the construction of a representative anatomical configuration for a population. Such a representative is important, for example, when investigating clinical hypotheses related to the shape of structures indicated in brain development or mental disorders. It can also be used to define a common spatial coordinate system so that spatial information from across a population can be accumulated, statistically analyzed, and presented in a single frame of reference.

One common approach for generating representative anatomical images is to simply select an individual image from the population. However, unless there is an *a priori* reason to choose one individual over the rest, this choice will bias the analysis for the particular anatomical configuration presented by the individual.¹ This method is also commonly used to establish spatial correspondence between image coordinate systems:

¹For example, this effect is described in (Rohlfing et al., 2004).

one image is chosen as a reference and all other images are spatially aligned with this reference image. Unfortunately, this is biased as well, since the image-to-image correspondences rely on the choice of the reference image.

A sensible approach for generating an anatomical representative image that is not biased by any particular individual is to generate an anatomical configuration that is in some sense centrally located—in terms of the configuration of anatomy within an image—with respect to the population under study. The mean in statistics captures this notion of “central location,” and the arithmetic mean is easily computed from data that are elements of a vector space. However, the shapes and shape changes evident in anatomy are not well described by vector spaces and, in particular, the naive approach of averaging an image voxel-wise clearly neither produces a realistic anatomical image nor captures the notion of mean *geometric* configuration.

Recently, the notion of Fréchet mean has been used to define mean shapes in non-linear shape spaces that have a metric space structure. For example, Fletcher et al. (2004) extend concepts such as averaging and principal component analysis to manifolds representing anatomical shape variability. This chapter presents a novel method for generating a Fréchet mean anatomical image, using large-deformation diffeomorphic transformations, from a collection of images representing a population. The resulting image serves as a natural average representation for the anatomy contained within the images and is not biased by the choice of any particular image or the ordering of the population images. The Fréchet mean image also provides a natural coordinate system for the collection of images. Because the method uses large-deformation diffeomorphic registration to establish correspondence between the image coordinate systems during the process of forming the mean, spatial data is easily mapped between images within the collection and, in particular, easily accumulated within the mean image coordinate system.

The strategy is to first define a notion of mean that is based on geometric change.

Next, a method for measuring geometric change between images is defined. These ideas are combined to define an algorithm for computing a mean image. The mean image is generated automatically from the input images using an iterative procedure of continuous joint alignment: the input images iteratively deform toward the evolving mean.

In the rest of this chapter, these methods are described and examples are presented for synthetic data. Related work is summarized in Section 4.4.

This work was developed in collaboration with Dr. Sarang Joshi. It has been published, with applications to medical image data, in Davis et al. (2004); Joshi et al. (2004); Xu et al. (2006); Goodlett et al. (2006); Gerig et al. (2006). Drs. Peter Lorenzen and Sarang Joshi led the effort to extend this methodology to generate an elegant multi-modal image set registration and averaging framework in Lorenzen et al. (2004a,b, 2006).

4.2 Methods

In this section I describe a method for computing a mean image that takes geometric change into account. The mean is computed automatically from a cohort of images. Section 4.2.1 reviews the definition of the arithmetic mean in Hilbert spaces. Section 4.2.2 describes why this definition is not applicable in the presence of large-scale geometric change, and it presents a new definition of mean—based on distances—that *is* applicable. Section 4.2.3 describes an algorithm that combines the diffeomorphisms described in Chapter 2 with the notion of Fréchet mean to generate mean images. Section 4.2.4 describes several properties of this algorithm. Finally, Section 4.2.5 describes implementation details.

4.2.1 Arithmetic mean in Hilbert spaces

A *Hilbert space* is a vector space H together with an inner product $\langle \cdot, \cdot \rangle$ such that H is complete relative to the metric $d(x, y) = \|x - y\| = \sqrt{\langle x - y, x - y \rangle}$ on H . Given N elements x^i of a Hilbert space H , the arithmetic mean is defined as the minimizer of the sum-of-squared-distances

$$\mu = \operatorname{argmin}_{x \in H} \frac{1}{N} \sum_{i=1}^N \langle x - x^i, x - x^i \rangle. \quad (4.1)$$

This definition depends on the vector space structure of H . In particular, it relies on well defined subtraction operations. In Euclidean space, the solution to Equation (4.1) is the familiar arithmetic mean

$$\mu = \frac{1}{N} \sum_{i=1}^N x^i. \quad (4.2)$$

Consider a collection of N 3D images $\{I^i\}_{i=1}^n$. Each image $I \in \mathcal{I}$ can be formally defined as an L^2 function from $\Omega \equiv [0, 1] \times [0, 1] \times [0, 1] \subset \mathbb{R}$ to the reals. In this case the mean image is defined using the L^2 inner product norm

$$\mu = \operatorname{argmin}_{I \in \mathcal{I}} \frac{1}{N} \sum_{i=1}^N \int_{\Omega} (I(x) - I^i(x))^2 dx. \quad (4.3)$$

In this case the mean image is

$$\mu(x) = \frac{1}{N} \sum_{i=1}^N I^i(x). \quad (4.4)$$

4.2.2 Fréchet mean image

We cannot rely on the frequently used vector space structure² of images themselves to generate satisfactory mean images. Although images can be added voxel-wise, the result is a loss of any identification with the geometric configuration of the image contents. Also, the L^2 norm does not take into account any notion of *geometric differences*. On the other hand, this thesis is based on the premise that geometric differences can be accounted for by transformations of the underlying image coordinates.

In this work, geometric change is represented as the action of the group of diffeomorphisms on images. This allows for the representation of the large geometric variability evident in anatomical images. Let $\text{Diff}_V(\Omega)$ be the group of diffeomorphisms that are isotopic to the identity (see Chapter 2). Each element $\phi : \Omega \rightarrow \Omega$ in $\text{Diff}_V(\Omega)$ deforms an image according to the following rule

$$\phi I = I_\phi = I \circ \phi^{-1}. \quad (4.5)$$

In the group of diffeomorphisms, *composition* is defined, but addition and subtraction are not. Therefore, Equation (4.1) is not defined for averaging diffeomorphisms. However, the idea of expected value of real random variables can be generalized to manifold-valued random variables via Fréchet expectation (Fréchet, 1948; Karcher, 1977).

Let f be a probability density on a Riemannian manifold \mathcal{M} . The Fréchet expectation is defined as

$$\mathbb{E}_f[p] \equiv \operatorname{argmin}_{q \in \mathcal{M}} \int_{\mathcal{M}} d(q, p)^2 f(p) dp \quad (4.6)$$

where d is the metric on the manifold \mathcal{M} . This definition is motivated by a minimum variance characterization of the mean, where variance is defined in terms of the metric.

²E.g, $L^2(\Omega)$

Fréchet expectation might not be unique (Karcher, 1977). Using this definition, an empirical estimate of the Fréchet mean, given a collection of observations $\{p_i, i = 1 \dots N\}$ on a manifold \mathcal{M} , is defined by

$$\mu = \operatorname{argmin}_{q \in \mathcal{M}} \frac{1}{N} \sum_i^N d(q, p^i)^2. \quad (4.7)$$

Figure 4.1 (a) illustrates a Fréchet mean on a curved manifold. This equation is similar to Equation (4.1) except that vector space structure is not required for the Fréchet mean. Indeed, plugging the Euclidean distance into Equation (4.7) leads directly to the arithmetic mean (Equation (4.2)).

The sum-of-squared distances in Equation (4.7) can be interpreted as a measure of variance for the observations with respect to the metric d , that is,

$$\operatorname{var} = \min_{q \in \mathcal{M}} \frac{1}{N} \sum_i^N d(q, p^i)^2. \quad (4.8)$$

In order to apply the Fréchet mean to images (see Figure 4.1 (b)), a satisfactory notion of distance for images must be established; in this dissertation the deformation-based metric described in Chapter 2 is used. That is,

$$d^2(I^F, I^M) = \operatorname{argmin}_{v: \dot{\phi}_t = v_t(\phi_t)} \int_0^1 \|v_t\|_v^2 dt + \frac{1}{\sigma^2} \|I_\phi^M - I^F\|_{L^2}^2. \quad (4.9)$$

Recall that the first term defines a metric on the space of diffeomorphisms that are generated by integrating velocity fields v . These diffeomorphisms are used to deform a moving image I^M to match the a fixed image I^F . The second term penalizes residual image mismatch. The free parameter σ governs the relative weight of these terms. See Section 2.3 for more details.

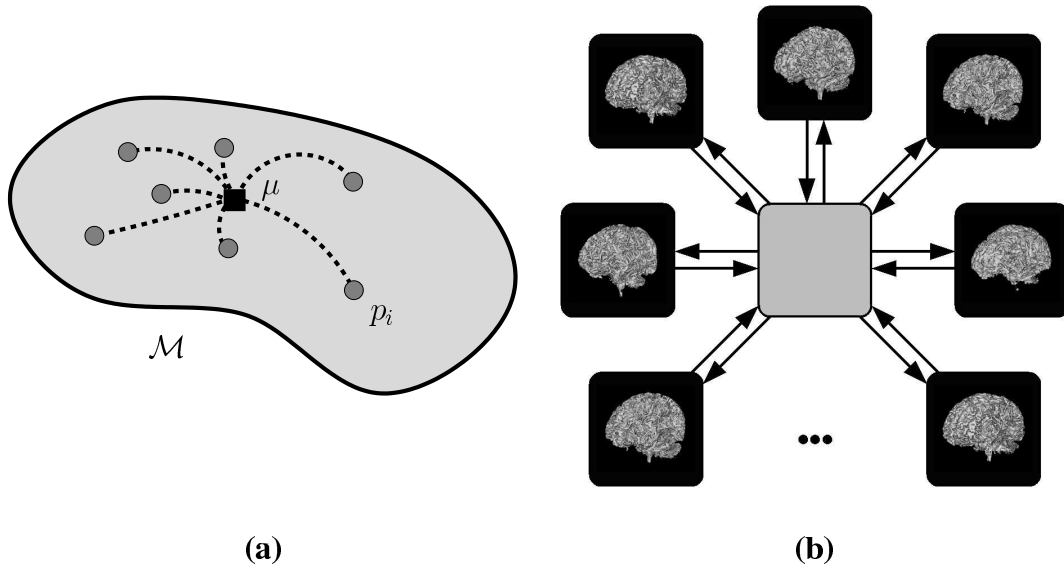


Figure 4.1: **Fréchet mean image schematic.** (a) Observed points p^i are represented by filled circles on the manifold \mathcal{M} . The Fréchet mean (filled square) is the point μ on the manifold that minimizes the sum of squared distances to the observations. Distances are measured along the manifold. (b) Fréchet mean image construction framework for 3D brain images. The mean image minimizes the sum of squared deformation distance (Equation (4.9)) required to match input images. The algorithm presented in this chapter automatically generates this mean image, as well as the deformations ϕ^i that map the mean image onto each input image I^i .

4.2.3 Computing Fréchet mean images with LDDMM and greedy solutions

Recall that given N points p^i on a Riemannian manifold \mathcal{M} , the Fréchet mean is defined as

$$\mu = \operatorname{argmin}_{q \in \mathcal{M}} \sum_{i=1}^N d(q, p^i)^2. \quad (4.10)$$

Given N images $I^i \in \mathcal{I}$, the goal is to compute the mean image \hat{I} that minimizes the sum of squared distances to the given images according to the metric defined in the previous section (Equation (4.9)):

$$\hat{I} = \operatorname{argmin}_{I \in \mathcal{I}} \sum_{i=1}^N d(I, I^i)^2. \quad (4.11)$$

Figure 4.1 (b) illustrates this idea.

This optimization problem (4.11) depends on deformations that deform the mean image \hat{I} to match each input image. Equivalently, it depends on deformations ϕ^i that deform each input image I^i to match the mean image \hat{I} .³ Plugging in the metric from Equation (4.9) results in

$$\hat{I}, \hat{\phi}^i = \operatorname{argmin}_{I, \phi^i \in \mathcal{I} \times \operatorname{Diff}_V(\Omega)^N} \sum_{i=1}^N \left[\int_0^1 \|v_t^i\|_V^2 dt + \frac{1}{\sigma^2} \|I - I_{\phi^i}^i\|_{L^2}^2 \right] \quad (4.12)$$

subject to

$$\phi_0^i = \operatorname{Id} \quad \phi_1^i = \phi^i \quad \dot{\phi}_t^i = v_t^i \circ \phi_t^i \quad \phi^i(x) = x + \int_0^1 v_t^i(\phi_t^i(x)) dt. \quad (4.13)$$

These equations express the following intuitive idea: a population of images can be

³Recall that $d_{\operatorname{Diff}_V(\Omega)}(\operatorname{Id}_{\operatorname{Diff}_V(\Omega)}, \phi) = d_{\operatorname{Diff}_V(\Omega)}(\operatorname{Id}_{\operatorname{Diff}_V(\Omega)}, \phi^{-1})$.

represented by the image that is centrally located, according to the metric d , among the observations. Since d is based on a metric for deformations, this implies that \hat{I} is the image that requires the “least amount of deformation” required to map onto the input images.

Equation (4.12) can be simplified by noticing that for fixed transformations the optimal mean image is

$$\hat{I} = \frac{1}{N} \sum_{i=1}^N I_{\phi^i}^i. \quad (4.14)$$

That is, \hat{I} is the voxel-wise arithmetic mean of the deformed images $I_{\phi^i}^i$. Note that this method for computing \hat{I} is determined by the L^2 image dissimilarity metric (recall Section 4.2.1). Other metrics would imply different methods for computing \hat{I} .

Combining equations (4.12) and (4.14) results in an optimization problem over the deformations ϕ^i alone:

$$\hat{\phi}^i = \underset{\phi^i \in \text{Diff}_V(\Omega)^N}{\text{argmin}} \sum_{i=1}^N \left[\int_0^1 \|v_t^i\|_V^2 dt + \frac{1}{\sigma^2} \left\| \frac{1}{N} \sum_{j=1}^N I_{\phi^j}^j - I_{\phi^i}^i \right\|_{L^2}^2 \right] \quad (4.15)$$

subject to

$$\phi_0^i = \text{Id} \quad \phi_1^i = \phi^i \quad \dot{\phi}_t^i = v_t^i \circ \phi_t^i \quad \phi^i(x) = x + \int_0^1 v_t^i(\phi_t^i(x)) dt. \quad (4.16)$$

For a fixed mean image \hat{I} , this equation can be written as N independent optimization problems

$$\hat{\phi}^i = \underset{\phi \in \text{Diff}_V(\Omega)}{\text{argmin}} \int_0^1 \|v_t^i\|_V^2 dt + \frac{1}{\sigma^2} \|\hat{I} - I_{\phi^i}^i\|_{L^2}^2 \quad (4.17)$$

subject to

$$\phi_0^i = \text{Id} \quad \phi_1^i = \phi^i \quad \dot{\phi}_t^i = v_t^i \circ \phi_t^i \quad \phi^i(x) = x + \int_0^1 v_t^i(\phi_t^i(x)) dt. \quad (4.18)$$

Equation (4.15) can be solved iteratively. First, the deformations are initialized to the identity transformation. Then the following two steps are iterated until convergence: (1) the deformations are fixed and the mean image is updated according to (4.14); (2) the mean image is fixed and the N deformations are updated by taking *incremental* steps along the gradient of the energy functionals in Equation (4.17). These gradients are given by

$$\nabla_{v_t^i} E_t = 2v_t^i - K \left(\frac{2}{\sigma^2} |D\phi_{t,1}^i| \nabla I_{\phi_{0,t}^i}^i (I_{\phi_{0,t}^i}^i - \hat{I}_{\phi_{1,t}^i}^i) \right). \quad (4.19)$$

This equation follows immediately from the derivation of the gradient of the image-to-image registration functional (Equation (2.19)) in Beg et al. (2005), which is derived in Appendix B of this document; the difference is that in this case the image I^i flows toward the evolving mean image and not toward an observed image. Convergence is achieved when incremental changes in the energy functional (Equation (4.15)) become negligible.⁴ This procedure is summarized in Algorithm 4.1.

Greedy solution

For the greedy solution, Equation (4.15) is solved using a similar iterative algorithm. First, the deformations are initialized to the identity transformation. Then the following two steps are iterated until convergence: (1) the deformations are fixed and the mean image is updated according to (4.14); (2) the mean image is fixed and the N deformations are updated by taking *incremental* steps according to the velocities v^i . Each v^i gives the locally optimal (greedy) solution given that the current estimate of the deformation

⁴Of course, a maximum number of iterations could also be used.

Algorithm 4.1 Intrinsic mean image: LDDMM solution

Input: A collection of N images: $\{I^i\}_{i=1}^N$

Output: A Fréchet mean image \hat{I} and a collection of diffeomorphic transformations $\{\phi^i\}_{i=1}^N$ such that ϕ^i maps I^i onto \hat{I}

```
1: // Initialize deformations with identity
2: for  $i = 1 : N$  do
3:    $\phi^i \leftarrow \text{Id}_{\text{Diff}_V(\Omega)}$ 
4: end for
5: // Perform gradient descent optimization
6: repeat
7:    $\hat{I} = \frac{1}{N} \sum_{i=1}^N I_{\phi^i}^i$  // Update mean image
8:   for  $i = 1 : N$  do
9:     // Compute gradient of functional at each discretized timepoint  $s$ 
10:    for  $j = 1 : T$  do
11:       $s \leftarrow \frac{j-1}{T}$ 
12:       $\nabla_{v_t^i} E_t \leftarrow 2v_t^i - K \left( \frac{2}{\sigma^2} |D\phi_{s,1}^i| \nabla I_{\phi_{0,s}^i}^i (I_{\phi_{0,s}^i}^i - \hat{I}_{\phi_{1,s}^i}) \right)$ 
13:    end for
14:     $v^i \leftarrow v^i - \epsilon \nabla_{v^i} E$  // Step along gradient
15:  end for
16: until convergence
```

is fixed

$$v^i = K \left(\frac{2}{\sigma^2} \nabla I_{\phi_{0,1}^i}^i (I_{\phi_{0,1}^i}^i - \hat{I}) \right) \quad (4.20)$$

Note that in contrast to the LDDMM case the velocity is only computed at time $t = 1$. The convergence criterion are identical to the LDDMM solution. This procedure is summarized in Algorithm 4.2.

4.2.4 Fréchet mean image algorithm properties

Groupwise registration

Groupwise registration is the process of mapping a collection of images into a common coordinate system. When transformations are invertible, they can be composed to map every input image onto every other input image. The algorithm described in this chapter

Algorithm 4.2 Intrinsic mean image: Greedy solution

Input: A collection of N images: $\{I^i\}_{i=1}^N$

Output: A Fréchet mean image \hat{I} and a collection of diffeomorphic transformations $\{\phi^i\}_{i=1}^N$ such that ϕ^i maps I^i onto \hat{I}

```
1: // Initialize deformations with identity
2: for  $i = 1$  : number of images do
3:    $\phi^i \leftarrow \text{Id}_{\text{Diff}_V(\Omega)}$ 
4: end for
5: repeat
6:    $\hat{I} = \frac{1}{M} \sum_{i=1}^M I_{\phi^i}^i$  // Update mean image
7:   for  $i = 1$  : number of images do
8:      $v_{\text{inc}} \leftarrow K \left( \frac{2}{\sigma^2} \nabla I_{\phi_{0,1}^i}^i (I_{\phi_{0,1}^i}^i - \hat{I}) \right)$  // Compute locally optimal velocity
9:      $\phi^i \leftarrow \text{Exp}_{\phi^i}(\epsilon v_{\text{inc}})$  // Step in direction of incremental velocity
10:  end for
11: until convergence
```

produces a groupwise registration method that is inherently unbiased in the sense that the results do not depend on the ordering of the input images or the arbitrary selection of a reference coordinate system.

In the solution to Equation (4.15), the transformations ϕ^i align the input images I^i in the common space Ω of the mean image. The solution to this minimization problem is independent of the ordering of the N images. Since each ϕ^i is a diffeomorphism, its inverse exists and can be calculated. Transformations mapping one image to another are defined by the composition rule

$$\phi_{i,j} : \Omega^i \rightarrow \Omega^j \equiv (\phi^j)^{-1} \circ \phi^i. \quad (4.21)$$

Inverse consistent registration

When this Fréchet mean image framework is applied to two images, the result is an inherently inverse consistent image registration algorithm—no correction term or penalty for consistency is required.

A registration framework is inverse consistent if image ordering does not affect the registration result. Many image registration algorithms are not inverse consistent be-

cause their image dissimilarity metrics are computed in the coordinate system of a single input image; the choice of such a reference image biases the result of the registration. Inverse consistent registration is desired when there is no *a priori* reason to choose one image over another as a reference image. Previous work (e.g., He and Christensen (2003); Christensen and Johnson (2001); Johnson and Christensen (2002)) has introduced methods for computing approximate inverse consistent registrations by applying inverse consistency constraints on intermediate incremental transformations.

The framework presented in this chapter leads to an inherently inverse consistent image registration algorithm. For two images I^1 and I^2 , Equation (4.15) becomes

$$\begin{aligned} \{\hat{\phi}^1, \hat{\phi}^2\} = \operatorname{argmin}_{\phi_1, \phi_2 \in \operatorname{Diff}_V(\Omega)^2} & \frac{1}{2} |I_{\phi^1}^1 - I_{\phi^2}^2|_{L^2}^2 \\ & + \int_0^1 \|v_t^1\|_V^2 dt + \int_0^1 \|v_t^2\|_V^2 dt. \end{aligned} \quad (4.22)$$

The transformations $\hat{\phi}^1$ and $\hat{\phi}^2$ map I^1 and I^2 onto the mean image. Composition is used to create the image-to-image mappings $\hat{\phi}_{1,2} = (\hat{\phi}^2)^{-1} \circ \hat{\phi}^1$ and $\hat{\phi}_{2,1} = (\hat{\phi}^1)^{-1} \circ \hat{\phi}^2$. This method is inverse consistent since $\hat{\phi}_{1,2} \circ \hat{\phi}_{2,1} = \hat{\phi}_{2,1} \circ \hat{\phi}_{1,2} = \operatorname{Id}_{\operatorname{Diff}_V(\Omega)}$.

4.2.5 Implementation notes

A multi-resolution approach is used to compute a Fréchet mean image at progressively higher resolutions. At each level the deformation is initialized by the results at the next coarsest scale level. This strategy has the dual benefits of (a) addressing the large scale shape changes first and (b) speeding algorithm convergence.

In this work, two differential operators are used in the Riemannian metric to measure squared distance on the space of diffeomorphisms. The first is motivated by the Navier-Stokes equations for a compressible viscous fluid with a very low Reynolds number

$$L = -\mu \nabla^2 - (\mu + \lambda) \nabla(\nabla \cdot) + \gamma, \quad (4.23)$$

where μ and λ are viscosity coefficients and γ ensures that L is invertible. This type of operator was first used for large-deformation image registration by Christensen et al. (1996). The second operator used is

$$L = -\mu\nabla^2 + \gamma, \quad (4.24)$$

which is used by Beg et al. (2005).

In order to solve Equation (4.19), the Green’s function of L must be applied. As in Christensen et al. (1996) and Beg et al. (2005), it is applied in a computationally efficient manner by inverting the operator L in the Fourier domain. Details are provided in Appendix D.

The dominating computation at each iteration is a Fast Fourier Transform. The order of the algorithm is $MNTn \log n$ where M is the number of iterations, N is the number of images, T is the number of discretized time steps, and n is the number of voxels along the largest dimension of the image. Therefore, the complexity grows linearly with the number of observations, making this algorithm suitable for application to large data sets. The greedy algorithm is more efficient than the LDDMM algorithm since it only updates one velocity field per iteration (that is, $T = 1$ for the greedy algorithm as opposed to T on the order of 20 for the LDDMM solution).

4.3 Algorithm demonstration with synthetic data


This section demonstrates the Fréchet mean image algorithm for 2D synthetic data. For each experiment a cohort of 256×256 images is generated from a known, underlying geometric model. Two mean images are generated for each cohort: a voxel-wise mean image and a Fréchet mean image (using the LDDMM solution Algorithm 4.1).

4.3.1 Ovals

In this experiment the cohorts consist of ovals in the image foreground that are either rotated (Figure 4.2 (a)) or deformed by bending (Figure 4.2 (b)). In the voxel-wise mean images (panels (c) and (d)) the geometric variation of these cohorts is evidenced by blurry edges: the average *shape* is not captured. In contrast, the Fréchet mean images (panels (e) and (f)) contain crisp shapes that match at least one notion of a “mean shape”.


For the “rotated ovals” experiment, the result is a disk—the long axis of each input oval has been deformed toward the center while the short axis is deformed slightly out. On the one hand, it can be argued that the mean shape *should* be something different, such as an ellipse with an average rotational pose. However, the diffeomorphic transformations used here are not intended to account for global pose. Instead, they are intended to account for residual deformations that occur *after* accounting for pose. Therefore, if the collection of input images is considered as a disk template that is deformed (elongated) to create each image then this result—the mean image is the initial disk template—is reasonable.

For the “deformed ovals” experiment, the results have a similar interpretation. The corners of the deformed ovals are midway between the corners of the fully deformed and the undeformed ovals. When the Fréchet mean image is considered as an initial template that is deformed to generate the input images, this result is reasonable. However, as in the previous case, there are other reasonable interpretations of mean shape for this example. The mean shape might be an oval (long axis horizontal) with no deformation whatsoever; it has been shown that Fréchet mean m-reps give this result when applied to this example (Merck et al., 2008).




(a)

(b)



(c)

(d)



(e)

(f)

Figure 4.2: **Synthetic data experiment: Fréchet mean ellipse images.** (a)-(b) Oval image samples of size 16. Each panel shows a tiled sample of individual 256×256 images (scaled down for display). Each image contains an ellipse that is rotated (a) or deformed (b). (c)-(d) Pixel-wise means for each sample. (e)-(f) Fréchet mean images computed from each sample.

4.3.2 Bullseyes

The bullseye experiments are shown in Figure 4.3. Three cohorts of images are shown: containing 4, 16, and 32 images, respectively. Each input image contains three disks—inner, middle, and outer—with radii determined by the functions r_{inner} , r_{middle} , and r_{outer} . These functions allow a wide variation in bullseye sizes and a random, nonlinear relationship between the disk radii while explicitly providing the constraint $r_{\text{inner}} < r_{\text{middle}} < r_{\text{outer}}$. They are

$$\begin{aligned} r_{\text{outer}} &= r_{\min} + \epsilon_{\text{outer}}(r_{\max} - r_{\min}) \\ r_{\text{middle}} &= r_{\text{outer}} \frac{3}{5} + \epsilon_{\text{middle}} r_{\text{outer}} \frac{1}{5} \\ r_{\text{inner}} &= r_{\text{outer}} \frac{1}{5} + \epsilon_{\text{inner}} r_{\text{outer}} \frac{1}{5} \end{aligned} \tag{4.25}$$

where each ϵ is a uniform distribution on $[0, 1]$. The outer radius is bounded by $r_{\min} = \frac{1}{7}$ th and $r_{\max} = \frac{1}{3}$ rd of the total image width. Once the disk geometries are fixed, independent, identically distributed Gaussian noise is added to the image intensities.

For each cohort, two means are computed: an *extrinsic*, pixel-wise mean (middle row of figure) and an *intrinsic*, Fréchet mean (bottom row) using the algorithm presented in this chapter (Algorithm 4.1). It is clear that the pixel-wise mean image does not maintain the bullseye geometric configuration. The large variation in sizes of the input bullseyes leads to a number of extra, blurry bands. The Fréchet mean images, on the other hand, do maintain the bullseye geometry. The red curves overlaid on each of these images represent the geometric mean of the known radii from the input images for each corresponding cohort. The close agreement of these curves with the mean images suggests that an intuitive, geometric notion of mean is captured by this algorithm.




Figure 4.3: **Synthetic data experiment: Fréchet means of bullseye images.** (a)-(c) Image samples of size 4, 16, and 32. Each panel shows a tiled sample of individual 256×256 bullseye images (scaled down for display). Outer band radii vary randomly between images; inner band radii vary randomly within each image. (d)-(f) Pixel-wise means for each sample. These images do not capture a geometric notion of a mean bullseye—the three-band structure for the bullseyes is not maintained. (g)-(i) Fréchet mean images computed from each sample. The banded structure is maintained; each band is approximately aligned with the geometric mean radii for each sample (identified by red overlays).

4.4 Related work

For low dimensional transformation groups, the Procrustes method (Goodall, 1991) has been used to align shapes. In the small deformation, high-dimensional setting, Miller et al. (1997) show how to construct average images by averaging transformation maps.

Fletcher et al. (2004) describe a method of applying principal component analysis (PCA) in curved Riemannian metric spaces in order to study anatomical shape variability. This work is applied to shape analysis for populations of m-reps (Fletcher et al., 2004; Pizer et al., 2003), which form a symmetric space, and diffusion tensors (Fletcher and Joshi, 2007), which form the group of symmetric positive definite matrices. Much of the developments in this chapter are built upon and continue this work.

Ramsay and Li (1998) define a method for the alignment of curves through re-parameterization of the curves (i.e., deformation of the “x” axis). They minimize the sum of squared distances from the deforming curves to a to-be-determined mean curve while penalizing deformations with high curvature. This method is the same as the algorithm presented in this chapter except that they use a different class of transformations, their deformation metrics differ, and their work is limited to a single dimension.

An information-theoretic approach to the large-deformation groupwise image registration problem is used in Twining et al. (2005); Cootes et al. (2004) and Davies et al. (2002). In this work, b-splines are used to align images into a common coordinate system. The b-spline parameters are selected in order to minimize the total description length for the collection of images. This description length is determined by the b-spline parameters themselves as well as the residual image differences after the mean image is deformed to each individual image.

Recently Rohlfing and Maurer (2007) presented a method for shape-based averaging of images that uses distance maps. An important difference between this method and groupwise registration methods is that no transformation is used to align images. Each voxel of the average image is the label (or intensity value) that is nearest to the voxels

spatial position (according to the distance maps) in all of the input images.

Craene et al. (2004) present a method for the joint alignment of label-maps. Their algorithm iteratively (1) updates the atlas labels based on the estimates of the current transformations and (2) updates the transformations in order to align each individual segmentation with the current mean segmentation. They use rigid, scaling, and non-rigid transformations based on finite element modeling.

Rueckert et al. (2003) present a method they call *statistical deformation models* that applies linear statistical methods to compute average anatomical images and modes of anatomical variation. They perform PCA on the b-spline control points of free-form deformations on images. For a collection of images, this method is used to compute a mean image and describe eigenmodes of variation by perturbing the control points using the eigenvectors of the PCA covariance matrix.

Charpiat et al. (2005) describe a method for computing a mean image based on continuously deforming each image toward an evolving mean image. They argue that the use of the mean image leads the image to be trapped in local optima. Instead they factor it out of the optimization: each image is incrementally deformed so that it best matches all of the other deformed images. At the end of this iterative process, the mean image can be recovered.

Learned-Miller (2006) describes a “congealing” method for continuous joint alignment of 2D images using affine transformations. Input images are incrementally aligned in a common coordinate space so that the integrated pixel-wise entropy of the aligned images is minimized. This idea was applied for the groupwise registration of 3D medical images in Zöllei et al. (2005).

A number of authors have proposed methods that specifically use diffeomorphic transformations; these methods are summarized below.

Lorenzen et al. (2006) extended the methodology presented in this chapter to generate a multi-modal image set registration and averaging framework. This allows com-

plementary information from different modalities (e.g., T1 and T2 MR images) to inform the averaging process. Spatial tissue probability maps are generated from each multi-modal image set. These probability maps are aligned using the Kullback-Leibler divergence as an image dissimilarity metric.

Avants and Gee (2004) use a tangent-space (velocity field) averaging approach in the large-deformation diffeomorphic setting. In their work an initial reference image flows toward the (unknown) mean in an iterative algorithm. The flow is induced by the velocity fields that are the average of the velocity fields used to deform the current mean to each image observation.

Vaillant et al. (2004) demonstrate principal component analysis (PCA) for sets of landmarks in the large-deformation diffeomorphic setting. This work is based on conservation of momentum in the tangent space.

Glaunès et al. (2004) describe a method for computing the average of unlabeled point sets in the large-deformation diffeomorphic setting. Their framework builds upon the work in this chapter (published in (Joshi et al., 2004)): the difference is that instead of an image match criterion, a metric on currents is used to measure the quality of the point-set match.

4.5 Conclusion

This chapter presents a new concept for unbiased construction of average anatomical images that is based on the Fréchet mean in metric spaces. This approach results in an algorithm where population images simultaneously deformation into a new average image that evolves iteratively. This technique avoids the systematic bias introduced by selecting a template but also the combinatorial problem of deformation of a large number of datasets into each new subject. Additionally, the new technique produces a population average image.

The method was demonstrated by computing Fréchet mean images from several collections of synthetic 2D images. In the next chapter, this method is extended in order to define a *regression* methodology for images.

Chapter 5

Manifold kernel regression

5.1 Introduction

An important area of medical image analysis is the development of methods for automated and computer-assisted assessment of anatomical change over time. For example, the analysis of structural brain change over time is important for understanding healthy aging. These methods also provide markers for understanding disease progression.

A number of longitudinal growth models have been developed to provide this type of analysis to time-series imagery of a single subject (e.g., Beg (2003); Clatz et al. (2005); Miller (2004); Thompson et al. (2000); Xu et al. (2008)). While these methods provide important results, their use is limited by their reliance on longitudinal data, which can be impractical to obtain for many medical studies. Also, while these methods allow for the study of an *individual's* anatomy over time, they do not apply when the *average* growth for a population is of interest.

Random design data sets, which contain anatomical data from many different individuals, provide a rich environment for addressing these problems. However, in order to detect time-related trends in such data, two distinct aspects of anatomical variation must be separated: individual variation and time effect. For measurements that naturally form Euclidean vector spaces, this separation can be achieved by *regressing* a

representative value over time from the data.

For example, in Figure 5.1 kernel regression is applied to measurements reported in a study by Mortamet et al. (2005) on the effect of aging on gray matter and ventricle volume in the brain. The regression curves demonstrate the average volume, as a function of subject age, of these structures. These trends—on average there is a loss of gray matter and expansion of the ventricles—have been widely reported in the medical literature on aging (Guttmann et al., 1998; Matsumae et al., 1996; Mortamet et al., 2005). While volume-based regression analysis is important, it does not provide any information about the detailed *shape* changes that occur in the brain, on average, as a function of age. This has motivated us to study regression of shapes.

Recent work has suggested that representing the geometry of shapes in *flat Euclidean vector spaces* limits our ability to represent natural variability in populations (Fletcher et al., 2004; Grenander and Miller, 1998; Miller, 2004). For example, Figure 5.2 demonstrates the amazing *nonlinear* variability in brain shape among a population of healthy adults. The analysis of transformation groups that describe shape change are essential to understanding this shape variability. These groups vary in dimensionality from simple rigid rotations to the infinite-dimensional group of diffeomorphisms (Miller and Younes, 2001). These groups are not generally vector spaces and are instead naturally represented as manifolds.

Hutton et al. (2003) apply kernel regression to estimate average growth trajectories within a linear PCA shape space. They apply this method to generate average face shapes as a function of age from a database of triangulated surface models of faces.

A number of authors have contributed to the field of statistical analysis on manifolds (see Pennec (2006) for a detailed history). Early work on manifold statistics includes directional statistics (Bingham, 1974; Jupp and Mardia, 1989) and statistics of point-set shape spaces (Kendall, 1984; Le and Kendall, 1993). The large sample properties of sample means on manifolds are developed in Bhattacharya and Patrangenaru (2002,




Figure 5.1: **Illustration of univariate kernel regression:** the effect of aging on gray matter (a) and ventricle volume (b) in the brain. Circles represent volume measurements relative to total brain volume. Kernel regression is used to estimate the relationship between subject age and structure volume (filled lines).

2003). Jupp and Kent (1987) describe a method of regression of spherical data that ‘unwraps’ the data onto a tangent plane, where standard curve fitting methods can be applied. Chung et al. (2005) and Spira et al. (2007) apply smoothing to scalar data that is defined on a manifold. Nilsson et al. (2007) apply regression to scalar functions on manifolds. In Fletcher et al. (2004); Joshi et al. (2004); Pennec (2006), statistical concepts such as averaging and principal components analysis were extended to manifolds representing anatomical shape variability. Many of the ideas are based on the method of averaging in metric spaces proposed by Fréchet (1948).

In this chapter the notion of Fréchet expectation is used to generalize regression to manifold-valued data. In particular, this methodology is applied to study geometric change, as a function of a predictor variable t such as time, in a random design database of images. This method generalizes Nadaraya-Watson kernel regression, using a weighted Fréchet mean for images, in order to compute representative images of this population for any value of t . To determine the local geometric change in the population with respect to t , we extend a dynamic growth model (Miller, 2004) to this population-representative image. The methodology is demonstrated on a database of synthetic images.




Figure 5.2: **Brain image database.** To demonstrate the extent of natural brain shape variability, a mid-axial image slice is presented for a sample of healthy subjects. The images are arranged in order of increasing subject age from 30 (top left) to 60 (bottom right). Because of the complexity of the shapes and the high level of natural shape variability, it is extremely difficult to visually discern any patterns within these data.

This represents joint work with Drs. Sarang Joshi, Tom Fletcher, and Elizabeth Bullitt. Versions of this chapter were published in (Davis et al., 2007) and (Davis et al., 2008).

5.2 Methods

5.2.1 Review of univariate kernel regression

Univariate kernel regression is a non-parametric method used to estimate the relationship, on average, between an independent random variable T and a dependent random variable Y (Hardle, 1990; Wand and Jones, 1995). The estimation is based on a set of observations $\{t_i, y_i\}_{i=1}^N$ drawn from the joint distribution of T and Y . This relationship between T and Y can be modeled as $y_i = m(t_i) + \epsilon_i$, where ϵ_i describes the random error of the model for the i th observation and m is the unknown function that is to be estimated.

In this setting, $m(t)$ is defined by the conditional expectation

$$m(t) \equiv E(Y|T = t) = \int y \frac{f(t, y)}{f_T(t)} dy \quad (5.1)$$

where $f_T(t)$ is the marginal density of T and $f(t, y)$ is the joint density function of T and Y . For random design data, both $f(t, y)$ and $f_T(t)$ are unknown and so m has no closed-form solution. A number of estimators for m have been proposed in the kernel regression literature.

One such estimator—the Nadaraya-Watson kernel regression estimator (Nadaraya, 1964; Watson, 1964)—can be derived from Equation (5.1) by replacing the unknown densities with their kernel density estimates

$$\hat{f}_T^h(t) \equiv \frac{1}{N} \sum_{i=1}^N K_h(t - t_i) \quad \text{and} \quad \hat{f}^{h,g}(t, y) \equiv \frac{1}{N} \sum_{i=1}^N K_h(t - t_i) K_g(y - y_i). \quad (5.2)$$

In these equations, K is a function that satisfies $\int_{\mathbb{R}} K(t) dt = 1$. $K_h(t) \equiv \frac{1}{h} K(\frac{t}{h})$ and $K_g(t) \equiv \frac{1}{g} K(\frac{t}{g})$ are kernel functions with bandwidths h and g respectively.

Plugging these density estimates into Equation (5.1) gives

$$\hat{m}_{h,g}(t) = \int y \frac{\frac{1}{N} \sum_{i=1}^N K_h(t - t_i) K_g(y - y_i)}{\frac{1}{N} \sum_{i=1}^N K_h(t - t_i)} dy. \quad (5.3)$$

Finally, assuming that K is symmetric about the origin, integration of the numerator leads to

$$\hat{m}_h(t) = \frac{\sum_{i=1}^N K_h(t - t_i) y_i}{\sum_{i=1}^N K_h(t - t_i)}. \quad (5.4)$$

Intuitively, the Nadaraya-Watson estimator returns the weighted average of the observations y_i , with the weighting determined by the kernel. Note that $\hat{f}^{h,g}(t, y)$ is factored out of the estimator—the weights only depend on the values t_i .

In Figure 5.1 we illustrate univariate kernel regression by applying it to demonstrate the effect of aging on ventricle volume and gray matter volume in the brain. This illustration is based on data collected by Mortamet et al. (2005). Each point represents a volume measurement, relative to total brain volume, for a particular subject. These measurements were derived from 3D MR images of 50 healthy adults ranging in age from 20 to 72 using an expectation-maximization based automatic segmentation method (Leemput et al., 1999). We used kernel regression to estimate the relationship, on average, between volume and subject age (filled lines). A Nadaraya-Watson kernel estimator with a Gaussian kernel of width $\sigma = 6$ years was used.

5.2.2 Kernel regression on Riemannian manifolds

In this section, the regression problem is defined in the more general setting of manifold-valued observations. Let $\{t_i, p_i\}_{i=1}^N$ be a collection of observations where the t_i are drawn from a univariate random variable T , but where p_i are points on a Riemannian manifold \mathcal{M} . The classical kernel regression methods presented in Section 5.2.1 are not applicable in this setting because they rely on the vector space structure of the observations. In particular, the addition operator in Equation (5.4) is not well defined.

The goal is to determine the relationship, on average, between the independent variable T and the distribution of the points $\{p_i\}$ on the manifold. This relationship can be modeled by

$$p_i = \text{Exp}_{m(t_i)}(\epsilon_i) \tag{5.5}$$

where $m : \mathbb{R} \rightarrow \mathcal{M}$ defines a curve on \mathcal{M} . The error term $\epsilon_i \in T_{m(t_i)}\mathcal{M}$ is a tangent vector that is interpreted as the displacement along the manifold of each observation p_i from the curve $m(t)$. The exponential mapping, Exp , returns the point on \mathcal{M} at time one along the geodesic flow beginning at $m(t_i)$ with initial velocity ϵ_i .

Following the univariate case, the regression function $m(t)$ can be defined in terms of expectation. However, in this case the idea of expectation of real random variables is generalized to manifold-valued random variables via Fréchet expectation, which is described in the previous Chapter 4. Recall (see Equation (4.7)) that the empirical estimate of the Fréchet mean, given a collection of observations $\{p_i, i = 1 \cdots N\}$ on a manifold \mathcal{M} , is defined by

$$\mu = \operatorname{argmin}_{q \in \mathcal{M}} \frac{1}{N} \sum_i^N d(q, p_i)^2. \quad (5.6)$$

Motivated by the definition of the Nadaraya-Watson estimator as a weighted averaging, we define a *manifold kernel regression estimator* using the weighted Fréchet empirical mean estimator as

$$\hat{m}_h(t) = \operatorname{argmin}_{q \in \mathcal{M}} \left(\frac{\sum_{i=1}^N K_h(t - t_i) d(q, p_i)^2}{\sum_{i=1}^N K_h(t - t_i)} \right). \quad (5.7)$$

This estimator is illustrated in Figure 5.3. Notice that when the manifold under study is a Euclidean vector space, equipped with the standard Euclidean norm, the above minimization results in the Nadaraya-Watson estimator.

5.2.3 Bandwidth selection

Within the kernel regression literature, it is well known that kernel width plays a crucial role in determining regression results (Wand and Jones, 1995). It is important to select a bandwidth that captures relevant population-wide changes without either oversmoothing and missing relevant changes or undersmoothing and biasing the results based on individual noisy data points. The ‘Goldilocks’ method of tuning the bandwidth until the results are most pleasing is a common subjective method for bandwidth selection. However, non-subjective methods may be required, for example, when kernel regression is part of a larger statistical study. A number of automatic kernel bandwidth selection




Figure 5.3: **Manifold kernel regression schematic.** (a) For any value of the predictor variable t , the manifold-valued observations p_i are summarized by the weighted Fréchet mean point $\hat{m}_h(t)$. (b) As in the univariate case, the weights are determined by the predictor values t_i and the kernel K_h .

techniques have been proposed for this purpose (for example Wand and Jones, 1995; Jones et al., 1996; Loader, 1999).

One classic method for automatic bandwidth selection is based on least squares cross-validation. This method is easily extended to the manifold regression setting in the following way. For observations $\{t_i, p_i\}_{i=1}^N$, with $t_i \in \mathbb{R}$ and $p_i \in \mathcal{M}$, the least squares cross-validation estimate for the optimal bandwidth h is defined as

$$\hat{h}_{\text{LSCV}} \equiv \operatorname{argmin}_{h \in \mathbb{R}^+} \frac{1}{N} \sum_{i=1}^N d(\hat{m}_h^{i-}(t_i), p_i)^2 \quad (5.8)$$

where

$$\hat{m}_h^{i-}(t) \equiv \operatorname{argmin}_{q \in \mathcal{M}} \left(\frac{\sum_{j=1, j \neq i}^N K_h(t - t_j) d(q, p_j)^2}{\sum_{j=1, j \neq i}^N K_h(t - t_j)} \right) \quad (5.9)$$

is the manifold kernel regression estimator with the i -th observation left out.

It is important to note that Equation (5.8) may achieve multiple local minima; this

is true even in Euclidean space (Hall and Marron, 1991).

5.2.4 Kernel regression for populations of anatomical images

In this section the manifold-regression methodology is applied to study random design anatomical image data. The database contains observations of the form $\{t_i, I_i\}_{i=1}^N$ where t_i is the age of subject i and I_i is a three-dimensional image that is identified with the anatomical configuration of subject i . We seek the unknown function m that associates a representative anatomical configuration, and its associated image \hat{I} , with each age.

Let $\Omega \subset \mathbb{R}^3$ be the underlying coordinate system of the observed images I_i . Each image $I \in \mathcal{I}$ can be formally defined as an L^2 function from Ω to the reals. As in previous chapters, anatomical differences are represented by diffeomorphic transformations of the underlying image coordinates. Let \mathcal{H} be the group of diffeomorphisms that are isotopic to the identity. Each element $\phi : \Omega \rightarrow \Omega$ in \mathcal{H} deforms an image according to the following rule

$$I_\phi(x) = I(\phi^{-1}(x)). \quad (5.10)$$

Using the metric on \mathcal{H} defined in Chapter 2, a distance between images is defined by¹

$$d_{\mathcal{I}}(I_1, I_2)^2 \equiv \min_{v: \dot{\phi}_s = v_s(\phi_s)} \left[\int_0^1 \|v_s\|_V^2 ds + \frac{1}{\sigma^2} \|I_1(\phi^{-1}) - I_2\|_{L_2}^2 \right] \quad (5.11)$$

where the second term accounts for the noise model of the image. While this construction is motivated by the metric on \mathcal{H} , it does not strictly define a Riemannian metric on the space of anatomical images (because of the second term). In the future we plan to define distance in terms of the elegant construction described in Trouvé and Younes (2005).

¹See Chapter 4 for more details

Having defined a metric on the space of images that accommodates anatomical variability, we can apply that metric to regress a representative anatomical configuration, with associated image, from our observations $\{t_i, I_i\}$

$$\hat{I}_h(t) = \operatorname{argmin}_{I \in \mathcal{I}} \left(\frac{\sum_{i=1}^N K_h(t - t_i) d_{\mathcal{I}}(I, I_i)^2}{\sum_{i=1}^N K_h(t - t_i)} \right). \quad (5.12)$$

Equation (5.12) expresses the following intuitive idea: For any age t , the population can be represented by the anatomical configuration that is centrally located, according to $d_{\mathcal{I}}$, among the observations that occur near in time to t . As in the univariate case, the weights are determined by the kernel K .

5.2.5 Diffeomorphic population growth model

Given a population-representative anatomical image, as a function of age, we consider the local shape changes evident—for the population—as a function of age. In this section the manifold kernel regression estimator is used to extend a single-subject longitudinal growth model in order to study *population-average* geometric change. In particular, an algorithm is presented for estimating the age-indexed diffeomorphism that quantifies the fine scale anatomical shape change of the population representative \hat{I} .

Single-subject growth model

The dynamic growth model described in Miller (2004) associates a single subject with a collection of image observations $J_t \in \mathcal{I}$, which are acquired over a period of time $t \in [0, 1]$. The goal is to determine the diffeomorphic flow g_t that deforms an exemplar image J_α through time in such a way that it matches these image observations. In practice J_0 is used as the exemplar image. This methodology has been applied, for example, to measure growth or atrophy of structures within the brain.

The formalization of the growth problem is similar to the definition of the image




Figure 5.4: **Population growth model schematic.** (a) The diffeomorphism g_t quantifies the geometric change of \hat{I} throughout the growth period. (b) The velocity field that is identified with the tangent vector $v_t = \dot{g}_t$ is overlaid on the underlying anatomical image $\hat{I}(t)$. The arrows indicate instantaneous shape change at age t . (c) This colormap identifies regions of local expansion and contraction of the underlying anatomy at time t . Red indicates expansion; blue indicates contraction. See text for details.

metric $d_{\mathcal{I}}$ (see Equation (5.11)) in that it is defined as a minimization problem that seeks to find a solution g_t that requires the least amount of deformation according to the metric $d_{\mathcal{H}}$ on the space of diffeomorphisms:

$$\operatorname{argmin}_{v: \dot{g}_t = v_t(g_t)} \left[\int_0^1 \|v_t\|_V^2 dt + \frac{1}{\sigma^2} \int_0^1 \|J_\alpha(g_t^{-1}) - J_t\|_{L_2}^2 dt \right]. \quad (5.13)$$

A primary difference between this equation and Equation (5.11) is that in the case of the growth model the second term is integrated over time. This enforces the requirement that the deforming exemplar image $J_\alpha(g_t^{-1})$ match the observed imagery J_t throughout the growth period.

It has been shown using the calculus of variations (Miller et al., 2002) that the solution to Equation (5.13) satisfies

$$L^\dagger L v_t = -\frac{1}{2\sigma^2} \nabla (J_\alpha \circ g_t^{-1}) \int_t^1 (J_u(g_u \circ g_t^{-1}) - J_\alpha(g_t^{-1})) |D(g_u \circ g_t^{-1})| du \quad (5.14)$$

where $\nabla(J_\alpha \circ g_t^{-1})$ is the gradient of the deformed exemplar image and $D(g_u \circ g_t^{-1})$ is the Jacobian of the diffeomorphic transformation that maps the anatomical configuration at time t to the configuration at time u . The discrete version of this equation is used to construct an iterative solution for v_t . Initially, g_t is set to the identity map for all t . At each iteration $v_t, t \in [0, 1]$ is updated according to the observed images J_t and the current estimate of $g_t, t \in [0, 1]$.

Population growth model

In order to extend this growth model to apply to a population of subjects, the subject-specific collection of observed imagery J_t is replaced with the *expected* observed imagery, as a function of time, for the population (cf. Figure 5.4). This is achieved by combining the manifold kernel regression estimator (Equation (5.12)) with the growth model (Equation (5.13)):

$$\begin{aligned} \operatorname{argmin}_{v: \dot{g}_t = v_t(g_t)} \int_0^1 \|v_t\|_V^2 dt + \\ \frac{1}{\sigma^2} \int_0^1 \|I_\alpha(g_t^{-1}) - \operatorname{argmin}_{I \in \mathcal{I}} \left(\frac{\sum_{i=1}^N K_h(t - t_i) d_{\mathcal{I}}(I, I_i)^2}{\sum_{i=1}^N K_h(t - t_i)} \right)\|_{L_2}^2 dt. \end{aligned} \quad (5.15)$$

In this way the population-representative images serve as a collection of *population average* time-sequence imagery.

In order to solve Equation (5.15) the solution for the interior minimization problem is first solved for a discrete collection of time points. This is legitimate since this problem does not depend on the growth deformation g_t . Once these population-representative images are computed, the time-indexed deformation g_t is computed using the iterative method based on Equation (5.14). The population exemplar image is taken to be $I_\alpha \equiv \hat{I}(0)$. In order to speed convergence, the algorithm is applied within a three-level multi-resolution framework where initial solutions at coarser scale levels are used to initialize the optimization procedure at finer scale levels.

Once g_t is computed, it can be analyzed to determine local, age-indexed geometric change for the population. For example, instantaneous local growth and atrophy can be measured via the log-determinant of the Jacobian of the velocity field defined by

$$\log \begin{vmatrix} \frac{\partial \dot{g}_t^1}{\partial x^1}(x) & \frac{\partial \dot{g}_t^1}{\partial x^2}(x) & \frac{\partial \dot{g}_t^1}{\partial x^3}(x) \\ \frac{\partial \dot{g}_t^2}{\partial x^1}(x) & \frac{\partial \dot{g}_t^2}{\partial x^2}(x) & \frac{\partial \dot{g}_t^2}{\partial x^3}(x) \\ \frac{\partial \dot{g}_t^3}{\partial x^1}(x) & \frac{\partial \dot{g}_t^3}{\partial x^2}(x) & \frac{\partial \dot{g}_t^3}{\partial x^3}(x) \end{vmatrix}. \quad (5.16)$$

Values of the log-Jacobian greater than zero indicate local expansion; values less than zero indicate local contraction.

5.3 Synthetic data experiment

This section demonstrates the manifold kernel regression methodology with a proof of concept experiment based on synthetic data. The manifold regression method is applied to a database of synthetic 2D images that were generated from a known, underlying geometric process. The goal is to recover, from the imagery alone, the underlying geometric change.

The database consists of two cohorts that each contain 100 256×256 2D bullseye images. The cohorts, B_1 and B_2 , differ by the amount of *random* geometric variation present. Each image is associated with a particular value of the synthetic predictor variable $t \in [0, 1]$; the values of t for the database were drawn from a uniform random distribution on $[0, 1]$. For the i^{th} image there are three disks that independently change in radii according to

$$\begin{aligned} r_1(t_i) &= f_1(t_i) + \epsilon_i + \epsilon_{i,1} \\ r_2(t_i) &= f_2(t_i) + \epsilon_i + \epsilon_{i,2} \\ r_3(t_i) &= f_3(t_i) + \epsilon_i + \epsilon_{i,3} \end{aligned} \quad (5.17)$$

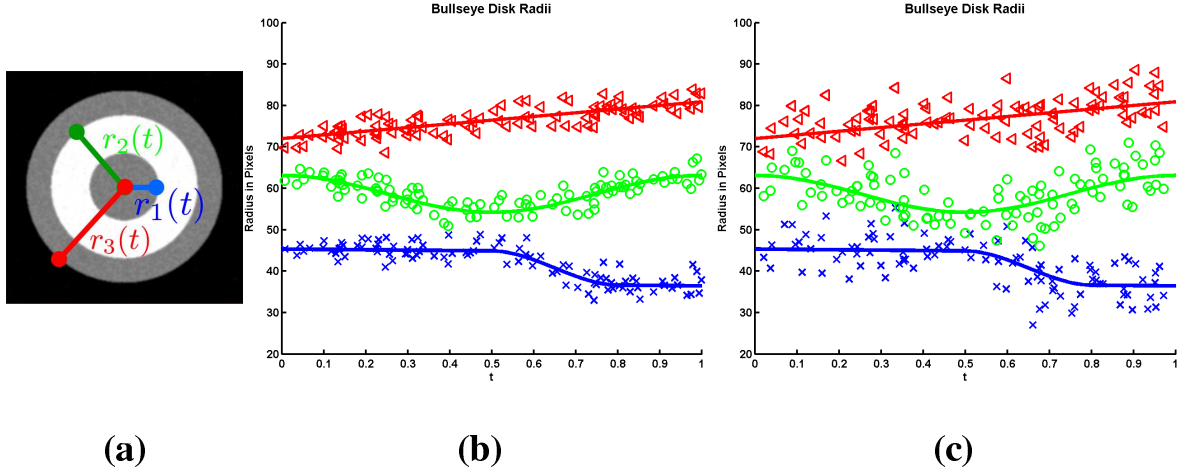


Figure 5.5: **Synthetic bullseye data set construction.** The bullseye database contains 200 total 2D images; each image is associated with a value of t drawn from a uniform distribution on $[0, 1]$. (a) Images are generated from three independent, noisy radius values: r_1, r_2 , and r_3 . (b) Cohort B_1 : each observed radius value (markers) is a function of t and is determined by adding random noise to the ground truth functions f_1, f_2 , and f_3 , which are depicted by the solid curves. (c) The second cohort, B_2 , was generated using a higher level of random geometric variation.

subject to

$$r_1(t) < r_2(t) < r_3(t) \text{ for all } t \in [0, 1]. \quad (5.18)$$

The functions f_1, f_2 , and f_3 are known; they *define* the noise-free, ground-truth geometric change as a function of t . Noise is added to these radius functions via the zero mean Gaussian random variables $\epsilon_i, \epsilon_{i,1}, \epsilon_{i,2}$ and $\epsilon_{i,3}$. For cohort B_1 , $\epsilon_i \sim N(\mu = 0, \sigma^2 = 4 \text{ pixels})$ and $\epsilon_{i,1}, \epsilon_{i,2}, \epsilon_{i,3} \sim N(\mu = 0, \sigma^2 = 1 \text{ pixels})$. For cohort B_2 , $\epsilon_i \sim N(\mu = 0, \sigma^2 = 16 \text{ pixels})$ and $\epsilon_{i,1}, \epsilon_{i,2}, \epsilon_{i,3} \sim N(\mu = 0, \sigma^2 = 4 \text{ pixels})$. Once the image geometries are fixed independent identically-distributed Gaussian noise is added to the image intensities. Figure 5.5 contains a schematic of the image generation process. Figure 5.6 displays a sample of the images from this database.

For each cohort, the manifold kernel regression algorithm was applied in order to regress a population-representative bullseye image for eight equally spaced values of




Figure 5.6: **Random design database of 2D bullseye images.** These images are taken from cohort B_2 . Associated time measurements increase from left to right and from top to bottom. Inner, middle, and outer disk radii are generated by adding noise to the underlying curves depicted in Figure 5.5 (b) and (c).

t. A kernel bandwidth of $\sigma = 0.045$ was used. For this experiment, the solutions to Equation (5.11) were computed using MATLAB codes based Algorithm 4.1 from Chapter 4.

Figure 5.7 contains the results of this experiment. The regressed images are shown in the background. The ground truth radii values, f_1 , f_2 , and f_3 , are depicted as colored overlays. The close agreement with the regressed images and the overlays indicates that the underlying geometric process was recovered from the image database—that is, the underlying time effect was separated from the random geometric variation. Comparing the results for the two cohorts, the regression of the geometries is rather robust with regard to the level of geometrical noise. Only a slight degradation in accuracy of the estimate is seen with a four-fold increase in the variance of the radii.




Figure 5.7: **Regression results for synthetic bullseye database.** These images show the regression results for the bullseye database at eight equally spaced time points for cohorts B_1 and B_2 . Colored overlays denote the ground-truth radii as determined by the underlying curves in Figure 5.6 (b) and (c).

5.4 Conclusion

This chapter introduces a method for *manifold kernel regression* that enables novel analysis of population shape and growth from random design data when the underlying shape model is non-Euclidean. While this method is quite general (for example, it is applied to rotational pose in Davis et al. (2007)), this chapter focuses on shape changes from populations of images. A population-representative image, indexed by an independent variable t , is regressed from a database of images. Finally, a longitudinal growth model is used to study the detailed local shape change that occurs, on average, as a function of t .

The definition of the manifold kernel regression estimator raises a number of important theoretical questions. Many of these questions have direct parallels in the Euclidean kernel regression setting: How much data do we need in order to achieve a stable regression estimate? How can we validate our results? How can we deal with edge effects? Can we apply robust regression in order to minimize the effects of outliers? How do we choose the best bandwidth? Should we use kernel regression or instead apply a parametric model? After applying regression, can we reliably solve the inverse problem (determine the predictor value from a new image)? Several questions are specific to the manifold regression setting: What are the properties of the Fréchet mean for the manifold under study? What are the convergence properties?

These questions serve as road map for future work on this project. Another important direction is the application of the manifold kernel regression methodology to other shape descriptions that are naturally elements of a shape manifold. Examples include m-reps (Pizer et al., 2003; Fletcher et al., 2004), which form a symmetric space; point-sets, curves, and surfaces within the diffeomorphic matching framework described in Glaunès et al. (2004); Glaunès and Joshi (2006); and diffusion tensor data, which is described by the space of 3×3 symmetric positive-definite matrices (Pennec et al., 2006; Fletcher and Joshi, 2007).

In the next chapter, manifold kernel regression is applied to study effect of aging on the brain by applying this methodology to a random design database of 3-dimensional MR brain images from healthy subjects.

Chapter 6

The aging brain: Measuring change using regression

6.1 Introduction

In this chapter the manifold kernel regression methodology is applied to study the aging brain. The observations are a collection of 3D magnetic resonance (MR) brain images of healthy subjects, with associated age measurements. In Section 6.2 this database is described in detail. In Section 6.3 results are presented for the regression of a representative image, as a function of age, from this database. Section 6.4 contains results for the measurement of local, age-indexed expansion and contraction in the brain. Finally, Section 6.5 describes how the kernel bandwidth was determined for these results.

6.2 Brain image database

The image database that was used in these experiments was collected by Dr. Elizabeth Bullitt at the University of North Carolina, Chapel Hill (Lorenzen et al., 2006). The database contains MRA, T1-FLASH, T1-MPRAGE, and T2-weighted images for over 97 healthy adults ranging in age from 20 to 79. For this study, only the T1-FLASH images were used; these images were acquired at a spatial resolution of $1\text{ mm} \times 1\text{ mm} \times$

1 mm using a 3 Tesla head-only scanner. The tissue exterior to the brain was removed using a mask generated by a brain segmentation tool described in Prastawa et al. (2004). This tool was also used for bias correction. In the final preprocessing step, all of the images were spatially aligned to an atlas using rigid registration.

Manifold kernel regression was applied for male-only, female-only, and combined cohorts. The final size of the male cohort is 38 subjects ranging in age from 22 to 72; the final size of the female cohort is 46 subjects ranging in age from 20 to 66. The final size of the combined cohort is 84 subjects ranging in age from 20 to 72. Midaxial slices of all subjects are shown in Figure 6.1. Figure 6.2 illustrates the spacing in time of the image data and the shape of several representative kernels used to compute the representative images.

6.3 Manifold kernel regression for MR brain images

The manifold kernel regression estimator was applied to compute representative anatomical images for each cohort. Images were computed for ages 30 to 60 at increments of one year using a Gaussian kernel with $\sigma = 6$ years which was truncated at 3 standard deviations. This bandwidth was subjectively determined.¹ High-resolution 2D slices from these representative images are shown in figures 6.3 through 6.5. Noticeable expansion of the lateral ventricles is clearly captured in this image data.

Figures 6.6 through 6.14 contain slices from these representative images for each cohort. These are not images from any particular patient—they are computed using the regression method proposed in Chapter 5.

¹See Section 6.5 for details.




Figure 6.1: **Healthy subject MRI brain image database.** To demonstrate the extent of natural brain shape variability within a population of healthy subjects, a mid-axial slice is presented for each image used in this study. The images are arranged, by row, in order of increasing patient age from 20 (top left) to 72 (bottom right). Because of the complexity of the shapes and the high level of natural shape variability, it is extremely difficult to visually discern any patterns within these data.




Figure 6.2: **Subject age distribution and kernel weights.** Illustration of kernel weights for regression of anatomical imagery. The crosses indicate time points where image data are available (the random vertical placement of these markers is used only for the purpose of visual differentiation). Atlases were created at one year time intervals starting at age 30 and continuing to age 60. The dashed lines graph the kernel weights that were used to compute each atlas (see Equation (5.7)).




Figure 6.3: **Regressed anatomical images: female cohort.** Representative anatomical images for the female cohort at ages 30 (left) and 60 (right). These images were generated from the random design 3D MR database using the shape regression method described in Chapter 5.




Figure 6.4: **Regressed anatomical images: male cohort.** Representative anatomical images for the male cohort at ages 30 (left) and 60 (right). These images were generated from the random design 3D MR database using the shape regression method described in Chapter 5.




Figure 6.5: **Regressed anatomical images: combined cohort.** Representative anatomical images for the combined cohort at ages 30 (left) and 60 (right). These images were generated from the random design 3D MR database using the shape regression method described in Chapter 5.




Figure 6.6: Average, aging brain: female cohort, axial view.

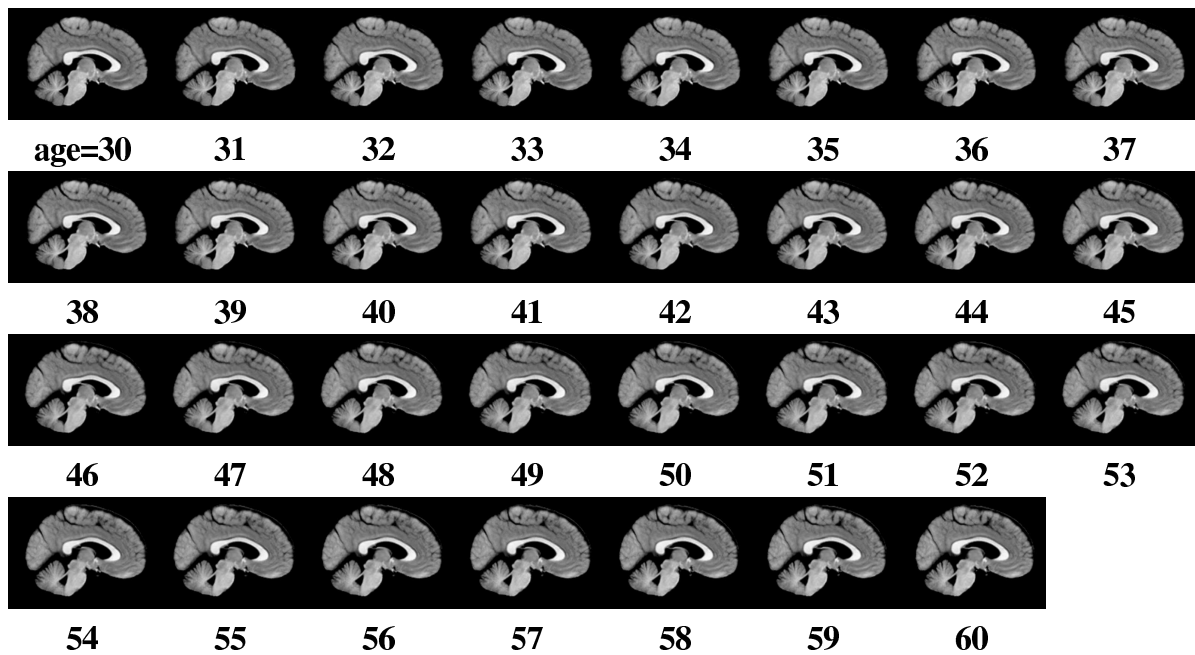


Figure 6.7: Average, aging brain: female cohort, sagittal view.




Figure 6.8: Average, aging brain: female cohort, coronal view.




Figure 6.9: Average, aging brain: male cohort, axial view.

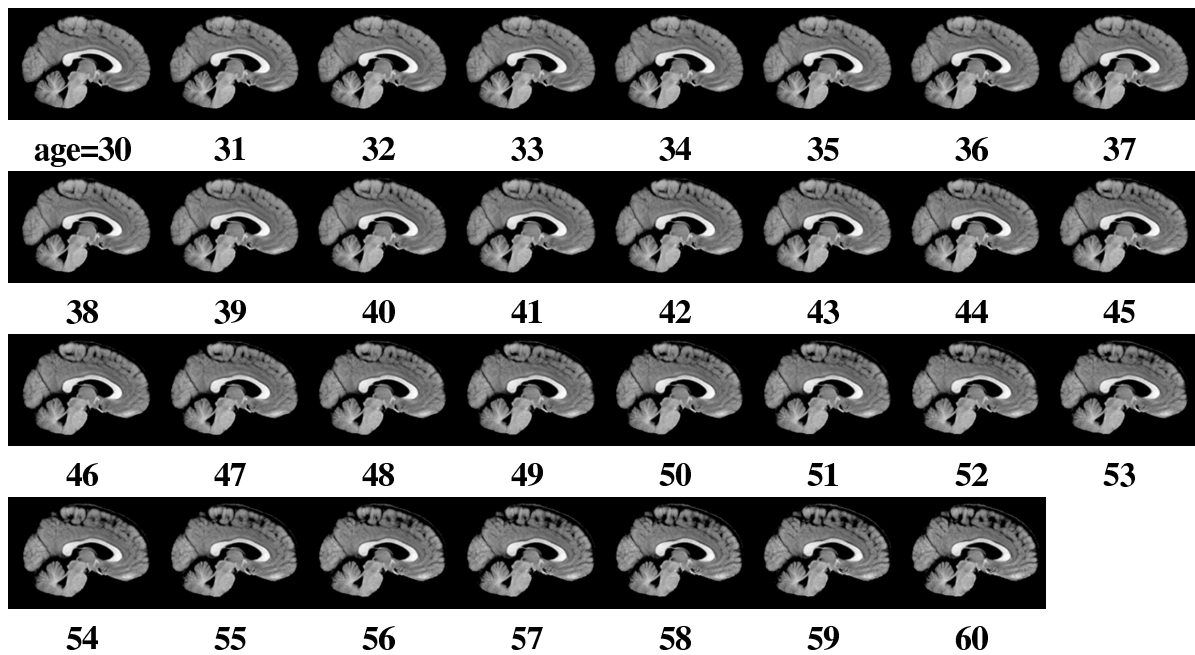


Figure 6.10: Average, aging brain: male cohort, sagittal view.




Figure 6.11: Average, aging brain: male cohort, coronal view.




Figure 6.12: Average, aging brain: combined cohort, axial view.

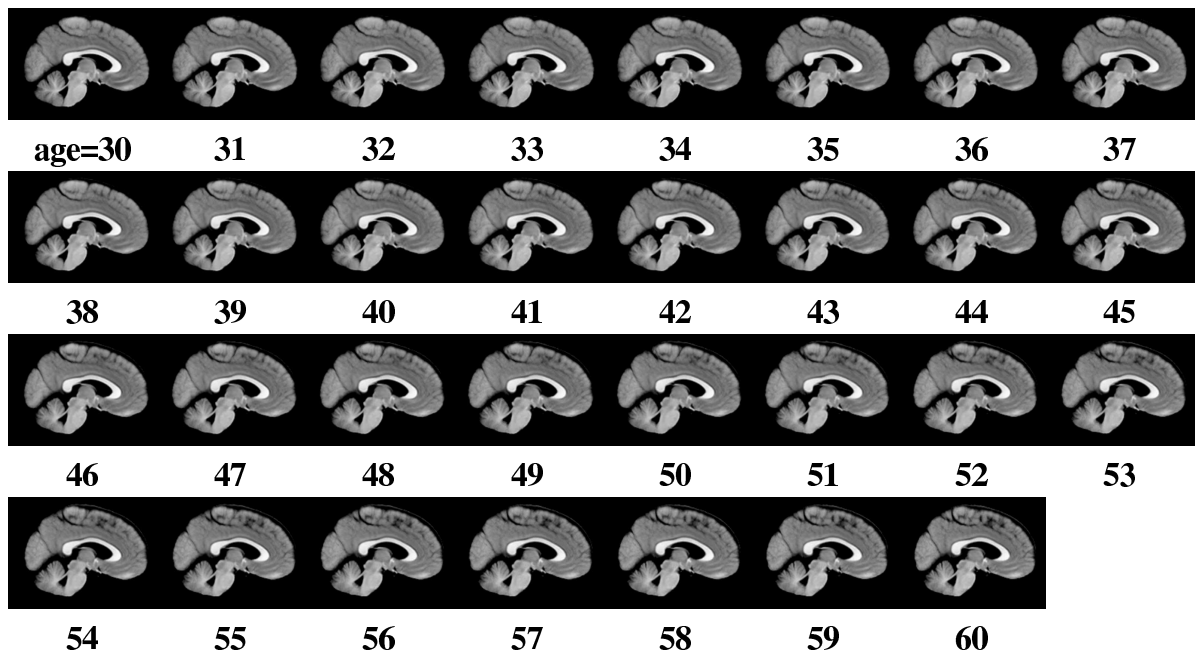


Figure 6.13: Average, aging brain: combined cohort, sagittal view.




Figure 6.14: Average, aging brain: combined cohort, coronal view.

6.4 Local, instantaneous volumetric change

The diffeomorphic growth estimation algorithm described in Chapter 5 was applied to determine the anatomical shape change over time for each cohort. Figures 6.15 through 6.17 illustrate the instantaneous change in the deformation every other year from age 30 to age 60 for each cohort. More precisely, the figures show the log-determinant of the Jacobian of the time-derivative of the deformation (see Equation (5.16) from Chapter 5). That is, in these images, bright pixels indicate *expansion* of the underlying tissue, at the given age, while dark pixels indicate *contraction*. According to these determinant maps, expansion of the ventricles is evident for each age group. However, the expansion is accelerated for ages 45 to 53. Note that this finding agrees well with volume-based regression analysis in Chapter 5 (Figure 5.1).

6.5 Investigation of kernel widths

As described in the previous chapter, one of the most important considerations in kernel regression is the choice of a kernel width that is appropriate for the data. In this case, it is important to select a bandwidth that captures relevant population-wide changes in the brain without biasing the results towards individual anatomy. The simplest *exploratory* strategy of choosing the bandwidth that is visually most pleasing was used for this study (see Figure 6.18). There are a number of automatic bandwidth selection techniques that could be applied; these techniques will be an important direction for future research.

The manifold kernel regression scheme was applied to each of the female, male, and combined cohorts using a wide range of kernel bandwidths: $\sigma = 1, 2, 4, 6, 8, 16$, and 32 years. Results for the female cohort are presented in Figure 6.19. For smaller bandwidths, rapid geometrical changes—with respect to age—are evident. These changes are the result of a bias to individual anatomy and do not represent population-wide change. For larger bandwidths, no changes are evident; especially for $\sigma = 32$ the kernel




Figure 6.15: **Aging brain shape change: female cohort.** Local brain shape *change* as a function of age for the female cohort. Bright pixels indicate instantaneous volumetric expansion; dark pixels indicate contraction.

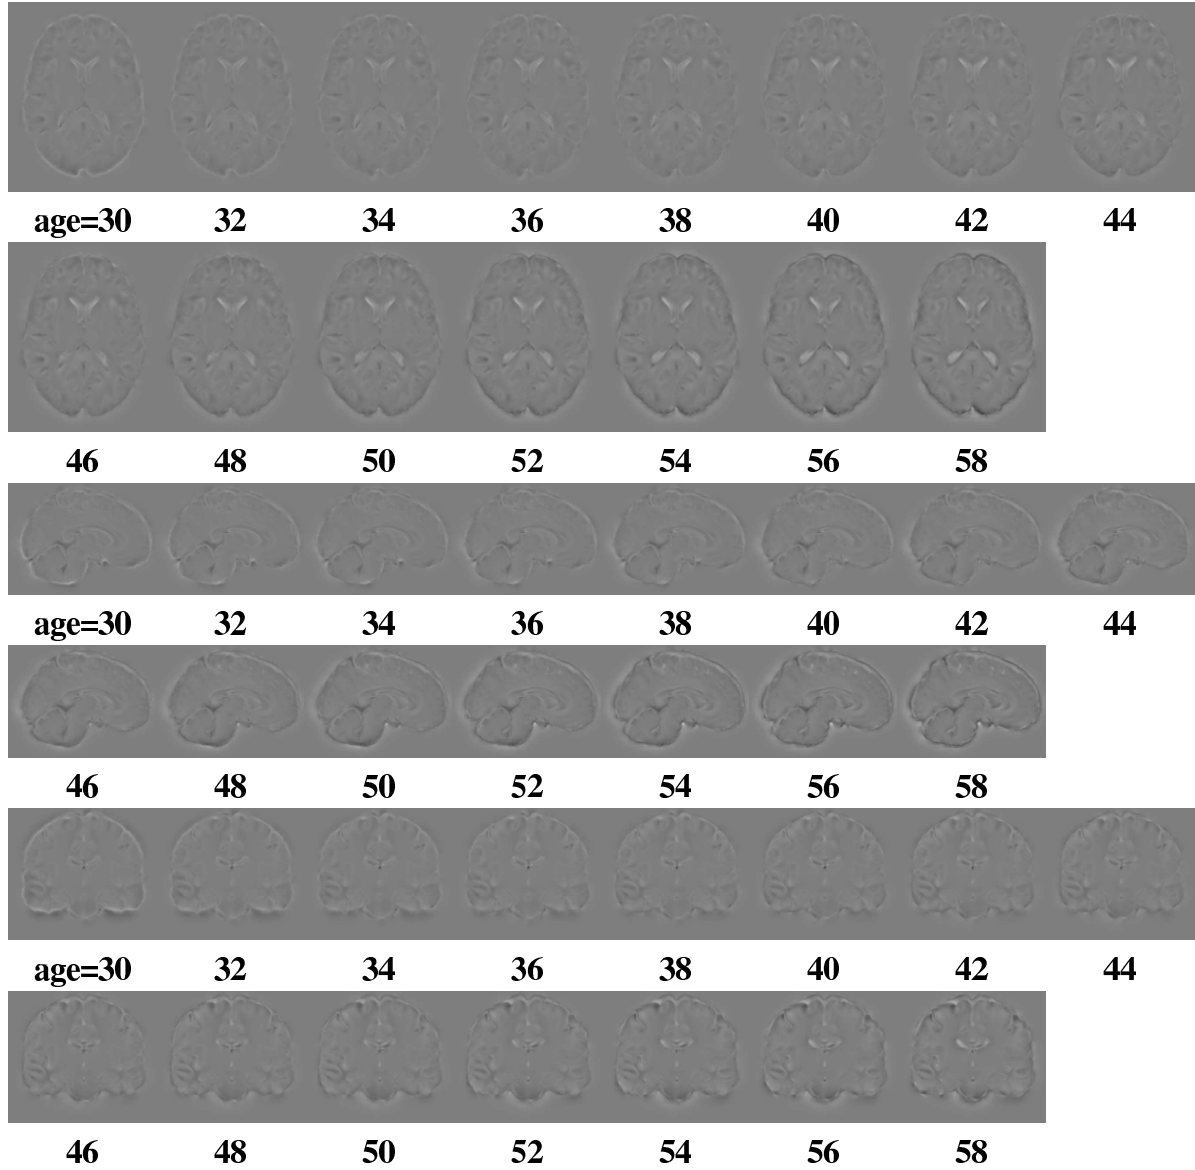


Figure 6.16: **Aging brain shape change: male cohort.** Local brain shape *change* as a function of age for the male cohort. Bright pixels indicate instantaneous volumetric expansion; dark pixels indicate contraction.




Figure 6.17: **Aging brain shape change: combined cohort.** Local brain shape *change* as a function of age for the combined cohort. Bright pixels indicate instantaneous volumetric expansion; dark pixels indicate contraction.




Figure 6.18: **Bandwidth selection example.** This figure demonstrates the bandwidth selection problem for kernel regression. A sinusoidal function, with noise added, is used to produce the underlying observations. Kernel regression, using three different bandwidths, is then used to recover the underlying function. The blue curve uses the narrowest kernel and seems to follow the noise in the data. The red curve uses the broadest kernel and clearly over-smooths the data. Visually, the green curve appears to follow the data while smoothing out small scale variation that is due to noise. In this sense, it is the ‘most pleasing’ bandwidth.

weights are nearly the same regardless of age. Qualitatively, it was determined that the bandwidth of six years provides the best balance between these extremes.

6.6 Computational details

For this study solutions for Equation (5.12) were approximated using an iterative greedy algorithm that is similar to Algorithm 4.2 from Chapter 4. Results were computed using a multi-threaded C++ implementation on an eight-processor (16-core) 2.4GHz system with approximately 64 gigabytes of RAM. Processing time averaged 116 minutes per $256 \times 256 \times 256$ regressed image volume.²

Figure 6.20 contains experimental runtime measurements for several problem sizes. All results are computed with $256 \times 256 \times 256$ floating-point (4 bytes per voxel) images.

The threaded implementation uses a fixed number (τ) of threads to solve Equation (4.19)—each thread is assigned to process a particular subset of input images. These threads are also utilized to compute the mean image at each iteration. In this case the mean image is partitioned into τ distinct geometrical regions and each thread is assigned to compute the mean values for a single region. Figure 6.21 shows a graph of the iteration speed as a function of processors for the input data size and system described above.

When computing each representative image $\hat{I}(x)$, a multi-resolution approach is applied which generates images at progressively higher resolutions, where each level is initialized by the results at the next coarsest scale. This strategy has the dual benefits of (a) addressing the large scale shape changes first and (b) speeding algorithm convergence.

The dominating computation at each iteration is a Fast Fourier Transform. The order of the algorithm is $MNn \log n$ where M is the number of iterations, N is the number

²Individual times measured 5 to 80 minutes and varied linearly according to the sampling (in age) density.






Figure 6.19: **Varying kernel width.** This figure demonstrates the result of modifying the kernel width within the image regression framework. Each row shows the regressed image as a function of age for a particular kernel width: $\sigma = 1, 2, 4, 6, 8, 16, 32$ years.



(a)



(b)

Figure 6.20: **Fréchet mean image experimental timing and memory results.** (a) Iteration time as a function of cohort size. (b) Memory usage as a function of cohort size. (c) Speedup as a function of processor count.




Figure 6.21: **Fréchet mean image experimental threading results.** Relative iteration speed as a function of processor count. The two curves denote two different solution strategies. For the green curve the mean image is updated once per iteration: after all the images have been deformed. For the blue curve, the mean image is update a every sub-iteration: every time an individual image is deformed. The second update strategy is much slower since the threads must wait while the mean is regenerated.

of images, and n is the number of voxels along the largest dimension of the images. Therefore, the complexity grows linearly with the number of observations, making this algorithm suitable for application to large data sets.

6.7 Conclusion

In this chapter the manifold kernel regression methodology was demonstrated for the study of the aging brain. It should be emphasized that, while these results are promising, more validation of this methodology must take place before these results should inform clinical decisions. In the future, I plan to apply this methodology on a structure-by-structure basis for substructures in the brain including the ventricles, caudate, hippocampus, and thalamus. Aside from generating regressed images of these structures, it will be informative to compare the results from the full image regression and the structure-by-structure regression.

Chapter 7

Discussion

7.1 Summary of contributions and thesis

This section revisits the thesis statement and claims that are presented in Chapter 1. Following each claim is a brief statement that summarizes how that claim is addressed in this dissertation.

The contributions of this dissertation are as follows:

(1) *I present a novel method, called manifold kernel regression, that enables regression analysis of manifold-valued data.*

Manifold kernel regression is presented in Chapter 5. The manifold kernel regression estimator is defined as a generalization of univariate kernel regression; it enables the study of systematic, population-wide trends from collections of manifold-valued observations. Manifold kernel regression is based on Fréchet expectation, which can be used to define averages on manifolds. A method for kernel bandwidth selection—a critical aspect of kernel regression—is presented in Section 5.2.3. Extensions and future directions for manifold kernel regression are presented at the end of Chapter 5 and later in this chapter.

(2) *I apply manifold kernel regression to the study of anatomical change from a random design database of medical images. In particular, it is defined for images using the*

framework of large-deformation diffeomorphic image mapping.

In this dissertation the general concept of manifold kernel regression is applied to the study of anatomy via diffeomorphisms of medical images. The large-deformation diffeomorphic image matching framework is reviewed in Chapter 2. Chapter 4 describes how Fréchet mean anatomical configurations are computed using distances measured on the diffeomorphic transformations. Chapter 5 defines a manifold kernel regression estimator using this Fréchet mean. Algorithms for computing the Fréchet mean images are presented in Chapter 4. The application of manifold kernel regression to 2D synthetic bullseye images is demonstrated in Section 5.3. Extensions and future directions for manifold kernel regression of images are presented at the end of Chapter 5 and later in this chapter.

(3) I demonstrate manifold kernel regression by measuring average geometric change in the aging brain from a random design dataset of 3D MR images. The effect of regression kernel width on the regressed shape is explored.

In Chapter 6 manifold kernel regression is used to study the aging brain. A manifold kernel regression estimator is applied to measure changes in the brain as a function of age from a random design database of 89 MR images taken from healthy adults ranging in age from 22 to 79 years. This data was collected by Dr. Elizabeth Bullitt at the University of North Carolina at Chapel Hill. A representative brain shape, as a function of age, is regressed from this population. The regression method is also applied to the male and female subpopulations separately.

Sections 5.2.5 and 6.4 develop a population growth model that is used to quantify local, instantaneous shape change for the population using the regression model. The growth model is applied to determine patterns of expansion and contraction for the average aging brain.

A preliminary, subjective study of the effect of kernel bandwidth on the regression of these images is presented in Section 6.5. In this study regressed images are produced

using a broad range of kernel bandwidths.

(4) I present a novel method for computing a Fréchet mean image from a collection of images.

Chapter 4 defines a Fréchet mean image using the large-deformation diffeomorphic framework, which is reviewed in Chapter 2. The Fréchet mean image is the solution to an optimization problem; it is the minimizer of the sum of squared distances to the input images. Squared distance is defined by the combination of a squared metric on the space of diffeomorphisms and a term that penalizes residual image mismatch. The mean image is generated automatically from the input images using an iterative procedure of continuous joint alignment: the input images iteratively deform toward the evolving mean.

The mean image defines a natural, unbiased average anatomical configuration and coordinate system from the collection of medical images. The mean is natural in the sense that it is based on intrinsic distances in the space of the diffeomorphic transforms that encode shape change. It is unbiased in the sense that (1) it does not depend on a choice of a final or initial geometrical coordinate system, and (2) it does not depend on the ordering of the input data. These properties are further described in Section 4.2.4.

Algorithms for computing Fréchet mean images are presented in Section 4.2.3. Section 4.3 demonstrates the Fréchet mean image methodology for collections of synthetic 2D images.

As described in Chapter 5, the Fréchet mean plays a central role in the definition of the manifold kernel regression estimator.

(5) I describe a program for efficiently computing Fréchet mean images and applying the manifold kernel regression analysis method to 2D and 3D images on shared-memory, multi-processor machines. Performance measurements are included.

Algorithms for computing a Fréchet mean image are described in Section 4.2.3. Sec-

tion 6.6 describes implementation details for a threaded, multiresolution, C++ implementation. This section includes a theoretical analysis of the algorithmic order and empirical measurements of the runtime and memory requirements for a modern shared-memory multi-processor computer.

(6) *I present a novel method for extending diffeomorphic image mapping to accommodate certain topological changes. The method is applied to track the changing position of the prostate relative to the pelvis in the context of transient bowel gas. The effectiveness of this method is tested in a retrospective study involving 40 3D computed tomography images from 3 patients undergoing adaptive radiation therapy.*

A method for accommodating extraneous structures in diffeomorphic image registration is presented in Chapter 3. Section 3.1 describes how registration is used to measure organ motion in adaptive radiation therapy of the prostate; this section also describes how transient gas in the rectum is not adequately handled using diffeomorphic image mapping. Section 3.2 provides a precise definition of topological equivalence for images and extraneous structures. This section also describes an algorithm for applying diffeomorphic image registration in the presence of extraneous structures by deflating the extraneous structures before applying registration. Diffeomorphic correspondence is determined for regions that do not contain extraneous structures by composing the deflation and registration transformations.

The application of this method to compute of Fréchet mean images is briefly discussed later in this chapter.

Thesis: *Manifold kernel regression is a natural generalization of kernel regression that enables regression analysis for points on a manifold. It extends classical kernel regression in order to estimate, from a collection of observations, the relationship—on average—between an independent predictor variable, such as time, and a dependent variable represented by points on a manifold. In particular, this method is useful for determining*

population average anatomical shape change over time from a random design database of medical images. Because it provides a quantitative link between a predictor variable and anatomical structure, manifold kernel regression is an effective tool for improving our understanding of anatomical changes within populations.

In Chapter 5, manifold kernel regression is described as a generalization of Nadaraya-Watson kernel regression. This generalization is natural in the sense that both Nadaraya-Watson and manifold kernel regression are expressed as weighted means of the response data. In both cases the weights are determined by the associated predictor data. Whereas the Nadaraya-Watson regression estimator is defined for vector space response data, the manifold kernel regression estimator is defined in terms of intrinsic operations on a manifold. In particular, it is based on the Fréchet mean—the minimizer of the sum of squared distances.

The manifold kernel regression estimator can be used for regression analysis—just like any classical regression estimator—to estimate the unknown relationship between two random variables: a predictor and response. The key feature of the manifold kernel regression estimator is that it relies only on intrinsic properties such as distance along the manifold. Therefore, it can be applied when the responses are elements of a curved manifold.

Chapter 4 describes algorithms for computing a Fréchet mean image within the large-deformation diffeomorphic image registration framework. Chapter 6 demonstrates the use of manifold kernel regression for the study of local shape changes in the average aging brain. This application demonstrates manifold kernel regression as a methodological tool for the study of anatomical change of populations. Patterns of expansion and contraction are produced using a population-based growth algorithm that is based on manifold kernel regression.

7.2 Benefits and limitations of diffeomorphisms for the study of shape

Diffeomorphic image matching is a powerful tool for aligning anatomical images and studying anatomical shape. In particular, it provides a rigorous mathematical framework with a composition operator, smoothness guarantees, and the ability to measure complex patterns of local geometric deformations. Furthermore, a metric can be defined on the space of these deformations, and this metric is useful for the analysis of shapes; it allows for the formulation of statistical methods like the Fréchet mean and manifold kernel regression estimator. The elastic and viscous fluid mechanical models used to generate diffeomorphisms provide an intuitive relationship between forces and tissue motion. Finally, a wide variety of geometric representations including multiple image channels, point sets, landmarks, curves, and surfaces can be used to generate diffeomorphisms.

There are also drawbacks to this diffeomorphic framework. First, deformation frameworks should only be applied after an initial, rigid alignment of structures. In practice this is usually handled by a cascading set of transformations starting with a rigid alignment and ending with a full deformation. Second, diffeomorphic transformations are not representative of all anatomical motion—for example, diffeomorphisms cannot model two structures sliding past each other. Third, it is clear from the study of anatomy that building diffeomorphic transformations between different individuals is not possible at all scales. For example, a particular sulcal fold in the brain may exist for one person and be absent for another. An investigation into this problem would be interesting; while multiresolution techniques are often used to compute transformations between individuals, they have not been used (to my knowledge) to determine scales at which a diffeomorphic match is possible for anatomical images. Similarly, one could investigate, for a fixed scale, which structures are likely to correspond diffeomorphically and which

do not.

Finally, diffeomorphic transformations are very high dimensional; theoretically, they are infinite dimensional. In practice, they are approximated by a displacement field that contains a vector for every pixel in the image. This leads to problems with high dimension and low sample size (HDLSS) for the study of anatomical variability. For example, Muller et al. (2008) have shown that the computation of shape spaces based on PCA on covariance is unstable in HDLSS settings. A number of methods are available to reduce the dimensionality of these transformations including smoothness constraints, parameterizations (such as b-splines), tangent space representations, and multi-resolution strategies. The computation of variance using these transformations will be more stable if one can aggressively reduce the number of parameters, use a more restrictive¹, lower dimensional nonlinear shape model, such as m-reps (Pizer et al., 2003), or use a combination of these strategies.

7.3 Hypothesis testing with manifold kernel regression

This section describes methods for using regression to answer the following questions: (1) are a predictor variable and a response variable independent, and (2) given two sets of observations, do their associated regression curves agree?

The first question is important because, while regression can be applied to any pair of predictor and response variables, the result is not useful—and might be misleading—if these variables are in fact independent; it is important to know if this is the case. The second question gives us the ability to compare, for example, regression curves for a healthy and diseased population in order to determine if they are significantly different.

The answer to these questions is simplified for univariate data; one can look at

¹in the sense that not all diffeomorphic transformations are admissible

a regression curve to subjectively determine if it captures the underlying relationship between the predictor and response in a reasonable way. Similarly, one can tell if the regression estimator only follows noise in the data. For higher dimensional data, such as the diffeomorphisms used in this dissertation, a visual assessment is difficult since the regressed data can change in complicated ways that are difficult to interpret. Therefore, purely analytic methods are needed. Also, in order to be applicable in the manifold setting, these methods must be based on operations that are valid for manifold-valued data.

The next two sections sketch methods that I have been developing to answer the questions posed above. These methods are based on Monte Carlo permutation testing (also called Fisher’s method of randomization, see (Bradley, 1968)) and are motivated by Pitman’s test for correlation, which uses randomization to determine if two variables are correlated (Bradley, 1968) given a collection of observations.

7.3.1 Hypothesis testing for independence of predictor and response

This section describes a method for testing the null hypothesis, H_0 , that a predictor variable and response variable are independent. The strategy is to define a permutation test that uses residual-sum-of-squares (RSS) as a test statistic. Under the null hypothesis, any observed predictor value x_i is equally likely to be assigned to any observed response value y_i . Therefore, a null distribution of RSS values can be determined empirically by randomly permuting the x_i observations, as is common for permutation tests.

Let r denote a random permutation of the N observation indices $1, 2, \dots, N$. Let r_i denote the i th index in r , i.e., $r_i \in \{1, 2, \dots, N\}$. The residual-sum-of-squares is defined

for each permutation using the regression estimator \hat{m}_h by

$$RSS = \sum_{i=1}^N (y_i - \hat{m}_h(x_{r_i}))^2. \quad (7.1)$$

Therefore, a null distribution can be constructed by repeating the following process: (1) permute the x_i observations, (2) apply regression using \hat{m}_h , (3) compute the resulting RSS value, and (4) store the RSS value in a histogram.

This method makes use of the kernel regression estimator \hat{m}_h , which is dependent on the scale (i.e., kernel bandwidth) parameter h . Therefore, separate null distributions can be constructed for a different bandwidths, and the question of the independence of the predictor and response can be analyzed across a broad range of scales (regression bandwidths).

On average, RSS should not change with randomization if x and y are independent. Therefore, H_0 can be rejected if the RSS of the observed data is in the lowest α percent of the distribution of RSS values. Using the null distribution, a significance level (p-value) may be assigned to the rejection of the H_0 , and we can say: “We reject the null hypothesis of the independence of x and y , at the scale h , at the α significance level.”

This method of hypothesis testing is applicable to manifold-valued data because it is based solely on intrinsic RSS (i.e., distance) measurements. Given a collection of manifold-valued response data, $p^i \in \mathcal{M}$, with associated predictors $t^i \in \mathbb{R}$, a manifold kernel regression estimator \hat{m}_h can be used to regress a continuous manifold-valued response as a function of t at scale h . A null distribution of residual-sum-of-squares can be created by repeatedly permuting the predictor values t_i , applying \hat{m}_h , and measuring the RSS value. This null distribution provides a confidence value for rejecting the null hypothesis that the observed predictor and manifold-valued response are independent at the scale h .

7.3.2 Hypothesis testing to determine if population trends coincide

A permutation test can also be used to test whether regression results from two populations are similar. This is important, for example, in determining if healthy and diseased populations differ.

Consider observations from populations A and B , given by predictor and response variables $\{x_i^A, y_i^A\}$ and $\{x_i^B, y_i^B\}$. Let the null hypothesis be that the observations are drawn from identical joint distributions. Under the null hypothesis, it should make no difference if observations from population A are randomly exchanged with observations from population B . Thus we can define a hypothesis test based on randomly assigning observations to the different populations.

For this test, the data are repeatedly randomized (observations are exchanged between A and B), regression is applied for each population, and RSS values are measured and tabulated in a histogram. Once the null distribution is determined in this way, it can be used to assign a significance level for rejecting the null hypothesis that the observed populations A and B were produced from identical joint distributions.

7.4 Future work

There are a number of questions and directions for future research related to this dissertation; many of these are mentioned in the conclusions of chapters 3, 4, 5, and 6. In this section a collection of these topics is briefly reviewed.

7.4.1 Properties of means and kernel regression estimators on manifolds

There are a number of unanswered questions regarding the application of the Fréchet mean and kernel regression in curved spaces. Many of these questions have direct parallels in the Euclidean kernel regression setting: How much data do we need in order to achieve a stable mean? How can we validate our results? How can we deal with edge effects? Can we apply robust regression in order to minimize the effects of outliers? How do we choose the best bandwidth? Should we use kernel regression or instead apply a parametric model? After applying regression, can we reliably solve the inverse problem (determine the predictor value from a new image)? Several questions are specific to the manifold regression setting: What are the properties of the Fréchet mean for the manifold under study? What are the convergence properties? These questions serve as a road map for future work on this project.

7.4.2 Parametric regression

This dissertation extends *kernel* regression to manifold valued data. *Parametric* regression is another common regression strategy, and there are theoretical and practical tradeoffs between these methods. Kernel regression is easier to apply and allows for an unbiased (by a parametric model) view of the data at varying scales. Parametric regression is useful when a model of the underlying process (and thus regression curve) is known; it also requires the estimation of fewer parameters.

It would be interesting to investigate parametric regression in the curved manifold setting where the parametric model corresponds to the initial momentum (tangent space representation) of the regression curve along the manifold. In the diffeomorphisms case, recent work on geodesic shooting and tangent space representatives of diffeomorphic flows (Miller et al., 2006; Vaillant et al., 2004; Beg and Khan, 2006), Ja-

cobi fields (Younes, 2007), and stationary velocity fields (Hernandez et al., 2007) provide the technical pieces that one could use to define a parametric manifold regression framework. Using geodesic shooting, for example, the initial momentum would parameterize the regression curve, i.e., the flow of diffeomorphisms.

7.4.3 Application to other shape representations and metrics

This dissertation has been focused on applying manifold kernel regression in the context of diffeomorphisms of medical images. However, manifold kernel regression can also be applied to a number of other shape descriptions and data representations that are popular within medical image analysis.

For example, m-reps, which are elements of a curved manifold with an associated Riemannian metric, have been used extensively for the statistical analysis of shapes (Pizer et al., 2003). Manifold kernel regression could be applied to m-reps using the techniques for computing Fréchet mean m-reps described in Fletcher et al. (2004).

Within the diffeomorphic framework, a number of anatomical object representations—e.g., unlabeled point sets, labeled landmarks, curves, and surfaces—have been used to drive diffeomorphic registration of anatomical data (Miller et al., 1997; Joshi and Miller, 2000; Miller et al., 2002; Glaunès et al., 2004; Vaillant and Glaunès, 2005). It is straightforward to compute Fréchet means from these data representations by replacing the image-matching term with a more general geometry-matching term. For example, Glaunès and Joshi (2006) demonstrate how to compute Fréchet means from unlabeled point sets. Combinations of these representations can be used in order to combine their various strengths. For example, landmarks can be manually added to an image-based registration problem in order drive the registration where the intensity information is vague or insufficient. Within the viscous fluid algorithm, the contributions of these different representations can be pooled by computing a body force from each representation separately and then summing the forces.

Another interesting possibility for applying regression to shapes is the combination of the diffeomorphisms of images with parametric shape descriptions, such as m-reps. Such combinations could combine the flexibility of deformable registration with the statistical strengths² of parametric shape models. The models can be combined sequentially by using one model to initialize another, e.g., to use deformable registration to capture small-scale, residual geometric change not captured by an m-rep model. These models could also be combined into a single framework by noticing that incremental geometric changes in one model can be mapped to incremental geometric changes in the other. Therefore, for example, the diffeomorphic registration procedure could be constrained by m-rep based shape statistics.

7.4.4 The inverse problem

The inverse problem is an important and useful extension to any regression methodology and a principal direction of future work for this dissertation. It allows for analysis of new response data that are collected without associated predictor values. The idea is to turn the regression problem on its head and ask: what predictor value best corresponds to a given response value? The inverse problem is important for regression of medical data because it is a method for gaining information about a new patient that can be used for classification, diagnosis, and evaluation. Within the context of Chapter 6, where the predictors are ages and the responses are brain shapes, the inverse problem is this: given a brain image from a new patient, what is the age (or more interestingly, the biological age) of the patient?

Any method for solving the inverse problem for manifold regression must take into account the curved structure of the manifold. In the rest of this section I sketch a naive solution to this problem. It remains to be seen if this solution is effective.

Given a collection of N observations (predictor and response pairs) $\{x_i, y_i\}$, manifold

²Such as the ability to stably compute variation.

kernel regression can be applied using the manifold kernel regression estimator $\hat{m}_h(x)$. Given a new predictor variable y' , the distance between y' and the regression curve is given by the function

$$f(x) = d(y', \hat{m}_h(x)). \quad (7.2)$$

The analysis of the function f leads to insight into how close y' is to any part of the regression curve. One way to look at f is that $\frac{1}{f}$ (after normalization) is a probability distribution on x where small distances between y' and $\hat{m}_h(x')$ imply that x' is a likely candidate for being the predictor that is associated with y' .

On the one hand, one can solve the inverse problem by choosing the value of x that maximizes $\frac{1}{f(x)}$. However, it seems unlikely that this procedure will be stable for high dimensional data. On the other hand, it is possible to graph $\frac{1}{f}$ and—optimistically—analyze its modes. Is there a clear mode that indicates a unique solution to the inverse problem? Are there a few modes? Are there many modes? Smoothing will undoubtedly help when analyzing this graph.

7.4.5 Managing outliers and topological change

Outliers and robust regression

Outliers cause a problem for the Fréchet mean by pulling the it toward outliers and away from the true, underlying mean. Therefore it would be useful to study methods for computing a robust manifold-valued means. One approach is to apply classical robust definitions of mean such as the trimmed mean and winzorized mean, which eliminate outliers based on orderings of the data. A similar approach is taken by Fletcher et al. (2008), who use the geometric median for computing means on manifolds rather than the Fréchet mean. The geometric median is defined as the minimizer of the sum of distances (rather than squared distances) and so it is less sensitive to the influence of

outliers.

Since manifold kernel regression is based on the computing means on a manifold, any method for computing a robust mean would translate to a mechanism for robust manifold kernel regression.

Topological changes

In the diffeomorphisms of image framework used in this dissertation, some outliers may occur because of topological changes (for example, organs sliding past each other or transient gas in the rectum) that lead to a high cost of deformation in order to match an image to the mean. In some sense, the fact that these images are outliers at all can be regarded as a failing of the diffeomorphic framework to faithfully describe anatomical motion and variability.

However, the effect of these “false” outliers on mean and regression results could be minimized using a method similar to the registration procedure described in Chapter 3. For example, a deflation transformation could be used to remove extraneous structures from input images—thus making them less like an outlier—before applying the Fréchet mean algorithm.

Longitudinal Data

Neither manifold kernel regression, as defined in this dissertation, nor individual-based diffeomorphic growth models (e.g., (Miller, 2004)) are well suited to the study of longitudinal data from multiple individuals.³ In this dissertation, manifold kernel regression is applied to random design data, where observations are assumed to be independent. In order to apply manifold kernel regression to longitudinal data, one could ignore this assumption by discarding information about repeated measures and treating all of the data as independent. However, in this case information is lost, and the underlying in-

³This is equally true for vector-space data.

dependence assumptions of the kernel regression method are ignored. A better strategy is to redefine manifold kernel regression in such a way that the longitudinal information is used. Some work along these lines is described in (Wand and Jones, 1995). Further investigation into the use of kernel regression for manifold-valued longitudinal data is an interesting direction for future work.

7.5 Other application areas

The Fréchet mean is not a new idea—it is a familiar mathematical concept that has received recent attention in the medical image analysis community, particularly for shape analysis. However, to my knowledge, a formulation of regression via the Fréchet mean where the response variable is an element of a general curved space is a novel concept. Therefore, it seems reasonable that manifold kernel regression, as a general approach (and not related to diffeomorphisms of images), will be a useful tool for scientists who deal with observed data that are elements of curved spaces. Examples of manifold-valued data can be found in many scientific fields such as physics, engineering, robotics, computer graphics, and scientific visualization. For example, curved configuration spaces are of interest in kinetic and kinematic problems. Also, the analysis of matrix valued data representing stress tensors, rotations, diffusion, etc., is very common throughout the sciences, and matrix groups often form curved manifolds.

Appendix A: Mathematical Preliminaries

A-1 Introduction

This appendix contains a brief review of the mathematical concepts used throughout this dissertation. Sections A-2 through A-7 review the elementary mathematical background of topological spaces, metric spaces, vector spaces, L^2 function spaces, groups, and differentiable geometry. Many of the definitions and propositions in these review sections are adapted from (Lee, 2003), (Boothby, 2003), and (Conway, 1990).

A-2 Topological spaces

Definition A-2.1 (Topological space (Lee p. 540)). A *topology* on a set X is a collection τ of subsets of X , called *open sets*, satisfying

- X and \emptyset (the null set) are open
- the union of any family of open sets is open
- the intersection of any finite family of open sets is open.

A pair (τ, X) consisting of a set X together with a topology τ on X is called a topological space.

A-3 Metric spaces

A metric allows one to measure distances between points in a space.

Definition A-3.1 (Metric Space). A *metric space* is a set X along with a *distance function* (or *metric*) $d : X \rightarrow \mathbb{R}$ that satisfies, for all $x, y, z \in X$,

- $d(x, y) \geq 0$
- $d(x, y) = 0$ if and only if $x = y$
- $d(x, y) = d(y, x)$
- $d(x, y) + d(y, z) \geq d(x, z)$.

The usual distance function in Euclidean space is an example: $d(x, y) = \sqrt{(x - y)^2}$. However, metrics exist for many more complicated spaces, such as Hilbert spaces and curved manifolds, which are defined below.

A-4 Vector spaces

Definition A-4.1 (Vector Space (Lee p. 558, Wikipedia)). A real *vector space* is a set endowed with two operations: *vector addition* $V \times V \rightarrow V$, denoted by $x + y$ for $x, y \in V$, and *scalar multiplication* $\mathbb{R} \times V \rightarrow V$, denoted by αx for $\alpha \in \mathbb{R}$. These operations are required to satisfy (for all $x, y, z \in V$ and $a, b \in \mathbb{R}$)

- Vector addition is associative: $x + (y + z) = (x + y) + z$
- Vector addition is commutative: $x + y = y + x$
- Vector addition has an identity element: $\text{Id}_V + x = x$
- Vector addition has an inverse element: there exists $x^{-1} \in V$ such that $x + x^{-1} = \text{Id}_V$
- Scalar multiplication satisfies $a(bx) = (ab)x$
- Scalar multiplication has an identity element: $1x = x$ where 1 denotes the multiplicative identity

- Scalar multiplication and vector addition are related by the following distributive laws:

$$(a + b)x = ax + bx \qquad a(x + y) = ax + ay$$

The terminology “vector space” will be used to denote a real vector space unless otherwise indicated.

A-4.1 Hilbert space

Definition A-4.2 (Hilbert space (Conway p. 4)). A *Hilbert space* is a vector space \mathcal{H} with an inner product $\langle \cdot, \cdot \rangle$ such that relative to the metric induced by the inner product, $d(x, y) = \|x - y\| = \sqrt{\langle x - y, x - y \rangle}$, \mathcal{H} is a complete metric space. That is, every Cauchy sequence of points in \mathcal{H} has a limit that is also in \mathcal{H} .

A-4.2 Images as L^2 functions

In this dissertation, images are modeled as functions that have a Hilbert space structure. Given a domain $\Omega \equiv [0, 1] \times [0, 1] \times [0, 1] \subset \mathbb{R}^3$, the functions $L^2(\Omega) : \Omega \rightarrow \mathbb{R}$ are the square integrable functions on Ω . That is, for $f \in L^2(\Omega)$,

$$\int_{\Omega} |f|^2 d\mu < \infty. \tag{A-1}$$

There are no further constraints on L^2 functions such as continuity or smoothness. $L^2(\Omega)$ forms a Hilbert space with the inner product given by

$$\langle f, g \rangle_{L^2(\Omega)} \equiv \int_{\Omega} f \bar{g} dx. \tag{A-2}$$

An image is formally defined as an element of $\mathcal{I} \equiv L^2(\Omega)$.

A-5 Properties of mappings

Let $f : U \rightarrow V$ be a function that maps points in the topological space U to points in the topological space V . The function f is a *bijection* if it is one-to-one and onto. The function f is a *homeomorphism* if (1) f is a bijection, (2) f is continuous⁴, and (3) f^{-1} exists and is continuous. The function f is a *C^r -diffeomorphism* if it is a homeomorphism and both f and f^{-1} are r -times differentiable.

A-6 Group theory

The notion of a group is central to this dissertation: group theory provides the abstract mathematical language used to describe collections of spatial transformations that can be used to deform anatomical structures. In this section the basic definitions are reviewed. See any introductory algebra text for more details (e.g., (Clark, 1984)).

Definition A-6.1 (Group). A *group* (G, \cdot) is a nonempty set G together with a binary operation $\cdot : G \times G \rightarrow G$ satisfying for all a , b , and c in G :

Associativity : $(a \cdot b) \cdot c = a \cdot (b \cdot c)$,

Identity Element : There exists $e \in G$ such that $e \cdot a = a \cdot e = a$, and

Inverse Element : There exists $a^{-1} \in G$ such that $a \cdot a^{-1} = a^{-1} \cdot a = e$.

The group operation is often abbreviated by concatenation, i.e., $ab \equiv a \cdot b$.

The group operation is not generally commutative; a group G where $ab = ba$ for all a and b in G is called *Abelian*. The following basic theorems follow from these axioms

- a group has exactly one identity element,
- every group element has a unique inverse,

⁴That is, for every open set $A \subset V$, $B = f^{-1}(A)$ is open in U .

- $(a^{-1})^{-1} = a$,
- $(ab)^{-1} = b^{-1}a^{-1}$,
- For all $a, b \in G$, there exists a unique $x \in G$ such that $ax = b$,
- $ax = ay \Rightarrow x = y$ and $xa = ya \Rightarrow x = y$.

Definition A-6.2 (Subgroup). A subset H of G is called a *subgroup* of (G, \cdot) , written $H \subseteq G$, if H satisfies the axioms of a group, using \cdot , restricted to the subset H . A *proper subgroup* of G is a subgroup which is not identical to G . A subgroup is *non-trivial* if it contains elements other than the identity element.

The following theorems relate to subgroups:

- the identity in H is identical to the identity in G ,
- the inverse of h in H is identical to the inverse of h in G ,
- a non-empty subset S of G is a subgroup iff $a \cdot b^{-1} \in S$ for all $a, b \in S$,
- the intersection of two or more subgroups is again a subgroup.

A-7 Differential geometry

In mathematics a manifold is a space that is locally Euclidean (i.e., locally flat). In this section the basic definitions and properties of a smooth manifold are reviewed. Manifolds are important for this dissertation since many transformation groups, such as the group of diffeomorphisms, form manifolds.

An n -dimensional manifold \mathcal{M} is a topological space along with a collection of coordinate charts $\mathcal{U} = \{U_\alpha, \phi_\alpha\}$ that map open subsets of \mathcal{M} (the U_α) into \mathbb{R}^n via the homeomorphic functions $\phi_\alpha : U_\alpha \rightarrow \mathbb{R}^n$. This dissertation deals exclusively with manifolds for which any change of coordinates for overlapping charts must be smooth (e.g., for $\phi_\alpha \circ \phi_\beta^{-1}$). Detailed definitions of these concepts follow.

Definition A-7.1 (Manifold of dimension n (Boothby p. 6)). A *manifold* \mathcal{M} of *dimension* n is a topological space with the following properties:

- every point $p \in \mathcal{M}$ has a neighborhood U which is homeomorphic to an open subset U' of \mathbb{R}^n
- \mathcal{M} is Hausdorff
- \mathcal{M} has a countable basis of open sets.

The first property gives a precise meaning to the intuitive idea of a *locally flat* space. It is important to emphasize that there is no requirement that \mathcal{M} be *globally* homeomorphic to any flat space. The second and third properties are not explicitly relevant to this discussion; they serve to exclude poorly behaved topological spaces that are not considered in this dissertation.

Suppose that U is an open subset of \mathcal{M} and ϕ is a homeomorphism from U to an open subset of \mathbb{R}^n . The pair U, ϕ is called a *coordinate chart*. For two overlapping charts U, ϕ and V, ψ there exist the homeomorphic *change of coordinates*

$$\phi \circ \psi^{-1} : \psi(U \cap V) \rightarrow \phi(U \cap V) \quad \text{and} \quad \psi \circ \phi^{-1} : \phi(U \cap V) \rightarrow \psi(U \cap V). \quad (\text{A-3})$$

When these change-of-coordinate functions satisfy additional smoothness constraints, \mathcal{M} is known as a smooth manifold. The following definitions make this idea precise.

Definition A-7.2 (C^∞ -compatible Charts (Boothby p. 52)). Two charts U, ϕ and V, ψ are C^∞ -compatible if $U \cap V$ nonempty implies that $\phi \circ \psi^{-1}$ and $\psi \circ \phi^{-1}$ are diffeomorphisms of the open subsets $\phi(U \cap V)$ and $\psi(U \cap V)$ of \mathbb{R}^n .

Definition A-7.3 (Smooth structure (Boothby p. 53)). A *smooth structure* on a topological manifold \mathcal{M} is a family $\mathcal{U} = \{U_\alpha, \phi_\alpha\}$ of coordinate charts such that

- the U_α cover \mathcal{M}

- U_α, ϕ_α and U_β, ϕ_β are C^∞ -compatible for any α, β
- any coordinate neighborhood V, ψ compatible with every $U_\alpha, \phi_\alpha \in \mathcal{U}$ is itself in \mathcal{U} .

Definition A-7.4 (Smooth Manifold (Boothby p. 53)). A *smooth manifold* is a topological manifold together with a smooth differentiable structure.

A tangent space can be attached to every point p on a smooth manifold. The tangent space is a vector space that contains the possible “directions” in which one can pass while traveling on a path along \mathcal{M} through p . In \mathbb{R}^n tangent vectors have a simple geometric interpretation through their natural identification with elements of \mathbb{R}^n . On general manifolds tangent vectors are defined as directional derivative operators on smooth functions.

Definition A-7.5 (Tangent Space (Boothby p. 104)). Let $C^\infty(p)$ be the algebra of C^∞ functions whose domain of definition includes some open neighborhood of p , with functions identified if they agree on any neighborhood of p . A *tangent space* $T_p(\mathcal{M})$ to \mathcal{M} at p is the set of all mappings $X_p : C^\infty(p) \rightarrow \mathbb{R}$ such that for all $\alpha, \beta \in \mathbb{R}$ and $f, g \in C^\infty(p)$ the mappings X_p are linear and Leibniz. That is,

- $X_p(\alpha f + \beta g) = \alpha(X_p f) + \beta(X_p g)$
- $X_p(fg) = (X_p f)g(p) + f(p)(X_p g)$.

A *tangent vector to \mathcal{M} at p* is any $X_p \in T_p(\mathcal{M})$.

For a manifold of dimension n , elements of $T_p(\mathcal{M})$ can be identified with elements of a vector space V^n . The particular identification depends on the choice of coordinates.

A-7.1 Measuring distances on manifolds

Metrics can be defined on curved spaces such as differentiable manifolds. Instead of measuring along straight lines, these distance functions measure distance along paths on a manifold. This construction requires additional structure for the manifold \mathcal{M} .

Definition A-7.6 (Riemannian manifold). A *Riemannian manifold* is a smooth manifold \mathcal{M} along with a smoothly varying inner product J_p defined at every point $p \in \mathcal{M}$. The length, L_R , of a path $\gamma : [0, 1] \rightarrow \mathcal{M}$ is defined by the integral equation

$$L_R(\gamma) = \int_0^1 \sqrt{\langle \dot{\gamma}(\tau), \dot{\gamma}(\tau) \rangle_{J_p}} d\tau \quad (\text{A-4})$$

where $\dot{\gamma}(\tau) \in T_{\gamma(\tau)}(\mathcal{M})$ is interpreted as the velocity of the path γ at the point $p = \gamma(\tau)$.

The *Riemannian distance*, $d_R(\cdot, \cdot)$, between any two points p and q on \mathcal{M} is defined as the infimum of this integral over all piecewise smooth curves γ that connect p and q . That is, for any points p, q on a Riemannian manifold \mathcal{M} , the Riemannian distance is defined by

$$d_R(p, q) = \inf_{\gamma: [0,1] \rightarrow \mathcal{M}, \gamma(0)=p, \gamma(1)=q} \int_0^1 \sqrt{\langle \dot{\gamma}(\tau), \dot{\gamma}(\tau) \rangle_{J_p}} d\tau. \quad (\text{A-5})$$

Note that the minimizing curves may not be unique.⁵

Proposition A-7.1 (Adapted from Lee proposition 11.20). *A connected Riemannian manifold, endowed with its Riemannian metric, is a metric space.*

Proof. The proof is sketched here. See (Lee, 2003) Proposition 11.20 for more details.

It must be shown that the Riemannian distance function d_R satisfies the properties of a metric (see section A-3).

- That $d_R(p, q) \geq 0$ for all $p, q \in \mathcal{M}$ is immediate from the definition of Riemannian distance.
- $d_R(p, p) = 0$ because any constant curve segment has length zero.
- Symmetry: Because any curve from p to q can be reparameterized to go from q to p , it follows that $d_R(p, q) = d_R(q, p)$.

⁵Consider the curves connecting the north and south pole of a sphere.

- Triangle inequality: Suppose γ_1 and γ_2 are piecewise smooth curves segments from p to q and from q to r , respectively. Let γ be a piecewise smooth curve segment that first follows γ_1 and then follows γ_2 . In this case

$$d_R(p, r) \leq L(\gamma) = L(\gamma_1) + L(\gamma_2). \quad (\text{A-6})$$

Since the distance is taken as the infimum over all such γ_1 and γ_2 , it follows that $d_R(p, r) \leq d_R(p, q) + d_R(q, r)$.

- Positivity: the last property to check is $d_R(p, q) > 0$ if $p \neq q$. A nonzero lower bound for $d_R(p, q)$ can be defined using a coordinate chart containing an open neighborhood U of p . Let p_2 be a point on the curve γ such that p_2 is also inside U and the Euclidean distance (in coordinates) between p and p_2 is ϵ . According to Lemma 11.19 of Lee, $d_R(p, q) \geq c\epsilon$ where c is a positive, nonzero constant. Therefore, (p, q) is strictly positive.

□

Appendix B: Euler-Lagrange equation for LDDMM image mapping

This appendix derives the Euler-Lagrange equation that is used to construct a gradient descent algorithm for the large-deformations diffeomorphic metric mapping problem. See Chapter 2 for appropriate context and notation. This derivation was published in (Beg, 2003) and (Beg et al., 2005).

Theorem B-0.3 provides a detailed derivation of the Euler-Lagrange equation. This theorem depends on the following lemma, which describes the effect of a perturbation of a velocity field on a deformation.

Let the deformation $\phi_{s,t}$ be defined by the composition of ϕ_t and ϕ_s^{-1} , i.e.,

$$\phi_{s,t} \equiv \phi_t \circ \phi_s^{-1}. \quad (\text{B-1})$$

Lemma B-0.2 (Beg (2003); Beg et al. (2005)). *The variation of the mapping $\phi_{s,t}^v$ when $v \in L^2([0, 1], V)$ is perturbed along $h \in L^2([0, 1], V)$ is given by*

$$\partial_h \phi_{s,t}^v = \lim_{\epsilon \rightarrow 0} \frac{\phi_{s,t}^{v+\epsilon h} - \phi_{s,t}^v}{\epsilon} \quad (\text{B-2})$$

$$= D\phi_{s,t}^v \int_s^t (D\phi_{s,u}^v)^{-1} h_u \circ \phi_{s,u}^v du, \quad (\text{B-3})$$

where $D\phi_{s,t}^v$ is the Jacobian of the mapping $\phi_{s,t}^v$, i.e.,

$$(D\phi_{s,t}^v)_{i,j} = \frac{\partial(\phi_{s,t}^v)_i}{\partial x_j}. \quad (\text{B-4})$$

Proof. In Beg (2003), this lemma is proved by solving for $\frac{d}{dt} (\partial_h \phi_{s,t}^v)$ and integrating it over time. The derivation of $\frac{d}{dt} (\partial_h \phi_{s,t}^v)$ starts with the computation of $\frac{d}{dt} (\phi_{s,t}^{v+\epsilon h})$ and

depends on the definition of flow and Taylor series expansions of $\phi_{s,t}^{v+\epsilon h}$, v_t , and h_t . \square

The intuition behind Equation (B-3) is the following: perturbations $h_u \circ \phi_{s,u}^v$ act on a particle traveling through time and space at times $u \in [s, t]$. These perturbations are accumulated by pulling them back to time s (via the inverse Jacobian matrix $(D\phi_{s,u}^v)^{-1}$). Finally, the accumulated perturbations are pushed forward to time t using the Jacobian matrix $D\phi_{s,t}^v$.

Theorem B-0.3 (Beg (2003)). *Given a continuously differentiable idealized template image I_0 and a noisy observation of anatomy I_1 , then $\hat{v} \in L^2([0, 1], V)$ for inexact matching of I_0 and I_1 is given by*

$$\hat{v} = \underset{v \in L^2([0, 1], V)}{\operatorname{arginf}} E(v) \equiv \int_0^1 \|v_t\|_V^2 dt + \frac{1}{\sigma^2} \|I_0 \circ \phi_{1,0} - I_1\|_{L^2}^2. \quad (\text{B-5})$$

Equation (B-5) satisfies the Euler-Lagrange equation given by

$$2\hat{v} - K \left(\frac{2}{\sigma^2} |D\phi_{t,1}^v| (J_t^0 - J_t^1) \nabla J_t^0 \right) = 0 \quad (\text{B-6})$$

where $J_t^0 \equiv I_0 \circ \phi_{t,0}^v$ and $J_t^1 \equiv I_1 \circ \phi_{t,1}^v$.

Proof. The derivation is based on the computation of the Gateaux variation for each term in equation (B-5) for an arbitrary perturbation $h \in L^2([0, 1], V)$. Since the variation is linear, these terms can be computed separately and later combined.

The Gateaux variation of the first term

$$E_1(v) = \int_0^1 \|v_t\|_V^2 dt \quad (\text{B-7})$$

is given by

$$\partial_h E_1(v) = \lim_{\epsilon \rightarrow 0} \frac{E_1(v + \epsilon h) - E_1(v)}{\epsilon} \quad (\text{B-8})$$

$$= \lim_{\epsilon \rightarrow 0} \frac{\int_0^1 \|v_t + \epsilon h_t\|_V^2 dt - \int_0^1 \|v_t\|_V^2 dt}{\epsilon} \quad (\text{B-9})$$

$$= \lim_{\epsilon \rightarrow 0} \frac{\int_0^1 \langle v_t + \epsilon h_t, v_t + \epsilon h_t \rangle_V dt - \int_0^1 \langle v_t, v_t \rangle_V dt}{\epsilon} \quad (\text{B-10})$$

$$= \lim_{\epsilon \rightarrow 0} \frac{\epsilon^2 \int_0^1 \|h_t\|_V^2 dt + 2\epsilon \int_0^1 \langle v_t, h_t \rangle_V dt}{\epsilon} \quad (\text{B-11})$$

$$= 2 \int_0^1 \langle v_t, h_t \rangle_V dt. \quad (\text{B-12})$$

The Gateaux variation of the second term

$$E_2(v) = \frac{1}{\sigma^2} \|I_0 \circ \phi_{1,0} - I_1\|_{L^2}^2 \quad (\text{B-13})$$

is given by

$$\lim_{\epsilon \rightarrow 0} \frac{E_2(v + \epsilon h) - E_2(v)}{\epsilon} \quad (\text{B-14})$$

$$= \frac{1}{\sigma^2} \lim_{\epsilon \rightarrow 0} \frac{\|I_0 \circ \phi_{1,0}^{v+\epsilon h} - I_1\|_{L^2}^2 - \|I_0 \circ \phi_{1,0}^v - I_1\|_{L^2}^2}{\epsilon} \quad (\text{B-15})$$

$$= \frac{1}{\sigma^2} \lim_{\epsilon \rightarrow 0} \frac{\langle I_0 \circ \phi_{1,0}^{v+\epsilon h} - I_1, I_0 \circ \phi_{1,0}^{v+\epsilon h} - I_1 \rangle_{L^2} - \langle I_0 \circ \phi_{1,0}^v - I_1, I_0 \circ \phi_{1,0}^v - I_1 \rangle_{L^2}}{\epsilon} \quad (\text{B-16})$$

$$= \frac{1}{\sigma^2} \lim_{\epsilon \rightarrow 0} \frac{\langle I_0 \circ \phi_{1,0}^v + DI_0 \circ \phi_{1,0}^v \epsilon \partial_h \phi_{1,0}^v - I_1, I_0 \circ \phi_{1,0}^v + DI_0 \circ \phi_{1,0}^v \epsilon \partial_h \phi_{1,0}^v - I_1 \rangle_{L^2}}{\epsilon} \quad (\text{B-17})$$

$$= \frac{1}{\sigma^2} \lim_{\epsilon \rightarrow 0} \frac{\epsilon^2 \langle DI_0 \circ \phi_{1,0}^v \partial_h \phi_{1,0}^v, DI_0 \circ \phi_{1,0}^v \partial_h \phi_{1,0}^v \rangle_{L^2} + 2\epsilon \langle I_0 \circ \phi_{1,0}^v, DI_0 \circ \phi_{1,0}^v \partial_h \phi_{1,0}^v \rangle_{L^2}}{\epsilon} \quad (\text{B-18})$$

$$= \frac{2}{\sigma^2} \langle I_0 \circ \phi_{1,0}^v - I_1, DI_0 \circ \phi_{1,0}^v \partial_h \phi_{1,0}^v \rangle_{L^2} \quad (\text{B-19})$$

where (B-17) follows from the first order expansion $I_0 \circ \phi_{1,0}^{v+\epsilon h} = I_0 \circ \phi_{1,0}^v + DI_0 \circ \phi_{1,0}^v \epsilon \partial_h \phi_{1,0}^v$.

Next, applying Lemma B-0.2 we get

$$= \frac{2}{\sigma_2} \left\langle I_0 \circ \phi_{1,0}^v - I_1, DI_0 \circ \phi_{1,0}^v \left(D\phi_{1,0}^v \int_1^0 (D\phi_{1,t}^v)^{-1} h_{t \circ \phi_{1,t}^v} dt \right) \right\rangle_{L^2} \quad (\text{B-20})$$

$$= \frac{-2}{\sigma_2} \left\langle I_0 \circ \phi_{1,0}^v - I_1, DI_0 \circ \phi_{1,0}^v \left(D\phi_{1,0}^v \int_0^1 (D\phi_{1,t}^v)^{-1} h_{t \circ \phi_{1,t}^v} dt \right) \right\rangle_{L^2} \quad (\text{B-21})$$

$$= \frac{-2}{\sigma_2} \int_0^1 \left\langle I_0 \circ \phi_{1,0}^v - I_1, D(I_0 \circ \phi_{1,0}^v) (D\phi_{1,t}^v)^{-1} h_{t \circ \phi_{1,t}^v} \right\rangle_{L^2} dt \quad (\text{B-22})$$

$$= \frac{-2}{\sigma_2} \int_0^1 \left\langle |D\phi_{t,1}^v| I_0 \circ \phi_{t,0}^v - I_1 \phi_{t,1}^v, D(I_0 \circ \phi_{t,0}^v) h_t \right\rangle_{L^2} dt \quad (\text{B-23})$$

$$= \frac{-2}{\sigma_2} \int_0^1 \left\langle |D\phi_{t,1}^v| (I_0 \circ \phi_{t,0}^v - I_1 \phi_{t,1}^v) D(I_0 \circ \phi_{t,0}^v), h_t \right\rangle_{L^2} dt \quad (\text{B-24})$$

$$= - \int_0^1 \left\langle K \left(\frac{2}{\sigma_2} |D\phi_{t,1}^v| (I_0 \circ \phi_{t,0}^v - I_1 \phi_{t,1}^v) D(I_0 \circ \phi_{t,0}^v) \right), h_t \right\rangle_V dt \quad (\text{B-25})$$

$$= - \int_0^1 \left\langle K \left(\frac{2}{\sigma_2} |D\phi_{t,1}^v| (J_t^0 - J_t^1) \nabla J_t^0 \right), h_t \right\rangle_V dt \quad (\text{B-26})$$

where (B-22) follows from the identity $D(I_0 \circ \phi_{1,0}^v) = DI_0 \circ \phi_{1,0} D\phi_{1,0}$; (B-23) follows from a change of variable for L^2 inner product $y = \phi_{t,1}^v(x)$; (B-25) follows from the definition of the inner product on V (see Chapter 2); and (B-26) follows from a relabeling $J_t^0 \equiv I_0 \circ \phi_{t,0}^v$ and $J_t^1 \equiv I_1 \circ \phi_{t,1}^v$.

Since the variation is linear, these terms can be combined. Thus, the optimizing velocity field satisfies

$$\partial_h E(\hat{v}) = \int_0^1 \left\langle 2\hat{v} - K \left(\frac{2}{\sigma_2} |D\phi_{t,1}^v| (J_t^0 - J_t^1) \nabla J_t^0 \right), h_t \right\rangle dt = 0. \quad (\text{B-27})$$

Since h is arbitrary in $L^2([0, 1], V)$, by the fundamental lemma of calculus of variations (see, for example, Luenberger (1969)), we get Equation (B-6). \square

Appendix C: Derivation of the differential operator for “fluid” image registration

C-1 Introduction

This appendix summarizes the relationship between the physics of fluid flow and differential operator

$$L = -\alpha\Delta - \beta\nabla(\nabla\cdot) + \gamma I \quad (\text{C-1})$$

that is used in “fluid” image registration. This operator relates a known body force term b , which is generated from image intensity and gradient information, with a velocity term v , which describes the instantaneous motion of the deforming continuum, according to

$$L^\dagger L \vec{v} = \vec{b}. \quad (\text{C-2})$$

This type of fluid registration was introduced by Christensen et al. (1994, 1996). This and other operators used for continuum-mechanics based operators in image registration are review in Holden (2008). This appendix summarizes the explanations from both of these sources.

C-2 The derivation of L

First, the conservation of momentum with a mass source, defining the relationship between body force and the resulting material deformation is

$$\rho \frac{d\vec{v}}{dt} + \vec{v}\eta - \nabla \cdot T - \vec{b} = 0 \quad (\text{C-3})$$

where ρ is the density of the material, v is the velocity of the material, η measures mass insertion per unit volume, T is the Cauchy stress tensor, and b is the body force. Now we must specify the constitutive behavior of the material by specifying the form of the stress tensor T .

The rate of deformation tensor is $D = \frac{1}{2}(\nabla \vec{v} + (\nabla \vec{v})^T)$. For a Navier-Poisson Newtonian fluid, the Cauchy stress tensor is related to D by

$$T = (\lambda(\text{tr} D) - p)I + 2\mu D \quad (\text{C-4})$$

where μ and λ are viscosity constants and p is pressure.

Now, substitute equation C-4 into $\nabla \cdot T$ from equation C-3 to get (note that divergence applies row-wise to a matrix):

$$\begin{aligned} \nabla \cdot T &= \nabla \cdot [(\lambda(\text{tr} D) - p)I + 2\mu D] \\ &= \nabla \cdot \left[\left(\lambda \left[\frac{\partial v_1}{\partial x} + \frac{\partial v_2}{\partial y} + \frac{\partial v_3}{\partial z} \right] - p \right) I + \mu \begin{pmatrix} 2\frac{\partial v_1}{\partial x} & \frac{\partial v_1}{\partial y} + \frac{\partial v_2}{\partial x} & \frac{\partial v_1}{\partial z} + \frac{\partial v_3}{\partial x} \\ \frac{\partial v_2}{\partial x} + \frac{\partial v_1}{\partial y} & 2\frac{\partial v_2}{\partial y} & \frac{\partial v_2}{\partial z} + \frac{\partial v_3}{\partial y} \\ \frac{\partial v_3}{\partial x} + \frac{\partial v_1}{\partial z} & \frac{\partial v_3}{\partial y} + \frac{\partial v_2}{\partial z} & 2\frac{\partial v_3}{\partial z} \end{pmatrix} \right] \\ &= \lambda \nabla (\nabla \cdot \vec{v}) - \nabla p + \mu \begin{pmatrix} 2\frac{\partial^2 v_1}{\partial x^2} + \frac{\partial^2 v_1}{\partial y^2} + \frac{\partial^2 v_2}{\partial xy} + \frac{\partial^2 v_1}{\partial z^2} + \frac{\partial^2 v_3}{\partial xz} \\ \frac{\partial^2 v_2}{\partial x^2} + \frac{\partial^2 v_1}{\partial yx} + 2\frac{\partial^2 v_2}{\partial y^2} + \frac{\partial^2 v_2}{\partial z^2} + \frac{\partial^2 v_3}{\partial yz} \\ \frac{\partial^2 v_3}{\partial x^2} + \frac{\partial^2 v_1}{\partial zx} + \frac{\partial^2 v_3}{\partial y^2} + \frac{\partial^2 v_2}{\partial zy} + 2\frac{\partial^2 v_3}{\partial z^2} \end{pmatrix} \\ &= (\lambda + \mu) \nabla (\nabla \cdot \vec{v}) - \nabla p + \mu \begin{pmatrix} \frac{\partial^2 v_1}{\partial x^2} + \frac{\partial^2 v_1}{\partial y^2} + \frac{\partial^2 v_1}{\partial z^2} \\ \frac{\partial^2 v_2}{\partial x^2} + \frac{\partial^2 v_2}{\partial y^2} + \frac{\partial^2 v_2}{\partial z^2} \\ \frac{\partial^2 v_3}{\partial x^2} + \frac{\partial^2 v_3}{\partial y^2} + \frac{\partial^2 v_3}{\partial z^2} \end{pmatrix} \\ &= (\lambda + \mu) \nabla (\nabla \cdot \vec{v}) - \nabla p + \mu \Delta \vec{v} \end{aligned}$$

Substituting this result into equation C-3 gives

$$\mu \Delta \vec{v} + (\lambda + \mu) \nabla (\nabla \cdot \vec{v}) + \vec{b} = \nabla p + \rho \frac{d\vec{v}}{dt} + \vec{v} \eta. \quad (\text{C-5})$$

Finally, we simplify the model by assuming a very low Reynold's number (where viscous forces dominate) where the pressure gradient ∇p and the inertial terms $\rho \frac{d\vec{v}}{dt}$ and $\vec{v}\eta$ are neglected, so that we get

$$\mu\Delta\vec{v} + (\lambda + \mu)\nabla(\nabla \cdot \vec{v}) + \vec{b} = 0. \quad (\text{C-6})$$

After substituting $\alpha \equiv \mu$ and $\beta = \mu + \lambda$, our operator is then $L = -\mu\Delta - (\lambda + \mu)\nabla(\nabla \cdot)$.

Finally, the term γI is added so that the operator L is guaranteed to be invertible. This is important since the differential equation (C-2) is solved in the Fourier domain by inverting L . See Appendix D for more details.

Appendix D: Numerical solution for velocity fields in fluid registration

D-1 Introduction

All of the image registration methods used in this dissertation are based on fluid registration. This appendix shows how to numerically solve for the velocity fields at each time step in the fluid image registration algorithm. The objective is to solve the equation

$$L^\dagger L v = b \tag{D-1}$$

where L is a differential operator (with adjoint L^\dagger), v is a velocity field, and b is a body force used to drive the registration.

D-2 Operators

Two differential operators are used in this dissertation. The first is motivated by the Navier-Stokes equation for compressible fluids with a very low Reynolds number

$$L = -\alpha \Delta - \beta \nabla(\nabla \cdot) + \gamma I. \tag{D-2}$$

In this equation $\nabla = \left[\frac{\partial}{\partial x} \frac{\partial}{\partial y} \frac{\partial}{\partial z} \right]^T$ is the gradient operator and $\Delta = \nabla \cdot \nabla = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}$ is the Laplacian operator (applied component-wise). The viscosity of the fluid is determined by α and β . The γ term, which can be thought of as a “body friction” term, ensures that the operator is positive definite, and hence invertible (Joshi and Miller, 2000). This operator was introduced for fluid image registration in Christensen et al.

(1996). However, Christensen did not use the γ term. The second operator is similar except that the β term is neglected

$$L_B = -\alpha\Delta + \gamma I. \quad (\text{D-3})$$

This operator is used in Beg et al. (2005). Holden (2008) provides a good review of these operators.

In expanded notation, these operators are

$$Lv = (-\alpha\Delta - \beta\nabla(\nabla\cdot) + \gamma I) v \quad (\text{D-4})$$

$$= -\alpha\Delta \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} - \beta\nabla \left(\nabla \cdot \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} \right) + \gamma \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} \quad (\text{D-5})$$

$$= -\alpha \begin{bmatrix} \frac{\partial^2 v_1}{\partial x^2} + \frac{\partial^2 v_1}{\partial y^2} + \frac{\partial^2 v_1}{\partial z^2} \\ \frac{\partial^2 v_2}{\partial x^2} + \frac{\partial^2 v_2}{\partial y^2} + \frac{\partial^2 v_2}{\partial z^2} \\ \frac{\partial^2 v_3}{\partial x^2} + \frac{\partial^2 v_3}{\partial y^2} + \frac{\partial^2 v_3}{\partial z^2} \end{bmatrix} - \beta\nabla \left(\frac{\partial v_1}{\partial x} + \frac{\partial v_2}{\partial y} + \frac{\partial v_3}{\partial z} \right) + \gamma \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} \quad (\text{D-6})$$

$$= -\alpha \begin{bmatrix} \frac{\partial^2 v_1}{\partial x^2} + \frac{\partial^2 v_1}{\partial y^2} + \frac{\partial^2 v_1}{\partial z^2} \\ \frac{\partial^2 v_2}{\partial x^2} + \frac{\partial^2 v_2}{\partial y^2} + \frac{\partial^2 v_2}{\partial z^2} \\ \frac{\partial^2 v_3}{\partial x^2} + \frac{\partial^2 v_3}{\partial y^2} + \frac{\partial^2 v_3}{\partial z^2} \end{bmatrix} - \beta \begin{bmatrix} \frac{\partial^2 v_1}{\partial x^2} + \frac{\partial^2 v_2}{\partial yx} + \frac{\partial^2 v_3}{\partial zx} \\ \frac{\partial^2 v_1}{\partial xy} + \frac{\partial^2 v_2}{\partial y^2} + \frac{\partial^2 v_3}{\partial zy} \\ \frac{\partial^2 v_1}{\partial xz} + \frac{\partial^2 v_2}{\partial yz} + \frac{\partial^2 v_3}{\partial z^2} \end{bmatrix} + \gamma \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} = \quad (\text{D-7})$$

$$\begin{bmatrix} -(\alpha+\beta)\frac{\partial^2}{\partial x^2} - \alpha\frac{\partial^2}{\partial y^2} - \alpha\frac{\partial^2}{\partial z^2} + \gamma & -\beta\frac{\partial^2}{\partial xy} & -\beta\frac{\partial^2}{\partial xz} \\ -\beta\frac{\partial^2}{\partial xy} & -\alpha\frac{\partial^2}{\partial x^2} - (\alpha+\beta)\frac{\partial^2}{\partial y^2} - \alpha\frac{\partial^2}{\partial z^2} + \gamma & -\beta\frac{\partial^2}{\partial yz} \\ -\beta\frac{\partial^2}{\partial xz} & -\beta\frac{\partial^2}{\partial yz} & -\alpha\frac{\partial^2}{\partial x^2} - \alpha\frac{\partial^2}{\partial y^2} - (\alpha+\beta)\frac{\partial^2}{\partial z^2} + \gamma \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} \quad (\text{D-8})$$

and

$$L_B v = (-\alpha \Delta + \gamma I) v \quad (\text{D-9})$$

$$= -\alpha \Delta \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} + \gamma \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} \quad (\text{D-10})$$

$$= -\alpha \begin{bmatrix} \frac{\partial^2 v_1}{\partial x^2} + \frac{\partial^2 v_1}{\partial y^2} + \frac{\partial^2 v_1}{\partial z^2} \\ \frac{\partial^2 v_2}{\partial x^2} + \frac{\partial^2 v_2}{\partial y^2} + \frac{\partial^2 v_2}{\partial z^2} \\ \frac{\partial^2 v_3}{\partial x^2} + \frac{\partial^2 v_3}{\partial y^2} + \frac{\partial^2 v_3}{\partial z^2} \end{bmatrix} + \gamma \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} = \quad (\text{D-11})$$

$$\begin{bmatrix} -\alpha(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}) + \gamma & 0 & 0 \\ 0 & -\alpha(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}) + \gamma & 0 \\ 0 & 0 & -\alpha(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}) + \gamma \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix}. \quad (\text{D-12})$$

The equations for two dimensions are similar

$$L v = \begin{bmatrix} -(\alpha + \beta) \frac{\partial^2}{\partial x^2} - \alpha \frac{\partial^2}{\partial y^2} + \gamma & -\beta \frac{\partial^2}{\partial xy} \\ -\beta \frac{\partial^2}{\partial xy} & -\alpha \frac{\partial^2}{\partial x^2} - (\alpha + \beta) \frac{\partial^2}{\partial y^2} + \gamma \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} \quad (\text{D-13})$$

and

$$L_B v = \begin{bmatrix} -\alpha(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}) + \gamma & 0 \\ 0 & -\alpha(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}) + \gamma \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \end{bmatrix}. \quad (\text{D-14})$$

D-3 Finite difference approximations

In this section, the finite difference approximations of the operators L and L_B are derived. Second order approximations are used; for some function $u : \mathbb{R}^3 \rightarrow \mathbb{R}$ they are

$$\frac{\partial^2 u}{\partial x^2}_{x,y,z} \approx u_{x+1,y,z} - 2u_{x,y,z} + u_{x-1,y,z} \quad (\text{D-15})$$

$$\frac{\partial^2 u}{\partial y^2}_{x,y,z} \approx u_{x,y+1,z} - 2u_{x,y,z} + u_{x,y-1,z} \quad (\text{D-16})$$

$$\frac{\partial^2 u}{\partial z^2}_{x,y,z} \approx u_{x,y,z+1} - 2u_{x,y,z} + u_{x,y,z-1} \quad (\text{D-17})$$

$$\frac{\partial^2 u}{\partial xy}_{x,y,z} = \frac{\partial^2 u}{\partial yx}_{x,y,z} \approx \frac{1}{4} [u_{x-1,y-1,z} + u_{x+1,y+1,z} - u_{x+1,y-1,z} - u_{x-1,y+1,z}] \quad (\text{D-18})$$

$$\frac{\partial^2 u}{\partial xz}_{x,y,z} = \frac{\partial^2 u}{\partial zx}_{x,y,z} \approx \frac{1}{4} [u_{x-1,y,z-1} + u_{x+1,y,z+1} - u_{x+1,y,z-1} - u_{x-1,y,z+1}] \quad (\text{D-19})$$

$$\frac{\partial^2 u}{\partial yz}_{x,y,z} = \frac{\partial^2 u}{\partial zy}_{x,y,z} \approx \frac{1}{4} [u_{x,y-1,z-1} + u_{x,y+1,z+1} - u_{x,y+1,z-1} - u_{x,y-1,z+1}]. \quad (\text{D-20})$$

These approximations are applied element-wise to

$$L \equiv \begin{bmatrix} l_{1,1}^1 & l_{1,2}^1 & l_{1,3}^1 \\ l_{2,1}^1 & l_{2,2}^1 & l_{2,3}^1 \\ l_{3,1}^1 & l_{3,2}^1 & l_{3,3}^1 \end{bmatrix} \quad \text{and} \quad L_B \equiv \begin{bmatrix} l_{1,1}^2 & l_{1,2}^2 & l_{1,3}^2 \\ l_{2,1}^2 & l_{2,2}^2 & l_{2,3}^2 \\ l_{3,1}^2 & l_{3,2}^2 & l_{3,3}^2 \end{bmatrix}. \quad (\text{D-21})$$

The resulting approximations for the differential operators are

$$\begin{aligned}\hat{l}_{1,1}^1 = & -(\alpha + \beta)(u_{x+1,y,z} - 2u_{x,y,z} + u_{x-1,y,z}) - \alpha(u_{x,y+1,z} - 2u_{x,y,z} + u_{x,y-1,z}) \\ & - \alpha(u_{x,y,z+1} - 2u_{x,y,z} + u_{x,y,z-1}) + \gamma u_{x,y,z}\end{aligned}\quad (\text{D-22})$$

$$\begin{aligned}\hat{l}_{2,2}^1 = & -\alpha(u_{x+1,y,z} - 2u_{x,y,z} + u_{x-1,y,z}) - (\alpha + \beta)(u_{x,y+1,z} - 2u_{x,y,z} + u_{x,y-1,z}) \\ & - \alpha(u_{x,y,z+1} - 2u_{x,y,z} + u_{x,y,z-1}) + \gamma u_{x,y,z}\end{aligned}\quad (\text{D-23})$$

$$\begin{aligned}\hat{l}_{3,3}^1 = & -\alpha(u_{x+1,y,z} - 2u_{x,y,z} + u_{x-1,y,z}) - \alpha(u_{x,y+1,z} - 2u_{x,y,z} + u_{x,y-1,z}) \\ & - (\alpha + \beta)(u_{x,y,z+1} - 2u_{x,y,z} + u_{x,y,z-1}) + \gamma u_{x,y,z}\end{aligned}\quad (\text{D-24})$$

$$\hat{l}_{1,2}^1 = \hat{l}_{2,1}^1 = -\frac{\beta}{4}(u_{x+1,y+1,z} + u_{x-1,y-1,z} - u_{x-1,y+1,z} - u_{x+1,y-1,z})\quad (\text{D-25})$$

$$\hat{l}_{1,3}^1 = \hat{l}_{3,1}^1 = -\frac{\beta}{4}(u_{x+1,y,z+1} + u_{x-1,y,z-1} - u_{x-1,y,z+1} - u_{x+1,y,z-1})\quad (\text{D-26})$$

$$\hat{l}_{2,3}^1 = \hat{l}_{3,2}^1 = -\frac{\beta}{4}(u_{x,y+1,z+1} + u_{x,y-1,z-1} - u_{x,y-1,z+1} - u_{x,y+1,z-1})\quad (\text{D-27})$$

and

$$\hat{l}_{i,i}^2 = -\alpha(u_{x+1,y,z} + u_{x-1,y,z} + u_{x,y+1,z} + u_{x,y-1,z} + u_{x,y,z+1} + u_{x,y,z-1} - 6u_{x,y,z}) + \gamma u_{x,y,z}\quad (\text{D-28})$$

$$\hat{l}_{i,j}^2 = 0 \text{ for } i \neq j.\quad (\text{D-29})$$

Applied component-wise to a particular spatial location (x, y, z) , these operators expand to

$$\begin{aligned}
(Lv_{x,y,z})^1 &\approx (6\alpha + 2\beta + \gamma)v_{x,y,z}^1 - (\alpha + \beta)v_{x+1,y,z}^1 - (\alpha + \beta)v_{x-1,y,z}^1 \\
&\quad - \alpha v_{x1,y-1,z}^1 - \alpha v_{x1,y+1,z}^1 - \alpha v_{x1,y,z-1}^1 - \alpha v_{x1,y,z+1}^1 \\
&\quad - \frac{\beta}{4}v_{x+1,y+1,z}^2 - \frac{\beta}{4}v_{x-1,y-1,z}^2 + \frac{\beta}{4}v_{x-1,y+1,z}^2 + \frac{\beta}{4}v_{x+1,y-1,z}^2 \\
&\quad - \frac{\beta}{4}v_{x+1,y,z+1}^3 - \frac{\beta}{4}v_{x-1,y,z-1}^3 + \frac{\beta}{4}v_{x-1,y,z+1}^3 + \frac{\beta}{4}v_{x+1,y,z-1}^3
\end{aligned} \tag{D-30}$$

$$\begin{aligned}
(Lv_{x,y,z})^2 &\approx -\frac{\beta}{4}v_{x+1,y+1,z}^1 - \frac{\beta}{4}v_{x-1,y-1,z}^1 + \frac{\beta}{4}v_{x-1,y+1,z}^1 + \frac{\beta}{4}v_{x+1,y-1,z}^1 \\
&\quad + (6\alpha + 2\beta + \gamma)v_{x,y,z}^2 - (\alpha + \beta)v_{x+1,y,z}^2 - (\alpha + \beta)v_{x-1,y,z}^2 \\
&\quad - \alpha v_{x1,y-1,z}^2 - \alpha v_{x1,y+1,z}^2 - \alpha v_{x1,y,z-1}^2 - \alpha v_{x1,y,z+1}^2 \\
&\quad - \frac{\beta}{4}v_{x,y+1,z+1}^3 - \frac{\beta}{4}v_{x,y-1,z-1}^3 + \frac{\beta}{4}v_{x,y-1,z+1}^3 + \frac{\beta}{4}v_{x,y+1,z-1}^3
\end{aligned} \tag{D-31}$$

$$\begin{aligned}
(Lv_{x,y,z})^3 &\approx -\frac{\beta}{4}v_{x+1,y,z+1}^1 - \frac{\beta}{4}v_{x-1,y,z-1}^1 + \frac{\beta}{4}v_{x-1,y,z+1}^1 + \frac{\beta}{4}v_{x+1,y,z-1}^1 \\
&\quad - \frac{\beta}{4}v_{x,y+1,z+1}^2 - \frac{\beta}{4}v_{x,y-1,z-1}^2 + \frac{\beta}{4}v_{x,y-1,z+1}^2 + \frac{\beta}{4}v_{x,y+1,z-1}^2 \\
&\quad + (6\alpha + 2\beta + \gamma)v_{x,y,z}^3 - (\alpha + \beta)v_{x+1,y,z}^3 - (\alpha + \beta)v_{x-1,y,z}^3 \\
&\quad - \alpha v_{x1,y-1,z}^3 - \alpha v_{x1,y+1,z}^3 - \alpha v_{x1,y,z-1}^3 - \alpha v_{x1,y,z+1}^3
\end{aligned} \tag{D-32}$$

and

$$(L_B v_{x,y,z})^1 \approx -\alpha (v_{x+1,y,z}^1 + v_{x-1,y,z}^1 + v_{x,y+1,z}^1 + v_{x,y-1,z}^1 + v_{x,y,z+1}^1 + v_{x,y,z-1}^1 - 6v_{x,y,z}^1) + \gamma v_{x,y,z}^1 \tag{D-33}$$

$$(L_B v_{x,y,z})^2 \approx -\alpha (v_{x+1,y,z}^2 + v_{x-1,y,z}^2 + v_{x,y+1,z}^2 + v_{x,y-1,z}^2 + v_{x,y,z+1}^2 + v_{x,y,z-1}^2 - 6v_{x,y,z}^2) + \gamma v_{x,y,z}^2 \tag{D-34}$$

$$(L_B v_{x,y,z})^3 \approx -\alpha (v_{x+1,y,z}^3 + v_{x-1,y,z}^3 + v_{x,y+1,z}^3 + v_{x,y-1,z}^3 + v_{x,y,z+1}^3 + v_{x,y,z-1}^3 - 6v_{x,y,z}^3) + \gamma v_{x,y,z}^3. \tag{D-35}$$

L is strictly diagonally dominate if $\gamma > 2\beta$. L_B is strictly diagonally dominate if

$\alpha > 0$ and $\gamma > 0$.

D-4 Fourier domain solution for L

As in the work of Beg et al. (2005) and Christensen et al. (1996), the numerical solution to $Av = b$ is determined using a Fourier transform. The strategy is to apply the Fast Fourier Transform (FFT) to the body force b , apply the inverse of the (discretized) differential operator A , and apply the inverse FFT to get the resulting velocity v .

Suppose that the body force is periodic and as a discretized vector field it has dimensions $N \times M \times L$. We can write v , $\hat{A}v$, and b in terms of their Fourier coefficients:

$$v(x, y, z) = \sum_{w=0}^{L-1} \sum_{v=0}^{M-1} \sum_{u=0}^{N-1} V(u, v, w) e^{2j\pi\left(\frac{ux}{N} + \frac{vy}{M} + \frac{wz}{L}\right)} \quad (\text{D-36})$$

$$\hat{A}v(x, y, z) = \sum_{w=0}^{L-1} \sum_{v=0}^{M-1} \sum_{u=0}^{N-1} \hat{L}V(u, v, w) e^{2j\pi\left(\frac{ux}{N} + \frac{vy}{M} + \frac{wz}{L}\right)} \quad (\text{D-37})$$

$$b(x, y, z) = \sum_{w=0}^{L-1} \sum_{v=0}^{M-1} \sum_{u=0}^{N-1} B(u, v, w) e^{2j\pi\left(\frac{ux}{N} + \frac{vy}{M} + \frac{wz}{L}\right)}. \quad (\text{D-38})$$

To solve $Av = b$ for v , the invertible matrix⁶ Λ must be found such that

$$\Lambda V(u, v, w) = B(u, v, w). \quad (\text{D-39})$$

The solution for v can then be written as

$$v(x, y, z) = \sum_{w=0}^{L-1} \sum_{v=0}^{M-1} \sum_{u=0}^{N-1} \Lambda^{-1} B(u, v, w) e^{2j\pi\left(\frac{ux}{N} + \frac{vy}{M} + \frac{wz}{L}\right)}. \quad (\text{D-40})$$

The strategy to solve for Λ is to find the intermediate matrix Ψ such that

$$\hat{A}V e^{2j\pi\left(\frac{ux}{N} + \frac{vy}{M} + \frac{wz}{L}\right)} = \Psi V \quad (\text{D-41})$$

⁶Not differential operator, a 3×3 matrix valued function of u , v , and w .

and then factor out $e^{2j\pi(\frac{ux}{N}+\frac{vy}{M}+\frac{wz}{L})}$ from Ψ . In other words,

$$\hat{A}V e^{2j\pi(\frac{ux}{N}+\frac{vy}{M}+\frac{wz}{L})} = \Psi V = \Lambda V e^{2j\pi(\frac{ux}{N}+\frac{vy}{M}+\frac{wz}{L})}. \quad (\text{D-42})$$

Ψ is simply the matrix $\hat{A}e^{2j\pi(\frac{ux}{N}+\frac{vy}{M}+\frac{wz}{L})}$ since V is constant with respect to the spatial domain variables x , y , and z . \hat{A} is symmetric, so Ψ and Λ will also be symmetric. The elements of Λ are derived below. Recall that $e^{\pm i\theta} = \cos \theta \pm i \sin \theta$ implies that $2 \cos \theta = e^{i\theta} + e^{-i\theta}$ and $2 \sin \theta = e^{i\theta} - e^{-i\theta}$. Also, $\cos(a+b) = \cos a \cos b - \sin a \sin b$.

$$\begin{aligned} \psi_{1,1} = & -(\alpha + \beta) \left(e^{2j\pi(\frac{u(x+1)}{N}+\frac{vy}{M}+\frac{wz}{L})} - 2e^{2j\pi(\frac{ux}{N}+\frac{vy}{M}+\frac{wz}{L})} + e^{2j\pi(\frac{u(x-1)}{N}+\frac{vy}{M}+\frac{wz}{L})} \right) \\ & - \alpha \left(e^{2j\pi(\frac{ux}{N}+\frac{v(y+1)}{M}+\frac{wz}{L})} - 2e^{2j\pi(\frac{ux}{N}+\frac{vy}{M}+\frac{wz}{L})} + e^{2j\pi(\frac{ux}{N}+\frac{v(y-1)}{M}+\frac{wz}{L})} \right) \\ & - \alpha \left(e^{2j\pi(\frac{ux}{N}+\frac{vy}{M}+\frac{w(z+1)}{L})} - 2e^{2j\pi(\frac{ux}{N}+\frac{vy}{M}+\frac{wz}{L})} + e^{2j\pi(\frac{ux}{N}+\frac{vy}{M}+\frac{w(z-1)}{L})} \right) \\ & + \gamma e^{2j\pi(\frac{ux}{N}+\frac{vy}{M}+\frac{wz}{L})} \end{aligned} \quad (\text{D-43})$$

which implies that

$$\begin{aligned} \lambda_{1,1} = & -(\alpha+\beta) \left(e^{2j\pi\frac{u}{N}} - 2 + e^{-2j\pi\frac{u}{N}} \right) - \alpha \left(e^{2j\pi\frac{v}{M}} - 2 + e^{-2j\pi\frac{v}{M}} \right) \\ & - \alpha \left(e^{2j\pi\frac{w}{L}} - 2 + e^{-2j\pi\frac{w}{L}} \right) + \gamma e^{2j\pi(\frac{ux}{N}+\frac{vy}{M}+\frac{wz}{L})} \end{aligned} \quad (\text{D-44})$$

$$= -(\alpha+\beta) \left(2 \cos \left[2\pi \frac{u}{N} \right] - 2 \right) - \alpha \left(2 \cos \left[2\pi \frac{v}{M} \right] - 2 \right) - \alpha \left(2 \cos \left[2\pi \frac{w}{L} \right] - 2 \right) + \gamma \quad (\text{D-45})$$

$$= -2(\alpha+\beta) \cos \left[2\pi \frac{u}{N} \right] - 2\alpha \cos \left[2\pi \frac{v}{M} \right] - 2\alpha \cos \left[2\pi \frac{w}{L} \right] + 6\alpha + 2\beta + \gamma. \quad (\text{D-46})$$

Similarly,

$$\lambda_{2,2} = -2\alpha \cos \left[2\pi \frac{u}{N} \right] - 2(\alpha + \beta) \cos \left[2\pi \frac{v}{M} \right] - 2\alpha \cos \left[2\pi \frac{w}{L} \right] + 6\alpha + 2\beta + \gamma \quad (\text{D-47})$$

$$\begin{aligned} \lambda_{3,3} = & -2\alpha \cos \left[2\pi \frac{u}{N} \right] - 2\alpha \cos \left[2\pi \frac{v}{M} \right] - 2(\alpha + \beta) \cos \left[2\pi \frac{w}{L} \right] + 6\alpha + 2\beta + \gamma. \\ & (\text{D-48}) \end{aligned}$$

For the cross terms,

$$\begin{aligned} \psi_{1,2} = & -\frac{\beta}{4} \left(e^{2j\pi\left(\frac{u(x+1)}{N} + \frac{v(y+1)}{M} + \frac{wz}{L}\right)} + e^{2j\pi\left(\frac{u(x-1)}{N} + \frac{v(y-1)}{M} + \frac{wz}{L}\right)} - e^{2j\pi\left(\frac{u(x-1)}{N} + \frac{v(y+1)}{M} + \frac{wz}{L}\right)} - e^{2j\pi\left(\frac{u(x+1)}{N} + \frac{v(y-1)}{M} + \frac{wz}{L}\right)} \right) \end{aligned} \quad (\text{D-49})$$

which implies that

$$\lambda_{1,2} = -\frac{\beta}{4} \left(e^{2j\pi\left(\frac{u}{N} + \frac{v}{M}\right)} + e^{2j\pi\left(-\frac{u}{N} - \frac{v}{M}\right)} - e^{2j\pi\left(-\frac{u}{N} + \frac{v}{M}\right)} - e^{2j\pi\left(\frac{u}{N} - \frac{v}{M}\right)} \right) \quad (\text{D-50})$$

$$= -\frac{\beta}{4} \left(2 \cos \left[2\pi \left(\frac{u}{N} + \frac{v}{M} \right) \right] - 2 \cos \left[2\pi \left(\frac{u}{N} - \frac{v}{M} \right) \right] \right) \quad (\text{D-51})$$

$$= \beta \sin \left[2\pi \frac{u}{N} \right] \sin \left[2\pi \frac{v}{M} \right]. \quad (\text{D-52})$$

Similarly,

$$\lambda_{1,3} = \beta \sin \left[2\pi \frac{u}{N} \right] \sin \left[2\pi \frac{w}{L} \right] \quad (\text{D-53})$$

$$\lambda_{2,3} = \beta \sin \left[2\pi \frac{v}{M} \right] \sin \left[2\pi \frac{w}{L} \right]. \quad (\text{D-54})$$

The solution to $\Lambda V = B$ is based on the Cholesky decomposition of Λ since Λ is symmetric and positive definite. First, the upper triangular matrix U is computed such that $UU^T = \Lambda$. This process is detailed in Algorithm 7.1. Then, forward and back substitution are used to solve for V .

Algorithm 7.1 Cholesky Decomposition for 3-by3 matrix Λ

Input: Λ is a 3-by3 symmetric positive definite matrix

Output: U is an upper triangular matrix such that $UU^T = \Lambda$

- 1: $U_{1,1} = \sqrt{\Lambda_{1,1}}$
 - 2: $U_{1,2} = \Lambda_{1,2}/U_{1,1}$
 - 3: $U_{1,3} = \Lambda_{1,3}/U_{1,1}$
 - 4: $U_{2,2} = \sqrt{\Lambda_{2,2} - U_{1,2}^2}$
 - 5: $U_{2,3} = (\Lambda_{2,3} - U_{1,2}U_{1,3})/U_{2,2}$
 - 6: $U_{3,3} = \sqrt{\Lambda_{3,3} - U_{1,3}^2 - U_{2,3}^2}$.
-

D-5 Fourier domain solution for L_B

The solution for L_B follows the same strategy as the solution for L . In this case, the off-diagonal elements of Ψ and Λ are zero. The diagonal elements are

$$\begin{aligned}
\psi_{i,i} = & -\alpha \left(e^{2j\pi \left(\frac{u(x+1)}{L} + \frac{vy}{M} + \frac{wz}{N} \right)} - 2e^{2j\pi \left(\frac{ux}{N} + \frac{vy}{M} + \frac{wz}{L} \right)} + e^{2j\pi \left(\frac{u(x-1)}{L} + \frac{vy}{M} + \frac{wz}{N} \right)} \right) \\
& - \alpha \left(e^{2j\pi \left(\frac{ux}{L} + \frac{v(y+1)}{M} + \frac{wz}{N} \right)} - 2e^{2j\pi \left(\frac{ux}{L} + \frac{vy}{M} + \frac{wz}{N} \right)} + e^{2j\pi \left(\frac{ux}{L} + \frac{v(y-1)}{M} + \frac{wz}{N} \right)} \right) \\
& - \alpha \left(e^{2j\pi \left(\frac{ux}{L} + \frac{vy}{M} + \frac{w(z+1)}{N} \right)} - 2e^{2j\pi \left(\frac{ux}{L} + \frac{vy}{M} + \frac{wz}{N} \right)} + e^{2j\pi \left(\frac{ux}{L} + \frac{vy}{M} + \frac{w(z-1)}{N} \right)} \right) \\
& + \gamma e^{2j\pi \left(\frac{ux}{L} + \frac{vy}{M} + \frac{wz}{N} \right)}
\end{aligned} \tag{D-55}$$

which implies that

$$\lambda_{i,i} = -\alpha \left(e^{2j\pi \frac{u}{L}} - 2 + e^{-2j\pi \frac{u}{L}} \right) - \alpha \left(e^{2j\pi \frac{v}{M}} - 2 + e^{-2j\pi \frac{v}{M}} \right) - \alpha \left(e^{2j\pi \frac{w}{N}} - 2 + e^{-2j\pi \frac{w}{N}} \right) + \gamma \tag{D-56}$$

$$= -\alpha \left(e^{2j\pi \frac{u}{L}} + e^{-2j\pi \frac{u}{L}} + e^{2j\pi \frac{v}{M}} + e^{-2j\pi \frac{v}{M}} + e^{2j\pi \frac{w}{N}} + e^{-2j\pi \frac{w}{N}} \right) + 6\alpha + \gamma \tag{D-57}$$

$$= -2\alpha \left(\cos \left[2\pi \frac{u}{L} \right] + \cos \left[2\pi \frac{v}{M} \right] + \cos \left[2\pi \frac{w}{N} \right] \right) + 6\alpha + \gamma. \tag{D-58}$$

Since Λ is diagonal the solution to $\Lambda V = B$ is trivial.

D-5.1 Note on interpretation of α and γ

For the Fourier transform, application of $-\alpha\Delta$ in the spatial domain is equivalent to multiplication by $\alpha((2\pi f_1)^2 + (2\pi f_2)^2 + (2\pi f_3)^2)$ in the frequency domain. The relationship between this solution and the discrete solution is apparent since $\cos \theta \approx 1 - \frac{\theta^2}{2}$ when θ is near zero. That is,

$$-2\alpha \left(\cos \left[2\pi \frac{u}{L} \right] + \cos \left[2\pi \frac{v}{M} \right] + \cos \left[2\pi \frac{w}{N} \right] \right) + 6\alpha \tag{D-59}$$

$$\approx -2\alpha \left(3 - \frac{(2\pi \frac{u}{L})^2}{2} - \frac{(2\pi \frac{v}{M})^2}{2} - \frac{(2\pi \frac{w}{N})^2}{2} \right) + 6\alpha \tag{D-60}$$

$$= \alpha \left(\left(2\pi \frac{u}{L} \right)^2 + \left(2\pi \frac{v}{M} \right)^2 + \left(2\pi \frac{w}{N} \right)^2 \right). \tag{D-61}$$

Thus, you can reason about the parameters by thinking about multiplying in the frequency domain by $\frac{1}{\alpha f^2 + \gamma}$. Increasing α decreases the width of the kernel in the Fourier domain, resulting in *more* smoothing of the body force. Decreasing α allows more high frequency components to remain. Increasing γ decreases the height of the kernel in the Fourier domain, resulting in smaller magnitudes for velocity relative to the body force (recall that the FFT is just a linear transformation). A value of $\gamma = 1$ will result in a kernel with height 1 in the Fourier domain. Thus $\alpha = 0$ and $\gamma = 1$ results in the identity operator.

Bibliography

- Agresti, A. (1990). *Categorical Data Analysis*. Wiley Interscience. 45
- Amenta, N., Choi, S., and Kolluri, R. K. (2001). The power crust. In *ACM Symposium on Solid Modeling and Applications*, pages 249–260. 36
- Avants, B. and Gee, J. C. (2004). Geodesic estimation for large deformation anatomical shape averaging and interpolation. *NeuroImage (Supplemental issue on Mathematics in Brain Imaging)*, 23:S139–S150. 68
- Beckmann, N., Laurent, D., Tigani, B., Panizzutti, R., and Rudin, M. (2004). Magnetic resonance imaging in drug discovery: lessons from disease areas. *Drug Discovery Today*, 9(1):1–49. 2
- Beg, M. and Khan, A. (2006). Computing an average anatomical atlas using lddmm and geodesic shooting. In *Proceedings of the 2006 IEEE International Symposium of Biomedical Imaging (ISBI 06)*. 121
- Beg, M., Miller, M., Trouvé, A., and Younes, L. (2005). Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International Journal of Computer Vision*, 61(2). 13, 18, 57, 61, 136, 144, 149
- Beg, M. F. (2003). Variational and computational methods for flows of diffeomorphisms in image matching and growth in computational anatomy. *PhD thesis, The Johns Hopkins University*. 7, 70, 136, 137
- Bhattacharya, R. and Patrangenaru, V. (2002). Nonparametric estimation of location and dispersion on Riemannian manifolds. *Journal of Statistical Planning and Inference*, 108:23–36. 71
- Bhattacharya, R. and Patrangenaru, V. (2003). Large sample theory of intrinsic and extrinsic sample means on manifolds i. *Annals of Statistics*, 31(1):1–29. 72
- Bingham, C. (1974). An antipodally symmetric distribution on the sphere. *The Annals of Statistics*, 2(6):1201–1225. 71
- Boothby, W. (2003). *An Introduction to Differentiable Manifolds*. Academic Press, second edition. 127
- Bradley, J. V. (1968). *Distribution-Free Statistical Tests*. Prentice Hall. 118
- Charpiat, G., Faugeras, O., and Keriven, R. (2005). Image statistics based on diffeomorphic matching. In *Proceedings of the 10th International Conference on Computer Vision (ICCV 2005)*. 67

- Christensen, G. E. and Johnson, H. J. (2001). Consistent image registration. *IEEE Transactions on Medical Imaging*, 20(7):568–582. 60
- Christensen, G. E., Rabbitt, R. D., and Miller, M. I. (1994). 3D brain mapping using a deformable neuroanatomy. *Physics in Medicine and Biology*, 39:609–618. 140
- Christensen, G. E., Rabbitt, R. D., and Miller, M. I. (1996). Deformable templates using large deformation kinematics. *IEEE Transactions On Image Processing*, 5(10):1435–1447. 12, 14, 18, 31, 35, 61, 140, 143, 149
- Chung, M. K., Robbins, S. M., Dalton, K. M., Davidson, R. J., Alexander, A. L., and Evans, A. C. (2005). Cortical thickness analysis in autism with heat kernel smoothing. *NeuroImage*, 25:1256–1265. 72
- Clark, A. (1984). *Elements of Abstract Algebra*. Dover. 130
- Clatz, O., Sermesant, M., Bondiau, P.-Y., Delingette, H., Warfield, S. K., Malandain, G., and Ayache, N. (2005). Realistic simulation of the 3D growth of brain tumors in MR images coupling diffusion with mass effect. *IEEE Transactions on Medical Imaging*, 24(10):1334–1346. 7, 70
- Conway, J. B. (1990). *A Course in Functional Analysis*. Graduate Texts in Mathematics. Springer, second edition. 127
- Cootes, T. F., Marsland, S., Twining, C. J., Smith, K., and Taylor, C. J. (2004). Group-wise diffeomorphic non-rigid registration for automatic model building. In Pajdla, T. and Matas, J., editors, *Proceedings of the 8th European Conference on Computer Vision (ECCV 2004)*, volume 3022 of *Lecture Notes in Computer Science*. Springer. 66
- Cootes, T. F., Taylor, C. J., Cooper, D. H., and Graham, J. (1995). Active shape models—their training and application. *Computer Vision and Image Understanding*, 61(1):38–59. 3
- Craene, M. D., du Bois d’Aische, A., Macq, B., and Warfield, S. K. (2004). Multi-subject registration for unbiased statistical atlas construction. In Barillot, C., Haynor, D. R., and Hellier, P., editors, *Proceedings of the 7th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI 2004)*, volume 3216 of *Lecture Notes in Computer Science*. Springer. 67
- Davies, R. H., Twining, C. J., Cootes, T. F., Waterton, J. C., and Taylor, C. J. (2002). A minimum description length approach to statistical shape modeling. *IEEE Transactions on Medical Imaging*, 21(5):525–537. 66
- Davis, B. C., Fletcher, P. T., Bullitt, E., and Joshi, S. C. (2007). Population shape regression from random design data. In *Proceedings of the 10th International Conference on Computer Vision (ICCV 2007)*. 9, 73, 86

- Davis, B. C., Fletcher, P. T., Bullitt, E., and Joshi, S. C. (Submitted 2008). Population shape regression from random design data. *International Journal of Computer Vision*. 73
- Davis, B. C., Foskey, M., Rosenman, J., Goyal, L., Chang, S., and Joshi, S. C. (2005). Automatic segmentation of intra-treatment CT images for adaptive radiation therapy. In Duncan, J. and Gerig, G., editors, *Proceedings of the 8th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI 2005)*, Lecture Notes in Computer Science. Springer. 21
- Davis, B. C., Lorenzen, P., and Joshi, S. C. (2004). Large deformation minimum mean squared error template estimation for computational anatomy. In *Proceedings of the 2004 IEEE International Symposium of Biomedical Imaging (ISBI 04)*. 50
- DeVita, V., Hellman, S., and Rosenberg, S., editors (2004). *Cancer: Principles and Practice of Oncology*, chapter Cancer of the Prostate. Lippincott Williams & Wilkins, 7th edition. 21
- Dice, L. R. (1945). Measures of the amount of ecologic association between species. *Ecology*, 26(3):297–302. 44
- Dupuis, P., Grenander, U., and Miller, M. (1998). Variational problems on flows of diffeomorphisms for image matching. *Quarterly of Applied Mathematics*, 56(3):587–600. 12, 13, 14, 15, 17, 18, 31
- Fletcher, P. T. and Joshi, S. (2007). Riemannian geometry for the statistical analysis of diffusion tensor data. *Signal Processing*, 87(2):250–262. 66, 86
- Fletcher, P. T., Joshi, S., Lu, C., and Pizer, S. M. (2004). Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Transactions on Medical Imaging*, 23(8):995–1005. 2, 5, 49, 66, 71, 72, 86, 122
- Fletcher, P. T., Venkatasubramanian, S., and Joshi, S. C. (2008). Robust statistics on Riemannian manifolds via the geometric median. In *Proceedings of the 2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 08)*. 124
- Foskey, M., Davis, B. C., Goyal, L., Chang, S., Chaney, E., Strehl, N., Tomei, S., Rosenman, J., and Joshi, S. (2005). Large deformation 3D image registration in image-guided radiation therapy. *Physics in Medicine and Biology*, 50:5869–5892. 21
- Fréchet, M. (1948). Les elements aleatoires de nature quelconque dans un espace distance. *Annales De L’Institut Henri Poincare*, 10:215–310. 52, 72
- Gerig, G., Davis, B. C., Lorenzen, P., Xu, S., Jomier, M., Piven, J., and Joshi, S. C. (2006). Computational anatomy to assess longitudinal trajectory of brain growth. In *Proceedings of the Third International Symposium on 3D Data Processing, Visualization and Transmission (3DPVT 2006)*, volume 3565. 50

- Glaunès, J. and Joshi, S. (2006). Template estimation from unlabeled point set data and surfaces for computational anatomy. In Pennec, X. and Joshi, S., editors, *Medical Image Computing and Computer Assisted Intervention Workshop: Mathematical Foundations of Computational Anatomy*, pages 29–39. 86, 122
- Glaunès, J., Trouvé, A., and Younes, L. (2004). Diffeomorphic matching of distributions: a new approach for unlabeled point-sets and sub-manifolds matching. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 04)*. 68, 86, 122
- Goodall, C. (1991). Procrustes methods in the statistical analysis of shape. *Journal of the Royal Statistical Society (Series B)*, 53(2):285–339. 66
- Goodlett, C., Davis, B. C., Jean, R., Gilmore, J., and Gerig, G. (2006). Improved correspondence for DTI population studies via unbiased atlas building. In Frangi, A. F. and Delingette, H., editors, *Proceedings of the 9th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI 2006)*, volume 3216 of *Lecture Notes in Computer Science*. Springer. 50
- Grenander, U. (1996). *Elements of Pattern Theory*. Johns Hopkins Studies in the Mathematical Sciences. The Johns Hopkins University Press. 3, 12
- Grenander, U. and Miller, M. I. (1998). Computational anatomy: An emerging discipline. *Quarterly of Applied Mathematics*, 56(4):617–694. 2, 3, 12, 71
- Guttmann, C., Jolesz, F., Kikinis, R., Killiany, R., Moss, M., Sandor, T., and Albert, M. (1998). White matter changes with normal aging. *Neurology*, 50(4):972–978. 71
- Hall, P. and Marron, J. S. (1991). Local minima in cross-validation functions. *Journal of the Royal Statistical Society (Series B)*, 53(1):245–252. 78
- Halperin, E. C., Perez, C. A., and Brady, L. W. (2007). *Perez and Brady’s Principles and Practice of Radiation Oncology*. Lippincott Williams & Wilkins. 1, 21
- Hardle, W. (1990). *Applied Nonparametric Regression*. Cambridge University Press. 73
- He, J. and Christensen, G. E. (2003). Large deformation inverse consistent elastic image registration. In Taylor, C. J. and Noble, J. A., editors, *IPMI3003, LNCS 2732*, pages 438–449. IPMI, Springer-Verlag. 60
- Hernandez, M., Bossa, M. N., and Olmos, S. (2007). Registration of anatomical images using geodesic paths of diffeomorphisms parameterized with stationary vector fields. In Nielsen, M., Niessen, W., and Westin, C., editors, *Proceedings of the 2007 IEEE Computer Society Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA 2007)*. 122
- Holden, M. (2008). A review of geometric transformations for nonrigid body registration. *IEEE Transactions on Medical Imaging*, 27(1):111–128. 140, 144

- Hutton, T. J., Buxton, B. F., Hammon, P., and Potts, H. W. W. (2003). Estimating average growth trajectories in shape-space using kernel smoothing. *IEEE Transactions on Medical Imaging*, 22(6):747–753. 71
- Johnson, H. J. and Christensen, G. E. (2002). Consistent landmark and intensity based registration. *IEEE Transactions on Medical Imaging*, 21(5):450–461. 60
- Jones, M. C., Marron, J. S., and Sheather, S. J. (1996). A brief survey of bandwidth selection for density estimation. *Journal of the American Statistical Association*, 91(433):401–407. 77
- Joshi, S., Lorenzen, P., Gerig, G., and Bullitt, E. (2003). Structural and radiometric asymmetry in brain images. *Medical Image Analysis*, 7(2):155–170. 34
- Joshi, S. C., Davis, B. C., Jomier, M., and Gerig, G. (2004). Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage (Supplemental issue on Mathematics in Brain Imaging)*, 23:S151–S160. 50, 68, 72
- Joshi, S. C. and Miller, M. I. (2000). Landmark matching via large deformation diffeomorphisms. *IEEE Transactions On Image Processing*, 9(8):1357–1370. 122, 143
- Jupp, P. and Mardia, K. (1989). A unified view of the theory of directional statistics, 1975–1988. *International Statistical Review*, 57(3):261–294. 71
- Jupp, P. E. and Kent, J. T. (1987). Fitting smooth paths to spherical data. *Applied Statistics*, 36(1):34–46. 72
- Karcher, H. (1977). Riemannian center of mass and mollifier smoothing. *Communications on Pure and Applied Mathematics*, 30:509–541. 52, 53
- Kendall, D. G. (1984). Shape manifolds, Procrustean metrics, and complex projective spaces. *Bulletin of the London Mathematical Society*, 16:18–121. 71
- Ketting, C. H., Austin-Seymour, M., Kalet, I., Unger, J., Hummel, S., and Jacky, J. (1997). Consistency of three-dimensional planning target volumes across physicians and institutions. *International Journal of Radiation Oncology*Biophysics*, 37:445–453. 23, 39
- Le, H. and Kendall, D. (1993). The Riemannian structure of Euclidean shape spaces: A novel environment for statistics. *The Annals of Statistics*, 21(3):1225–1271. 71
- Learned-Miller, E. G. (2006). Data driven image models through continuous joint alignment. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(2):236–250. 67
- Lee, J. M. (2003). *Introduction to Smooth Manifolds*. Graduate Texts in Mathematics. Springer. 29, 127, 134

- Leemput, K. V., Maes, F., Vandermeulen, D., and Suetens, P. (1999). Automated model-based tissue classification of MR images of the brain. *IEEE Transactions on Medical Imaging*, 18(10):897–908. 75
- Leunens, G., Menten, J., Weltens, C., Verstraete, J., and van der Scheuren, E. (1993). Quality assessment of medical decision making in radiation oncology: Variability in target volume delineation for brain tumours. *Radiotherapy and Oncology*, 29:169–175. 39
- Loader, C. R. (1999). Bandwidth selection: Classical or plug-in? *The Annals of Statistics*, 27(2):415–438. 77
- Lorenzen, P., Prastawa, M., Davis, B., Gerig, G., Bullitt, E., and Joshi, S. (2006). Multi-modal image set registration and atlas formation. *Medical Image Analysis*, 10(3):440–451. 17, 50, 67, 88
- Lorenzen, P. J., Davis, B. C., and Joshi, S. C. (2004a). Model based symmetric information theoretic large deformation multi-modal image registration. In *Proceedings of the 2004 IEEE International Symposium of Biomedical Imaging (ISBI 04)*. 50
- Lorenzen, P. J., Davis, B. C., and Joshi, S. C. (2004b). Multi-class posterior atlas formation via unbiased kullback-leibler template estimation. In Barillot, C., Haynor, D. R., and Hellier, P., editors, *Proceedings of the 7th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI 2004)*, volume 3216 of *Lecture Notes in Computer Science*. Springer. 50
- Luenberger, D. G. (1969). *Optimization by vector space methods*. Series In Decision and Control. Wiley-Interscience. 139
- Malladi, R., Sethian, J., and Vemuri, B. (1995). Shape modeling with front propagation: A level set approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17:158–175. 3
- Matsumae, M., Kikinis, R., Mórocz, I., Lorenzo, A., Sándor, T., Sándor, T., Albert, M., Black, P., and Jolesz, F. (1996). Age-related changes in intracranial compartment volumes in normal adults assessed by magnetic resonance imaging. *Journal of Neurosurgery*, 84:982–991. 71
- Merck, D., Tracton, G., Saboo, R., Levy, J., Chaney, E. L., Pizer, S., and Joshi, S. C. (2008). Training models of anatomic shape variability. *Medical Physics (Under Review)*. 62
- Miller, M. I. (2004). Computational anatomy: shape, growth, and atrophy comparison via diffeomorphisms. *NeuroImage*, 23:S19–S33. 2, 7, 70, 71, 72, 79, 125
- Miller, M. I., Banerjee, A., Christensen, G. E., Joshi, S. C., Khaneja, N., Grenander, U., and Matejic, L. (1997). Statistical methods in computational anatomy. *Statistical Methods in Medical Research*, 6:267–299. 66, 122

- Miller, M. I., Joshi, S. C., and Christensen, G. E. (1999). Large deformation fluid diffeomorphisms for landmark and image matching. In Toga, A. W., editor, *Brain Warping*, chapter 7. Academic Press. 35
- Miller, M. I., Trounev, A., and Younes, L. (2002). On the metrics and Euler-Lagrange equations of computational anatomy. *Annual Review of Biomedical Engineering*, 4:375–405. 13, 80, 122
- Miller, M. I., Trouné, A., and Younes, L. (2006). Geodesic shooting for computational anatomy. *Journal of Mathematical Imaging and Vision*, 24(2). 13, 121
- Miller, M. I. and Younes, L. (2001). Group actions, homeomorphisms, and matching: A general framework. *International Journal of Computer Vision*, 41:61–84. 3, 13, 71
- Mori, S., Wakana, S., van Zijl, P. C. M., and Nagae-Poetscher, L. M. (2005). *MRI Atlas of Human White Matter*. Elsevier Science. 2
- Mortamet, B., Zeng, D., Gerig, G., Prastawa, M., and Bullitt, E. (2005). Effects of healthy aging measured by intracranial compartment volumes using a designed MR brain database. In *Lecture Notes in Computer Science (LNCS)*, volume 3749, pages 383–391. Medical Image Computing and Computer Assisted Intervention (MICCAI). 71, 75
- Muller, K. E., Chi, Y., Ahn, J., and Marron, J. S. (2008). Limitations of high dimension, low sample size principal components for gaussian data. *Under Review*. 117
- Nadaraya, E. A. (1964). On estimating regression. *Theory of Probability and its Applications*, 10:186–190. 74
- Nilsson, J., Sha, F., and Jordan, M. I. (2007). Regression on manifolds using kernel dimension reduction. In *Proceedings of the 24th International Conference on Machine Learning (ICML 2007)*. 72
- Pennecc, X. (2006). Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements. *Journal of Mathematical Imaging and Vision*, 25:127–154. 5, 71, 72
- Pennecc, X., Fillard, P., and Ayache, N. (2006). A Riemannian framework for tensor computing. *International Journal of Computer Vision*, 66(1):41–66. 86
- Pizer, S., Fletcher, P., Joshi, S., Thall, A., Chen, J., Fridman, Y., Fritsch, D., Gash, A., Glotzer, J., Jiroutek, M., Lu, C., Muller, K., Tracton, G., Yushkevich, P., and Chaney, E. (2003). Deformable m-reps for 3D medical image segmentation. *International Journal of Computer Vision*, 55(2):85–106. 2, 3, 66, 86, 117, 122
- Pizer, S. M., Broadhurst, R. E., Levy, J., Liu, X., Jeong, J., Stough, J., Tracton, G., and Chaney, E. L. (2008). Segmentation by posterior optimization of m-reps: Strategy and results. *Under Review*. 45

- Prastawa, M., Bullitt, E., Ho, S., and Gerig, G. (2004). A brain tumor segmentation framework based on outlier detection. *Medical Image Analysis*, 8(3):275–283. 89
- Ramsay, J. O. and Li, X. (1998). Curve registration. *Journal of the Royal Statistical Society (Series B)*, 60(2):351–363. 66
- Rasch, C., Barillot, I., Remeijer, P., Touw, A., van Herk, M., and Lebesque, J. V. (1999). Definition of the prostate in CT and MRI: a multi-observer study. *International Journal of Radiation Oncology*Biophysics*, 43(1):57–66. 39
- Rohlfing, T., Brandt, R., Menzel, R., and Maurer, C. R. (2004). Evaluation of atlas selection strategies for atlas-based image segmentation with application to confocal microscopy images of bee brains. *NeuroImage*, 21(4):1428–1442. 48
- Rohlfing, T. and Maurer, C. R. (2007). Shape-based averaging. *IEEE Transactions On Image Processing*, 16(1):153–161. 66
- Rueckert, D., Frangi, A. F., and Schnabel, J. A. (2003). Automatic construction of 3-D statistical deformation models of the brain using nonrigid registration. *IEEE Transactions on Medical Imaging*, 22(8):1014–1024. 67
- Spira, A., Kimmel, R., and Sochen, N. (2007). A short-time Beltrami kernel for smoothing images and manifolds. *IEEE Transactions On Image Processing*, 16(6):1628–1636. 72
- Styner, M., Lieberman, J. A., McClure, R., Weinberger, D., Jones, D., and Gerig, G. (2005). Morphometric analysis of lateral ventricles in schizophrenia and healthy controls regarding genetic and disease-specific factors. *Proceedings of the National Academy of Sciences*, 102(13):4872–4877. 2
- Tewell, M. A. and Adams, R. (2004). The PLUNC 3D treatment planning system: a dynamic alternative to commercially available systems. *Medical Dosimetry*, 29(2):134–138. 2
- Thompson, P. M., Giedd, J. N., Woods, R. P., MacDonald, D., Evans, A. C., and Toga, A. W. (2000). Growth patterns in the developing brain detected by using continuum mechanical tensor maps. *Nature*, 404(6774):190–193. 7, 70
- Trouvé, A. (1998). Diffeomorphisms group and pattern matching in image analysis. *International Journal of Computer Vision*, 28(3):213–221. 13
- Trouvé, A. and Younes, L. (2005). Metamorphoses through lie group action. *Foundations of Computational Mathematics*, 5(2):173–198. 78
- Twining, C. J., Cootes, T. F., Marsland, S., Petrovic, V., Schestowitz, R., and Taylor, C. J. (2005). A unified information-theoretic approach to groupwise non-rigid registration and model building. In Christensen, G. E. and Sonka, M., editors, *Proceedings of the 19th International Conference on Information Processing in Medical Imaging (IPMI 2005)*, volume 3565 of *Lecture Notes in Computer Science*. Springer. 66

- Vaillant, M. and Glaunés, J. (2005). Surface matching via currents. In Christensen, G. E. and Sonka, M., editors, *Proceedings of the 19th International Conference on Information Processing in Medical Imaging (IPMI 2005)*, volume 3565 of *Lecture Notes in Computer Science*. Springer. 122
- Vaillant, M., Miller, M. I., Younes, L., and Trouvé, A. (2004). Statistics on diffeomorphisms via tangent space representations. *NeuroImage (Supplemental issue on Mathematics in Brain Imaging)*, 23:S161–S169. 68, 121
- Valley, J.-F. and Mirimanoff, R.-O. (1993). Comparison of treatment techniques for lung cancer. *Radiotherapy and Oncology*, 28:168–173. 39
- van Herk, M., Bruce, A., Kroes, A. P. G., Shouman, T., Touw, A., and Lebesque, J. V. (1995). Quantification of organ motion during conformal radiotherapy of the prostate by three dimensional image registration. *International Journal of Radiation Oncology*Biology*Physics*, 33(5):1311–1320. 23
- Wand, M. P. and Jones, M. C. (1995). *Kernel Smoothing*. Number 60 in Monographs on Statistics and Applied Probability. Chapman & Hall/CRC. 73, 76, 77, 126
- Warach, S., Gaa, J., Siewert, B., Wielopolski, P., and Edelman, R. R. (1995). Acute human stroke studied by whole brain echo planar diffusion-weighted magnetic resonance imaging. *Annals of Neurology*, 37(2):231–241. 2
- Watson, G. S. (1964). Smooth regression analysis. *Sankhya*, 26:101–116. 74
- Xu, S., Styner, M., Davis, B. C., Joshi, S. C., and Gerig, G. (2006). Group mean differences of voxel and surface objects via nonlinear averaging. In *Proceedings of the 2006 IEEE International Symposium of Biomedical Imaging (ISBI 06)*. 50
- Xu, S., Styner, M., Gilmore, J., and Gerig, G. (2008). Multivariate longitudinal statistics for neonatal-pediatric brain tissue development. In *Proceedings of the SPIE Conference on Medical Imaging (SPIE 2008)*. SPIE. 70
- Yan, D., Lockman, D., Brabbins, D., Tyburski, L., and Martinez, A. (2000). An off-line strategy for constructing a patient-specific planning target volume in adaptive treatment process for prostate cancer. *International Journal of Radiation Oncology*Biology*Physics*, 48(1):289–302. 21
- Younes, L. (2005). Deformation analysis for shape and image processing. *Lecture Notes*. 13
- Younes, L. (2007). Jacobi fields in groups of diffeomorphisms and applications. *Quarterly of Applied Mathematics*, 65:113–134. 122
- Zijdenbos, A. P., Dawant, B. M., Margolin, R. A., and Palmer, A. C. (1994). Morphometric analysis of white matter lesions in MR images: Method and validation. *IEEE Transactions on Medical Imaging*, 13(4):716–724. 44

- Zöllei, L., Learned-Miller, E., Grimson, E., and Wells, W. M. (2005). Efficient population registration of 3D data. In *Proceedings of the 1st International Workshop on Computer Vision for Biomedical Image Applications (CVBIA 05)*, volume 3765/2005 of *Lecture Notes in Computer Science*, pages 291–301. Springer. 67
- Zou, K. H., Warfield, S. K., Baharatha, A., Tempany, C., Kaus, M. R., Haker, S. J., Wells, W. M., Jolesz, F. A., and Kikinis, R. (2004). Statistical validation of image segmentation quality based on a spatial overlap index. *Academic Radiology*, 11:178–189. 39, 45