WASHINGTON UNIVERSITY

SEVER INSTITUTE OF TECHNOLOGY

DEFORMABLE SHAPE MODELS FOR ANATOMY

by

Gary Edward Christensen

Prepared under the direction of Dr. Michael I. Miller

A dissertation presented to the Sever Institute of
Washington University in partial fulfillment
of the requirements for the degree of

DOCTOR OF SCIENCE

August, 1994

Saint Louis, Missouri

WASHINGTON UNIVERSITY

SEVER INSTITUTE OF TECHNOLOGY

ABSTRACT

DEFORMABLE SHAPE MODELS FOR ANATOMY

by Gary Edward Christensen

ADVISOR : Dr. Michael I. Miller

August, 1994

Saint Louis, Missouri

Medical imaging modalities, such as magnetic resonance (MR), histological images, and positron emission tomography (PET), enable study of anatomy and function in animals and humans. The technology to collect such data greatly exceeds tools to analyze it. This research seeks to address this issue by developing methods that automatically synthesize labeled electronic atlases tailored to individuals.

The approach taken is based on Grenander's global shape models. It consists of defining a deformable textbook that mathematically represents the shape and variability of an anatomic region or object at the resolution provided by the sensor. The textbook consists of anatomical information in the form of digitized image volumes (such as isotropic MR or blockface histological 3D data sets) and descriptive information such as structure names, locations, shapes, and functions. Shape variability is modeled by applying random transformations to the textbook coordinate system.

The deformable textbook is used to generate individualized ones. An individualized textbook provides subject specific information (like structure location, volume, shape, etc.) normally only known for the textbook anatomy. It is generated by deforming a generalized textbook into the shape of a particular individual's anatomy. Topological relationships between textbook structures are preserved during deformation by constraining the transformation to obey linear-elastic or viscous-fluid continuum mechanical models.

The deformation techniques developed in this research generalize and extend existing registration methods by accommodating both global and local shape variability. This is accomplished by discretizing the transformation into a dense volume of displacement vectors and using the subject specific image volume for parameter estimation. The estimation procedures are formulated in a Bayesian framework.

The deformable textbook serves as a standard reference frame for comparing function between individuals. Functional data (such as PET images) is mapped from the individual to the textbook coordinate system. This removes anatomical shape variability between subjects and allows precise comparison.

© Copyright by Gary Edward Christensen 1994.
All Rights Reserved.

No part of this dissertation may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the author.

Table of Contents

Page

List of Tables

Table Page

List of Figures

DEFORMABLE SHAPE MODELS FOR ANATOMY

1. Introduction

Imaging modalities such as magnetic resonance (MR), x-ray computed tomography (CT), and positron emission tomography (PET) provide exquisitely detailed *in vivo* information regarding the anatomy and physiological function of specific subjects. However, the interpretation of the data has been hindered by the inability to expeditiously relate such information to specific anatomical regions. The difficulty lies in two areas: images from atlases and other modalities must be registered; but more fundamentally, even when registered, normal variation in anatomy makes interpretation extremely difficult. This dissertation provides algorithms for representing *normal neuroanatomies* by precisely specifying the global anatomical relationships between structures and how they can vary from one individual to another. The goal is to provide representations which allow for the generalization of a single *ideal electronic anatomical textbook* to the *individual*.

The method of Grenander's shape models is used to represent the structure and variability of the anatomy. The structure of a typical normal anatomy is represented by constructing an anatomical textbook, and the shape variability is accommodated by defining probabilistic transformations on the coordinate system of the textbook. By applying the transformations to the textbook, a rich family of anatomies may be represented using a single textbook.

The anatomical textbook is a registered set of images generated from a typical healthy anatomy within the population. It consists of anatomic, symbolic, and database modalities. Examples of anatomical modalities are MR images, CT images, digitized photographs of histological sections, etc. The symbolic modalities contain information about the anatomy, such as structure names, sizes, locations, functionality, dependence on other structures, etc. The database modalities of the textbook refer to data sets that have been registered with the textbook and that have become part of the textbook. Examples of textbook-registered database modalities might be collections of functional images such as PET or SPECT images.

The textbook is individualized by estimating the transformation that accommodates the shape differences between the anatomical modalities of the textbook and those taken of a subject. Once this transformation is known, all the information contained in the symbolic images of the textbook are mapped onto the anatomical images of the subject creating an individualized textbook.

The transformations on the textbook need to maintain the topology of the textbook, and also be of high spatial dimension to accommodate the complex anatomical details. This is achieved by constraining the set of transformations to be consistent with physical deformations of real materials such as elastic solids and viscous fluids. Transformations such as those based on the theory of elasticity develop restoring forces which are proportional to the square of the deformed distance. Except for small deformations, such elastic transformations may prevent the textbook from being fully deformed into the shape of the subject's brain. The shortcomings of the elasticity model can be overcome by using a viscous fluid model which allows the restoring forces to relax over time.

1.1 Description of Deformable Textbook

The deformable-textbook approach for representing the complex structure and variability of anatomy is based on Grenander's global shape models. In this framework the textbook serves as a template that is a representation of a typical anatomy from a set of normal anatomies. The neuroanatomy textbook is a set of a registered anatomical, symbolic, and database images that correspond to a typical healthy brain within the population.

The anatomical modalities of the textbook are used to deform the textbook into the shape of the individual. Examples of the anatomical modalities include

- Magnetic Resonance (MR),

- X-ray Computed Tomography (CT),
- spiral CT,
- ultrasound,
- X-rays, and
- digitized images of histological sections.

The symbolic information contained in the deformable textbook augment and enhance the shape and sensor information contained in the anatomic modalities of the textbook. It is the symbolic data that describes an image scene. For example, the symbolic information describes the attributes of a point in the textbook domain such as:

- structure name,
- function,
- imaging characteristics,
- segmentation label,
- references to research papers,
- visualization parameters,
- textural description,
- probability of a lesion, and
- surgical procedures.

In addition, the textbook information describes the relationships between different structures such as:

- part of,
- next to,
- supplied by, and

- involved in function.

Still other information can be inferred using the above information such as

- structure volume,
- structure shape, and
- shape statistics.

For other examples of symbolic information in an electronic atlases, see the excellent work of Karl Hahns Höhne, et al. [2, 3]. Karl Höhne and colleagues have been pioneers in the field of electronic anatomical atlases over the last decade, and have set the standard for anatomical atlases with the Voxelman software.

The deformable textbook contains a third class of modalities which are called database modalities. Database modalities differ from the anatomical and symbolic modalities of the textbook in that they are acquired as the deformable textbook is used. Examples of database modalities are collections of functional modalities such as

- Positron-Emission Tomography (PET),
- Single-Photon-Emission Tomography (SPECT),
- and stained histological sections,

that have been registered to the textbook. Once registered to the textbook, these modalities become a part of the textbook forming a database of functional information. These database modalities may be used to generate composite images by averaging or subtracting images in the database. Another example is the database of textbook transformations. Each time a data set is registered to the textbook the transformation is saved in the transformation database. This database can be used to generate the average shape of an anatomy or can be used to generate shape statistics.

1.2 Topology of Transformations

It is important that template transformations have properties that preserve the topology of the template, i.e., critical points get mapped to critical points and level curves get mapped to level curves. The template represents our prior knowledge of the shape and composition of an object from a set of topologically similar objects. We assume that any object from this set can be represented by the template by applying a topology preserving transformation to the template coordinate system. The template is used to draw inferences about a particular object by estimating the transformation that deforms the template into the shape of the object and then associating our prior knowledge of the template topology with that of the unknown object.

In the present case, the template is the deformable neuroanatomy textbook. In order to preserve the topology of the textbook anatomy, we restrict the set of textbook transformations to the set of homeomorphic mappings of the textbook coordinate system. A transformation is said to be homeomorphic if it is continuous, globally 1-to-1, and onto (see Chapter 3).

A homeomorphic transformation has the property that connected regions in the textbook are mapped to connected regions in the deformed textbook. This property ensures us that a connected structure in the textbook such as the thalamus or a ventricle gets mapped to a connected structure. In addition, each structure in the textbook gets mapped to only a single structure in the deformed textbook, i.e., structures are not broken up into multiple structures.

Another important property of a homeomorphic transformation is that it has an inverse. An inverse map allows functional information associated with the individual, such as PET, to be mapped from the individual back to the coordinate system of the textbook anatomy. Using the inverse transformation in this manner for multiple individuals removes individual anatomical shape differences from functional modalities.

1.3 Deformable Textbook Applications

The material presented in this section is an overview of a large number of applications of the deformable textbook. These applications are a research topics in their own right and will supply many researchers work for decades if not lifetimes. The important breakthrough that the deformable textbook offers is the ability to rigorously incorporate vast amounts of *a priori* knowledge about the anatomy into the solution of a particular problem. The applications of the deformable textbook are limited only by one's imagination of how to incorporate information into a particular problem.

1.3.1 Comparing Function between Individuals

Mapping function from one anatomy to another is an area of great promise. Currently, a physician or researcher has two options for comparing functional information between two individuals. In the first method, the investigator uses his/her own expertise to map function from one anatomy to the other by eye; i.e., the investigator locates the function in the one image and marks this function in the other based on locating similarities of the anatomical areas. This method is relatively robust because human experts are good at recognition and identification but it is extremely time consuming and tedious. This method also suffers from the problem that no two experts perform the mapping exactly the same way. It also has the limitation that the function that is mapped must be very localized; i.e., small points of activity are easily mapped from one anatomy to another, but large, non-regularly shaped regions are very hard to map by hand.

The second standard method for comparing brain function between individuals consists of averaging functional images across the individuals after the images have been corrected for rigid body translation, rotation, and non-rigid scale. This method has seen success in the area of determining function from PET scans. However, these simple global transformations do not correct for the complex shape differences between different anatomies.

On the other hand, the deformable textbook offers a natural method for mapping functional information from one individual to another. The deformable textbook can be used to remove anatomical shape differences between two functional images by first registering the functional images to anatomical images and then deforming the textbook anatomical image into the study anatomical image. Because the textbook transformation is a homeomorphism, the inverse mapping exists which allows the study functional image to be mapped into the coordinates of the textbook. This procedure has the effect of removing the global and local shape differences between the individual and the textbook.

In order to make the above idea clearer, assume that functional PET volumes are rigidly registered to MR volumes. Once the textbook MR volume is mapped to the study MR volume, the study PET volume can be projected onto the textbook anatomy. Repeating this procedure for multiple subjects will give data sets that can be be averaged together to give better localization of brain activity as well as improving the signal-to-noise ratio. In the results chapter (Chapter 10), an example is presented to show the viability of using the deformable textbook to remove individual anatomical shape differences in PET images. This work is an on going collaborative project with Dr. M. Raichle and Dr. T. Videen of the Neurology Department of Washington University.

A similar application involves mapping the functional information from one monkey visual cortex to another. In a collaborative project Dr. D. Van Essen and Dr. T. Coogan of the Anatomy and Neurobiology Department at Washington University, we have been able to map the visual cortex of one monkey to another (see Chapter 10). By mapping back to the textbook the function data generated from many different experiments, it will be possible to generate a functional database of the monkey visual cortex. Someday it may even be possible to map the functional database of the monkey visual cortex to the human visual cortex in a principled manner.

Eventually the deformable textbook has the potential to become a powerful tool for the diagnosis of abnormal function. For example in PET, a current method used to diagnose abnormal function is to subtract the PET images of normal patients from those of abnormal patients. What is not known is whether this functional difference is due to abnormal function or is due to abnormal anatomical shape differences. The deformable textbook should be able to give researchers a new tool for investigating such questions.

1.3.2 Automatic Segmentation/Neuromorphometrics

Image segmentation is another very useful application of the deformable textbook. It is assumed that the textbook contains an optimal segmentation, in some sense, of the textbook anatomy and what is desired is a similar segmentation of a study anatomy. Once the textbook is deformed into the study, the textbook segmentation is projected (using the estimated transformation) into the coordinate system of the study generating an automatic segmentation of the study. Since the textbook transformation is homeomorphic, the local neighborhood structure of the segmentation is retained in the automatic segmentation. This means that connected structures are mapped to connected structures and adjacent objects remain adjacent.

Closely related to image segmentation is neuromorphometrics, which is the study of the shape of neuroanatomical structures. Currently, researchers trace by hand the structure shapes using data sets from normal and abnormal subjects. Physicians then use these segmentations to diagnosis normal and abnormal structure shapes. Because this approach is very tedious, only a small number of regions are segmented. In sharp contrast, the deformable textbook is both automatic and can segment a larger number of regions than is currently feasible by hand. The number of segmented regions generated by the deformable textbook is limited only by the number of segmented regions in the textbook.

Chapter 10 shows some preliminary results in which the deformable textbook was used to automatically segment MR volumes. In this collaborative project with Dr. M.W. Vannier et al. of the Mallinckrodt Institute of Radiology at Washington University, we are investigating this technique for the diagnosis of Alzheimer's disease. In this project, it is important to measure the atrophy of the gray and white matter and the shape change of the hippocampus.

This technique may prove useful also in the diagnosis of schizophrenia, Huntington's disease, and Parkinson's disease. For example, the structural variability associated with schizophrenia includes abnormally large cerebral ventricles (specific spaces in the brain which contain cerebrospinal fluid) and an apparent thinning of certain areas of the cortex that are involved in emotional expression [4]. In Huntington's disease and Parkinson's disease, it is essential to identify the exact size and position of the head of the caudate nucleus and to distinguish it from the adjacent ventricle [5].

1.3.3 Statistical Characterization of Shape Variability

Statistical characterization of complex shape is a new and important area of research. With the advent of high-resolution 2D and 3D imaging modalities, researchers are now trying to find new ways of describing 3D structures [1, 6], other than by the simple measures of volume and center of mass, in order to characterize normal and abnormal anatomical shapes. Prosthesis designers are particularly interested in characterizing shape so they need to build only a small number optimally shaped parts that will fit most of the population without having to custom manufacture each part.

The deformable textbook offers a new unified approach for statistically characterizing shape. First, parameterized shape models are naturally included into the textbook as part of the *a priori* knowledge. Second, these models give a sample of the underlying shape statistics each time the textbook is deformed into a new study. The textbook is used to classify normal shape by calculating the statistics of the shape model parameters when the textbook is deformed into a large set of normal anatomies. Likewise, abnormal shape is characterized by the statistics of the parameters when the textbook is deformed into a large number of diseased anatomies.

The deformable-textbook idea is well suited for designing generically shaped prosthesis parts or equipment because it can be deformed into the average shape of a set of anatomies. The average anatomy of a set of anatomies is generated in a two step process. The average transformation is generated by averaging the transformations of the textbook to each anatomy in the set of anatomies and then applying the average transformation to the deformable textbook. Once the average anatomy is determined, it is easy to design the particular part or piece of equipment to fit this average anatomy.

1.3.4 Anatomical and Shape Database

One important use of the deformable textbook is for the construction of a database of anatomical and shape information. Each time the textbook is deformed into a new anatomy, we gain important new information such as new shape information or function location in the coordinates of the textbook. Often only part of the information contained in the deformed textbook is useful to an investigator and the rest is discarded. The construction of the database allows other researchers access all of the information generated by the deformations of the textbook and to use the discarded information.

1.3.5 Localization of Regions of Interest

Often times the sizes of the 3D data sets are too large to work with at full resolution while at the same time it is necessary to accurately identify only one or more small regions of the anatomy. This problem can be solved by quickly determining a tight region of interest using the deformable textbook that contains the object in a downsampled data volume. An accurate transformation can then be determined in this small region of interest at the full resolution much faster than could be done working with the entire volume. In addition, this multiresolution approach offers increased performance because less data means faster execution times. This focusing on the particular region of interest can also be used to automatically place active contours and active surfaces [7, 8, 9] for parametric shape representations.

Another possible application for region-of-interest localization is to apply Markov-random-field (MRF) texture-segmentation algorithms [10, 11, 12] in these smaller regions in the cases for which the textbook anatomy differs significantly from the study anatomy. By focusing on a region that contains only two or three structures, the MRF texture segmentation should give good results. Effectively, the focusing on a small region of interest is a simplification or reduction of the complexity of the image scene which will improve the performance of these types of pixel-based algorithms. A drawback of this approach is that it offers no mapping information between the textbook and study anatomies.

1.3.6 Radiation-Treatment Planning

The deformable textbook has a number of interesting possible applications in radiation-treatment planning (RTP). The goal of RTP is to destroy a tumor by radiating it with a focused beam of radiation while sparing nearby normal tissue. The approach is to radiate the tumor from multiple directions while avoiding vital structures that would not survive the radiation dose. The deformable textbook would be used to locate the tumor and to generate an automatic segmentation of nearby vital and nonvital tissues. The segmentation would then be used to generate the treatment dosages and projection paths for treatment of the tumor. Additionally, the deformable textbook would be able to track the shape changes of the tumor over time to evaluate whether or not, additional treatment is required.

Another RTP application is the real-time tracking of a patient's head during treatment. A group at Stanford University [13] is using x-ray projection data of the head collected during treatment to locate the position and orientation of the patients head in real-time. After rigid registration of the x-ray data with previously collected CT data, the path of the radiation beam can be corrected for patient movement. The advantage of this approach is that it eliminates the need to attach a painful stereotactic frame to the skull of the patient for treatment. The individualized textbook used in the planning stage would be natural to use for this registration step and would also provide real-time visualization of the treatment.

1.3.7 Surgical Tool

A growing area of research is the application of computers in the operating room for planning, visualization, object location, and instrument guidance. With the availability of high-resolution MR and inexpensive visualization workstations, surgeons are now able to visualize a surgical procedure before performing the surgery. The deformable textbook will aid in this task by giving the surgeon an atlas that corresponds to the specific patient's anatomy. The deformable textbook could also provide the surgeon with descriptions of surgical procedures if queried. In addition, the deformable atlas can be used as a teaching tool for medical students while they learn about anatomy or practice surgeries.

Another exciting application of the deformable textbook is the fusion of a textbook-generated tumor surface with a real-time video of the patient on the operating table. Current methods [14] have been used by surgeons to trace the area of the tumor on the patient while looking at the fusion of live video and tumor surface. In addition, the region of the tumor is not always evident during the surgery, so that a composite video of the tumor surface and live video can aid the surgeon during removal of the tumor. Research has already begun to create virtual reality head sets for the surgeons that will fuse the textbook information with the live video [14].

Additionally, the individualized deformable textbook could aid in the alignment of surgical instruments. New techniques are being developed for using a stereotactic coordinate system in the operating room to pinpoint and guide surgical instruments [15, 16, 13, 17, 18, 19, 20]. Eventually, the deformable textbook could even be used to guide robots to perform precise surgical procedure such as drilling.

Still another application of the deformable textbook is to predict the movement of the brain based on its physical characteristics and the application of external forces. As an example, it would be helpful to predict how the brain changes its shape and position as a tumor grows and when it is removed. It may be possible to predict these movements by deforming the individualized textbook using the true elastic and fluid characteristics of the brain and the applied forces. Another example is modeling the change in position of the brain after the skull cap is removed and fluid is drained off during surgery.

1.3.8 Incorporating Anatomical Information in Reconstruction Algorithms

Using information from an individualized textbook for other applications is another potentially important use of the deformable textbook. To illustrate this concept, assume that we desired to generate a PET reconstruction algorithm that utilizes *a priori* information about the locations of structure boundaries. The deformable textbook can be used to generate the structure boundaries automatically, given some additional information such as a MR scan of the study region. These boundaries can then be used in the PET reconstruction algorithm to produce a better PET reconstruction than if the boundaries were not known. An example of this application is given in Chapter 10, Results.

1.3.9 Intrasubject Registration

Another use of the deformable textbook is for registering multiple-modality images taken of the same subject. These modalities might be CT, MR, and PET, or they could be the same modality taken at different times. Alignment of different imaging modalities can be accomplished in a number of different ways; one of which is to align each modality to its corresponding modality in the textbook and another is to use corresponding features extracted from two different modalities to drive the registration.

An example of when it is necessary to register a time sequence of images is tracking tumor growth over time in order to validate whether or not a treatment is working. In this case, the first image is considered as the textbook and it is deformed into the next image which is considered the study. This processes is then repeated by viewing the second image in time as the textbook and deforming it into the third, and so on. Other examples include tracking growth or atrophy of structures over time or tracking healing after reconstructive surgery. The deformable textbook idea can even be extended to track motion in movies by viewing the current movie frame as the textbook and the subsequent frame as the study.

1.4 Organization of Dissertation

Chapter 2 reviews the previous research in the field of anatomical atlas registration.

In Chapter 3, the topological properties of the textbook transformations are discussed. Theorems from calculus are presented that describe the properties of homeomorphic (continuous, onto, and one-to-one) transformations. The two most important of theorems are the theorems that give conditions for which a transformation is globally one-to-one. These theorems are used throughout this work to test whether or not a particular transformation is globally one-to-one. Finally, this chapter concludes with simple 2D examples which illustrate how non-homeomorphic transformations destroy the topology of the template.

Chapter 4 defines the deformable textbook mathematically and presents a general approach for individualizing the textbook. This approach is formulated in a Bayesian framework and is based on Grenander's global shape models.

Chapter 5 describes a general approach for inducing prior models by driving linear differential operators with random fields.

Chapter 6 describes the minimum mean-squared error estimation procedure used to individualize the deformable textbook. This method is based on stochastic differential equations which are used to draw samples from the posterior distribution of textbook transformations.

Chapter 7 describes the first prior model investigated in this research. This model is based on the theory of linear elasticity and is induced using the methods described in Chapter 5.

Chapter 8 describes a second method for inducing prior models. This approach is based on using a spatial parameterization of the textbook transformation. Linear elasticity is again used as an example.

A third approach for imposing topological constraints on the textbook transformations is presented in Chapter 9. In this chapter a viscous fluid model is presented for constraining the textbook deformations. The viscous-fluid model has the advantage that it can accommodate long-distance, nonlinear deformations of the textbook. This is in sharp contrast to the small linear deformation assumptions of the linear elasticity, the Laplacian, and the biharmonic models.

Chapter 10 is the results chapter. Two- and three-dimensional results are shown for linear-elastic and fluid deformations of various simulated and real-data textbooks.

Chapters 11 and 12 present a discussion of various topics and the conclusions, respectively.

2. Review of the Literature

Over the last decade, there has been a lot of research in the area of deformable atlases (see for example the proceedings of a recent workshop [21] for a substantive introduction). The methods for accommodating individual shape variability based on deforming a single atlas vary from fully human interactive to semiautomatic to fully automatic. In addition, these methods can be classified as low-dimensional or high-dimensional transformations. Section B has summaries of the papers discussed in this section.

The methods of Fox et al. [22] , Evans et al. [5], and Greitz et al. [23], all fall under the full human intervention, low-dimensional category. The images are manually aligned to the atlas coordinate system using rigid transformations and nonrigid scales. In addition to rigid transformations, the Greitz et al. atlas can also be deformed manually with a small set of nonrigid transformations.

The methods of Bajcsy et al. [24], Pelizzari et al. [25], Undril et al. [26], Bookstein et al. [27], and Evans et al. [28] fall under the semiautomatic, low-dimensional category. All of these methods require the investigator locate landmarks in the data images that correspond with landmarks in the atlas. The first three of these methods accommodate individual shape variability by transforming the atlas with rigid translations, rotations, and scales. The atlas of Undril et al. also includes boundaries of certain structures which can be nonrigidly deformed after the original rigid alignment. The last two approaches match the landmarks in the atlas and the data using thin-plate splines.

The method described by Amit et al. [29] falls under the automatic, low-dimensional category. This method differs from the previous methods in that an image is deformed into another via non-rigid probabilistic transformations. The transformation is defined in terms of a displacement field which is modeled as a two-dimensional Gaussian process with covariance structure given by the Green's function of the Laplacian operator. This displacement field is represented by its Karhunen-Loève expansion where the orthogonal basis is given by the eigenfunctions of the

Laplacian operator. In practice, the Karhunen-Loève expansion was truncated to have under 200 coefficients and is therefore classified as low dimensional.

The methods of Bajcsy et al. [30] and Collins et al. [31] fall under the automatic, high-dimensional category because the transformation is estimated at each pixel in the data images. Both methods maximize the cross-correlation between the deformed atlas and the data and assume a small deformation of the atlas. The method of Bajcsy et al. was the first to use elasticity to regularize the high-dimensional spatial transformation of the anatomical atlas. The method of Collins et al. matched the atlas to the data using a multiresolution approach without regularization.

The work presented in this dissertation is most akin to the *physically based* modeling work of Terzoupolis and Waters [32] and Bajcsy and Kovacic [30] in which transformations are constructed to obey certain physical laws. The approach presented here, while identical in spirit to both, substantially differs from and extends the previous methods. First, the driving force acting throughout the continuum is the non-linear variation of the distortion between the deforming template and the multi-modality imaging data. The full non-linear optimization is solved without linearization of the driving force. Only under the condition that the template and data are sufficiently similar so that the deformations are small, do the linearized solutions give equivalent results (see [33]). Secondly, the deformation procedure is accomplished by solving a sequence of optimization problems from coarse to fine scale via parametrically defined deformation fields. This is analogous to multi-grid methods but here the notion of refinement from coarse to fine is accomplished by increasing the number of basis components. The final stage is to use the translation field over the entire continuum. Thirdly, and perhaps most important, the deformations are not restricted to small linear deformations. By tracking the nonlinear kinematics of the deforming continuum, large distance deformations are possible while maintaining the smoothness and homeomorphic properties of the transformation. Conventional linear approximations assure these properties only for very small linear deformations.

3. Topology in an Eulerian Reference Frame

3.1 Eulerian description of the deformable textbook

The task of individualizing the anatomical textbook is to track the course of every material particle of the textbook from its original to its final position. Each particle can be identified on the basis of either its initial or final coordinates. If a particle is tracked with respect to its original coordinates, the reference frame is called *Lagrangian*; otherwise, if the particle is tracked with respect to its final coordinates, the reference frame is called *Eulerian*.

In this work, an Eulerian reference frame is used to track the deformation of the textbook. This reference frame was chosen because the deformation of the textbook is tracked as a volume.

In an Eulerian reference frame, the state of the deformed textbook is tracked with respect to a fixed voxel lattice. Therefore, there is no extra processing required to display the deformed textbook or to compare it to a study image. However, in a Lagrangian coordinate system, the trajectories of individual material particles of the textbook are tracked over time. As the textbook deforms the particles being tracked move to locations that do not necessarily correspond to pixel locations. Therefore an extra step is required to project the deformed continuum onto a voxel lattice for comparison with a study and for display.

Another reason that an Eulerian reference frame is used is that it is more natural for tracking large-distance, nonlinear deformations. This is because the hard problem of tracking the nonlinear trajectories of individual particles does not have to be performed. Instead, the easier task of observing the continuum as it deforms through fixed observation points is performed.

The Eulerian displacement field $\vec{u}(\vec{X}, t)$ diagrammed in Fig. 3.1 is used to track the deformation of the textbook over time. It is defined as a map from points in the template to fixed observation points in the deforming continuum such that a mass particle instantaneously located at $\vec{X}$ at time t originated at point $\vec{X} - \vec{u}(\vec{X}, t)$. The time t can correspond to real or simulation time.

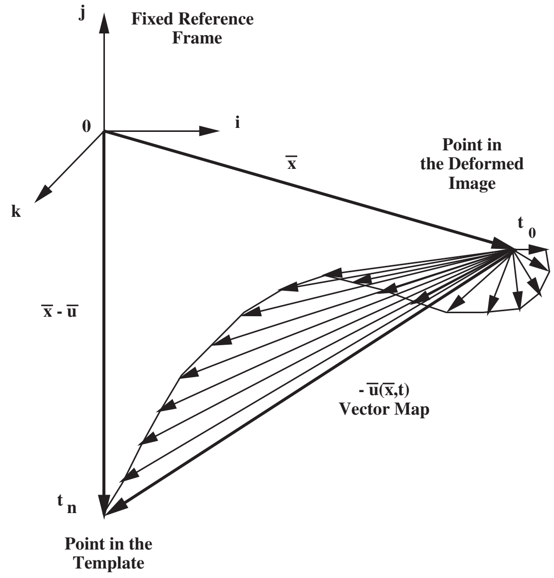

Figure 3.1: Eulerian frame showing coordinate system used to track the kinematics of the deformation. In the Eulerian frame, vector and scalar field variables are associated with fixed spatial positions through which particles move rather than with moving material particles. The point $\vec{X} - \vec{u}(\vec{X}, t_i)$ denotes the particle from the template which is resident at location $\vec{X}$ at time t_i.

3.2 Topological Properties of the Textbook Transformations

There are two kinds of variations which must be accommodated: normal variation between anatomies and diseased or abnormal states. Disease and abnormal variation other than those based on shape changes have not been addressed in this work. Focusing on normal human variation, let $\mathcal{H}$ be the set of homeomorphic transformations of the textbook such that $\vec{h} \; : \; \Omega \; \rightarrow \; \Omega$ for $\vec{h} \in \mathcal{H}$. A homeomorphic transformation is defined to be a continuous, one-to-one, onto mapping. These homeomorphisms are generated from translation groups applied to points $\vec{x} \in \Omega$:

$$\vec{h} \; : \; \vec{x} = (x_1, x_2, x_3) \; \rightarrow (x_1 - u_1(\vec{x}), x_2 - u_2(\vec{x}), x_3 - u_3(\vec{x})) \; . \tag{3.1}$$

The vector field $\vec{u}(\vec{x}) = [u_1(\vec{x}), u_2(\vec{x}), u_3(\vec{x})]^T$ is called the displacement field. The maps constructed from these transformations allow for the dilation, contraction, and warping of the underlying ideal coordinates of the template into the coordinates of the individual anatomy at a very local level. The set of normal anatomies generated from the textbook $(\Omega, T, \mathcal{T})$ becomes $\{T \circ \vec{h} \;:\; \vec{h} \in \mathcal{H}\}$, with $\circ$ the composition operator.

It is illustrative to break down the definition of a homeomorphism in order to understand its importance fully. The first requirement for a transformation to be homeomorphic is that it must be continuous. A transformation is said to be continuous if it satisfies the following definition.

Definition 3.1 ***(Buck [34], p. 332)*** *A transformation $\vec{h}$ defined on a set $D \subset \mathbb{R}^n$ is said to be continuous at a point $p_0 \in D$ if and only if for any $\epsilon > 0$ there is a $\delta > 0$ such that $|h(p) - h(p_0)| < \epsilon$ whenever $|p - p_0| < \delta$ and $p \in D$.*

The notation C is used to denote the class of continuous transformations, and C^n is used to denote the class of transformations which have the property that the n^{th} derivative of each component is continuous. In addition, the notation C' will be used to denote C^1 and C'' to denote C^2.

A continuous transformation of the textbook ensures that the local relationships between structures is maintained. That is, structures that are adjacent to each other before transformation are still adjacent after transformation even though they may have changed shape. An important property of continuous transformations is that the concatenation of continuous transformations is again continuous. Concatenation of transformations occurs when global and local transformations are combined and when templates are propagated in the fluid model (see Chapter 9). These ideas are made rigorous by the following two theorems.

Theorem 3.1 ***(Buck, p. 333)*** *Let $\vec{h}$ be continuous on a set D. Then, any compact set $C \subset D$ is carried by $\vec{h}$ into a compact set $h(C)$, and any connected set $S \subset D$ is carried into a connected set $h(S)$.*

Theorem 3.2 ***(Buck, p. 333)*** *Let the transformation $\vec{g}$ be continuous on a set A and $\vec{h}$ be continuous on a set B, and let $p_0 \in A$ and $g(p_0) = q_0 \in B$. Then, the product transformation $\vec{h}g$, defined by $\vec{h}g(p) = \vec{h}(\vec{g}(p))$, is continuous at p_0.*

The second requirement for a transformation to be a homeomorphism is that it must be globally 1-to-1. A necessary condition for a transformation to be globally 1-to-1 is that it must first be locally 1-to-1.

Definition 3.2 ***(Buck, p. 354)*** *A transformation which is defined in an open set D is said to be locally 1-to-1 in D if about any point $p \in D$ there is a neighborhood in which $\vec{h}$ is 1-to-1.*

A transformation may be tested to determine if it is locally 1-to-1 by computing its Jacobian. If the Jacobian of the transformation is positive at each point in a given set D, the transformation is said to be locally 1-to-1 over D. Additionally, a nonnegative Jacobian over a set D implies that open subsets of D are mapped to open subsets. The following definition and theorems makes the above statements rigorous.

Definition 3.3 ***(Buck, p.354)*** *If $\vec{h}$ is a transformation from R^n into R^n which is of class C' in a set D, then the Jacobian of $\vec{h}$ is the function J defined in D by*

$$J(p) = det\left(\vec{\nabla}\vec{h}|_p\right)$$

For example, if $\vec{h}$ is given by

$$\vec{h}(\vec{x}) = \begin{cases} h_1 = h_1(x_1, x_2, x_3) \\ h_2 = h_2(x_1, x_2, x_3) \\ h_3 = h_2(x_1, x_2, x_3) \end{cases}$$

then

$$J(p) = \begin{vmatrix} \frac{\partial h_1(p)}{\partial x_1} & \frac{\partial h_1(p)}{\partial x_2} & \frac{\partial h_1(p)}{\partial x_3} \\ \frac{\partial h_2(p)}{\partial x_1} & \frac{\partial h_2(p)}{\partial x_2} & \frac{\partial h_2(p)}{\partial x_3} \\ \frac{\partial h_3(p)}{\partial x_1} & \frac{\partial h_3(p)}{\partial x_2} & \frac{\partial h_3(p)}{\partial x_3} \end{vmatrix}.$$

Theorem 3.3 ***(Buck, p. 355)*** *Let $\vec{h}$ be a transformation from R^n into R^n which is of class C' in an open set D, and suppose that $J(p) \neq 0$ for each $p \in D$. Then, $\vec{h}$ is locally 1-to-1 in D.*

Theorem 3.4 ***(Buck, p. 356)*** *Let $\vec{h}$ be of class C' on an open set D in n space, taking values in n space; suppose that $J(p) \neq 0$ for all $p \in D$. Then, $h(D)$ is an open set; thus, $\vec{h}$ carries every open set in D into an open set.*

The next two theorems give the sufficient conditions to insure that a locally 1-to-1 transformation is globally 1-to-1. In practice, these two theorems are very useful for determining whether or not a particular transformation is globally 1-to-1.

Theorem 3.5 ***(Buck, p. 361)*** *If $\vec{h}$ maps a convex set D onto another convex set D^* in such a way that the boundary of D is mapped exactly onto the boundary of D^*, and if $\vec{h}$ is locally 1-to-1 in D, then $\vec{h}$ must be globally 1-to-1.*

Theorem 3.6 ***(Buck, p. 361)*** *If $\vec{h}$ is defined on a rectangular region D in the plane, and if neither the Jacobian of $\vec{h}$ nor any of the diagonal entries of $\vec{\nabla}\vec{h}$ has a zero value in D, then $\vec{h}$ must be globally 1-to-1 in D.*

An important property of a homeomorphic transformation is that a unique inverse transformation exists and is continuous. In addition, if the homeomorphic transformation is also of class C' then the inverse transformation is also of class C'. The existence of the inverse mapping ensures that the transformation can be used to map information not only from the textbook to the individual but also from the individual to the textbook. The next two theorems make these ideas rigorous.

Theorem 3.7 ***(Buck, p. 353)*** *If $\vec{h}$ is continuous and 1-to-1 on a compact set D, then $\vec{h}$ has a unique inverse $\vec{h}^{-1}$ which maps $\vec{h}(D) = D^*$ 1-to-1 onto D, and $\vec{h}^{-1}$ is continuous on D^*.*

Theorem 3.8 ***(Buck, p. 358)*** *Let $\vec{h}$ be of class C' in an open set D, with $J(p) \neq 0$ for all $p \in D$. Suppose also that $\vec{h}$ is globally 1-to-1 in D, so that there is an inverse transformation $\vec{h}^{-1}$ defined on the open set $\vec{h}(D) = D^*$. Then, $\vec{h}^{-1}$ is of class C' on D^*, and $\vec{\nabla}(\vec{h}^{-1})|_q = (\vec{\nabla}\vec{h}|_p)^{-1}$, where $q = \vec{h}(p)$.*

3.3 Example of a Transformation with Singular Jacobian

It is important that the textbook transformation $\vec{h}$ be a homeomorphism so that the deformed textbook is topologically equivalent to the original textbook. A necessary condition that $\vec{h}$ is a homeomorphism is that the Jacobian of $\vec{h}$ for all $\vec{x} \in \Omega$ is positive. The next example demonstrates what happens when this condition is violated in both a Lagrangian and Eulerian reference frame.

Assume that $\vec{h}(\vec{x}) = (x_1 + \frac{1}{4}\sin 2\pi x_1 \sin \pi x_2, x_2)$ for $\vec{x} \in \Omega = [0,1]^2$. Although this transformation is continuous and onto in Ω, it is singular on Ω because its Jacobian $J(\vec{x}) = 1 + \frac{\pi}{2}\cos 2\pi x_1 \sin \pi x_2$ is both positive and negative on Ω. The region where $J(\vec{x})$ is less than or equal to zero is shown as the black region in Fig. 3.2.

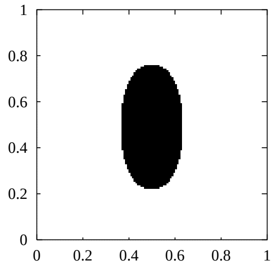

Figure 3.2: The black region is where the Jacobian of $\vec{h}(\vec{x}) = (x_1 + \frac{1}{4}\sin 2\pi x_1 \sin \pi x_2, x_2)$ for $\vec{x} \in \Omega = [0,1]^2$ is less than or equal to zero.

If we interpret $\vec{h}$ with respect to a Lagrangian reference frame, $\vec{h}$ maps points from the undeformed coordinate system to points in the deformed system as a function of points in the undeformed coordinate system. Thus, $\vec{h}(\vec{x})$ gives the location in the deformed coordinate system to which $\vec{x}$ in the undeformed coordinate system maps.

Fig. 3.3 demonstrates the effect of mapping the template shown in the left panel with $\vec{h}$ with respect to a Lagrangian reference. The fact that the Jacobian of $\vec{h}$ is singular on Ω implies that the $\vec{h}$ does not preserve the topology of the template. This is evident because the two vertical lines in the template are mapped to lines that cross over one another in the deformed coordinate system. Transformations with singular Jacobians are often referred to as a folding transformations because they "fold" the coordinate system causing the lines in this case to cross over one another.

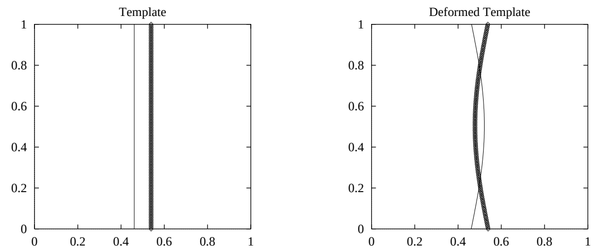

Figure 3.3: Two vertical lines before (left) and after (right) transformation with $\vec{h}(\vec{x}) = (x_1 + \frac{1}{4}\sin 2\pi x_1 \sin \pi x_2, x_2)$ when $\vec{x}$ is with respect to a Lagrangian reference frame.

This example also illustrates that when $\vec{h}$ has a singular Jacobian in a Lagrangian reference frame, $\vec{h}$ is a many-to-one mapping. The two points where the deformed lines intersect are examples of the many-to-one property of $\vec{h}$.

Next we interpret $\vec{h}$ in terms of a Eulerian reference frame. In this case, $\vec{h}$ maps points from the undeformed coordinate system to points in the deformed system as a function of points in the deformed coordinate system. Thus, $\vec{h}(\vec{x})$ is the location of the point in the undeformed coordinate system that maps to $\vec{x}$ in the deformed coordinate system.

Fig. 3.4 demonstrates the effect of mapping three different images with $\vec{h}$ in a Eulerian reference frame. Panels (a)-(c) show the three images before transformation and panels (d)-(f) show the corresponding images after transformation.

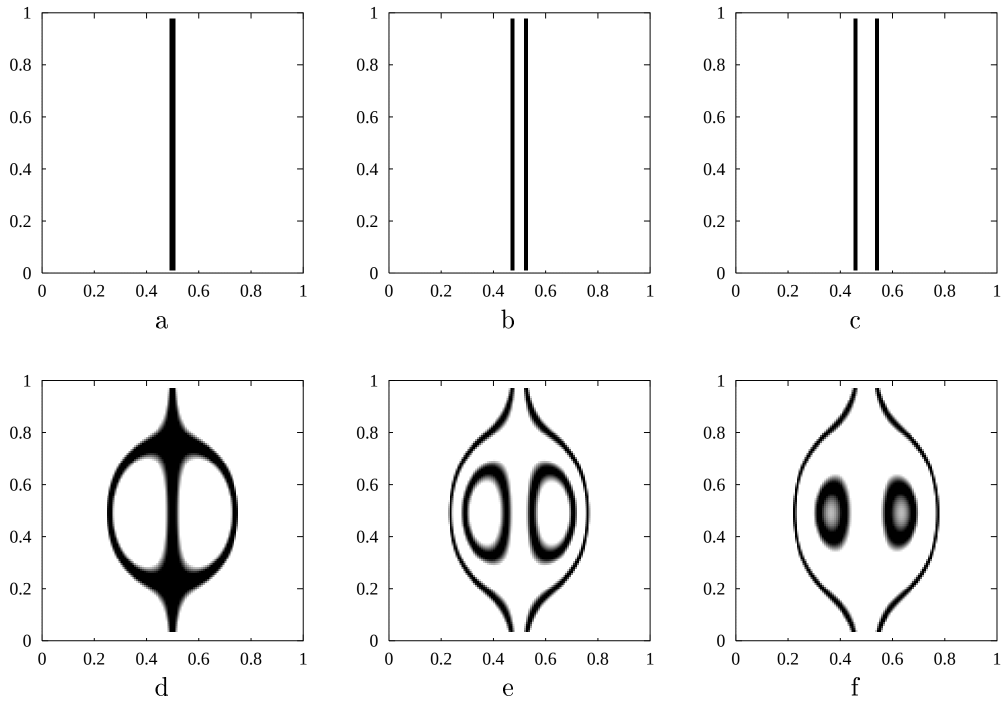

Figure 3.4: Top row shows three images before transformation with $\vec{h}(\vec{x}) = (x_1 + \frac{1}{4}\sin 2\pi x_1 \sin \pi x_2, x_2)$ and the bottom row shows the corresponding images after transformation when $\vec{x}$ is with respect to a Eulerian reference frame.

Similar to the Lagrangian case, the fact that the Jacobian of $\vec{h}$ in the Eulerian reference frame is singular implies that $\vec{h}$ does not preserve the topology of the mapped images. This is evidenced by the fact that vertical lines in the undeformed coordinate system are mapped to non-lines.

In contrast to the Lagrangian case, this example shows that a transformation $\vec{h}$ with a singular Jacobian (with respect to a Eulerian reference frame) is a one-to-many mapping. This is seen by the fact that the single line in panel (a) is mapped to the triple line in panel (d) and that each line in panels (b) and (c) is mapped to a disconnected line and circle.

In contrast to these examples, a homeomorphic mapping has a positive Jacobian. Therefore, this crossing of lines and mapping lines into nonlines cannot happen because topology is preserved.

4. Textbook Individualization Using Bayesian Estimation

4.1 Characterization of the Textbook and the Population

The goal of this work is to understand and compare data sets called *studies* from a large population of topologically equivalent anatomies $\mathbb{A}$. We do this by constructing an anatomical textbook which contains all of our current knowledge of a typical anatomy from $\mathbb{A}$. Understanding is achieved by associating the information in the textbook with the study. This is accomplished by homeomorhically deforming the textbook into the shape of the study, thereby creating an individualized textbook. Comparison of multiple studies is achieved by homeomorphically mapping each study into the common coordinate system of the textbook. This removes anatomical shape variations between the studies, thus allowing for precise comparisons.

Mathematically, the textbook is defined on a fixed coordinate system $\Omega \subset \mathbb{R}^3$ to be an N-valued mapping $T : \Omega \rightarrow \mathcal{T}$. The range space $\mathcal{T}$, is assumed to be an N-fold product of spaces $\mathcal{T}_1 \times \mathcal{T}_2 \times \ldots \mathcal{T}_N$. Each component T_n of T corresponds to a different modality of the textbook and has range space $\mathcal{T}_n$. T_n is real valued when it corresponds to a measurement such as an MR image and is symbolic in value when it corresponds to information such as structure names and functional descriptions. The triple $(\Omega, T, \mathcal{T})$ is termed the *anatomical textbook (template)*.

A study S is an M-valued vector function consisting of M-characterizing data sets $\{S_m\}_{1\leq m\leq M}$, or *sub-studies*. Each sub-study is an examination of the patient's tissue via some sensor. It is assumed that all of the study types already exist in the *ideal* textbook which implies $S_m : \Omega \rightarrow \mathcal{T}_{n_m}$, for some $n_m \in \{1, 2, \ldots, N\}$. For the work described here, the study modalities are assumed to be acquired in register; in general, a second processing step would be required to register modalities from a single patient.

We make the assumption that there exists a set of homeomorphic transformations $\mathcal{H}$ on Ω that transform the shape of the textbook into the shape of each study generated from an anatomy in $\mathbb{A}$. The transformation of the textbook into the shape of a study generates an *individualized* textbook. This individualized textbook contains the same information as the original textbook, but it is customized to the study anatomy.

4.2 Bayesian Formulation

The shape variability of the anatomies in $\mathbb{A}$ is assumed to be a random process. That is, we assume that there exists a function $p(\vec{h}|S)$ which describes the probabilistic dependence of $\vec{h} \in \mathcal{H}$ which individualizes the textbook for a given study S. Because the posterior $p(\vec{h}|S)$ is generally not known, Bayes rule [35]

$$p(\vec{h}|S) \propto p(S|\vec{h})p(\vec{h}) \tag{4.1}$$

is used to define the posterior in terms of a likelihood model $p(S|\vec{h})$ and *a prior* model $p(\vec{h})$. The likelihood models the statistics of the data-collection process while the prior models our a priori knowledge of the statistics of the transformation.

Often it is more convenient to work with the displacement field $\vec{u}(\vec{X}) = \vec{X} - \vec{h}(\vec{X})$ directly instead of $\vec{h}$. In this case, the posterior in terms of $\vec{h}$ is reformulated in terms of $\vec{u}$ by the change of variables $\vec{h}(\vec{X}) = \vec{X} - \vec{u}(\vec{X})$. Therefore, the likelihood and prior models become $p(S|\vec{u}) = p_{S|u}(S|\vec{u}) = p_{S|h}(S|\vec{h}) = p_{S|h}(S|\vec{X} - \vec{u}(\vec{X}))$ and $p(\vec{u}) = p_u(\vec{u}) = p_h(\vec{h}) = p_h(\vec{X} - \vec{u}(\vec{X}))$, respectively. Using Bayes rule, the posterior of $\vec{u}$ given S is

$$p(\vec{u}|S) \propto p(S|\vec{u})p(\vec{u}). \tag{4.2}$$

4.3 Likelihood Models

The likelihood model $p(S|\vec{h})$ defines the probabilistic dependence of the (observed) study images S on the transformation $\vec{h}$. It models the physical constraints of the imaging system, texture variability, and the noise present in the system.

Without loss of generality[1], we assume that the likelihood has a Gibbs' form $p(S|\vec{h}) = \frac{1}{Z}e^{-H(S|\vec{h})}$ where Z is a normalizing constant, and $H(S|\vec{h})$ is the Gibbs' potential. This form of the likelihood is a convenient representation because we will see in Chapter 6 that only the Gibbs' potential of the posterior is required to individualize the textbook.

An example of a Gibbs' potential is the Gaussian potential given by

$$H(S|\vec{h}) = \frac{1}{2\sigma^2}\int_\Omega |T(\vec{h}(\vec{X})) - S(\vec{X})|^2 d\vec{X}. \tag{4.3}$$

where σ is a constant. This potential is useful for systems that can be modeled by Gaussian data collection, structures with Gaussian texture models, and/or systems that contain additive Gaussian noise.

Another example is Csiszár's I-divergence [36] potential given by

$$H(S|\vec{h}) = \int_\Omega T(\vec{h}(\vec{X}))\log\left(\frac{T(\vec{h}(\vec{X}))}{S(\vec{X})}\right)d\vec{X} \; - \int_\Omega \left(T(\vec{h}(\vec{X})) - S(\vec{X})\right)d\vec{X}. \tag{4.4}$$

This potential can be used only when both the textbook and study modalities are positive. The Poisson potential [37] given by

$$H(S|\vec{h}) = \int_\Omega T(\vec{h}(\vec{X}))d\vec{X} - \int_\Omega \ln(T(\vec{h}(\vec{X})))S(d\vec{X}) \tag{4.5}$$

is more appropriate for PET or SPECT data.

In addition to being a function of sensor data, the likelihood can also be a function of features such as landmarks or fiduciary points [27, 28] and edges [30]. This type of likelihood occurs naturally when two different modalities—such as t1-weighted and t2-weighted MR images or MR and PET images—are registered. An example of a feature-based potential is the potential

$$H(S|\vec{h}) = \sum_{m=1}^{M} |\vec{t}_m - \vec{h}(\vec{s}_m)|^2 \tag{4.6}$$

where $\vec{t}_m$ and $\vec{s}_m$ correspond to equivalent landmarks in the textbook and study, respectively.

---

[1] A nonnegative function $f(\vec{x})$ can always be put in the Gibbs' form by using the identity $f(\vec{x}) = \exp(\ln f(\vec{x}))$.

In general, the study consists of many different modalities, each of which may possibly have a different likelihood model. In this case, the likelihood potential becomes a hybrid potential consisting of a weighted sum of potentials or loglikelihoods. To illustrate this, assume that T and S consist of M modalities each satisfying the Gaussian model Eq. (4.3). The likelihood potential is written as the weighted sum

$$H(S|\vec{h}) = \int_\Omega \sum_{m=1}^{M} w_m |T_m(\vec{h}(\vec{X})) - S_m(\vec{X})|^2 d\vec{X}. \tag{4.7}$$

The weights w_m are used to weight the likelihood models of each individual modality. When all of the modalities contain equal amounts of information, the weights are set equal. But, if one modality contains more information than the others, its weight is adjusted to have a larger magnitude than the others.

4.4 Prior Models

The prior model $p(\vec{h})$ measures how likely transformations are regardless of the study S. As mentioned in Chapter 3, we desire the set of textbook transformations to be homeomorphic. In the Bayesian formulation this assumption and others are incorporated into the prior model $p(\vec{h})$.

The approach used to define the prior model on $\vec{h}$ is to define a prior model on the displacement field $\vec{u}(\vec{X}) = \vec{X} - \vec{h}(\vec{X})$. We assume that the prior has the Gibbs' form $p(\vec{u}) = \frac{1}{Z} e^{-H(\vec{u})}$ where $H(\vec{u})$ is the prior potential and Z is a normalizing constant.

Some examples of prior potentials are the Laplacian or membrane model [29, 38]

$$H(\vec{u}) = \frac{1}{2} \int_\Omega \sum_{i=1}^{3} \left(\frac{\partial u_i(\vec{X})}{\partial X_i} \right)^2 d\vec{X}, \tag{4.8}$$

the biharmonic or thin-plate model [38, 39, 40]

$$H(\vec{u}) = \frac{1}{2} \int_\Omega \sum_{i=1}^{3} \sum_{j=1}^{3} \left(\frac{\partial^2 u_i(\vec{X})}{\partial X_i \partial X_j} \right)^2 d\vec{X}, \tag{4.9}$$

and the linear-elasticity model [41, 42]

$$H(\vec{u}) = \int_{\Omega} \sum_{i=1}^{3} \sum_{j=1}^{3} \frac{\lambda}{2} \left(\frac{\partial u_i(\vec{X})}{\partial X_i} \right) \left(\frac{\partial u_j(\vec{X})}{\partial X_j} \right) + \frac{\mu}{4} \left(\frac{\partial u_i(\vec{X})}{\partial X_j} + \frac{\partial u_j(\vec{X})}{\partial X_i} \right)^2 d\vec{X}. \tag{4.10}$$

The coefficients μ and λ are called the Lamé elasticity constants. The linear elasticity model is derived in Appendix A and will be discussed in more depth in the chapters that follow.

4.5 Bayesian Estimators

In order to determine the solution $\vec{h}$ which individualizes the textbook using the Bayesian formulation, we must compute a suitable statistic on the posterior $p(\vec{h}|S)$. Typical statistics are the *Maximum-a-Posteriori* (MAP) estimator, the *Maximum-Likelihood Estimate* (MLE), and the *Minimum-Variance Estimate* (MVE) [35].

The MAP estimate of $\vec{h}$ corresponds the maximum value of the posterior $p(\vec{h}|S)$. This statistic is in general hard to compute unless the functional form of the posterior is known. The MLE estimate is a special case of the MAP estimate which corresponds to the maximum value of the posterior when the prior $p(\vec{h})$ is uniform. In this case the prior assigns equally likely weights to all transformations $\vec{h}$.

In this work the MVE estimate is used to determine the optimal transformation $\vec{h}$ that individualizes the textbook given a study S. The MVE estimate is also referred to as the minimum mean squared error (MMSE) estimate and the *conditional mean* estimate. Chapter 6 discusses this statistic and describes a method to compute it using a stochastic differential equations (SDE).

5. Inducing Priors via PDEs Driven by Random Fields

This chapter describes a general approach for inducing prior distributions on the space of textbook transformations which consists of driving partial differential equations (PDEs) with random fields. The form of the PDEs are based on continuum mechanical models and are selected to enforce desired topological properties on the textbook transformations.

5.1 Linear PDEs Driven by Random Fields

The prior statistics of the displacement field $\vec{u}$ are induced by driving a linear partial differential equation with a random driving function. That is, the statistics of the displacement field $\vec{u}$ are induced from the partial differential equation

$$L\vec{u}(\vec{x}) = \vec{e}(\vec{x}) \tag{5.1}$$

where the operator L is assumed to be a local (bounded support) self-adjoint differential operator and $\vec{e}(\vec{x})$ is a vector-valued random field. Examples of operators are the Laplacian operator such that $L\vec{u}(\vec{x}) = \nabla^2\vec{u}(\vec{x})$ and the linear-elasticity operator such that $L\vec{u}(\vec{x}) = \mu\nabla^2\vec{u}(\vec{x}) + (\mu+\lambda)\vec{\nabla}(\vec{\nabla}\cdot\vec{u}(\vec{x}))$ (μ and λ are the Lamé elasticity constants).

We assume that the operator L has eigenvectors $\{\vec{\phi}_n(\vec{x})\}$ and eigenvalues $\{\kappa_n\}$ which are defined by the equation

$$L\vec{\phi}_n(\vec{x}) = \kappa_n\vec{\phi}_n(\vec{x}).$$

In addition, we assume that the random field $\vec{e}(\vec{x})$ has orthogonal expansion

$$\vec{e}(\vec{x}) = \sum_{n=0}^{\infty} \epsilon_n \vec{\phi}_n(\vec{x}) \tag{5.2}$$

where the coefficients $\{\epsilon_n\}$ are zero-mean independent Gaussian variables with variances σ_n^2 and $\sum_{n=0}^{\infty}\sigma_n^2/\kappa_n^2 < \infty$. The covariance structure of $\vec{e}(\vec{x})$ is therefore $K_e(\vec{x},\vec{y}) = \sum_{n=0}^{\infty}\sigma_n^2\vec{\phi}(\vec{x})\vec{\phi}^T(\vec{y})$.

The statistics of $\vec{u}$ are determined from Eq. 5.1 by solving for $\vec{u}$. Substituting Eq. (5.2) and

$$\vec{u}(\vec{x}) = \sum_{n=0}^{\infty}\mu_n\vec{\phi}_n(\vec{x}) \tag{5.3}$$

into Eq. (5.1) gives

$$\sum_{n=0}^{\infty}\kappa_n\mu_n\vec{\phi}_n(\vec{x}) = \sum_{n=0}^{\infty}\epsilon_n\vec{\phi}_n(\vec{x}). \tag{5.4}$$

Solving this equation for the coefficients $\{\mu_n\}$ of the displacement field gives

$$\mu_n = \frac{\epsilon_n}{\kappa_n} \quad \text{or} \quad \mu_n \sim N\left(0, \frac{\sigma_n^2}{\kappa_n^2}\right) \tag{5.5}$$

provided that $\kappa_n \neq 0$. If $\kappa_n = 0$ we define $\mu_n = 0$. The energy of $\vec{u}$ is given by the inner product

$$< \vec{u}, \vec{u} >= \int_{\Omega}\vec{u}(\vec{x})\cdot\vec{u}(\vec{x})d\vec{x} = \sum_{n=0}^{\infty}\frac{\sigma_n^2}{\kappa_n^2}.$$

We see that the assumption that $\sum_{n=0}^{\infty}\sigma_n^2/\kappa_n^2 < \infty$ insures that $\vec{u}$ has finite energy. Note also that $\vec{u}$ is a Gaussian random process (see Section 5.3) which has covariance structure

$$K_u(\vec{x},\vec{y}) = E[\vec{u}(\vec{x})\vec{u}^T(\vec{y})] = \sum_{n=0}^{\infty}\frac{\sigma_n^2}{\kappa_n^2}\vec{\phi}(\vec{x})\vec{\phi}^T(\vec{y}). \tag{5.6}$$

The infinite dimensional measure on $\vec{u}$ is determined from its finite N-dimensional densities

$$p_N(\mu_1,\ldots,\mu_N) = \left(\prod_{n=1}^{N}\frac{\kappa_n}{\sigma_n\sqrt{2\pi}}\right)\exp -\frac{1}{2}\sum_{n=1}^{N}\frac{\kappa_n^2\mu_n^2}{\sigma_n^2}. \tag{5.7}$$

5.2 Linear PDEs and the Gibbs' Potential Energy

Another way to express Eq. (5.7) is to express it in terms the linear-differential operator L.

Let $\vec{u}_N(\vec{x}) = \sum_{n=1}^{N} \mu_n \vec{\phi}_n(\vec{x})$ represent the N term truncation of $\vec{u}(\vec{x}) = \sum_{n=1}^{\infty} \mu_n \vec{\phi}_n(\vec{x})$ where $L\phi_n = \kappa_n \phi_n$. Let $< \vec{g}, \vec{f} >$ denote the standard inner product $\int_\Omega \vec{g}(\vec{x}) \cdot \vec{f}(\vec{x}) d\vec{x}$ and let $< \vec{g}, \vec{f} >_{K_e^{-1}}$ denote the inner product $\int_\Omega \vec{g}(\vec{x}) \cdot [\int_\Omega K_e^{-1}(\vec{x}, \vec{y}) \vec{f}(\vec{y}) d\vec{y}] d\vec{x}$ where $K_e^{-1}(\vec{x}, \vec{y}) = \sum_{n=0}^{\infty} \vec{\phi}(\vec{x}) \vec{\phi}^T(\vec{y}) / \sigma_n^2$.

Using Parseval's theorem, we can express the energy of the displacement field $\vec{u}_N$ by its inner product, that is,

(5.8)

$$\begin{aligned} < L\vec{u}_N, L\vec{u}_N >_{K_e^{-1}} &= < L \sum_{n=1}^{N} \mu_n \vec{\phi}_n, L \sum_{n=1}^{N} \mu_n \vec{\phi}_n >_{K_e^{-1}} = < \sum_{n=1}^{N} \kappa_n \mu_n \vec{\phi}_n, \sum_{n=1}^{N} \kappa_n \mu_n \vec{\phi}_n >_{K_e^{-1}} \\ &= < \sum_{n=1}^{N} \kappa_n \mu_n \vec{\phi}_n, K_e^{-1} \sum_{n=1}^{N} \kappa_n \mu_n \vec{\phi}_n > = < \sum_{n=1}^{N} \kappa_n \mu_n \vec{\phi}_n, \sum_{n=1}^{N} \frac{\kappa_n \mu_n}{\sigma^2} \vec{\phi}_n > \\ &= \int_\Omega \left(\sum_{n=1}^{N} \kappa_n \mu_n \vec{\phi}_n(\vec{x}) \right) \cdot \left(\sum_{n=1}^{N} \frac{\kappa_n \mu_n}{\sigma^2} \vec{\phi}_n(\vec{x}) \right) d\vec{x} = \sum_{n=1}^{N} \frac{\kappa_n^2 \mu_n^2}{\sigma^2}. \end{aligned}$$

Substituting Eq. (5.8) into Eq. (5.7) gives

$$p_N(\mu_1, \dots, \mu_N) = \left(\prod_{n=1}^{N} \frac{\kappa_n}{\sigma_n \sqrt{2\pi}} \right) \exp -\frac{1}{2} < L\vec{u}_N, L\vec{u}_N >_{K_e^{-1}}, \tag{5.9}$$

which means that the Gibbs' potential energy for the N parameters $\mu_1, \dots, \mu_N$ is given by $-1/2 < L\vec{u}_N, L\vec{u}_N >_{K_e^{-1}}$.

5.3 Orthogonal Expansion Properties

The next two theorems state that Eq. (5.1) induces a quadratic-mean-continuous Gaussian random field.

Definition 5.1 *Let $\{g_i\}$ denote an arbitrary set of single-valued, square-integrable functions with support on the compact set Ω, then a vector valued random field $\vec{u}(\vec{X}) = [u_1(\vec{X}), \dots, u_N(\vec{X})]$ with domain Ω is said to be a* ***Gaussian random field*** *if every linear functional*

$$y = \sum_{i=1}^{N} \int_\Omega g_i(\vec{X}) u_i(\vec{X}) d\vec{X} \tag{5.10}$$

is a Gaussian random variable.

Theorem 5.1 *Let $\{\mu_n\}$ be a sequence of independent, zero-mean, Gaussian random variables with variance σ_n^2 such that $\sum_{n=0}^{\infty} \sigma_n^2 < \infty$ and let $\{\vec{\phi}_n\}$ be a complete orthonormal basis on Ω with respect to finite-energy functions then $\vec{u}(\vec{x}) = \sum_{n=0}^{\infty} \mu_n \vec{\phi}_n(\vec{x})$ is a Gaussian random field.*

Proof. Because the basis functions $\vec{\phi}_n$ form a complete orthonormal basis, we can represent an arbitrary, finite-energy, vector function $\vec{g}(\vec{X})$ on Ω by

$$\vec{g}(\vec{x}) = \sum_{n=0}^{\infty} g_n \vec{\phi}_n(\vec{x}) \tag{5.11}$$

where

$$g_n = \int_{\Omega} \vec{g}(\vec{x}) \cdot \vec{\phi}_n(\vec{x}) d\vec{x}, \quad \text{and} \quad \sum_{n=0}^{\infty} g_n^2 < \infty.$$

Next notice that

$$y = \int_{\Omega} \vec{g}(\vec{x}) \cdot \vec{u}(\vec{x}) d\vec{x} = \int_{\Omega} \left[\sum_{n=0}^{\infty} g_n \vec{\phi}_n(\vec{x}) \right] \cdot \left[\sum_{m=0}^{\infty} \mu_m \vec{\phi}_m(\vec{x}) \right] d\vec{x} = \sum_{n=0}^{\infty} g_n \mu_n.$$

In addition,

$$E[y^2] = E\left[\left(\sum_{n=0}^{\infty} g_n \mu_n \right)^2 \right] \leq E\left[\sum_{n=0}^{\infty} g_n^2 \sum_{n=0}^{\infty} \mu_n^2 \right] = \sum_{n=0}^{\infty} g_n^2 \sum_{n=0}^{\infty} \sigma_n^2 < K \sum_{n=0}^{\infty} g_n^2 \leq \infty.$$

Thus, y is a Gaussian random variable because it is the weighted sum of Gaussian random variables and $E[y^2] < \infty$. This implies $\vec{u}(\vec{X})$ is a Gaussian random field by definition. □

Definition 5.2 *A vector random field $\vec{u}(\vec{X})$ is said to be* ***quadratic mean (q.m.) continuous*** *at $\vec{X} \in \Omega$ if*

$$E\{(\vec{u}(\vec{X}) - \vec{u}(\vec{Y}))^2\} \xrightarrow{|\vec{X}-\vec{Y}| \to 0} 0$$

for $\vec{Y} \in \Omega$. If a field is q.m. continuous at every $\vec{X} \in \Omega$, we shall say that it is a q.m. continuous field.

Theorem 5.2 *If $\{\vec{\phi}_n\}$ is a complete orthonormal basis on Ω and each $\phi_n(\vec{x})$ is continuous on Ω then $\vec{u}(\vec{x}) = \sum_{n=0}^{\infty} \mu_n \vec{\phi}_n(\vec{x})$ is q.m. continuous when $\mu_n \sim N(0, \sigma_n^2)$ and $\sum_{n=0}^{\infty} \sigma_n^2 < M < \infty$.*

Proof. Using the Schwarz inequality, we can write

$$\left(\sum_{n=0}^{\infty}\mu_n\vec{\phi}_n(\vec{x})-\sum_{n=0}^{\infty}\mu_n\vec{\phi}_n(\vec{y})\right)^2\leq\sum_{n=0}^{\infty}\mu_n^2\sum_{n=0}^{\infty}(\vec{\phi}_n(\vec{x})-\vec{\phi}_n(\vec{y}))^2. \tag{5.12}$$

Therefore,

$$E\{(\vec{u}(\vec{x})-\vec{u}(\vec{y}))^2\}=E\left\{\left(\sum_{n=0}^{\infty}\mu_n\vec{\phi}_n(\vec{x})-\sum_{n=0}^{\infty}\mu_n\vec{\phi}_n(\vec{y})\right)^2\right\}\leq E\left\{\sum_{n=0}^{\infty}\mu_n^2\sum_{n=0}^{\infty}(\vec{\phi}_n(\vec{x})-\vec{\phi}_n(\vec{y}))^2\right\}$$
$$\leq\sum_{n=0}^{\infty}\sigma_n^2\sum_{n=0}^{\infty}(\vec{\phi}_n(\vec{x})-\vec{\phi}_n(\vec{y}))^2<M\sum_{n=0}^{\infty}(\vec{\phi}_n(\vec{x})-\vec{\phi}_n(\vec{y}))^2\xrightarrow{|\vec{x}-\vec{y}|\to 0}0$$

The summation goes to zero due to the continuity of the basis functions. □

It is also possible to show under certain conditions [43] that the displacement field $\vec{u}$ satisfying Eq. (5.1) is *almost surely* (a.s.) continuous (see Amit and Piccioni [43] for details).

6. Gibbs' Sampling via Stochastic Differential Equations

The intent of this chapter is to give a background on the stochastic differential equation and to describe how it is used to draw samples from the posterior.

6.1 Stochastic Differential Equations

Let $b_i(t,x)$ and $\sigma_{ij}(t,x)$, $1 \le i,j \le r$, be Borel-measurable functions from $[0,\infty) \times \mathbb{R}^r$ into $\mathbb{R}$, and define the $(r \times 1)$ *drift vector* $b(t,x) = \{b_i(t,x)\}_{1 \le i \le r}$ and the $(r \times r)$ *dispersion matrix* $\sigma(t,x) = \{\sigma_{ij}(t,x)\}_{1 \le i,j, \le r}$. Then an r-dimensional stochastic differential equation (SDE) is given by

$$dX_t = b(t,X_t)dt + \sigma(t,X_t)dW_t \tag{6.1}$$

which is written in component form as

$$dX_t^{(i)} = b_i(t,X_t)dt + \sum_{j=1}^{r} \sigma_{ij}(t,X_t)dW_t^{(j)} \tag{6.2}$$

where $W = \{W_t, \mathcal{F}_t; 0 \le t < \infty\}$ is an r-dimensional Brownian motion with independent components [44]. Eq. (6.1) is a formal equation and only has meaning if interpreted as

$$X_t = X_0 + \int_0^t b(s,X_s)ds + \int_0^t \sigma(s,X_s)dW_s; \quad 0 \le t < \infty, \tag{6.3}$$

or equivalently,

$$X_t^{(i)} = X_0^{(i)} + \int_o^t b_i(s,X_s)ds + \sum_{j=1}^{r} \int_o^t \sigma_{ij}(s,X_s)dW_s^{(j)}; \quad 0 \le t < \infty, \ 1 \le i \le r. \tag{6.4}$$

The last integral in the previous two equations is interpreted as an Ito stochastic integral [45]. The next theorem specifies when the solution to Eq. (6.3) exists and is unique.

Definition 6.1 *Let $||\cdot||$ in the next theorem have the following interpretations:*

- *For every $(r \times 1)$ vector b, let $||b||^2 \equiv \sum_{i=1}^{r} b_i^2$.*
- *For every $(r \times d)$ matrix σ, let $||\sigma||^2 \equiv \sum_{i=1}^{r}\sum_{j=1}^{d} \sigma_{ij}^2$.*
- *For every r-dimensional process $\{X_t = (X_t^{(1)}, \dots, X_t^{(r)}), 0 \le t \le T\}$, let $||X_t||^2 \equiv \sum_{i=1}^{r} (X_t^{(i)})^2$.*

Theorem 6.1 *Let $(\mathbf{\Omega}, \mathcal{F}, \mathbf{P})$ be a fixed probability space and let $W = \{W_t, \mathcal{F}_t; 0 \le t < \infty\}$ be a separable, r-dimensional Brownian defined on this probability space. Let X be an $\mathbb{R}^r$-valued random vector measurable with respect to $\mathcal{F}_0$, independent of W, and satisfy $E||X||^2 < \infty$. Suppose that the coefficients $b(t,x)$ and $\sigma(t,x)$ satisfy the global Lipschitz and linear growth conditions*

$$||b(t,x) - b(t,y)|| + ||\sigma(t,x) - \sigma(t,y)|| \le K||x-y||, \tag{6.5}$$

$$||b(t,x)||^2 + ||\sigma(t,x)||^2 \le K^2(1 + ||x||^2), \tag{6.6}$$

for every $0 \le t < \infty$, $x, y \in \mathbb{R}^r$, where K is a positive constant. Under these hypotheses, there exists a separable and measurable process $\{X_t, \mathcal{F}_t; 0 \le t < \infty\}$ with the following properties:

1. *For each t in $[0,T]$ X_t is $\mathcal{F}_t$-measurable.*
2. *X_t is square-integrable, i.e., for every $T > 0$, there exists a constant C, depending only on K and T, such that*

   $$E||X_t||^2 \le C(1 + E||X||^2)e^{Ct}; \quad 0 \le t \le T. \tag{6.7}$$

3. *$\{X_t, 0 \le t \le T\}$ satisfies Eq. (6.3) with $X_0 = X$ almost surely.*
4. *With probability 1, $\{X_t, 0 \le t \le T\}$ is sample continuous.*
5. *$\{X_t, 0 \le t \le T\}$ is unique with probability 1.*

6. $\{X_t, 0 \le t \le T\}$ *is a Markov process, i.e.,* $P(X_{t+s} \in E|\mathcal{F}_t) = P(X_{t+s} \in E|X_t)$ *for all* $s \ge 0$.

Proof. Properties 1-4 are given by Thm 5.2.9, pg. 289 of Karatzas and Shreve [44]. Property 5 is given by Thm 5.2.5, pg. 287 Karatzas and Shreve. Property 6 is proved by extending the proof of property 6 of Proposition 4.4.1, pg.150 of Wong and Hajek to the r-dimensional case [45]. □

Definition 6.2 *A* r*-dimensional random process* $\{X_t, \mathcal{F}_t, 0 \le t \le T < \infty\}$ *is a diffusion process if it has the following two properties:*

1. X_t *is Markov.*
2. X_t *is sample path continuous.*

Definition 6.3 *The infinitesimal mean and variance of an* r*-dimensional process* X_t *is defined to be:*

$$\text{(6.8)} \qquad \textit{Infinitesimal mean:} \quad m(x,t) = \lim_{h \downarrow 0} \frac{1}{h} E\{X_{t+h} - X_t | X_t = x\}$$

$$\text{(6.9)} \qquad \textit{Infinitesimal variance:} \quad v(x,t) = \lim_{h \downarrow 0} \frac{1}{h} E\{(X_{t+h} - X_t)(X_{t+h} - X_t)^T | X_t = x\}.$$

The r-dimensional process X_t which satisfies Eq. (6.1) and the conditions of Thm. 6.1 is a diffusion. A diffusion process is characterized by its infinitesimal mean and variance.

Theorem 6.2 *The infinitesimal mean and variance of the* r*-dimensional process* $\{X_t, \mathcal{F}_t, 0 \le t \le T\}$ *generated from Eq. (6.1) which satisfies Thm. 6.1 has infinitesimal mean* $m(x,t) = b(x,t)$ *and infinitesimal variance* $v(x,t) = \sigma(x,t)\sigma^T(x,t)$*, or equivalently,* $m_i(x,t) = b_i(x,t)$ *and* $v_{ij}(x,t) = \sum_{k=1}^{r} \sigma_{ik}(x,t)\sigma_{jk}(x,t)$ *for* $1 \le i, j \le r$.

Proof. The infinitesimal mean is given by

$$\begin{aligned} E\{X_{t+h} - X_t | X_t = x\} &= E\left\{ \int_t^{t+h} b(s, X_s)ds + \int_t^{t+h} \sigma(s, X_s)dW_s \middle| X_t = x \right\} \\ &= E\left\{ \int_t^{t+h} b(s, X_s)ds \middle| X_t = x \right\}. \end{aligned}$$

The second integral is zero because $X_t^{(i)}$ and $dW_t^{(j)}$ for $1 \leq i, j \leq r$ are independent and the intervals of Brownian motion have zero expected value. Using the mean value theorem for integrals for $t \leq u \leq t+h$

$$\begin{aligned}
&= E\left\{ b(u, X_u) \int_t^{t+h} ds \middle| X_t = x \right\} \\
&= E\{b(u, X_u)|X_t = x\}h.
\end{aligned}$$

As $h \to 0^+$, we have $u \to t$ and $E\{b(u, X_u)|X_t = x\} \to b(t, X_t)$. This last statement assumes that $b(t, x)$ is continuous in both arguments and uses the fact that X_t is a.s. sample continuous. Thus, the infinitesimal mean is given by

$$m(x,t) = \lim_{h \downarrow 0} \frac{1}{h} E\{X_{t+h} - X_t | X_t = x\} = b(x,t). \tag{6.10}$$

The infinitesimal variance is given by

$$\begin{aligned}
E\{(X_{t+h} - X_t)^2 | X_t = x\} &= E\left\{(X_{t+h} - X_t)(X_{t+h} - X_t)^T | X_t = x\right\} \\
&= E\left\{ \left(\int_t^{t+h} b(s, X_s) ds + \int_t^{t+h} \sigma(s, X_s) dW_s \right)^2 \middle| X_t = x \right\} \\
&= E\Bigg\{ \left(\int_t^{t+h} b(s, X_s) ds \right)^2 + \int_t^{t+h} b(s, X_s) ds \int_t^{t+h} \sigma^T(s, X_s) dW_s \\
&+ \int_t^{t+h} b^T(s, X_s) ds \int_t^{t+h} \sigma(s, X_s) dW_s + \left(\int_t^{t+h} \sigma(s, X_s) dW_s \right)^2 \Bigg| X_t = x \Bigg\} \\
&= E\left\{ \left(\int_t^{t+h} b(s, X_s) ds \right)^2 + \left(\int_t^{t+h} \sigma(s, X_s) dW_s \right)^2 \middle| X_t = x \right\}
\end{aligned}$$

(The last equality is because $X_t^{(i)}$ and $dW_t^{(j)}$ for $1 \leq i, j \leq r$ are independent and the intervals of Brownian motion have zero expected value.)

$$= E\left\{ \left(b(u, X_u) \int_t^{t+h} ds \right)^2 + \left(\sigma(v, X_v) \int_t^{t+h} dW_s \right)^2 \middle| X_t = x \right\}$$

where the mean value theorem for integrals was used twice and $t \leq u, v \leq t+h$

$$= E\{b(u,X_u)b^T(u,X_u)|X_t = x\}h^2 + E\{\sigma(v,X_v)\sigma^t(v,X_v)|X_t = x\}h.$$

As $h \to 0^+$, we have $u \to t$, $v \to t$, $E\{b(u,X_u)b^T(u,X_u)|X_t = x\} \to b(t,X_t)b^T(t,X_t)$, and $E\{\sigma(u,X_u)\sigma^T(u,X_u)|X_t = x\} \to \sigma(t,X_t)\sigma^T(t,X_t)$. This last statement assumes that $b(t,x)$ and $\sigma(t,x)$ are continuous in both arguments and uses the fact that X_t is a.s. sample continuous. Thus, the infinitesimal variance is given by

$$v(t,x) = \lim_{h \downarrow 0} \frac{1}{h} E\{(X_{t+h} - X_t)^2 | X_t = x\} = \sigma(t,x)\sigma^T(t,x). \tag{6.11}$$

□

The next two theorems are presented to give the conditions under which the stochastic differential equation may be used to draw samples from a Gibbs distribution. The proofs of the next two theorems can be found in [46].

Theorem 6.3 *Let $H : \mathbb{R}^r \to \mathbb{R}$ be of class C^2 with gradient ∇H satisfying the global Lipschitz and linear growth conditions defined in Thm. 6.1. If W_t is a standard, $\mathbb{R}^r$-valued Brownian motion then*

$$dX_t = -\frac{1}{2}\nabla H(X_t)dt + dW_t \tag{6.12}$$

has a unique solution for every initial distribution on X_0, and it has the unique invariant measure

$$\mu(dx) = \frac{1}{Z} e^{-\frac{1}{2}H(x)} dx \quad \text{on } \mathcal{B}(\mathbb{R}^r) \tag{6.13}$$

is invariant for Eq. (6.12).

Theorem 6.4 *Let X_t be the r-dimensional Markov process satisfying Thm. 6.3. Then μ is the unique invariant measure of X_t, and the associated chain $X_{i\Delta}$, $\Delta > 0$, converges in total variation norm to μ the invariant measure.*

Proof. This theorem is a simplification of a Thm. 2. [1] which was proved by Grenander and Miller. □

In other words, Thm. 6.4 says that the transition measure of X_t satisfying Thm. 6.3 converges in variation norm to the stationary measure, i.e.,

$$\lim_{t \to \infty} P_t(x, dy) = \mu(dy)$$

with equality in the variation norm sense. The importance of this result is that the transition measure converges to the stationary measure no matter what the initial condition as $t \to \infty$. We use this to get at the empirical average of the time samples of X_t converge to the mean of the invariant measure μ, i.e.,

$$\lim_{N \to \infty} \frac{1}{N} \sum_{i=1}^{N} X_{i\Delta} = \int_{\mathbb{R}^r} x\mu(dx) = \int_{\mathbb{R}^r} x \frac{1}{Z} e^{-\frac{1}{2}H(x)} dx. \tag{6.14}$$

7. Inducing the Linear-Elastic-Basis Prior

This chapter describes the first prior model investigated in this research. This model is based on the theory of linear elasticity and is induced using the methods described in Chapter 5.

7.1 Boundary Conditions

Two different sets of boundary conditions were investigated for the domain $\Omega = [0,1]^3$. The first set are called the *bending* boundary conditions and are given by

$$\begin{aligned}
\left.\frac{\partial u_1}{\partial X_1}\right|_{(0,X_2,X_3)} = \left.\frac{\partial u_1}{\partial X_1}\right|_{(1,X_2,X_3)} = \left.\frac{\partial u_2}{\partial X_2}\right|_{(X_1,0,X_3)} = 0 \\
\left.\frac{\partial u_2}{\partial X_2}\right|_{(X_1,1,X_3)} = \left.\frac{\partial u_3}{\partial X_3}\right|_{(X_1,X_2,0)} = \left.\frac{\partial u_3}{\partial X_3}\right|_{(X_1,X_2,1)} = 0 \\
u_1(X_1,0,X_3) = u_1(X_1,1,X_3) = u_1(X_1,X_2,0) = u_1(X_1,X_2,1) = 0 \\
u_2(0,X_2,X_3) = u_2(1,X_2,X_3) = u_2(X_1,X_2,0) = u_2(X_1,X_2,1) = 0 \\
u_3(0,X_2,X_3) = u_3(1,X_2,X_3) = u_3(X_1,0,X_3) = u_3(X_1,1,X_3) = 0.
\end{aligned} \tag{7.1}$$

These boundary conditions allow the unit cube to be mapped to a warped cube in which the boundaries of the cube are free to bend inward or outward on all six faces of the cube while the corners of the the cube are fixed in their original position. These boundary conditions are useful for the case when either the object in the textbook and/or the study is close to the boundary. A 2D example of a transformation satisfying these boundary conditions is shown in the left panel of Fig. 7.1.

Figure 7.1: Left panel shows the 2D bending boundary conditions and the right panel shows the 2D sliding boundary conditions.

The second set of boundary conditions

$$
\begin{aligned}
&u_1(0,X_2,X_3) = u_1(1,X_2,X_3) = u_2(X_1,0,X_3) = 0 \\
&u_2(X_1,1,X_3) = u_3(X_1,X_2,0) = u_3(X_1,X_2,1) = 0 \\
&\frac{\partial u_1}{\partial X_2}\Big|_{(X_1,0,X_3)} = \frac{\partial u_1}{\partial X_2}\Big|_{(X_1,1,X_3)} = \frac{\partial u_1}{\partial X_3}\Big|_{(X_1,X_2,0)} = \frac{\partial u_1}{\partial X_3}\Big|_{(X_1,X_2,1)} = 0 \\
&\frac{\partial u_2}{\partial X_1}\Big|_{(0,X_2,X_3)} = \frac{\partial u_2}{\partial X_1}\Big|_{(1,X_2,X_3)} = \frac{\partial u_2}{\partial X_3}\Big|_{(X_1,X_2,0)} = \frac{\partial u_2}{\partial X_3}\Big|_{(X_1,X_2,1)} = 0 \\
&\frac{\partial u_3}{\partial X_1}\Big|_{(0,X_2,X_3)} = \frac{\partial u_3}{\partial X_1}\Big|_{(1,X_2,X_3)} = \frac{\partial u_3}{\partial X_2}\Big|_{(X_1,0,X_3)} = \frac{\partial u_3}{\partial X_2}\Big|_{(X_1,1,X_3)} = 0.
\end{aligned}
\tag{7.2}
$$

are called the *sliding* boundary conditions. These boundary conditions differs from the previous in that they map the unit cube into the unit cube. The term sliding refers to the fact that points located on the boundary of the cube are mapped onto the boundary of the cube, but they can slide along the side of the cube. The right panel of Fig. 7.1 shows a 2D example of a transformation which satisfies the sliding boundary conditions.

7.2 3D Eigenvectors of the Linear-elasticity Operator

The linear operator L associated with linear elasticity is given by

$$L\vec{u} = \mu\nabla^2\vec{u}(\vec{X}) + (\mu+\lambda)\vec{\nabla}(\vec{\nabla}\cdot\vec{u}(\vec{X})) \tag{7.3}$$

where λ, and μ are the Lamé elasticity constants. The eigenvectors of L were derived in Appendix A for the bending Eq. (7.1) and sliding Eq. (7.2) boundary conditions and correspond to

$$\vec{\phi}_{1,ijk}(\vec{X}) = \alpha_1 \begin{bmatrix} i\cos i\pi X_1 \sin j\pi X_2 \sin k\pi X_3 \\ j \sin i\pi X_1 \cos j\pi X_2 \sin k\pi X_3 \\ k \sin i\pi X_1 \sin j\pi X_2 \cos k\pi X_3 \end{bmatrix}, \quad \vec{\phi}_{2,ijk}(\vec{X}) = \alpha_2 \begin{bmatrix} -j\cos i\pi X_1 \sin j\pi X_2 \sin k\pi X_3 \\ i \sin i\pi X_1 \cos j\pi X_2 \sin k\pi X_3 \\ 0 \end{bmatrix},$$

$$\vec{\phi}_{3,ijk}(\vec{X}) = \alpha_3 \begin{bmatrix} ik\cos i\pi X_1 \sin j\pi X_2 \sin k\pi X_3 \\ jk \sin i\pi X_1 \cos j\pi X_2 \sin k\pi X_3 \\ -(i^2+j^2) \sin i\pi X_1 \sin j\pi X_2 \cos k\pi X_3 \end{bmatrix}, \tag{7.4}$$

and

$$\vec{\phi}_{1,ijk}(\vec{X}) = \alpha_1 \begin{bmatrix} i\sin i\pi X_1 \cos j\pi X_2 \cos k\pi X_3 \\ j \cos i\pi X_1 \sin j\pi X_2 \cos k\pi X_3 \\ k \cos i\pi X_1 \cos j\pi X_2 \sin k\pi X_3 \end{bmatrix}, \quad \vec{\phi}_{2,ijk}(\vec{X}) = \alpha_2 \begin{bmatrix} -j\sin i\pi X_1 \cos j\pi X_2 \cos k\pi X_3 \\ i \cos i\pi X_1 \sin j\pi X_2 \cos k\pi X_3 \\ 0 \end{bmatrix},$$

$$\vec{\phi}_{3,ijk}(\vec{X}) = \alpha_3 \begin{bmatrix} ik\sin i\pi X_1 \cos j\pi X_2 \cos k\pi X_3 \\ jk \cos i\pi X_1 \sin j\pi X_2 \cos k\pi X_3 \\ -(i^2+j^2) \cos i\pi X_1 \cos j\pi X_2 \sin k\pi X_3 \end{bmatrix}, \tag{7.5}$$

respectively. Both sets of eigenvectors have the eigenvalues

$$\kappa_{1,i,j,k} = -\pi^2(2\mu+\lambda)(i^2+j^2+k^2) \quad \text{and} \quad \kappa_{2,i,j,k} = \kappa_{3,i,j,k} = -\pi^2\mu(i^2+j^2+k^2). \tag{7.6}$$

and the coefficients (which scale each eigenvector to unit energy)

$$\alpha_1 = \sqrt{\frac{8}{i^2+j^2+k^2}}, \quad \alpha_2 = \sqrt{\frac{8}{i^2+j^2}}, \quad \text{and} \quad \alpha_3 = \sqrt{\frac{8}{(i^2+j^2)(i^2+j^2+k^2)}}. \tag{7.7}$$

7.3 3D Eigenvectors for the Laplacian Operator

The Laplacian model [29] is presented here to compare with the linear-elasticity model. The Laplacian operator is a simplification of the linear-elasticity operator and is given by

$$L = \nabla^2 = \frac{\partial^2}{\partial x_1^2} + \frac{\partial^2}{\partial x_2^2} + \frac{\partial^2}{\partial x_3^2}. \tag{7.8}$$

This operator has the eigenvectors

$$\begin{aligned}
\vec{\phi}_{1,ijk}(\vec{X}) &= \frac{1}{\sqrt{8}}[\cos i\pi X_1 \sin j\pi X_2 \sin k\pi X_3, 0, 0]^T \\
\vec{\phi}_{2,ijk}(\vec{X}) &= \frac{1}{\sqrt{8}}[0, \sin i\pi X_1 \cos j\pi X_2 \sin k\pi X_3, 0]^T \\
\vec{\phi}_{3,ijk}(\vec{X}) &= \frac{1}{\sqrt{8}}[0, 0, \sin i\pi X_1 \sin j\pi X_2 \cos k\pi X_3]^T
\end{aligned}$$

for the bending boundary conditions given in Eq. (7.1) and

$$\begin{aligned}
\vec{\phi}_{1,ijk}(\vec{X}) &= \frac{1}{\sqrt{8}}[\sin i\pi X_1 \cos j\pi X_2 \cos k\pi X_3, 0, 0]^T \\
\vec{\phi}_{2,ijk}(\vec{X}) &= \frac{1}{\sqrt{8}}[0, \cos i\pi X_1 \sin j\pi X_2 \cos k\pi X_3, 0]^T \\
\vec{\phi}_{3,ijk}(\vec{X}) &= \frac{1}{\sqrt{8}}[0, 0, \cos i\pi X_1 \cos j\pi X_2 \sin k\pi X_3]^T
\end{aligned}$$

for the sliding boundary conditions given in Eq. (7.2). The eigenvalues associated with these eigenvectors are

$$\kappa_{n,i,j,k} = -\pi^2(i^2 + j^2 + k^2) \qquad \text{for } n = 1, 2, 3.$$

7.4 Linear-Elastic-Basis Model

The model used to induce the linear-elastic-basis prior is derived as follows. Let $\vec{u}$ satisfy

$$L\vec{u}(\vec{x}) = \vec{e}(\vec{x}) \tag{7.9}$$

where L is given by Eq. (7.3) and

$$\vec{e}(\vec{x}) = \sum_{\substack{i,j,k=0 \\ i+j+k>0}}^{d} \sum_{n=1}^{3} \epsilon_{n,i,j,k} \vec{\phi}_{n,i,j,k}(\vec{x})$$

such that $\epsilon_{n,i,j,k} \sim N(0, \kappa_{n,i,j,k})$, the basis functions $\phi_{n,i,j,k}$ are given by Eq. (7.4) or (7.5), and $\kappa_{n,i,j,k}$ are given by Eq. (7.6). Setting

$$\vec{u}(\vec{x}) = \sum_{\substack{i,j,k=0 \\ i+j+k>0}}^{d} \sum_{n=1}^{3} \mu_{n,i,j,k} \vec{\phi}_{n,i,j,k}(\vec{x})$$

and solving for $\mu_{n,i,j,k}$ in Eq. (7.9) gives $\mu_{n,i,j,k} \sim N(0, \kappa_{n,i,j,k})$. Therefore, the Gibbs' potential associated with $\vec{u}$ is given by

$$H(\vec{u}) = \sum_{\substack{i,j,k=0 \\ i+j+k>0}}^{d} \sum_{n=1}^{3} \kappa_{n,i,j,k} \xi^2_{n,i,j,k}. \tag{7.10}$$

Alternatively, Eq. (7.10) can be derived using the continuous strain-energy potential for linear elasticity derived in Appendix A. Restated here, the strain-energy potential for linear elasticity is given by

$$H(\vec{u}) = \sum_{i=1}^{3} \sum_{j=1}^{3} \int_{\Omega} \frac{\lambda}{2} \left(\frac{\partial u_i(\vec{X})}{\partial X_i} \right) \left(\frac{\partial u_j(\vec{X})}{\partial X_j} \right) + \frac{\mu}{4} \left(\frac{\partial u_i(\vec{X})}{\partial X_j} + \frac{\partial u_j(\vec{X})}{\partial X_i} \right)^2 d\vec{X}. \tag{7.11}$$

We assume that the displacement field can be represented by

$$\vec{u}(\vec{X}) = \underset{d\to\infty}{\mathrm{l.i.m.}} \sum_{\substack{i,j,k=0 \\ i+j+k>0}}^{d} \sum_{n=1}^{3} \xi_{n,i,j,k} \vec{\phi}_{n,i,j,k}(\vec{X}). \tag{7.12}$$

where $\phi_{n,i,j,k}$ are given by Eq. (7.4) or (7.5) and $\xi_{n,i,j,k}$ are unknown random variables. The finite truncation of $\vec{u}$ is denoted by

$$\vec{u}_d(\vec{X}) = \sum_{\substack{i,j,k=0 \\ i+j+k>0}}^{d} \sum_{n=1}^{3} \xi_{n,i,j,k} \vec{\phi}_{n,i,j,k}(\vec{X}). \tag{7.13}$$

Inserting Eq. (7.13) into (7.11) gives the potential energy

$$H(\vec{u}) = \sum_{\substack{i,j,k=0 \\ i+j+k>0}}^{d} \sum_{n=1}^{3} \kappa_{n,i,j,k} \xi_{n,i,j,k}^2 \tag{7.14}$$

where $\kappa_{n,i,j,k}$ are given by Eq. (7.6). Thus, the coefficients $\xi_{n,i,j,k}$ are independent, zero mean, Gaussian random variables with variance $\kappa_{n,i,j,k}$.

7.5 Generalizations of the Linear-Elastic-Basis Model

The linear-elastic-basis model derived in Section 7.4 can be generalized in the following ways.

The first generalization is to change the statistics of the random driving function by changing the statistics of the basis coefficients of the random field $\vec{e}$. For example, suppose the statistics of the basis coefficients $\vec{e}$ are changed from $\epsilon_{n,i,j,k} \sim N(0, \kappa_{n,i,j,k})$ to $\epsilon_{n,i,j,k} \sim N(0,1)$. This results in a prior model on $\vec{u}$ such that the prior Gibbs' potential is given by

$$H(\vec{u}) = \sum_{\substack{i,j,k=0 \\ i+j+k>0}}^{d} \sum_{n=1}^{3} \kappa_{n,i,j,k}^2 \xi_{n,i,j,k}^2. \tag{7.15}$$

Notice that this model allows for less variation in the displacement field because the coefficients $\xi_{n,i,j,k}$ have smaller variances, that is, the parameters $\xi_{n,i,j,k}$ are now $N(0, \kappa_{n,i,j,k}^2)$ as opposed to $N(0, \kappa_{n,i,j,k})$.

Another generalization of the model $L\vec{u} = \vec{e}$ is to use a linear operator L' that is a polynomial function p of the operator L. For example, if $L = \nabla^2$ is the Laplacian operator and $p(x) = x^2$, then a new model can be induced on $\vec{u}$ using $L'\vec{u} = \vec{e}$ where $L' = p(L) = \nabla^4$ is the biharmonic operator.

Let the eigenvectors and eigenvalues of L be given by ϕ_n and λ_n, respectively, then the eigenvectors and eigenvalues of $L' = p(L)$ are given by ϕ_n and $p(\lambda_n)$, respectively. The statistical model induced by $L'\vec{u} = \vec{e}$ (where the basis coefficients of $\vec{e}$ are $N(0,1)$) is given by the Gibbs' distribution with potential

$$H(\vec{u}) = \sum_{\substack{i,j,k=0 \\ i+j+k>0}}^{d} \sum_{n=1}^{3} p^2(\kappa_{n,i,j,k})\xi_{n,i,j,k}^2. \tag{7.16}$$

7.6 Multiresolution Using the Orthogonal-Basis Approximation

We follow the multiresolution approach first presented by Amit, et al. [29, 33]. It is natural to estimate the large global shape changes between the textbook and the study first and then refine the transformation to account for smaller and smaller shape differences. This approach avoids many local minima in the estimation of the textbook transformation and also increases the rate of convergence of the estimation procedure.

Eq. (7.13) defines an indexed set of multiresolution displacement fields that can be naturally refined from coarse to fine representation. For notational convenience, Eq. (7.13) is restated as

$$\vec{u}^d(\vec{X}) = \sum_{\substack{i,j,k=0 \\ i+j+k>0}}^{d} \sum_{n=1}^{3} \xi_{n,i,j,k}^d \vec{\phi}_{n,i,j,k}(\vec{X})$$

to emphasize the dependence on the value of d. Notice that as d is increased this approximation allows for more variability by increasing the harmonic frequency of the basis functions used in the approximation. Due to orthogonality of the basis functions, notice that the approximation $\vec{u}^d$ is equal to the approximation $\vec{u}^{d+1}$ when

$$\xi_{n,i,j,k}^{d+1} = \begin{cases} \xi_{n,i,j,k}^d, & \text{if } 1 \le n \le 3 \text{ and } 0 \le i,j,k \le d, \\ 0, & \text{if } 1 \le n \le 3 \text{ and } i,j,k = d+1. \end{cases} \tag{7.17}$$

Thus, it is possible to estimate the displacement field first at a coarse resolution or small value of d and then to refine the estimate by increasing d.

The parameters of the $d+1$ approximation are initialized as follows. The parameters that are common between the $d+1$ and d approximations are set equal and the additional parameters are initialized in one of two following ways. The first approach is to initialize the additional parameters with their mean value. This corresponds to using Eq. (7.17) when the basis coefficients have zero mean. The second approach is to initialize the new parameters by drawing samples of the coefficients based on their statistics. In practice, we have used the first approach.

7.7 Topology of an Orthogonal Expansion

As mentioned in Chapter 3, the transformations that are applied to the textbook should be homeomorphic, i.e., continuous, globally 1-to-1, and onto. Unfortunately, the linear-elastic model is only globally 1-to-1 for small deformations and does not constrain transformations to be globally 1-to-1. This shortcoming is due to the linearization of the strain-displacement relationship in the derivation of the linear-elastic model (see Appendix A). The next theorem illustrates this fact by deriving the conditions for the orthogonal expansion to be globally 1-to-1 and showing that these conditions are not necessarily met.

Theorem 7.1 *Let $\{\mu_n\}$ be a sequence of independent, zero-mean, Gaussian random variables with variance σ_n^2 such that $\sum_{n=0}^{\infty} \sigma_n^2 < \infty$ and let $\{\vec{\phi}_n\}$ be a complete orthonormal basis on the compact set Ω such that each $\phi_n \in C^1$, then $\vec{h}(\vec{x}) = \vec{x} - \vec{u}(\vec{x})$ where $\vec{u}(\vec{x}) = \sum_{n=0}^{\infty} \mu_n \vec{\phi}_n(\vec{x})$ is not necessarily a globally 1-to-1 transformation.*

Proof. To prove this theorem, it is only necessary to show that the diagonal entries of $\nabla\vec{h}$ may be less than zero on Ω (see Theorem 3.6).

A necessary condition for $\vec{h}$ to be globally 1-to-1 is that the diagonal entries of $\nabla\vec{h}$ must be positive, that is,

$$\frac{\partial h_i}{\partial x_i} = 1 - \frac{\partial u_i}{\partial x_i} > 0 \quad \text{for } i\text{=1, 2, 3} \tag{7.18}$$

or

$$1 - \sum_{n=0}^{\infty} \mu_n \frac{\partial \phi_n^{(i)}(\vec{x})}{\partial x_i} > 0 \qquad \text{for } i = 1, 2, 3. \tag{7.19}$$

However, the coefficients are $\mu_n \sim N(0, \sigma_n^2)$ which does not ensure that the conditions in Eq. (7.19) are met. A trivial example can be constructed by assuming that all of the coefficients of the expansion are zero except for μ_n. For this case, Eq. (7.19) becomes

$$1 \pm \mu_n \frac{\partial \phi_n^{(i)}(\vec{x})}{\partial x_i} > 0 \qquad \text{for } i = 1, 2, 3.$$

or

$$|\mu_n| < \frac{1}{\partial \phi_n^{(i)}(\vec{x}) / \partial x_i} \qquad \text{for } i = 1, 2, 3.$$

Clearly there is a finite probability in which these conditions are violated because μ_n is a Gaussian random variable and $\partial \phi_n^{(i)}(\vec{x}) / \partial x_i$ is bounded on Ω. □

The implication of Thm. 7.1 is that when using the Laplacian operator or the linear-elasticity operator or any polynomial power of these operators, one must manually check to determine if the generated transformation is globally one-to-one. This can be done by checking that the Jacobian of the transformation is positive for all $\vec{x} \in \Omega$ (see Thm. 3.5). An example of a polynomial power of the the Laplacian operator $L = \nabla^2$ is the the biharmonic operator $L^2 = \nabla^4$.

7.8 Minimum Mean-Square Error Estimation

The conditional mean of the posterior

$$p(\vec{h}|S) \propto p(S|\vec{h}) p(\vec{h}) \tag{7.20}$$

(where $p(S|\vec{h})$ is the likelihood and $p(\vec{h})$ is the prior) is defined to be the transformation $\vec{h}$ that individualizes the textbook to the study S. The conditional mean estimate of $\vec{h}$ is also called the Minimum Mean Squared Error (MMSE) estimate. The method of solution is to generate the transformation field $\vec{h}$ which is the mean of the posterior distribution. For this, the stochastic differential equation (SDE) is used to draw samples from the posterior. Under conditions specified in Chapter 6, the averages of the parameter estimates generated from the SDE converge to the mean of the posterior. We note that because of finite the finite samples used to calculate the average, we are only assured that the mean of the posterior distribution is computed locally around the large energy minima.

The posterior for the linear elasticity basis model is constructed by combining the Gaussian likelihood model in Eq. (4.3) and the linear elastic basis model in Eq. (7.10). After discretizing the Gaussian potential, this posterior has Gibbs' potential

$$H(\vec{\xi}|S) = \frac{1}{\sigma^2}\sum_{n=0}^{N-1} |T(\vec{X}_n - \vec{u}(\vec{X}_n)) - S(\vec{X}_n)|^2 + \sum_{\substack{i,j,k=0\\ i+j+k>0}}^{d} \sum_{n=1}^{3} \kappa_{n,i,j,k}\xi^2_{n,i,j,k}, \tag{7.21}$$

where

$$\vec{u}(\vec{X}) = \sum_{\substack{i,j,k=0\\ i+j+k>0}}^{d} \sum_{n=1}^{3} \xi_{n,i,j,k}\vec{\phi}_{n,i,j,k}(\vec{X}), \tag{7.22}$$

$\phi_{n,i,j,k}$ is defined by Eqs. (7.4) or (7.5), and $\kappa_{n,i,j,k}$ is defined by Eq. (7.6). The stochastic differential equation that samples from the Gibbs' distribution with potential given in Eq. (7.21) is

$$d\xi_t^{(p,i,j,k)} = -\frac{1}{2}\frac{\partial H(\vec{u}_t|S)}{\partial \xi_t^{(p,i,j,k)}}dt + dW_t^{(p,i,j,k)} \qquad \text{for } p = 1,2,3 \tag{7.23}$$

where

$$\frac{\partial H(\vec{u})}{\partial \xi_{p,i,j,k}} = -\frac{1}{\sigma^2}\sum_{n=1}^{N}\Big(T(\vec{X}_n - \vec{u}(\vec{X}_n)) - S(\vec{X}_n)\Big)\nabla T(\vec{X}_n - \vec{u}(\vec{X}_n))\cdot \vec{\phi}_{p,i,j,k}(\vec{X}_n) + \beta_{p,i,j,k}\xi_{p,i,j,k}, \tag{7.24}$$

and $W_t^{(n,i,j,k)}$ is a $3 \times (d+1)^3$-dimensional Brownian motion with independent components.

8. Inducing a Spatially Sampled Elasticity Prior

This chapter presents a second method for inducing a prior model for our Bayesian framework. It is based on discretizing the strain-energy potential of the linear-elasticity model to generate a spatially sampled parameterization of the displacement field. This model is useful for accommodating local shape differences as opposed to global shape differences.

8.1 Motivation

The spatial parameterization of the textbook transformation is motivated by the following two factors:

1. There is a need to accommodate the complex shape and variation of anatomical structures not only at the global level but also at the local level.

2. A spatial parameterization of the transformation maps naturally onto a massively parallel computer.

The fact that anatomical and biological shapes are so complex and variable from one subject to another requires that transformations of the textbook be able to accommodate shape differences at a very fine level of detail. The spatially sampled parameterization of the the transformation allows for very local shape accommodation because the interactions between displacement vectors are local and not global. This is in contrast to the basis parameterizations presented in Chapter 7 which have global support over the whole domain Ω. With the former representations, a small change in one of the basis coefficients has an effect over the whole image. Thus, local shape accommodation is more difficult with global representations than with the local representations. As a side-note, wavelet bases [33] are somewhere in between the very local interactions of the spatially sampled parameterizations and the global parameterizations.

Spatial parameterizations of volume transformations are easily computed on massively parallel computers. This is due to the fact that spatially sampled volume transformations can be partitioned into 2D planes for storage and computation. In addition, spatial parameterizations transformations only have local interactions. In other words, most computations need either no-neighbor or nearest-neighbor communication. Massively parallel machines are designed for just such computations and are therefore well suited to calculate spatially sampled transformations.

8.2 Spatial Parameterization of Textbook Transformation

Partition the domain Ω into a set of ordered points

$$\vec{X}_n = \left[(n \mod N_1)\Delta_1, (n \mod (N_1 N_2))\Delta_2, \lfloor \frac{n}{N_1 N_2} \rfloor \Delta_3\right]^T, \tag{8.1}$$

where N_1, N_2, and N_3 are the numbers of partition points of Ω in the X_1, X_2, and X_3 directions, respectively. The separations between partition points in the X_1, X_2, and X_3 directions are defined to be

$$\Delta_1 = \frac{1}{N_1}, \quad \Delta_2 = \frac{1}{N_2}, \quad \text{and} \quad \Delta_3 = \frac{1}{N_3}, \tag{8.2}$$

respectively. A natural partitioning of Ω corresponds to the voxel locations in the study $\vec{S}$. The displacement field $\vec{u}$ is then approximated by its spatial samples given by the set

$$\mathbb{U} = \{\vec{u}_n = \vec{u}(\vec{X}_n) : 0 \leq n < N_1 N_2 N_3 = N\}. \tag{8.3}$$

8.3 Spatial-Linear-Elastic Prior

The model that is used to induce the spatial-linear-elastic prior starts with the continuous strain-energy potential for linear elasticity that is derived in Appendix A. Restated here, the continuous strain-energy potential for linear elasticity is given by

$$H(\vec{u}) = \sum_{i=1}^{3}\sum_{j=1}^{3} \int_\Omega \frac{\lambda}{2}\left(\frac{\partial u_i(\vec{X})}{\partial X_i}\right)\left(\frac{\partial u_j(\vec{X})}{\partial X_j}\right) + \frac{\mu}{4}\left(\frac{\partial u_i(\vec{X})}{\partial X_j} + \frac{\partial u_j(\vec{X})}{\partial X_i}\right)^2 d\vec{X}. \tag{8.4}$$

The spatially parameterized, linear-elastic, strain-energy potential is determined from Eq. (8.4) by discretizing the integrals and derivatives. Performing this discretization gives

$$(8.5)\quad H(\mathbb{U}) = \sum_{i,j=1}^{3} \sum_{n=0}^{N-1} \frac{\lambda}{2} \left(\frac{\mathbb{P}_i(u_{i,n}) - u_{i,n}}{\Delta_i} \right) \left(\frac{\mathbb{P}_j(u_{j,n}) - u_{j,n}}{\Delta_j} \right) + \frac{\mu}{4} \left(\frac{\mathbb{P}_j(u_{i,n}) - u_{i,n}}{\Delta_j} + \frac{\mathbb{P}_i(u_{j,n}) - u_{j,n}}{\Delta_i} \right)^2 ,$$

where

$$(8.6)\qquad \mathbb{P}_j(u_{i,n}) = u_i(\vec{x}_n + \Delta_j \vec{e}_j)$$

is the plus shift operator. The j^{th} unit vector $\vec{e}_j$ has all zero components except the j^{th} component which is one. It will also be convenient to define the minus shift operator as

$$(8.7)\qquad \mathbb{M}_j(u_{i,n}) = u_i(\vec{x}_n - \Delta_j \vec{e}_j).$$

Notice that the the parameters of Eq. (8.5) are just the elements of $\mathbb{U}$.

The parameterization of the displacement field $\vec{u}$ by $\mathbb{U}$ is very high-dimensional. Normally there is a one-to-one correspondence between grid points and voxels in the study. This means that the displacement field is parameterized by three times the number of voxels. The advantage of this high-dimensional approximation is that it can accommodate very local shape differences between the textbook and the study.

8.4 Minimum Mean-Squared-Error Estimation

This section presents a second algorithm for individualizing the deformable textbook using the linear elasticity model. This algorithm is used to accommodate local nonrigid shape differences between the textbook and the study. The potential of the posterior for the spatially sampled linear-elasticity model is generated by combining the Eqs. (4.3) and (8.5). Thus, the posterior is a measure on $\mathbb{U}$ and has Gibbs' form $\frac{1}{Z}e^{-H(\mathbb{U}|S)}$ with potential

$$
H(\mathbb{U}|S) = \frac{1}{\sigma^2}\sum_{n=1}^{N}|T(\vec{X}_n - \vec{u}_n) - S(\vec{X}_n)|^2 + \sum_{i,j=1}^{3}\sum_{n=1}^{N}\frac{\lambda}{2}\left(\frac{\mathbb{P}_i(u_{i,n}) - u_{i,n}}{\Delta_i}\right)\left(\frac{\mathbb{P}_j(u_{j,n}) - u_{j,n}}{\Delta_j}\right) + \frac{\mu}{4}\left(\frac{\mathbb{P}_j(u_{i,n}) - u_{i,n}}{\Delta_j} + \frac{\mathbb{P}_i(u_{j,n}) - u_{j,n}}{\Delta_i}\right)^2 \tag{8.8}
$$

where $\mathbb{P}_i$ is defined in Eq. (8.6).

The vector of displacement vectors $\mathbb{U}$ are estimated by drawing samples from the posterior distribution $p(\mathbb{U}|S) \propto \frac{1}{Z}e^{-H(\mathbb{U}|S)}$ where $H(\vec{u}|S)$ is given in Eq. (8.5) using the Stochastic Differential Equation (SDE)

$$
d\mathbb{U}_t^{(i,n)} = -\frac{1}{2}\frac{\partial H(\mathbb{U}_t|S)}{\partial u_{1,n}}dt + dW_t^{(i,n)} \quad 1 \le i \le 3,\ 1 \le n \le N, \tag{8.9}
$$

where W_t is a standard $3 \times N$-valued Brownian motion with independent components. These samples are then averaged in order to compute the conditional mean of the posterior.

The drift matrix $\partial H(\mathbb{U}_t|S)/\partial u_{1,n}$ of Eq. (8.9) is determined as follows. First, the partial derivatives of $H(\mathbb{U}|S)$ are computed with respect to $u_{1,m} = [\vec{u}_m]_1$, that is,

$$
\begin{aligned}
\frac{\partial H(\mathbb{U}|S)}{\partial u_{1,m}} &= -\frac{1}{\sigma^2}\Big(T(\vec{X}_m - \vec{u}_m) - S(\vec{X}_m)\Big)\frac{\partial T(\vec{X}_m - \vec{u}_m)}{\partial x_1} \\
&+ \mu\sum_{j=1}^{3}\frac{2u_{1,m} - \mathbb{P}_j(u_{1,m}) - \mathbb{M}_j(u_{1,m})}{\Delta_j^2} \\
&+ \lambda\sum_{j=1}^{3}\frac{u_{j,m} - \mathbb{P}_j(u_{j,m}) + \mathbb{M}_1(\mathbb{P}_j(u_{j,m})) - \mathbb{M}_1(u_{j,m})}{\Delta_1\Delta_j} \\
&+ \mu\sum_{j=1}^{3}\frac{u_{j,m} - \mathbb{P}_1(u_{j,m}) + \mathbb{M}_j(\mathbb{P}_1(u_{j,m})) - \mathbb{M}_j(u_{j,m})}{\Delta_1\Delta_j}.
\end{aligned} \tag{8.10}
$$

Next, Eq. (8.10) is made more accurate by replacing the terms

$$
\frac{u_{j,m} - \mathbb{P}_j(u_{j,m}) + \mathbb{M}_1(\mathbb{P}_j(u_{j,m})) - \mathbb{M}_1(u_{j,m})}{\Delta_1\Delta_j} \tag{8.11}
$$

and

$$\frac{u_{j,m} - \mathbb{P}_1(u_{j,m}) + \mathbb{M}_j(\mathbb{P}_1(u_{j,m})) - \mathbb{M}_j(u_{j,m})}{\Delta_1 \Delta_j} \tag{8.12}$$

with

$$\frac{\mathbb{P}_1(\mathbb{P}_j(u_{j,m}) - \mathbb{M}_j(u_{j,m})) - \mathbb{M}_1(\mathbb{P}_j(u_{j,m}) - \mathbb{M}_j(u_{j,m}))}{4\Delta_1 \Delta_j}, \quad j = 2, 3. \tag{8.13}$$

because the terms in Eqs. (8.11) and (8.12) are $o(\Delta_1 \Delta_j)$ discrete approximations of the $\frac{\partial^2 u_j}{\partial X_1 \partial X_j}$ while Eq. (8.13) is an $o(\Delta_1^2 \Delta_j^2)$ approximation for $j = 1, 2$. Rewriting Eq. (8.10) making the above substitutions gives

$$\begin{aligned} \frac{\partial H(\mathbb{U}|S)}{\partial u_{1,m}} &= -\frac{1}{\sigma^2}\Big(T(\vec{X}_m - \vec{u}_m) - S(\vec{X}_m)\Big)\frac{\partial T(\vec{X}_m - \vec{u}_m)}{\partial x_1} \\ &+ \mu \sum_{j=1}^{3} \frac{2u_{1,m} - \mathbb{P}_j(u_{1,m}) - \mathbb{M}_j(u_{1,m})}{\Delta_j^2} + (\lambda + \mu)\frac{2u_{1,m} - \mathbb{P}_1(u_{1,m}) - \mathbb{M}_1(u_{1,m})}{\Delta_1^2} \\ &+ (\lambda + \mu)\sum_{j=2}^{3} + \frac{\mathbb{P}_1(\mathbb{P}_j(u_{j,m}) - \mathbb{M}_j(u_{j,m})) - \mathbb{M}_1(\mathbb{P}_j(u_{j,m}) - \mathbb{M}_j(u_{j,m}))}{4\Delta_1 \Delta_j}. \end{aligned} \tag{8.14}$$

Next it is noticed that Eq. (8.14) is nothing other than the first component of the discretized version of the linear-elastic PDE, i.e.,

$$\frac{\partial H(\mathbb{U}|S)}{\partial u_{1,m}} = [\mu \nabla^2 \vec{u}(\vec{X}_m) + (\mu + \lambda)\vec{\nabla}(\vec{\nabla}^T \vec{u}(\vec{X}_m)) + \vec{b}(\vec{X}_m, \vec{u}(\vec{X}_m))]_1 \tag{8.15}$$

where

$$\vec{b}_i(\vec{X}_m, \vec{u}_m) = -\frac{1}{\sigma^2}\Big(T(\vec{X}_m - \vec{u}_m) - S(\vec{X}_m)\Big)\frac{\partial T(\vec{X}_m - \vec{u}_m)}{\partial x_i} \tag{8.16}$$

for $1 \le i \le 3$. Thus, the general formula for the drift matrix in Eq. (8.9) is given by

$$\frac{\partial H(\mathbb{U}|S)}{\partial u_{i,m}} = [\mu \nabla^2 \vec{u}(\vec{X}_m) + (\mu + \lambda)\vec{\nabla}(\vec{\nabla}^T \vec{u}(\vec{X}_m)) + \vec{b}(\vec{X}_m, \vec{u}_m)]_i \tag{8.17}$$

where the derivatives on the right side are interpreted as symmetric difference approximations.

9. Inducing a Fluid Prior Via PDE's

This chapter presents a method for individualizing the anatomical textbook based on viscous-fluid transformations. This approach follows that of Rabbitt [47]. The advantage of this method is that anatomical variation can be accommodated which require large deformations while maintaining continuity and smooth deformation of the template.

9.1 Motivation

The development of the viscous-fluid-continuum model presented herein was motivated by the following. (1) The complexity of the patterns which must be represented (ventricles, cortical folds, etc.) require high-dimensional transformations. These high-dimensional transformations require regularization (priors) to ensure that they transform the textbook continuously and smoothly. (2) The fact that large deformations must and will be accommodated, implies that models with small-deformation assumptions (upon which linear models and quadratic penalties are based) cannot be use solely. Examples of models with small-deformation assumptions are the membrane or Laplacian model [29, 38], the thin-plate or biharmonic model [38, 39, 40], and linear elasticity model (see Appendix A and [41, 42]).

9.2 Bayesian variational problem.

The approach herein is to derive partial differential equations (PDE's) which the transformation fields must satisfy. The PDE's are crafted using principles obtained directly from continuum mechanics for deformable bodies. The deformation process does not attempt to model the physical process of brain growth, but instead uses the mechanical analogies to enforce topological properties that adhere to large deformation kinematics. Towards this end, the classical conservation equations of mechanics have been modified to account for non-mass-conserving deformations. This allows for the growth and shrinkage of regions within the continuum. When combined with the various constitutive laws—for instance, a viscous-fluid model as in this chapter—we are provided with a mathematical framework embueing our deformations with smoothness properties of which we have complete control.

A crucial part of the specification of the PDE is the force which drives the deformation of the template to match the configuration of the data. This force is induced by adopting the Grenander/Bayesian view of the group transformations in which there exists a prior distribution on the transformation. The solution is generated as samples from the posterior distribution relating the transformation fields to the data as collected by the sensor. The posterior is taken in Gibbs' form with the total potential associated with the sum of the statistical model of the sensor and the prior distribution. The PDE formulation is then the variational maximizer (minimizer) of the Gibbs' potential. The forcing function in the PDE is then none other then the variation of the sensor statistical likelihood function with respect to the vector displacement field.

This view is based on the existence of a variational principle from which the field equations can be induced: such principles are well known for conservative systems [48, 49] (for linear elasticity for example where stress is only a function of instantaneous strain) and link the PDE formulation as the minimizer (maximizer) of some potential. Variational formulations have also been developed for numerous nonconservative systems including the Newtonial fluid model addressed herein [50, 51]. We rely upon existence of this variational formulation for interpretation of the method (see section 11, however for brevity the field equations are constructed via the Newtonian equations of motion to have physical and topological properties which we demand of the resulting transformations. The body force appearing in the PDE is found directly from the variation of the Gibbs' statistical model on the sensor data with respect to the transformation field as has been done previously [29, 52, 41] for conservative systems. The variational solution represented by the PDE then corresponds to a measure on space-time (see [50]).

9.3 Eulerian Velocity Field

As before, we use an Eulerian reference frame to describe the deformation of the textbook. But now we need to describe not only the displacement of the continuum, but also its velocity over time. The Eulerian velocity field $\vec{v}$ is related to the Eulerian displacement field by the equation

$$\vec{v}(\vec{X},t) = \frac{\partial \vec{u}(\vec{X},t)}{\partial t} + \sum_{k=1}^{3} v_k(\vec{X},t)\frac{\partial \vec{u}(\vec{X},t)}{\partial X_k}. \tag{9.1}$$

To derive Eq. (9.1), we first assume that there exists a sufficiently differentiable and single-valued function $\vec{X}(\vec{x},t)$ defining the motion of a particle, even though it may not be known. This function allows us to write the Eulerian displacement field in terms of its initial and current position as

$$\vec{u}(\vec{X}(\vec{x},t),t) = \vec{X}(\vec{x},t) - \vec{x}. \tag{9.2}$$

In addition, the Lagrangian velocity of a particle originally located at $\vec{x}$ at time t is given by

$$\vec{V}(\vec{x},t) = \frac{\partial \vec{X}(\vec{x},t)}{\partial t}. \tag{9.3}$$

The Eulerian velocity $\vec{v}(\vec{X},t)$ is related to the Lagrangian velocity by

$$\vec{v}(\vec{X}(\vec{x},t),t) = \vec{V}(\vec{x},t). \tag{9.4}$$

Combining Eqs. (9.2 – 9.4) gives

$$\vec{v}(\vec{X}(\vec{x},t),t) = \frac{\partial \vec{X}(\vec{x},t)}{\partial t} = \frac{d}{dt}(\vec{x} + \vec{u}(\vec{X}(\vec{x},t),t)) = \frac{d\vec{u}(\vec{X}(\vec{x},t),t)}{dt}. \tag{9.5}$$

Using the chain rule for differentiation, we see that

$$v_m(\vec{X},t) = \left.\frac{\partial u_m}{\partial t}\right|_{(\vec{X},t)} + \sum_{k=1}^{3} \left.\frac{\partial u_m}{\partial X_k}\right|_{(\vec{X},t)} \left.\frac{\partial X_k}{\partial t}\right|_{(\vec{x},t)} \qquad 1 \le m \le 3 \tag{9.6}$$

where $\vec{X} = \vec{X}(\vec{x},t)$. Rewriting this equation with the aid of Eqs. (9.3) and (9.4) gives the desired result

$$v_m(\vec{X},t) = \left.\frac{\partial u_m}{\partial t}\right|_{(\vec{X},t)} + \sum_{k=1}^{3} v_k(\vec{X},t)\left.\frac{\partial u_m}{\partial X_k}\right|_{(\vec{X},t)} \qquad 1 \le m \le 3. \tag{9.7}$$

Notice that each term in this last equation is a function of $\vec{X}$ and can be computed without the knowledge of the function $\vec{X}(\vec{x},t)$. The operator

$$\frac{d}{dt} = \frac{\partial}{\partial t} + \sum_{k=1}^{3} v_k(\vec{X},t)\frac{\partial}{\partial X_k} \tag{9.8}$$

is called the material derivative operator and it provides the time rate of change experienced by an element of material instantaneously at $\vec{X}$ at time t.

The term $\sum_{k=1}^{3} v_k(\vec{X},t)\frac{\partial \vec{u}(\vec{X},t)}{\partial X_k}$ in the material derivative Eq. (9.1) accounts for the kinematic nonlinearities of the displacement field $\vec{u}$. Without this term, particles tracked in a Eulerian frame cannot have curved trajectories over time because there is only a linear relationship between $\vec{u}$ and $\vec{v}$. The importance of this nonlinear term is illustrated in the following 2D example.

Example: Assume that particles of a 2D fluid have velocity $\vec{v}(\vec{X},t) = [-X_2 \;\; X_1]^T$ and initial displacement $\vec{u}(\vec{X},0) = [0 \;\; 0]^T$. This velocity field corresponds to the particles of a fluid rotating around the origin over time. Solving (9.1) for $\vec{u}(\vec{X},t)$ with the nonlinear term gives

$$\vec{u}(\vec{X},t) = \begin{bmatrix} X_1 - (X_1 \cos t + X_2 \sin t) \\ X_2 - (-X_1 \sin t + X_2 \cos t) \end{bmatrix} \tag{9.9}$$

and without the nonlinear term gives

$$\vec{u}(\vec{X},t) = \begin{bmatrix} -X_2 t \\ X_1 t \end{bmatrix} . \tag{9.10}$$

Notice that $\vec{u}(\vec{X},t)$ in (9.9) tracks the fluid particles as they travel in circular paths about the origin over time. In sharp contrast, the $\vec{u}(\vec{X},t)$ of (9.10) can track only particles moving with linear trajectories and cannot follow circular paths.

9.4 Formulation of the Viscous-Fluid PDE

This section describes how to formulate a PDE which governs the behavior of a material by combining the fundamental conservation laws of mechanics with a constitutive equation. The

PDE of most interest is that of a viscous fluid given by

$$\mu\nabla^2\vec{v} + (\lambda+\mu)\vec{\nabla}(\nabla\cdot\vec{v}) + \vec{b} = 0 \tag{9.11}$$

where the velocity $\vec{v}$ is related to the displacement $\vec{u}$ by Eq. (9.1). The body force $\vec{b}$ provides the connection between the sensor data and the deforming continuum.

9.4.1 Fundamental Conservation Laws of Mechanics

In the field of continuum mechanics it is generally assumed that a control volume can be selected such that mass within the volume is conserved. For deformation of image templates, however, it is often desirable to allow the addition of mass (growth) or the subtraction of mass (shrinkage) on a local level. In order to allow for this, Dr. R.D. Rabbitt reformulated the basic laws to account for a mass "creation" term. The following basic laws are derived as in Malvern [53] modified to include a spatially dependent mass source.

Conservation of mass

Consider a fixed volume Ω bounded by surface $\partial\Omega$. A source of mass supplies or extracts mass from the volume at a rate of η per unit volume. By conservation of mass the rate of mass change within the volume is equal to the rate of mass source minus the flux of mass through the surface. This is written as

$$\iiint_\Omega \frac{\partial\rho}{\partial t}dV = \iiint_\Omega \eta dV - \iint_{\partial\Omega} \rho\vec{v}\cdot\vec{n}ds, \tag{9.12}$$

where $\vec{n}$ is the unit outward normal to $\partial\Omega$ and ρ is the density of the material. Using the divergence theorem and noting that the volume integral is arbitrary we obtain the differential form of the mass conservation law with a mass source term:

$$\frac{\partial\rho}{\partial t} + \vec{\nabla}\cdot(\rho\vec{v}) = \eta. \tag{9.13}$$

This equation is expressed in terms of the material derivative by

$$\frac{d\rho}{dt} + \rho\vec{\nabla}\cdot\vec{v} = \eta, \tag{9.14}$$

where d/dt is defined in Eq. (9.8).

Reynolds transport theorem with a mass source

In determining the equations describing conservation of momentum it will be necessary to determine the rate of change of a field variable associated with the material. The rate of change of the product $\psi\rho$ associated with the material instantaneously in a fixed volume Ω is equal to the rate of increase of $\psi\rho$ inside Ω plus the rate outward flux of $\psi\rho$ carried by mass transport through the volume surface $\partial\Omega$. This is written as

$$\frac{d}{dt}\iiint_{\Omega}\rho\psi dV = \iiint_{\Omega}\frac{\partial\rho\psi}{\partial t}dV + \iint_{\partial\Omega}\rho\psi\vec{v}\cdot\vec{n}ds. \tag{9.15}$$

which can be rewritten as

$$\frac{d}{dt}\iiint_{\Omega}\rho\psi dV = \iiint_{\Omega}(\rho[\frac{\partial\psi}{\partial t} + \vec{v}\cdot\vec{\nabla}\psi] + \psi[\frac{\partial\rho}{\partial t} + \vec{\nabla}\cdot(\rho\vec{v})])dV, \tag{9.16}$$

using the divergence theorem. Applying the continuity equation with a mass source term in Eq. (9.13) to this last expression gives

$$\frac{d}{dt}\iiint_{\Omega}\rho\psi dV = \iiint_{\Omega}(\rho[\frac{\partial\psi}{\partial t} + \vec{v}\cdot\vec{\nabla}\psi] + \psi\eta)dV, \tag{9.17}$$

which is a generalized version of Reynolds Transport Theorem including a mass source term η.

Conservation of momentum

Consider a fixed control volume Ω with boundary $\partial\Omega$. A distributed traction force $\vec{\tau}$ acts on the surface and a distributed body force $\vec{b}$ per unit volume acts on all material elements within the body. Applying Newton's second law in the form of conservation of linear momentum yields

$$\frac{d}{dt}\iiint_{\Omega}\rho\vec{v}dV = \iint_{\partial\Omega}\vec{\tau}dA + \iiint_{\Omega}\vec{b}dV. \tag{9.18}$$

The traction force $\vec{\tau}$ is related to the Cauchy stress tensor $\mathbb{T}$ by $\vec{\tau} = \mathbb{T}\cdot\vec{n}$ where $\vec{n}$ is the outward unit normal from the surface. Using this and the divergence theorem provides

$$\frac{d}{dt}\iiint_{\Omega}\rho\vec{v}dV = \iiint_{\Omega}(\vec{\nabla}\cdot\mathbb{T}+\vec{b})dV. \tag{9.19}$$

Using the generalized Reynolds transport theorem including a mass source yields

$$\iiint_{\Omega}(\rho\frac{d\vec{v}}{dt}+\vec{v}\eta-\vec{\nabla}\cdot\mathbb{T}-\vec{b})dV = 0. \tag{9.20}$$

This is valid for all choices of control volumes so conservation of momentum with a mass source becomes

$$\rho\frac{d\vec{v}}{dt}+\vec{v}\eta-\vec{\nabla}\cdot\mathbb{T}-\vec{b} = 0. \tag{9.21}$$

Eq. (9.21) is called the momentum equation and provides the relationship between applied body forces and the resulting material deformation. The body force $\vec{b}$ provides the connection between the sensor data (study) and the deforming continua (textbook) and is discussed in the next section.

In order to complete the PDE formulation, we need to specify the constitutive behavior of the material by specifying the form of the stress tensor $\mathbb{T}$. Specific examples for a Newtonian viscous fluid and a linear elastic solid are provided in the next section.

9.4.2 The Viscous-Fluid Constitutive Law

We choose a constitutive law which allows for large, nonlinear deformations while maintaining a continuous homeomorphic map with smooth deformations of the template. This is an alternative constitutive law to Hooke's law used in linear elasticity, one corresponding to a viscous-fluid with the property that the stress within the deformed configuration is allowed to relax with time. This implies that the mechanical energy (penalty) does not necessarily increase with the magnitude of the pointwise deformation. A Navier-Poisson Newtonian fluid model is used for which the Cauchy stress tensor $\mathbb{T}$ is related to the rate of deformation tensor D by $\mathbb{T} = (\lambda(trD) - p)\vec{I} + 2\mu D$, where μ and λ are the viscosity constants, p is the pressure, and the rate of deformation tensor is $D = \frac{1}{2}(\vec{\nabla}\vec{v} + (\vec{\nabla}\vec{v})^T)$. Making these substitutions in Eq. (9.21) gives

$$\mu\nabla^2\vec{v} + (\lambda+\mu)\vec{\nabla}(\vec{\nabla}\cdot\vec{v}) + \vec{b}(\vec{u}) = \vec{\nabla}p + \rho\frac{d\vec{v}}{dt} + \vec{v}\eta \ . \tag{9.22}$$

A simplified model is obtained for very-low-Reynold's-number flow where the pressure gradient $\vec{n}ablap$ and the inertia terms $\rho\frac{d\vec{v}}{dt} + \vec{v}\eta$ are neglected, i.e., Eq. (9.21) becomes

$$\mu\nabla^2\vec{v} + (\lambda+\mu)\vec{\nabla}(\vec{\nabla}\cdot\vec{v}) + \vec{b}(\vec{u}) = \vec{0} \ , \tag{9.23}$$

where $\nabla^2 = \nabla^T\nabla$ is the Laplacian operator.

The template is deformed by solving Eq. (9.23) subject to a body force $\vec{b}$ that drives the template viscously into registration with the study. Two mechanisms are clearly identified. The first, represented by $\mu\nabla^2\vec{v}$, is associated with constant-volume viscous flow of the template. The second, represented by $(\lambda+\mu)\vec{\nabla}(\vec{\nabla}\cdot\vec{v})$, is non-zero when regions of the template grow or dissolve. Using continuity, (Eq. 9.14, Appendix A) we note that $\vec{\nabla}\cdot\vec{v} = \frac{\eta}{\rho} - \frac{d\rho}{dt}$, so the second term in Eq. (9.23) is related to dilatation caused by the mass injection per unit volume η and the rate of change of density $\frac{d\rho}{dt}$. In the present work we are specifically interested in allowing growth and shrinkage of local regions of the template and hence $\vec{\nabla}\cdot\vec{v}$ in general is non-zero. The viscosity coefficient λ appears in the fluid constitutive law multiplying the rate of volume dilatation $tr(D)$ and is adjusted to control the rate of growth or shrinkage of local regions within the deforming template. The viscosity coefficient μ is responsible for generation of shear stress in opposition to the rate of deformation D and is adjusted to control shearing between adjacent regions of the template.

9.4.3 The Body Force

Choosing the body force appearing in the PDE associated with the momentum Eq. (9.21) follows directly from the Bayesian view of the PDE as the variational minimizer of the posterior distribution. Then the body force is given directly from the potential term in the posterior distribution associated with the sensor, linking the data to the template. The body force becomes the gradient of the total Lagrangian, call it C, associated with the log-likelihood of the data source. This is consistent with our previous work as described in [29, 52, 41]. For a Gaussian sensor model [1], the cost function is

$$C(T(\vec{X}), S(\vec{X}), \vec{u}) = \frac{1}{2}\int_{\Omega} |T(\vec{X} - \vec{u}(\vec{X}, t)) - S(\vec{X})|^2 d\vec{X}, \tag{9.24}$$

yielding the body force

$$\vec{b}(\vec{X}, \vec{u}(\vec{X}, t)) = -\Big(T(\vec{X} - \vec{u}(\vec{X}, t)) - S(\vec{X})\Big)\vec{\nabla}T\Big|_{\vec{X}-\vec{u}(\vec{X},t)} . \tag{9.25}$$

9.5 Numerical Solution

The solution of the viscous-fluid PDE

$$\mu\nabla^2\vec{v} + (\lambda + \mu)\vec{\nabla}(\vec{\nabla}\cdot\vec{v}) + \vec{b}(\vec{u}) = 0 \ , \tag{9.26}$$

$$\frac{\partial \vec{u}}{\partial t} = \vec{v} - \sum_{i=1}^{3} v_i \frac{\partial \vec{u}}{\partial X_i} \ , \tag{9.27}$$

$$\vec{b}(\vec{X}, \vec{u}(\vec{X}, t)) = -\Big(T(\vec{X} - \vec{u}(\vec{X}, t)) - S(\vec{X})\Big)\Big(\vec{\nabla}T\Big|_{\vec{X}-\vec{u}(\vec{X},t)}\Big) \tag{9.28}$$

includes nonlinearities associated with the body force and kinematic nonlinearities associated with the material derivative. In order to solve these nonlinear equations, we use a classical method

[1] For MRI this appears to be reasonably appropriate [54, 55, 56]. Although for PET image reconstruction, Poisson models [57] and their associated variation would be more appropriate for the forcing function.

of solution via "real time" integration. We decompose the non-linear PDE into a set of linear equations on the velocity field $\vec{v}(\vec{X},t)$ for each instance of time t, and then forward integrate using the material derivative non-linear relationship between the u and v fields.

Algorithm 1:

1. *Initialize* $t = 0$ *and* $\vec{u}(\vec{X},0) = 0$.

2. *Using Eq. (9.28), calculate the body force* $\vec{b}(\vec{X},\vec{u}(\vec{X},t))$ *given* $\vec{u}(\vec{X},t)$.

3. *If* $\vec{b}(\vec{X},\vec{u}(\vec{X},t))$ *is below a threshold for all* $\vec{X} \in \Omega$, *then STOP.*

4. *Solve the linear PDE Eq. (9.26) for the instantaneous velocity* $\vec{v}(\vec{X},t)$ *at time* t *assuming that* $\vec{v}(\vec{X},t) = 0, \vec{X} \in \partial\Omega$.

5. *Perform real-time integration by calculating* $\vec{u}(\vec{X},t+\Delta)$ *from the discretized version of the material derivative Eq. (9.27) given by*

$$\vec{u}(\vec{X},t+\Delta) = \vec{u}(\vec{X},t) + \Delta\vec{v}(\vec{X},t) - \Delta\sum_{i=1}^{3} v_i(\vec{X},t)\frac{\partial\vec{u}(\vec{X},t)}{\partial X_i} \ . \tag{9.29}$$

6. *Let* $t = t + \Delta$ *and go to step 2.*

Notice in step 4 that the nonlinearity introduced by the body force term in Eq. (9.26) is handled by fixing $\vec{b}(\vec{X},\vec{u}(\vec{X},t))$ at time t. Since $\vec{b}(\vec{X},\vec{u}(\vec{X},t))$ is fixed for time t, solving for the the instantaneous velocity is a linear problem. Also note in step 4 that the assumption that the velocity is zero on the boundary of Ω implies that the displacement field is also zero on the boundary of Ω. This assumption ensures that the boundary of the deformed domain maps back onto the boundary of the original domain Ω through the identity mapping.

9.5.1 Solution on a Finite Spatial Grid

The viscous-fluid PDE is solved on a discrete lattice associated with Ω. For large curved deformations, the transformation evaluated at these spatial grid points becomes singular over time even though the transformation evaluated on the continuum would not. To circumvent this problem, a method of regridding the template is used by generating a new, or propagated template, whenever the discretized transformation approaches local singularity. This is accomplished as follows.

The discretized displacement field is propagated through time using the procedure in the previous section until the Jacobian of the transformation indicates that the transformation is approaching singularity. When the magnitude of the Jacobian drops below a threshold (0.5 for the results shown later), the computation is stopped and a new propagated template is generated equal to the deformed template at the previous instant. The algorithm is restarted using the new template. When mapped back to the original template, this is equivalent to defining a new nonhomogeneous computation grid on the undeformed continuum. Initial conditions for the propagated template are set to match the final state of the previous template, i.e., the instantaneous velocity remains the same and the displacement field corresponding to the new template is set to zero. The total transformation is continuously tracked via the concatenation of the displacement fields associated with all of the propagated templates.

The method of propagating templates (i.e. regridding) is formalized in the following way. Let $T^{(j)}, j = 0, 1, 2\ldots$ denote the sequence of propagated templates with associated displacement field $\vec{U}^{(j)}$, so that

$$T^{(j)}(\vec{X} - \vec{U}^{(j)}(\vec{X}, t_i)) \tag{9.30}$$

is the deformation of the j^{th} template. Next, define t_{p_j} as the time that template $T^{(j-1)}$ is propagated to template $T^{(j)}$ such that $t_{p_0} = t_0 = 0 < t_{p_1} < \ldots < t_{p_M}$. Each propagated template $T^{(j)}$, except $T^{(0)}$ which is equal to the original template T, is defined as the deformation of the previous template $T^{(j-1)}$ at time t_{p_j}:

$$T^{(j)}(\vec{X}) \triangleq \begin{cases} T(\vec{X}), & j = 0, \\ T^{(j-1)}(\vec{X} - \vec{U}^{(j-1)}(\vec{X}, t_{p_j})), & 0 < j \leq M. \end{cases} \tag{9.31}$$

The time progression of displacement field $U^{(j)}$ becomes

(9.32)

$$\vec{U}^{(j)}(\vec{X}, t_i) \triangleq \begin{cases} 0, & i = p_j, \\ \vec{U}^{(j)}(\vec{X}, t_{i-1}) + (t_i - t_{i-1})(\vec{v}(\vec{X}, t_{i-1}) - \Gamma_{i-1}\vec{U}^{(j)}(\vec{X}, t_{i-1})), & p_j < i \leq p_{j+1}, \end{cases}$$

where $\Gamma_i \triangleq \sum_{j=1}^{3} v_j(\vec{X}, t_i) \frac{\partial}{\partial X_j}$ for $0 \le i \le N$. This definition ensures that each propagated displacement field satisfies the discrete approximation of the material derivative Eq. (9.1) and that $T^{(j)}(\vec{X} - \vec{U}^{(j)}(\vec{X}, t_{p_j})) = T^{(j)}(\vec{X})$ at propagation time t_{p_j}. The displacement field $\vec{u}$ associated with the original template T is defined in terms of the propagated displacement fields $\vec{U}^{(j)}$ by

$$\vec{u}(\vec{X}, t_i) \triangleq \begin{cases} \vec{U}^{(0)}(\vec{X}, t_i), & 0 \le i \le p_1, \\ \vec{U}^{(j)}(\vec{X}, t_i) + \vec{U}^{(j-1)}(\vec{X} - \vec{U}^{(j)}(\vec{X}, t_i), t_{p_j}), & p_j < i \le p_{j+1}. \end{cases} \tag{9.33}$$

Eq. (9.33) is a consistency condition ensuring

$$\begin{aligned} T^{(0)}(\vec{X} - \vec{U}^{(0)}(\vec{X}, t_i)) &= T(\vec{X} - \vec{U}^{(0)}(\vec{X}, t_i)) \\ &= T(\vec{X} - \vec{u}(\vec{X}, t_i)), \quad 0 \le i \le p_1 \end{aligned}$$

and that

$$\begin{aligned} T^{(j)}(\vec{X} - \vec{U}^{(j)}(\vec{X}, t_i)) &= T(\vec{X} - \vec{U}^{(j)}(\vec{X}, t_i) - \vec{U}^{(j-1)}(\vec{X} - \vec{U}^{(j)}(\vec{X}, t_i), t_{p_j})) \\ &= T(\vec{X} - \vec{u}(\vec{X}, t_i)), \qquad p_j < i \le p_{j+1}. \end{aligned}$$

Incorporating Eq's (9.31) - (9.33) into Algorithm 1 gives the following algorithm:

Algorithm 2:

1. *Let* $t = 0$, $i = 0$, $T^{(0)}(\vec{X}) = T(\vec{X})$, *and* $\vec{U}^{(0)}(\vec{X}, 0) = 0$.

2. *Calculate the body force*

$$\vec{b}(\vec{X}, \vec{U}^{(i)}(\vec{X}, t)) = -\left(T^{(i)}(\vec{X} - \vec{U}^{(i)}(\vec{X}, t)) - S(\vec{X})\right) \vec{\nabla} T^{(i)}\Big|_{\vec{X} - \vec{U}^{(i)}(\vec{X}, t)}$$

3. *If* $\vec{b}(\vec{X}, \vec{U}^{(i)}(\vec{X}, t))$ *is below a threshold for all* $\vec{X} \in \Omega$, *then STOP.*

4. *Solve the linear PDE Eq. (9.26) for the instantaneous velocity* $\vec{v}(\vec{X}, t)$ *at time* t *assuming that* $\vec{v}(\vec{X}, t) = 0, \vec{X} \in \partial\Omega$ *using Successive overrelaxation (SOR) (see Section 9.5.2).*

5. *Calculate the perturbation of the displacement field*

$$\vec{R}(\vec{X}) = \vec{v}(\vec{X},t) - \sum_{i=1}^{3} v_i(\vec{X},t) \frac{\partial \vec{U}^{(i)}(\vec{X},t))}{\partial X_i} \ .$$

6. *Choose a real-time step size* Δ *which is a function of* $\gamma = \max_{\vec{X}\in\Omega} ||\vec{R}(\vec{X})||$.

7. *If the Jacobian of* $\vec{X} - \vec{U}^{(i)}(\vec{X},t) - \Delta\vec{R}(\vec{X})$ *is less than* 0.5 *then propagate to template* $i+1$

   - $T^{(i+1)}(\vec{X}) = T^{(i)}(\vec{X} - \vec{U}^{(i)}(\vec{X},t))$, $\vec{U}^{(i+1)}(\vec{X},t) = 0$, *and* $i = i+1$,
   - *go to step 5.*

   else update the i^{th} *displacement field*

   - $\vec{U}^{(i)}(\vec{X},t+\Delta) = \vec{U}^{(i)}(\vec{X},t) + \Delta\vec{R}(\vec{X})$ *and* $t = t+\Delta$,
   - *go to step 2.*

In practice, the real-time step size in item 6 above is chosen so that the maximum displacement at each instant of real-time is between 0.3 and 1.0 pixels.

The next theorem states that under sufficiently smooth conditions on the transformation and velocity fields, the solution generated by Algorithm 2 solves the fluid PDE given by Eq's (9.26 - 9.28) as $t_{i+1} - t_i$ goes to zero for all $0 \le i < N$. This of course proves that the regridding method allows large curved deformations to be calculated on a discrete grid while maintaining a diffeomorphism (i.e. avoiding a singular transformation).

Theorem 9.1 *Let* $\vec{U}^{(j)} \in C^2(\Omega)$ *and* $\vec{v} \in C^1(\Omega)$ *where* $\Omega \subset \Re^3$ *then Algorithm 2 satisfies the fluid PDE given by Eq's (9.26 - 9.28) in the sense that for times* t_i *with* $t_0 = 0 < t_1 < \ldots < t_N < T_{max}$, *the following system of equations are satisfied:*

$$\mu\nabla^2\vec{v}(\vec{X},t_i) + (\lambda+\mu)\vec{\nabla}(\vec{\nabla}\cdot\vec{v}(\vec{X},t_i)) + \vec{b}(\vec{X},\vec{u}(\vec{X},t_i)) = 0, \tag{9.34}$$

$$\vec{u}(\vec{X},t_i) = \begin{cases} 0, & i=0 \\ \vec{u}(\vec{X},t_{i-1}) + (t_i - t_{i-1})\Big(\vec{v}(\vec{X},t_{i-1}) - \Gamma_{i-1}\vec{u}(\vec{X},t_{i-1})\Big) + o(t_i - t_{i-1}), & 0 < i. \end{cases} \tag{9.35}$$

where $\Gamma_i \triangleq \sum_{j=1}^{3} v_j(\vec{X},t_i)\frac{\partial}{\partial X_j}$ *for* $0 \le i \le N$.

Proof. Eq. (9.34) is satisfied by Algorithm 2 if the real-time integration Eq. (9.35) is satisfied. Thus, it is only necessary to show that the real-time integration of Algorithm 2 given by Eq's (9.32) and (9.33) implies Eq. (9.35).

We now show by induction that Eq's (9.32) and (9.33) satisfy Eq. (9.35).

Basis:

Substituting Eq. (9.33) into (9.32) gives Eq. (9.35) for times $t_0 = t_{p_0} \leq t_i \leq t_{p_1}$.

Induction:

Assume that Algorithm 2 satisfies Eq. (9.35) for time t_i and show that Eq. (9.35) is satisfied for t_{i+1}. Substituting Eq. (9.32) into (9.33) gives

$$\begin{aligned}\vec{u}(\vec{X}, t_{i+1}) = \vec{U}^{(j)}(\vec{X}, t_i) + (t_{i+1} - t_i)(\vec{v}(\vec{X}, t_i) - \Gamma_i \vec{U}^{(j)}(\vec{X}, t_i)) \\ + \vec{U}^{(j-1)}(\vec{X} - \vec{U}^{(j)}(\vec{X}, t_i) - (t_{i+1} - t_i)(\vec{v}(\vec{X}, t_i) - \Gamma_i \vec{U}^{(j)}(\vec{X}, t_i)), t_{p_j}) \ .\end{aligned}$$

Using the fact that $\vec{U}^{(j-1)} \in C^2(\Omega)$ we can write the Taylor series expansion of $\vec{U}^{(j-1)}$ about the point $\vec{X} - \vec{U}^{(j)}(\vec{X}, t_i)$ as

$$\begin{aligned}\vec{u}(\vec{X}, t_{i+1}) =& \vec{U}^{(j)}(\vec{X}, t_i) + (t_{i+1} - t_i)(\vec{v}(\vec{X}, t_i) - \Gamma_i \vec{U}^{(j)}(\vec{X}, t_i)) + \vec{U}^{(j-1)}(\vec{X} - \vec{U}^{(j)}(\vec{X}, t_i), t_{p_j}) \\ & - (t_{i+1} - t_i) \vec{\nabla} \vec{U}^{(j-1)}\Big|_{(\vec{X} - \vec{U}^{(j)}(\vec{X}, t_i), t_{p_j})} \cdot \Big(I - \vec{\nabla} \vec{U}^{(j)}\Big|_{(\vec{X}, t_i)}\Big) \vec{v}(\vec{X}, t_i) \\ & + o\Big((t_{i+1} - t_i) || \vec{v}(\vec{X}, t_i) - \Gamma_i \vec{U}^{(j)}(\vec{X}, t_i) ||\Big) \ ,\end{aligned}$$

where the (3×3) matrix $\vec{\nabla} \vec{u}$ is defined to have elements $[\vec{\nabla} \vec{u}]_{ij} = \frac{\partial u_i}{\partial X_j}$. Rearranging terms and using the fact that $\vec{v}$ and $\vec{\nabla} \vec{U}^{(j)}$ are bounded because they are continuous on the compact set Ω gives

$$\begin{aligned}\vec{u}(\vec{X}, t_{i+1}) =& (t_{i+1} - t_i) \Big(\vec{v}(\vec{X}, t_i) - \Gamma_i \Big[\vec{U}^{(j)}(\vec{X}, t_i) + \vec{U}^{(j-1)}(\vec{X} - \vec{U}^{(j)}(\vec{X}, t_i), t_{p_j})\Big]\Big) \\ & + \vec{U}^{(j)}(\vec{X}, t_i) + \vec{U}^{(j-1)}(\vec{X} - \vec{U}^{(j)}(\vec{X}, t_i), t_{p_j}) + o(t_{i+1} - t_i) \ .\end{aligned}$$

Substituting Eq. (9.33) into this last expression gives Eq. (9.35) for time t_{i+1}. $\square$

9.5.2 Successive Overrelaxation (SOR) Algorithm

Successive overrelaxation (SOR) with checker-board update [58] is used to solve (9.23) for the instantaneous velocity $\vec{v}(\vec{X}, t_i)$ of the displacement field at each grid point $\vec{X} \in \Omega$. The SOR formula used for solving (9.23) in 2D at time t with grid spacing Δ is given by

$$\begin{aligned}
v_{ij}^{(1)n+1} &= v_{ij}^{(1)n-1} + \omega\Big(\frac{\Delta^2}{6}(b_{ij}^{(1)} + 2v_{i+1,j}^{(1)n-1} + 2v_{i-1,j}^{(1)n-1} + v_{ij+1}^{(1)n-1} + v_{ij-1}^{(1)n-1} \\
&\quad + \frac{1}{4}(v_{i+1,j+1}^{(2)n} + v_{i-1,j-1}^{(2)n} - v_{i+1,j-1}^{(2)n} - v_{i-1,j+1}^{(2)n})) - v_{ij}^{(1)n}\Big) \\
v_{ij}^{(2)n+1} &= v_{ij}^{(2)n-1} + \omega\Big(\frac{\Delta^2}{6}(b_{ij}^{(2)} + 2v_{ij+1}^{(2)n-1} + 2v_{ij-1}^{(2)n-1} + v_{i+1,j}^{(2)n-1} + v_{i-1,j}^{(2)n-1} \\
&\quad + \frac{1}{4}(v_{i+1,j+1}^{(1)n} + v_{i-1,j-1}^{(1)n} - v_{i+1,j-1}^{(1)n} - v_{i-1,j+1}^{(1)n})) - v_{ij}^{(2)n}\Big)
\end{aligned}$$

where $\mu = 1$ and $\lambda = 0$ to simplify the notation. The subscripts ij denote the ij^{th} pixel location in the study, the superscript n denotes the relaxation iteration, and the superscript (m) denotes the m^{th} component of a vector. Extending Eq. (9.36) to 3D is straightforward. All of the points with index ijk whose sum $i + j + k$ is even are updated first and all of the points with $i + j + k$ odd are updated second. In addition, there are three equations, one for each component of $\vec{v}$, to evaluate instead of two.

10. Results

10.1 Two-Dimensional Elasticity Experiments

10.1.1 2D Elastic MR-to-MR Experiment

A 2-dimensional textbook was constructed from an MR study collected at Duke University of a normal patient using standard magnetic resonance spin-echo sequences to generate a spin-density, t1-weighted, and t2-weighted series. These 2-D scans were hand segmented into three gray matter nuclei (thalamus, putamen, and head of caudate nucleus), ventricles, other brain matter, and background. The hand segmentation was generated using all three MR images, along with the detailed horizontal brain section on page 28 of the DeArmond et al. [59] anatomy atlas. This textbook is shown in Fig. 10.1 and consists of the three magnetic-resonance images, and the hand labeled segmentation. We shall refer to this textbook as textbook W.

Two different patient studies, A and B, were analyzed. Study A (shown in the top row of Fig. 10.2) consisted of $N = 3$ sub-studies (spin-density, t1-weighted, and t2-weighted MR images) and study B (shown in the left column of Fig. 10.3) consisted of $N = 2$ sub-studies (spin-density

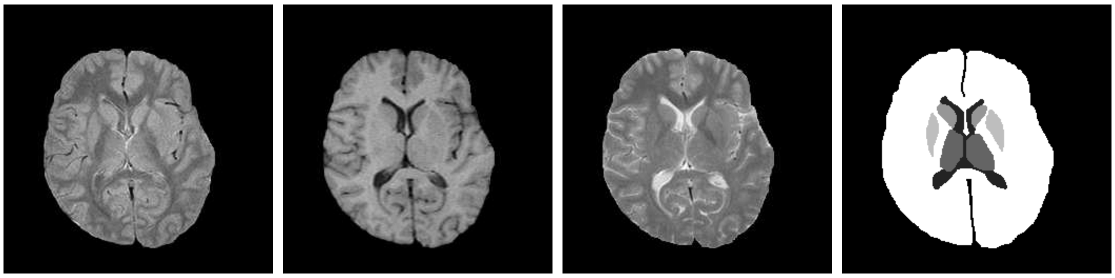

Figure 10.1: Textbook components: spin-density, t1-weighted, t2-weighted MR images, and a hand segmentation.

and t2-weighted MR images). These studies were selected because of their similar orientations and level in the brain and not for their closeness in brain size or similarity to the textbook.

The linear elastic transformation of textbook W into study A is shown in the middle row of Fig. 10.2. This transformation was estimated using all three modalities of the study. Each modality contributed different information that aided in the estimation of the transformation: the spin-density gave good contrast between the white and gray matter, and the t1 and t2 images gave good contrast between the CSF and other brain matter. The likelihood model that was used was a weighted sum of three equally weighted Gaussian likelihood models—one for each of the three modalities. Notice the remarkable correspondence between the transformed textbook (middle row) and the study (top row). In particular, notice the good correspondence between the ventricles and also, the outer contour of the sections. The bottom row shows the magnitude of the difference images between the transformed textbook and the sub-studies (top row minus the middle row) demonstrating the near perfect alignment.

Fig. 10.3 shows the linear elastic transformation of textbook W into study B in the middle column and the difference images in the right column. This transformation was estimated using both modalities of the study and the likelihood was the sum of two equally weighted Gaussian models. Again notice the good correspondence between the study and the deformed textbook.

The linear elastic transformations $\hat{h}_A$ and $\hat{h}_B$—of textbook W into studies A and B—were estimated using a multiresolution approach consisting of two steps. In the first step, the global shape differences between the textbook and the study were accommodated using the elastic-basis model described in Chapt. 7. This global transformation was then refined using the spatial-elastic model described in Chapt. 8 in order to accommodate the local shape differences.

Fig. 10.4 shows the t-2 weighted MR difference image for the linear-elastic deformation of W into A initially (left), after the basis transformation (middle), and after the spatial transformation (right). This middle panel demonstrates that the large global shape disparities were accommodated using the linear-elastic-basis model, and the right panel demonstrates that the local shape disparities were accommodated using the spatial linear elasticity model.

The bottom row of Fig. 10.5 shows the automatic segmentations of studies A (left) and B (right). These segmentations were generated by applying $\hat{h}_A$ and $\hat{h}_B$ to the gray matter nuclei and ventricle segmentation of textbook W. In order to judge the performance of the automatic

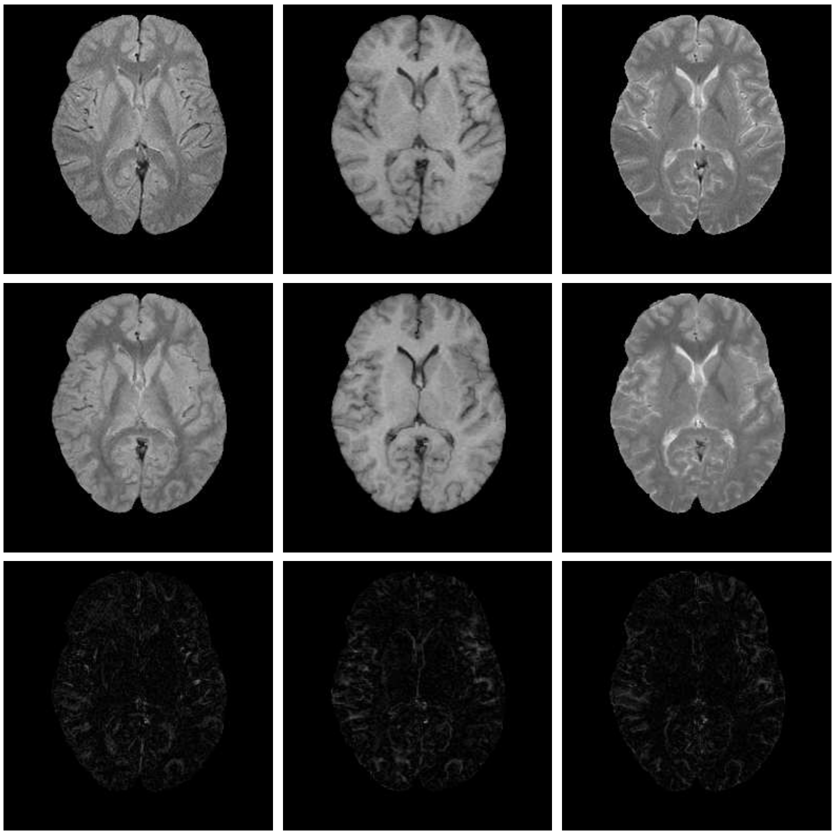

Figure 10.2: Top row shows the spin-density (left), t1-weighted (middle) and t2-weighted (right) images of patient A; middle row shows the transformed textbook; bottom row shows the magnitude difference images between the transformed textbook and the patient.

segmentations, the top row shows a hand labeling of the structures of studies A (left) and B (right). Comparing these images visually, we see that the automatic segmentations although not perfect are pretty good. Differences between the segmentations are due to errors in the original textbook segmentation, limitations of the algorithm, and errors in the study segmentations.

One thing in particular that should be noticed in the automatic segmentations in Fig. 10.5 is the fragmentation of the structures. This fragmentation occurred at places where the Jacobian of the transformation is singular (see Fig. 10.6) and corresponds to the fact that the transformation did not preserve the topology of the textbook. Fig. 10.6 shows the regions where the Jacobian of the transformations $\hat{h}_A$ (left) and $\hat{h}_B$ (right) are negative in white and positive in black. A

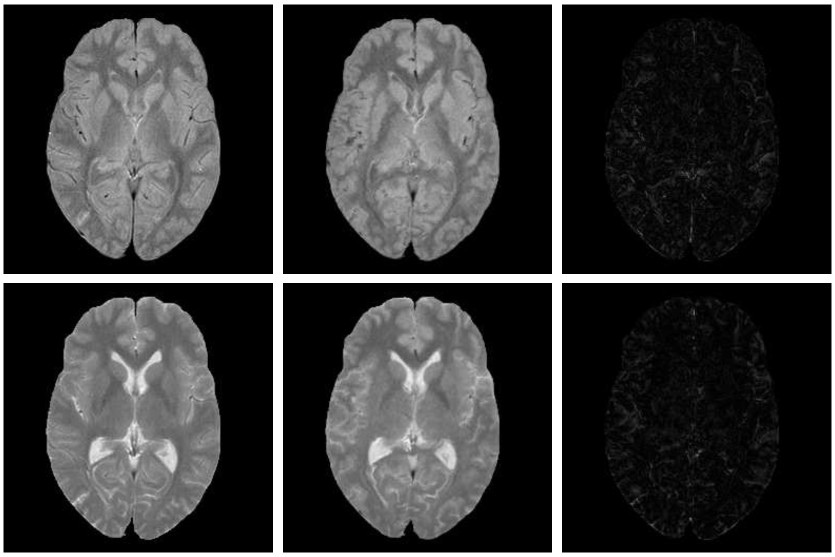

Figure 10.3: Left column shows the spin-density (top) and t2 (bottom) of patient B; middle column shows the transformed textbook; right column shows the magnitude difference images between the textbook and the patient.

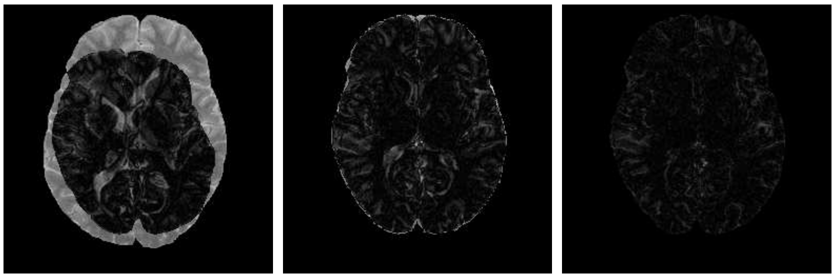

Figure 10.4: Magnitude difference images between the t2 component of patient A and the textbook (left), the globally deformed textbook (middle), and the locally deformed textbook (right).

transformation is singular at points where the Jacobian is zero which corresponds to the boundary of the black and white regions.

It is not surprising that the transformations $\hat{h}_A$ and $\hat{h}_B$ have regions where the Jacobian is singular. This is because the linear elasticity model does not ensure that the transformation is globally one-to-one, but rather assumes that transformation is globally one-to-one due to small linear deformations. We will see later in this chapter that the viscous-fluid model does not have this problem because the fluid transformation is globally one-to-one.

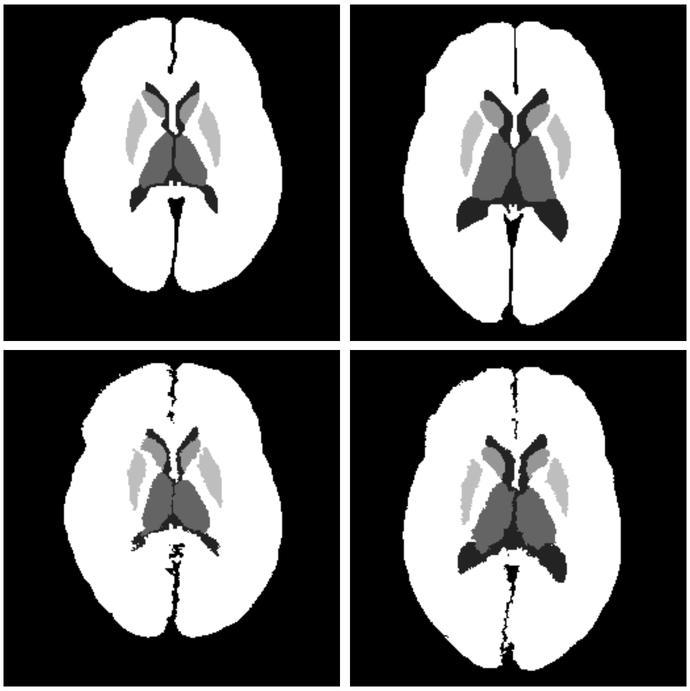

Figure 10.5: Top row shows the hand labeled structures of patients A (left) and B (right); the bottom row shows the automatic labeled structures of patients A (left) and B (right).

10.1.2 2D Elastic Visual-Cortex Experiment

The intent of this experiment was to test the ability of the 2D linear elastic model to accommodate the shape variability of the cortical folds of gray matter. In order to do this, we used high-resolution images of the cut surface taken during histological sectioning of the monkey visual cortex. These images are also referred to as cryosectioned blockface images because the images are taken of the frozen brain after each section is removed with a microtome. These data were provided courtesy of Dr. D. Van Essen and Dr. T. Coogan from the Department of Anatomy and Neurobiology at Washington University.

Fig. 10.7 shows two 2D visual cortex cryosectioned blockface images from two different macque monkeys: K (left) and L (right). These cryosection data are at approximately 10$\times$ the resolution of the MR data and show the exquisite detail of the cortical mantle. Notice that the data are much simpler than the 2D MR textbook because only two structures are included: white matter and gray matter.

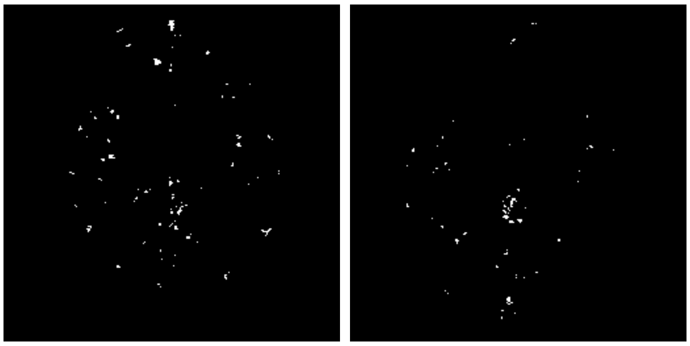

Figure 10.6: The regions where the Jacobian of $\hat{h}_A$ (left) and $\hat{h}_B$ (right) are negative are shown in white and positive in black.

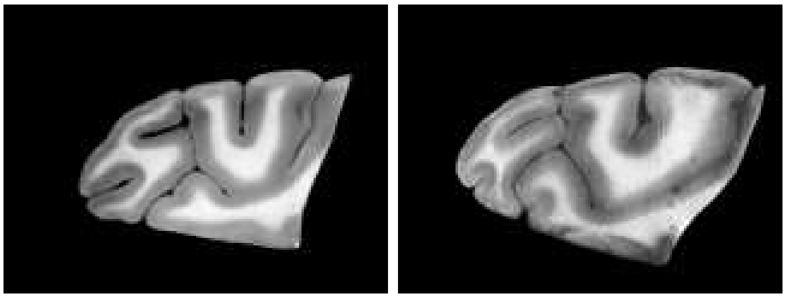

Figure 10.7: 2D Visual Cortex cryosectioned blockface images from monkey K (left) and L (right). Data courtesy of Dr. D. Van Essen, et al.

Two different sets of Lamé elasticity coefficients λ and μ were used in this experiment. The first set of of coefficients, $\mu = 0.1$ and $\lambda = 0$, corresponds to strong elasticity model and the second set, $\mu = 0.001$ and $\lambda = 0$, corresponds to a weak elasticity model.

The trade-off between the strong and weak coefficients are shown in Fig. 10.8. The left column shows the result of deforming the K template into L study using strong elasticity model in the top panel and the weak elasticity model in the bottom panel. The difference images between the deformed template and the study are shown in the right column. Notice that the strong elastic transformation (top) has a much larger difference than the weak elastic transformation (bottom). This is because the strong elasticity "held back" the template and prevented it from fully deforming into the study while the weak elasticity did not.

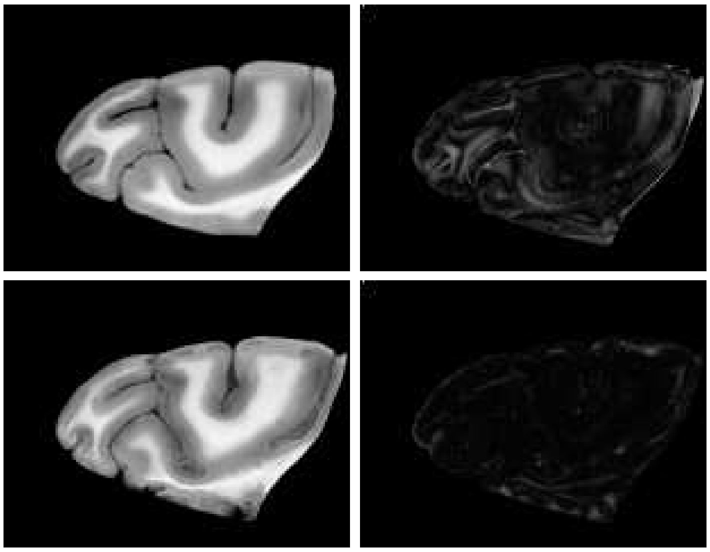

Figure 10.8: The left column shows a strong (top) and weak elastic (bottom) transformation of monkey visual cortex K into L. The right column shows the difference images between the deformed template and the study.

The top row of Fig. 10.9 shows in white the points where the Jacobian of each transformation was less than or equal to zero. The strong elastic transformation had no points where the Jacobian was negative while the weak elastic transformation had many points. Thus, we can conclude using the theorems in Chapter 3 that the strong elastic model produced a homeomorphic map while the weak elastic model did not.

The strong and weak transformations were applied to rectangular grids creating the deformed grids shown in the bottom row of Fig. 10.9. The left grid corresponds to the strong elastic transformation and shows very little deformation. The grid in the right panel corresponds to the weak elastic transformation and shows that there was very large deformations in this case. However, this later transformation did not preserve the topology of the grid in the regions where the Jacobian of the transformation is negative (see top-right panel of Fig. 10.8). The destruction of the original grid topology is evident in this deformed grid because of large knots that are present (see Section 3.3).

This experiment demonstrates that even though linear elasticity does not guarantee a homeomorphic transformation, in practice a homeomorphic transformation can be generated using an

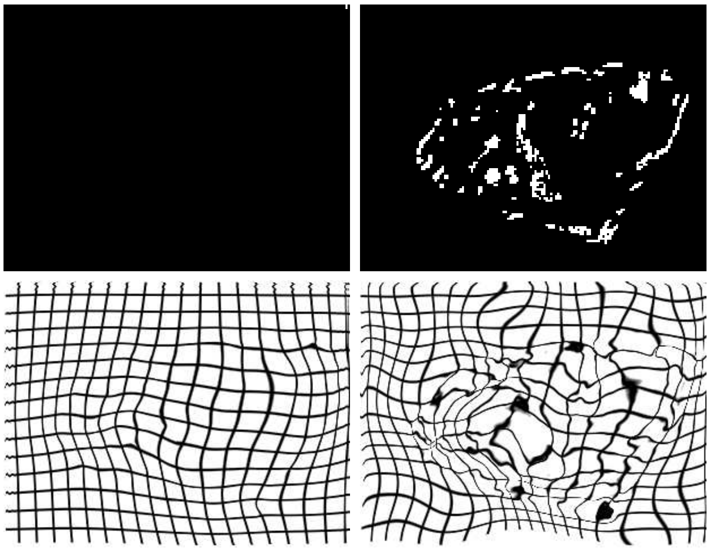

Figure 10.9: The white regions in the top two panels show where the the Jacobian of the strong (left) and weak (right) transformations of K to L is negative. The left and right panels of the bottom row shows the strong and weak linear elastic transformations applied to a rectangular grid, respectively.

elasticity model corresponding to strong elasticity. The trade-off is that only small deformations can be generated. This is consistent with the small deformation assumption used in the derivation of linear elasticity.

10.2 Two-Dimensional Fluid Experiments

10.2.1 2D Fluid Patch-to-"C" Experiment

One of the major advantages of the viscous-fluid model is the fact that it accommodates large-distance, nonlinear deformations. To demonstrate this, Fig. 10.10 shows the the wedge shaped template (left) and the "C" shaped study (right) used in the next experiment. Note that the template overlaps the lower right section of the study.

The left column of Fig. 10.11 shows the template before (top) and after (bottom) fluidly transforming the wedge shaped patch into the "C." The difference in area of these objects demonstrates the importance of the mass source term in Eq. (9.23), without which, the template could not

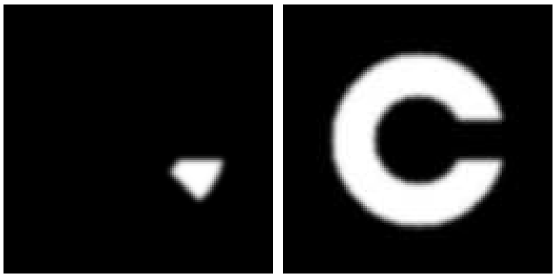

Figure 10.10: Template (left) and study (right) for patch-to-"C" experiment.

have grown in mass (area) to match the shape of the study. The right column of Fig. 10.11 shows three marker points (top) in the original template and their trajectory during the fluid transformation (bottom). Notice the extent of the curved trajectories. This solution emphasizes the ability of the kinematics as manifest via the material derivative to accommodate non-linear kinematic trajectories (see Section 9.3).

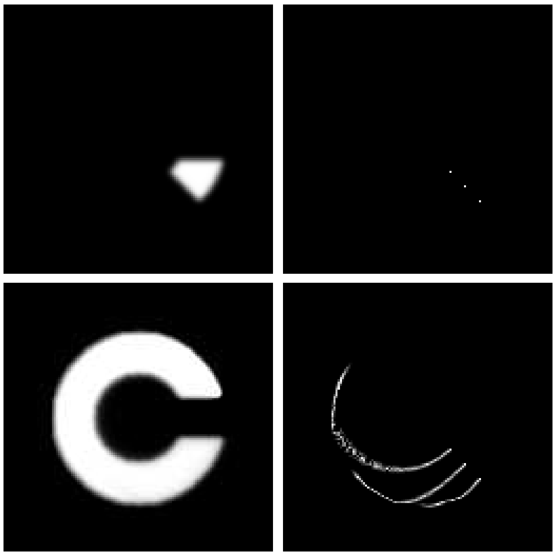

Figure 10.11: Top row shows the template (left) and three marker points (right) before transformation. The lower left panel shows the viscous-fluid transformed template. The lower right panel shows the trajectory that the three maker points traveled during the fluid transformation.

A time sequence of deformed grids which were generated by applying the fluid transformation—obtained by deforming the template (Fig. 10.10 left) into the study (Fig. 10.10 right)—to a rectangular grid and is shown in Fig. 10.12. Notice that because the fluid transformation is continuous and globally one-to-one that the grid lines are connected and continuous. The nonuniformity in the thickness of the deformed grid lines results from the fact that the original grid is expanded and compressed during the transformation.

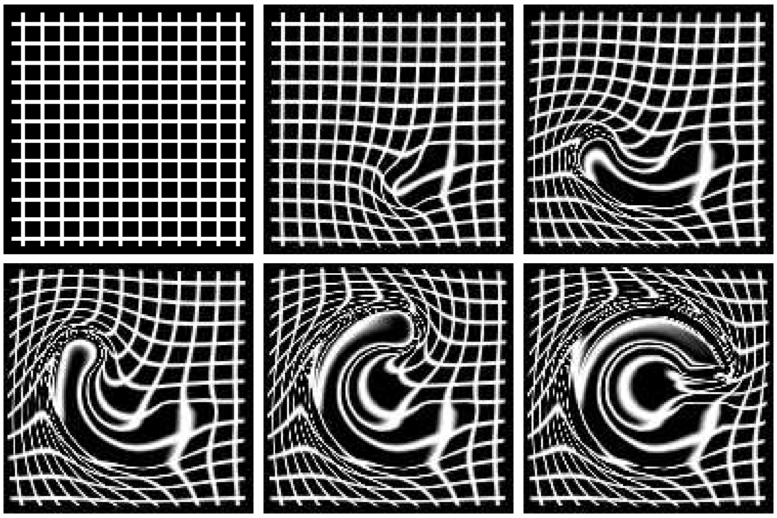

Figure 10.12: Time progression of the fluid transformation applied to a rectangular grid.

Fig. 10.13 shows a vector representation of the final displacement field for this experiment. Each displacement vector tracks the trajectory of a mass particle from its original location to its final location as a function of their final position. This is due to the Eulerian reference frame. Because the transformation is evaluated on a regular grid, these vectors terminate at regular grid locations but start at a very irregular set of corresponding locations in the textbook.

The method of propagating templates (see section 9.5) is demonstrated in Fig. 10.14. The top row shows propagated templates $T^{(6)}$, $T^{(12)}$, $T^{(18)}$, and $T^{(30)}$ (from left to right) that were used to fluidly deform the wedge shaped patch into the "C." The middle row shows the Jacobians of the transformations $\vec{x} - \vec{U}^{(i-1)}(\vec{x}, t_{p_i})$ that were used to determine the propagation times t_{p_i} for $i = 6$, 12, 18, and 30. Because the Jacobian of each propagated transformation $\vec{x} - \vec{U}^{(i-1)}(\vec{x}, t_{p_i})$

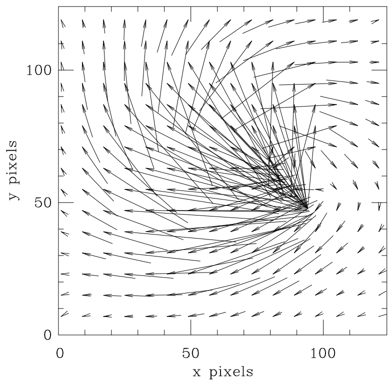

Figure 10.13: Vector map of the viscous-fluid transformation for the wedge shaped patch-to-"C" experiment.

is greater than 0.5 for all $x \in \Omega$, each propagated transformation is a continuous and globally one-to-one transformation.

The bottom row of Fig. 10.14 shows the Jacobians of the transformations $\vec{h}(\vec{x}, t_{p_i}) = \vec{x} - \vec{u}(\vec{x}, t_{p_i})$ at the propagation times $i = 6$, 12, 18, and 30 where $\vec{u}(\vec{x}, t_{p_i})$ is computed using Eq. (9.33). The Jacobian of $\vec{h}(\vec{x}, t_{p_i})$ is computed by multiplying the Jacobians of the first $i-1$ propagated transformations together. Thus, the Jacobian of $\vec{h}(\vec{x}, t_{p_i})$ is guaranteed to be positive for each $\vec{x} \in \Omega$ because the Jacobian of each of the $i-1$ propagated transformations are positive for $\vec{x} \in \Omega$.

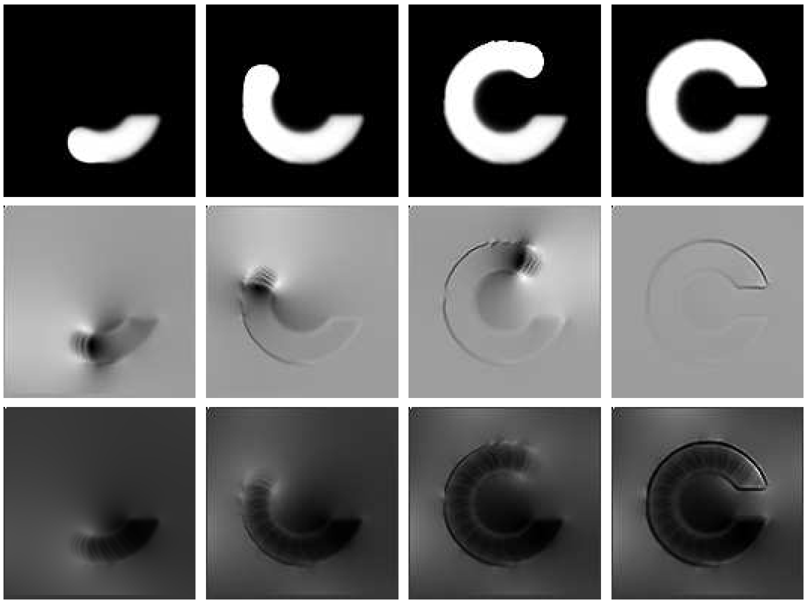

Figure 10.14: Top row shows propagated templates $T^{(i)}$, middle row shows the Jacobian of $\vec{x} - \vec{U}^{(i-1)}(\vec{x}, t_{p_i})$, and bottom row shows the Jacobian of $\vec{x} - \vec{u}(\vec{x}, t_{p_i})$ for i = 6, 12, 18, and 30.

10.2.2 2D Fluid MR-to-MR Experiment

We now return to the transformation of textbook W (see Fig.10.1) into studies A (top row Fig. 10.2) and B (left column Fig. 10.3). Fig. 10.15 shows the fluid deformation of W into A and B for the t2 modality. The top row shows the original template (left), study A (middle), and the fluid transformation of W into A (right). The bottom row shows study B (middle), and the fluid transformation of W into B (right).

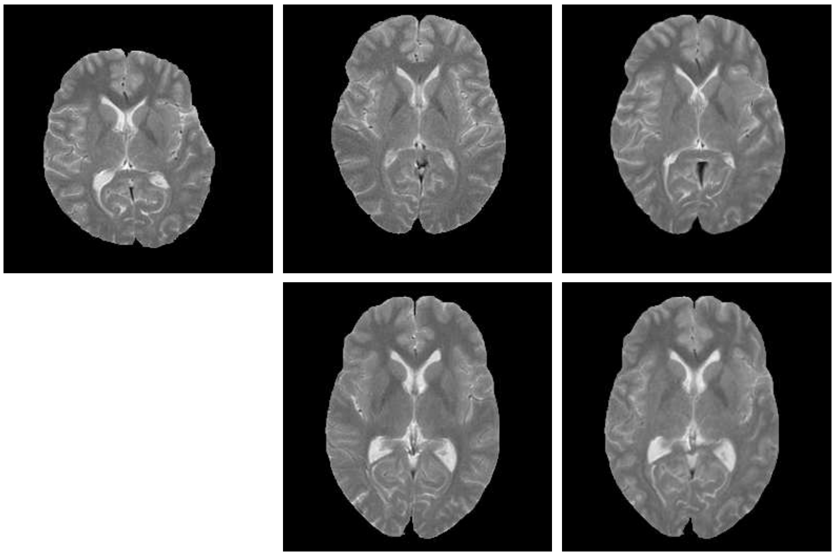

Figure 10.15: The top row shows the t2 modality of textbook W undeformed (left), study A (middle), and fluid deformation of W to A (right). The bottom row shows study B (middle) and the fluid deformation of W to B (right). Data courtesy of Scott Nadel.

A two-step multiresolution procedure was used to deform the textbook W. First, the global shape differences were accommodated using the linear-elastic-basis model and then the transformation was refined to accommodate the local shape differences with the viscous-fluid model. Fig. 10.16 shows t2 modality of textbook W before transformation (left), after basis transformation (middle), and after the fluid transformation (right). Below each image is its the difference image with the t2 modality of study A.

The $x-$ and $y-$components of the final displacement field for the fluid transformation of W to A are shown in the left and middle panels of Fig 10.17. Bright areas correspond to a positive displacement and dark areas to a negative displacement. The low-frequency components of the

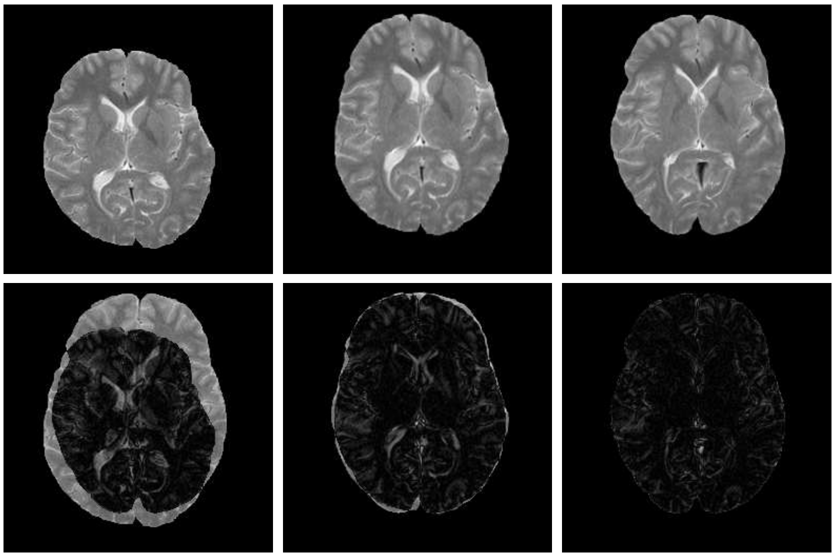

Figure 10.16: The top row shows the W template undeformed (left), globally deformed (middle), and locally deformed (right). Below each image is its the difference image with the study A. Data courtesy of Scott Nadel.

displacement fields accommodate the global shape differences between the template and study while the high-frequency components of the displacement fields accommodate the local or fine structure differences. The right panel shows the effect of the applying the transformation for this experiment to a rectangular grid.

The automatic segmentation of A generated by applying the fluid transformation to the segmentation of W is shown in Fig. 10.18. Notice how similar the hand segmentation of A (middle) is to the automatic segmentation of A (right). In addition, notice that all of the segmented regions in the deformed segmentation remained connected and did not fracture as did the automatic segmentation generated by the linear elasticity model (see right panel of Fig. 10.18). This is because the fluid algorithm generates a homeomorphic transformation while the linear elasticity model, in general, does not.

In order to quantify the performance of the automatic segmentation of study A generated from fluid transformation of the segmentation of W, the relative overlap measure will be used [60]. The relative overlap of a common structure in two different segmentations is defined to be the area of the intersection divided by the area of the union.

Figure 10.17: The left and middle panels show the x- and y- components of the displacement field and the right panel shows the W-to-A fluid transformation applied to a rectangular grid. Data courtesy of Scott Nadel.

Figure 10.18: The textbook segmentation of W (left), hand segmentation of A (middle), and the fluidly transformed segmentation of W to A (right).

The relative overlap is a useful measure for comparing two different segmentations, but it must be interpreted with caution. First, the relative overlap is not a linear measure of the overlap of structures. For instance, if two regions are of equal area and overlap by 50 percent of their area, their relative overlap is only 0.33. Second, the relative overlap of two manual segmentations of the same anatomy has an error due to the uncertainty of the segmentations. The uncertainty in the segmentations is because of the uncertainty of where edges are placed in the segmentation. This uncertainty is due to an accumulation of many different factors such as partial volume effects, discretization, and operator error. The error due to this uncertainty is larger for structures with a large ratio of surface area to volume than for structures with a small ratio; this is because the uncertainty occurs at the edges of a structure and not in the interior of the structure.

Table 10.1 shows the relative overlap of all the segmented regions before transformation, after the linear-elastic-basis transformation, and after the viscous-fluid transformation. This table shows that the relative overlap increased for each of the structure increased after transformation with the basis and again after the fluid transformation. Notice that the head of the caudate nucleus did not overlap initially in the textbook and study, but had a relative overlap of 0.78 after transformation. Also, the relative overlap of the background went from 0.79 to 0.99 which means that outer contour of the deformed textbook matched almost perfectly with the outer contour of the study.

Table 10.1: Relative overlap for the fluid deformation of W to A.

Structure	Before Transformation	After Elastic-Solid Basis Transformation	After Fluid Transformation
background	0.79	0.96	0.99
ventricles	0.017	0.43	0.63
thalamus	0.097	0.64	0.81
head of caudate nucleus	0.0	0.53	0.78
putamen	0.0	0.45	0.70
other brain matter	0.53	0.86	0.95

10.2.3 2D Fluid Visual-Cortex Experiment

The fluid transformation accommodates variation of tremendous detail. This is precisely the kind of detail and accommodation of variability which we require for the cryosection brain data shown in Fig. 10.19. The top-left and top-middle panels of Fig. 10.19 shows optical cryosection images of visual cortex from monkeys K and L, respectively. The result of fluidly deforming the template into the study is shown in the top-right panel.

Once again, the template was first elastically deformed into the study using the linear-elastic basis and then refined with the fluid model. The bottom row shows the deformed grid lines which have been smoothly deformed. The x- and y-components of the displacement field are shown in the bottom-left and bottom-right panels respectively. Intensities that are close to white correspond to positive displacements and dark intensities correspond to negative displacements. The low-frequency intensity changes in displacement images cause global shape changes and resulted from

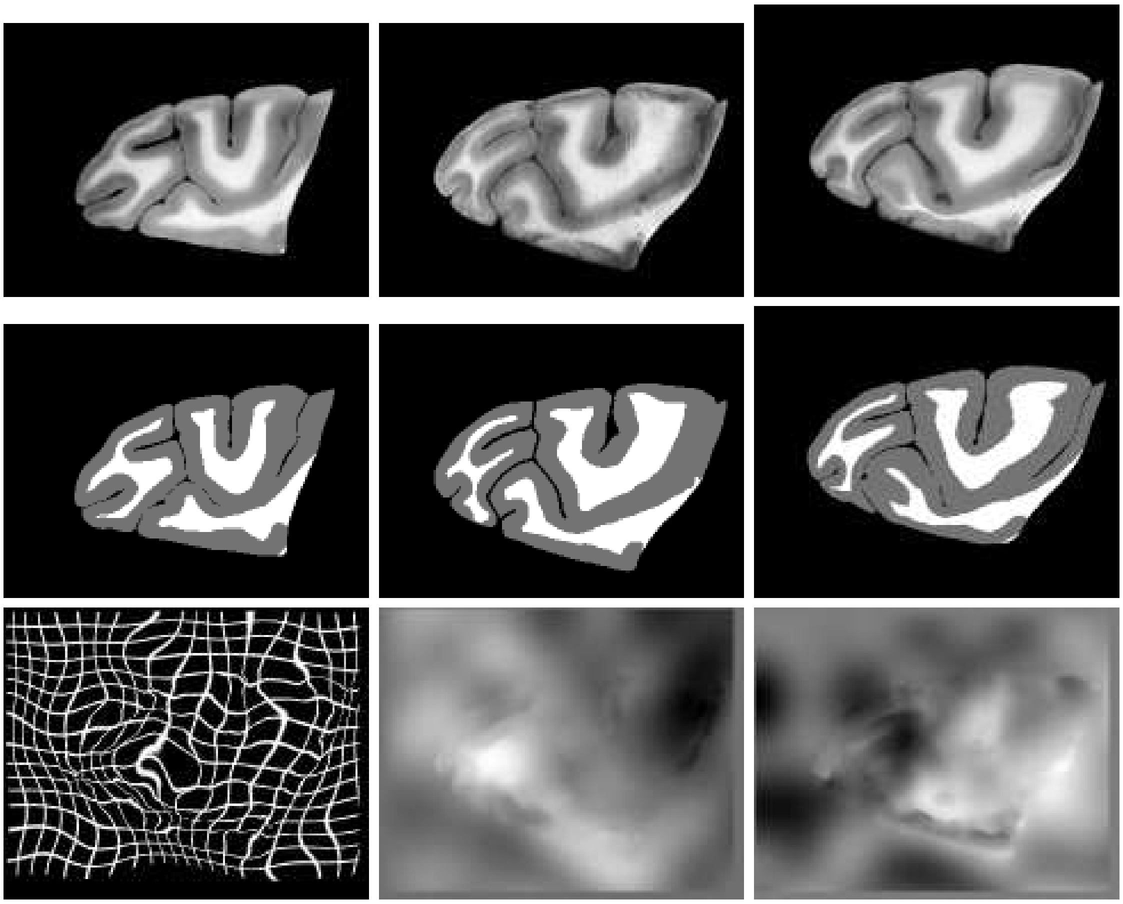

Figure 10.19: Top-left panel shows the monkey visual cortex cryosection template and the top-middle panel shows the second monkey slice. The top-right panel shows the result of fluidly deforming the top-left into the top-middle anatomy. The second row shows a gray/white matter hand segmentation of the template (left) and study (middle). Shown in the middle-right panel is an automatic segmentation of the study that was generated by applying the fluid transformation to the template segmentation. The bottom-left panel shows the fluid transformation applied to the rectangular grid. The x- and y-components of the estimated transformation are shown in the bottom-middle and bottom-right panels respectively. Data courtesy of Dr. D. Van Essen, et al.

the elastic-solid transformation. The high-frequency intensity changes correspond to local shape changes resulting from the fluid transformation.

This transformation was then applied to the template segmentation and grid to generate the right-middle and right-bottom panels, respectively. The automatic segmentation shows excellent agreement with the study (top-middle) and a hand segmentation of the study (middle). Differences between the hand and automatic segmentation of the study are due to both errors in the hand segmentation and errors with the automatic segmentation.

10.2.4 2D Fluid Circle-to-"C" Experiment

The next experiment illustrates how a template fluidly deforms into a study if the template differs locally from the study by a concave region. Shown in Fig. 10.20 is the template (left) and study (middle) for this experiment. Both the template and the study are 128×128 images that have intensity zero for the background and 255 for the foreground. The radius of the circle is 31 pixels and the inner and outer radii of the "C" are 21 and 41 pixels, respectively. The gap of the "C" is 20 pixels wide.

Figure 10.20: The template (left), study (middle), and fluidly deformed template (right) for the circle-to-"C" experiment.

The right panel of Fig. 10.20 shows the result of fluidly deforming the circle into the "C" and Fig. 10.21 shows a time evolution of the deformation visualized by the deformation of a rectangular grid. The deformed grids show that the circle deformed into the "C" by moving the region of the circle that overlapped the gap in the "C" to the center. Notice that the right side of the template that overlaps the gap in the "C" was stretched into the shape of the interior edge of the "C."

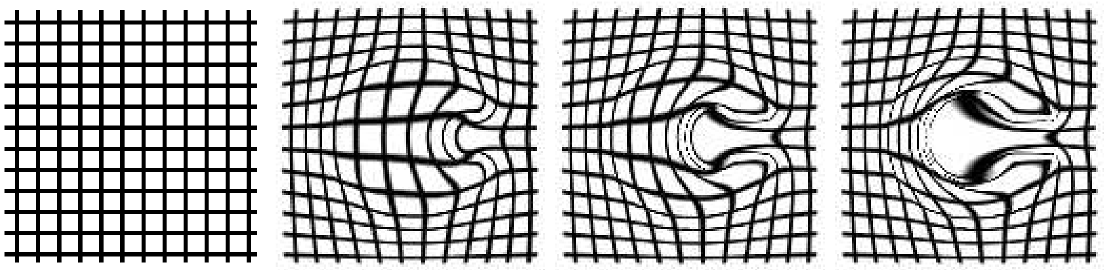

Figure 10.21: Time sequence of the fluid deformation of a circle to "C."

Fig. 10.22 shows that particles of material of the circle took large distance, nonlinear trajectories in order to deform the circle to the "C." The material of the circle that voided the center of the "C" moved first to the left into the gap of the"C" and then outward to match the interior contour of the "C." To give a feeling for the distance that some of the template particles traveled, particles at the right side of the circle moved a distance of 52 pixels (41% of the overall width of the image) to match the interior of the "C."

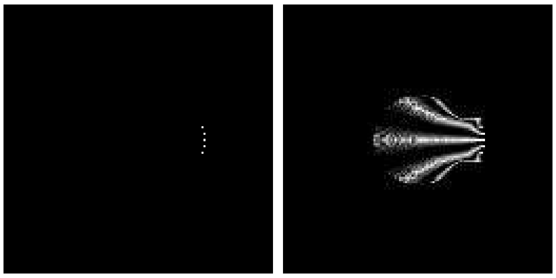

Figure 10.22: Initial positions of five marker points (left) and their trajectory throughout the transformation (right).

Notice that the trajectories of these five marker points spread apart as they move from their original to final positions. This is consistent with the fact that (1) the right side of the template that overlaps the gap in the "C" was stretched into the shape of the interior edge of the "C" (see above), and that (2) the perimeter of the circle increased as it deformed into the "C."

Fig. 10.23 shows a target segmentation of the circle before and after transformation. Notice that all of the rings are evident and connected in the deformed segmentation; this is due to the fact that the transformation is smooth and globally one-to-one. An interesting point to notice about this deformed segmentation is that the rings do not have equal thickness at points A and B. The width of the rings near point B appear to be more compressed than the rings at point A. This behavior will be explored in more depth in the next experiment.

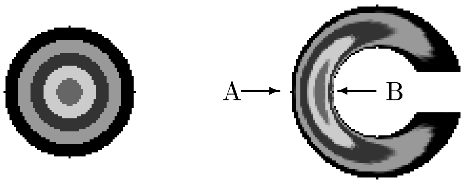

Figure 10.23: The template segmentation before (left) and after (right) transformation.

The result of this experiment is important for a number of reasons. First, the circle template was able to be fluidly deformed into the shape of the "C" even though large nonlinear deformations were required. Second, only a small portion of the template was stretched to account for the larger perimeter length of the study. This deformation is very different from deforming the circle by uniformly squeezing the circle into a large cigar shape and then bending it into the shape of the "C." And finally, the segmentation shows that the material of the template was not uniformly compressed between points A and B (see Fig. 10.23).

10.2.5 2D Fluid Square-to-Rectangle Experiment

This experiment illustrates how a region of constant intensity fluidly expands into a larger region of constant intensity. Shown in the top row of Fig. 10.25 is the template (left) and study (middle) for this experiment. Both the template and the study are 128×128 images that have intensity zero for the background and 255 for the foreground. The width of the square is 32 pixels and the width and height of the rectangle are 64 and 32 pixels, respectively.

Fig. 10.25 shows the result of fluidly deforming the square template into the rectangle study. The left column shows the template before (top) and after (bottom) fluid deformation. Notice that the deformed square matches exactly with the shape of the rectangle study. The middle and right columns show a rectangular grid and a template segmentation before (top) and after (bottom)

Figure 10.24: The template (left) and study (right) for the square to rectangle experiment.

deformation, respectively. The deformed grid and segmentation show that the constant intensity template was not uniformly stretched into the shape of the study. In fact, these images show that the deformed template was compressed on the right side and expanded on the left side.

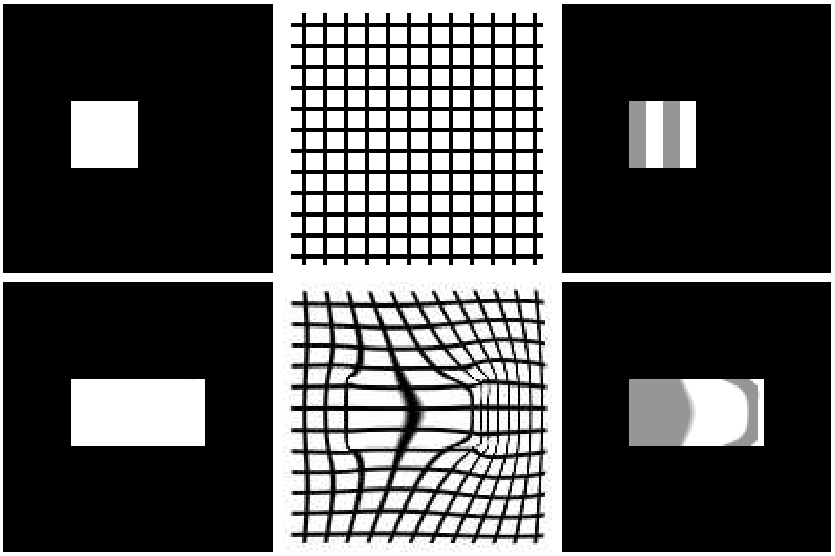

Figure 10.25: The columns show the template (left), grid (middle), and template segmentation (right) before and after deformation for the square-to-rectangle experiment.

The fact that the compression occurs on the right side of the deformed template and the expansion on the left side is a surprising result. Intuitively, one might guess that the expansion would occur on the right side of the template because that is the side that experiences the deforming force. However, this is not the case.

The reason that the square expands non-uniformly in the x-direction is because the fluid model is not solved in space and time simultaneously. The current implementation solves the model simultaneously in the three spatial coordinates (Eq. (9.26) is solved using SOR) and then uses forward integration (see Eq. (9.29)) to propagate the displacement field through time. The result is that the solution is able to relax over space but not over space and time together. Thus, the square is stretched uniformly over space, but not over time.

The reason that the interior edges of the segmentation are curved is due to the boundary conditions and viscous nature of the fluid model. The boundary conditions are zero velocity along the perimeter of the template domain. The effect of this boundary condition is that material along the edge of the template domain does not move from its original position while the material in the center of the domain is free to deform into the shape of the study (see the deformed grid). The viscous nature of the fluid model causes the material in near the edges of the domain to move slower than the material in the center of the domain. The varying speeds of the deforming continuum from the center to the edges of domain causes the interior edges of the segmentation to bend during deformation.

10.3 Three-Dimensional Experiments

10.3.1 3D Elastic MR-to-MR Experiment

The next experiment demonstrates the 3D elastic deformation of an MR textbook into a study. Fig. 10.26 shows the $128 \times 128 \times 148$ MR volume used in this experiment for the textbook (left) and the study (middle). The MR sequence of these data is called MPRAGE and was provided courtesy of Michael W. Vannier, et al. of the Mallinckrodt Institute of Radiology at Washington University and Marcus E. Raichle, et al. of the Department of Neurology at Washington University.

The original data were downsampled from $256 \times 256 \times 128$ to $128 \times 128 \times 128$ using trilinear interpolation and was padded by 10 planes of zero intensity background in the z-direction. The volumes displayed in this figure were generated by surface rendering the fatty tissue of the head and replacing the top surface of the head with the proper axial MR data slice. These volumes were visualized with the help of Sarang Joshi and Jing Wang. Notice that when the two MR volumes are cut at the same data slice that the anatomy differs greatly from the textbook and the study. Also notice the large difference in the appearance of the face of the textbook and the study.

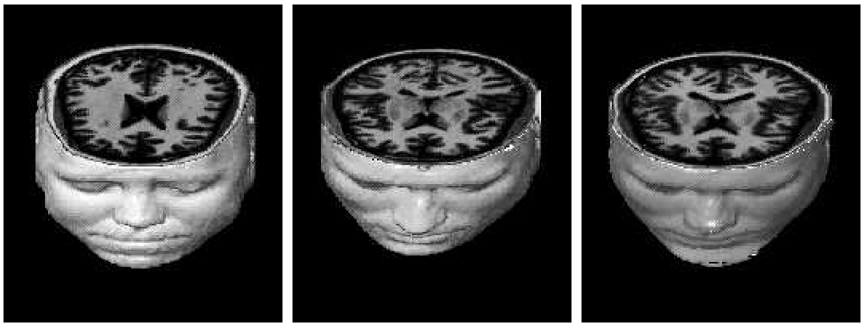

Figure 10.26: The MR textbook (left), study (middle), and elastically deformed textbook (right) for the 3D elastic MR-to-MR experiment.

The right panel of Fig. 10.26 shows the 3D elastic deformation of the textbook into the study. Notice how similar the anatomy of the deformed (individualized) textbook is to that of the study volume, and how similar the surface of the deformed textbook face is to the study face. The reason that deformed-textbook and study faces look almost identical is because the squared error forcing function minimizes the mismatch between the background and the skin. A small mismatch between the skin and the background produces a large body force at the point of mismatch because there is a large difference in intensity between the skin and the background.

The top and bottom rows of Fig. 10.27 show the sagittal MR slices 74 and 62 from the textbook (left), study (middle), and elastically deformed textbook (right), respectively. The differences between the textbook and study for these slices illustrate the need to accommodate anatomical shape differences in all three coordinate directions. For example, notice how different the shape of the corpus callosum of the study (top-middle) is from the textbook (top-left) and how similar it is to the deformed textbook (top-right). Likewise, the bottom row of Fig. 10.27 shows good agreement in the shape of the eye sockets, nasal cavity, ventricles, cortical folds, and cerebellum after deformation.

Fig. 10.28 shows the $x-$, $y-$, and $z-$components of displacement field $\vec{u}(\vec{x})$ for sagittal slice 74 and 62 of the elastically deformed textbook. The bright regions in these images correspond to positive displacements and dark regions correspond to negative displacements. The low-frequency

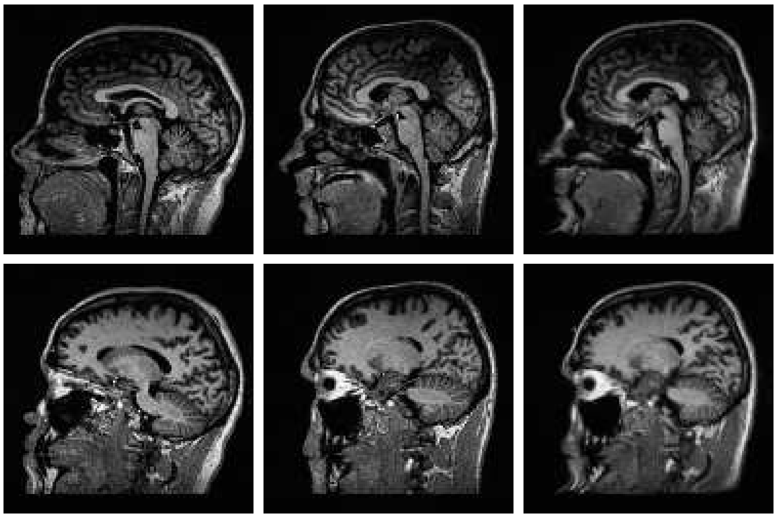

Figure 10.27: The top and bottom rows show sagittal MR slices 74 and 62 from the MR textbook (left), study (middle), and elastically deformed textbook (right), respectively.

components of the images in Fig. 10.28 correspond to global deformations of the textbook while the high-frequency components correspond to local deformations of the textbook.

This transformation was estimated using the 3D linear-elastic-basis model and was then refined using the 3D spatial linear elastic model. The elastic-basis transformation was parameterized by 192 basis coefficients and the spatial transformation was parameterized by $3 \times 128 \times 128 \times 148 = 7.3 \times 10^6$ floating-point numbers corresponding to the components of the $128 \times 128 \times 148 = 2.4 \times 10^6$ displacement vectors.

Fig. 10.29 shows the coordinates where the Jacobian of the 3D elastic transformation is less than or equal to zero in white. Because these images are not totally black, this 3D elastic transformation is not a homeomorphism. This outcome was expected because the linear elastic model does not constrain the transformation to be homeomorphic, but rather assumes that it is homeomorphic due to a small linear deformation. The fact that these images are not homeomorphic means that the textbook information is not smoothly and uniquely associated with the study anatomy. This point is not evident when viewing the deformed textbook anatomy and study in Fig. 10.27 and comparing similarities.

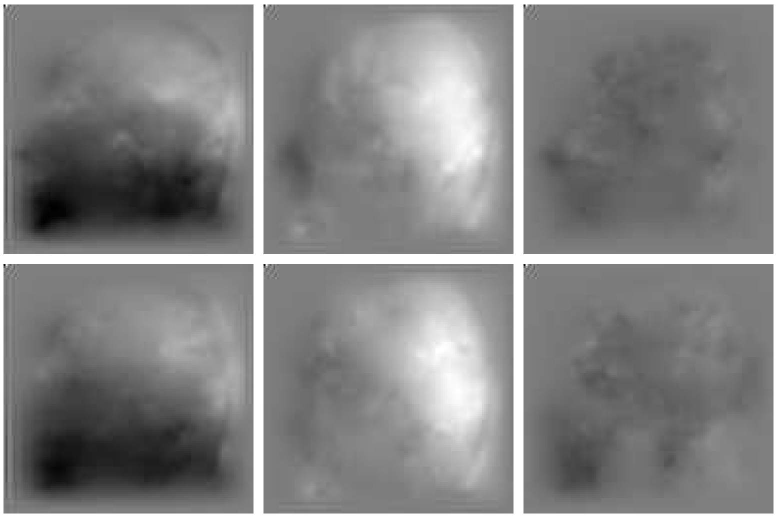

Figure 10.28: The top and bottom rows show the x- (left), y- (middle), and z-components (right) of the displacement field $\vec{U}(\vec{x})$ used to generate slices 74 and 62 of the elastically deformed textbook, respectively.

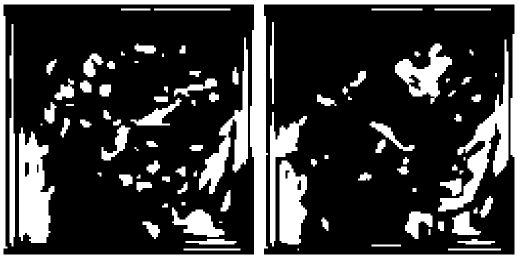

Figure 10.29: The white regions show where the Jacobian of the transformation from S to D is less than or equal to zero for sagittal slices 74 and 62, respectively.

10.3.2 3D Elastic and Fluid Monkey-Visual-Cortex Experiments

The next experiment was designed to test the performance of the 3D elasticity and 3D fluid models for accommodate the three dimensional shape variation of the monkey visual cortex.

In this experiment, two 3D volumes of the monkey visual cortex were constructed by downsampling 26 consecutive images from their original resolution of 640×480 to 128×96. These volumes were then zero padded to make the total volume $128 \times 128 \times 42$. Slices 14, 18, 22, 26, and 30, of data set K and L are shown in the top two rows of Fig. 10.30, respectively. These data were provided courtesy of Dr. D. Van Essen and Dr. T. Coogan, et al. of the Department of Anatomy and Neurobiology at Washington University.

The bottom row of Fig. 10.30 shows the 3D elastic deformation of K into L for slices 14, 18, 22, 26, and 30, respectively. This transformation was generated from a 3D spatial linear elastic refinement of a 3D linear-elastic-basis transformation. The top row shows the textbook, the middle row the study, and the bottom row the deformed textbook. As with the previous 3D elastic experiment, the deformed textbook looks almost identical to the study.

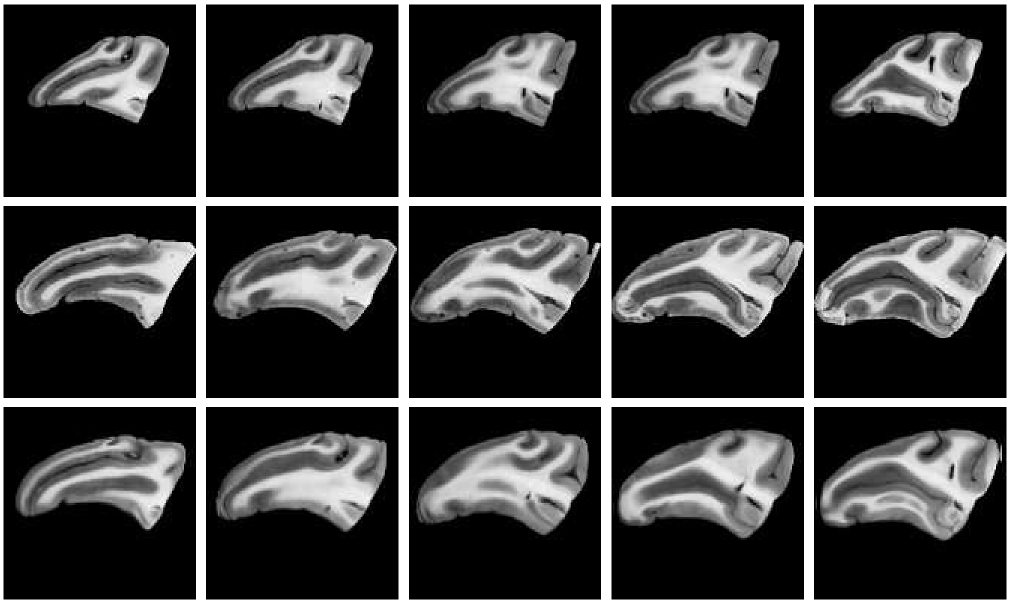

Figure 10.30: Slices 14, 18, 22, 26, and 30 (right to left) of the monkey cortex K template (top row), the monkey cortex L study (middle row), and the 3D elastically deformed template (bottom row).

The white points in Fig. 10.31 indicate where the Jacobian of the 3D elastic transformation of K to L is non-positive for slices 14, 18, 22, 26, and 30, respectively. As in the previous linear elasticity experiments, this figure shows that although the textbook can be deformed so that it looks like the study using the 3D elastic model, the estimated transformation is not homeomorphic.

Figure 10.31: The white points indicate where the Jacobian of the 3D elastic transformation of K to L is singular for slices 14, 18, 22, 26, and 30, respectively.

Fig. 10.32 shows slices 14, 18, 22, 26, and 30 of the 3D fluid deformation of visual cortex K into L. This transformation was generated from a 3D viscous-fluid refinement of a 3D linear-elastic-basis transformation. The top row shows the textbook, the middle row the study, and the bottom row the deformed textbook.

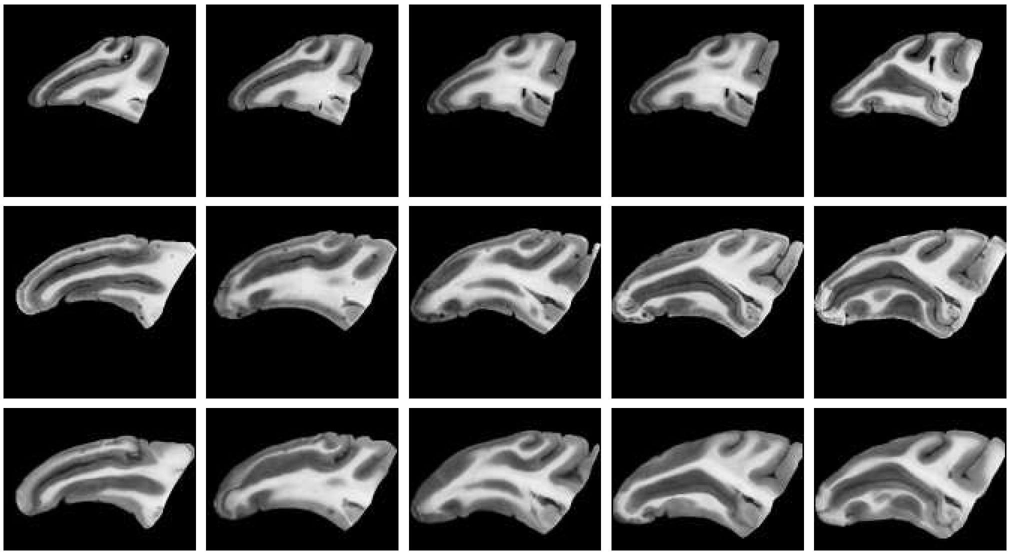

Figure 10.32: Slices 14, 18, 22, 26, and 30 (right to left) of the K monkey visual cortex template (top row), the L monkey visual cortex study (middle row), and the 3D fluidly deformed template (bottom row).

Visual comparison of the elastic and fluid deformations in Figs. 10.30 and 10.32 shows that the fluidly deformation of K to L is slightly better than the elastic deformation. However, the fluid transformation is clearly superior to the elastic transformation because it is a homeomorphic mapping while the elastic transformation is not.

The fluid transformation of K to L was used to generate an automatic segmentation of visual cortex L. The top row of Fig. 10.33 shows slices 14, 18, 22, 26, and 30 from the undeformed segmentation of K and the bottom row of Fig. 10.33 shows slices 14, 18, 22, 26, and 30 of the fluid deformation of K into L. A hand segmentation of L corresponding to these slices are shown in the middle row for comparison.

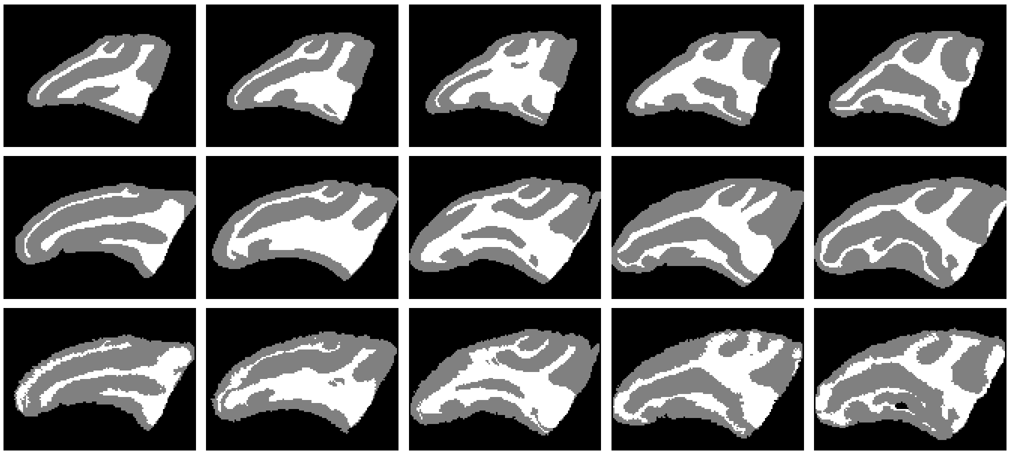

Figure 10.33: Slices 14, 18, 22, 26, and 30 (right to left) of a gray/white matter segmentation of template K (top row), a hand segmentation of study L (middle row), and the 3D fluid deformed template segmentation (bottom row).

In order to quantify the segmentation, the segmentation correspondence of high-confidence points (SCHCP) was computed. This statistic takes into account the uncertain nature of the original textbook segmentation. The uncertainty in the segmentations is because of the uncertainty of where edges are placed in the segmentation. This uncertainty is due to an accumulation of many different factors such as partial volume effects, discretization, and operator error. This uncertainty means that voxels near edges have a lower confidence of being labeled correctly then voxels farther away from the edges.

The SCHCP statistic is computed in the following manner. First, uniform random locations are generated in the deformed coordinate system. These points are then mapped back to the textbook. If the point maps back to a location that is near an edge in the segmentation it is discarded as being a point of low confidence. On the other hand if the point maps back to a point that is away from an edge, it is labeled as a point of high confidence. Next, the segmentation labels for the points of high confidence are compared to a hand segmentation. The SCHCP statistic is computed as the percentage of correctly labeled voxels (as determined by the hand segmentation) to the total number of high-confidence voxels used in each segmented region.

Table 10.2 shows the SCHCP statistic for the automatic segmentation visual cortex L. The threshold of certainty used in this table was that voxels had to be at least two voxels away from an edge. This table states that the gray matter and white matter regions increased from about 60 percent agreement initially to about 80 percent after the elastic basis to nearly 98-99 percent after the fluid transformation.

Table 10.2: Segmentation correspondence for points of high confidence in K-to-L 3D fluid experiment.

Structure	Before Transformation	After Elastic-Solid Basis Transformation	After Fluid Transformation
background	90.7 %	99.9 %	99.9 %
gray matter	62.6 %	86.6 %	99.4 %
white matter	61.4 %	80.2 %	97.5 %

10.3.3 3D Fluid MR-to-MR Experiment

Returning to human MR data, this section presents four different 3D fluid deformations. The data and the hand segmentations were provided courtesy of Michael W. Vannier, et al. of the Mallinckrodt Institute of Radiology at Washington University. The MR segmentations were generated by Barry Bruneston.

Figs. 10.34–10.37 show four different 3D fluid experiments. In each figure, the textbook is shown in the left panel, the study in the middle panel, and the deformed textbook in the right panel. All of these data sets were linearly interpolated from their original size of $256 \times 256 \times 128$ to $128 \times 128 \times 80$ and padded by 10 planes of background in the z-direction. This generated data sets of $128 \times 128 \times 100$ with voxel size $2mm^3$. The volumes were thresholded and surface rendered to create the faces in these panels. Notice how similar the faces of the deformed textbooks are to their corresponding target study.

All of the transformations shown in Figs. 10.34–10.37 were generated by using the 3D fluid model to refine the 3D linear-elastic-basis model. The 3D basis transformations were inspected

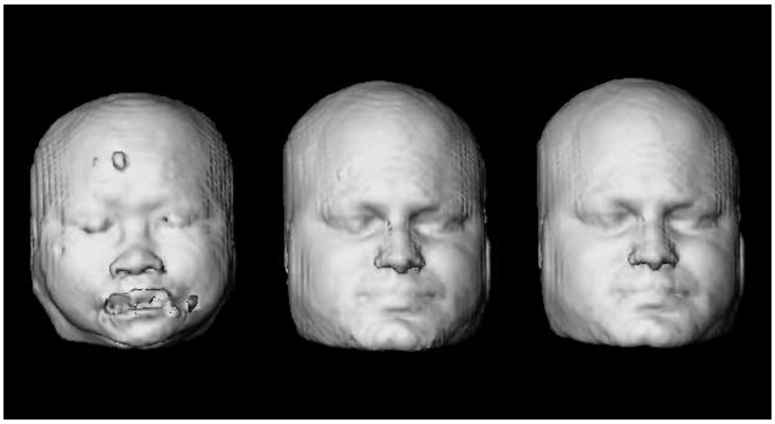

Figure 10.34: FJ textbook (left), P study (middle), and 3D fluid deformation of FJ to P (right).

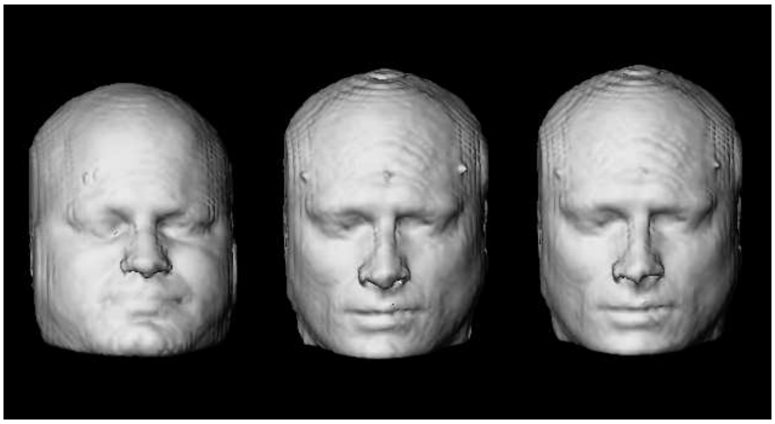

Figure 10.35: P textbook (left), D study (middle), and 3D fluid deformation of P to D (right).

to make sure that they were homeomorphic transformations. The elastic-basis transformation was parameterized by 192 basis coefficients and the spatial transformation was parameterized by $6 \times 128 \times 128 \times 100 = 9.8 \times 10^6$ floating-point numbers corresponding to the components of the $128 \times 128 \times 100 = 1.6 \times 10^6$ displacement and velocity vectors.

Figs. 10.38-10.40 show sagittal slices 32, 38, and 43, respectively, of the fluid transformation of P into FJ. The top row of each of these figures shows the undeformed textbook slice (left), the study slice (middle), and the deformed textbook slice (right). Under each image is its corresponding white/gray matter segmentation. Although these segmentations may not look perfect, they look

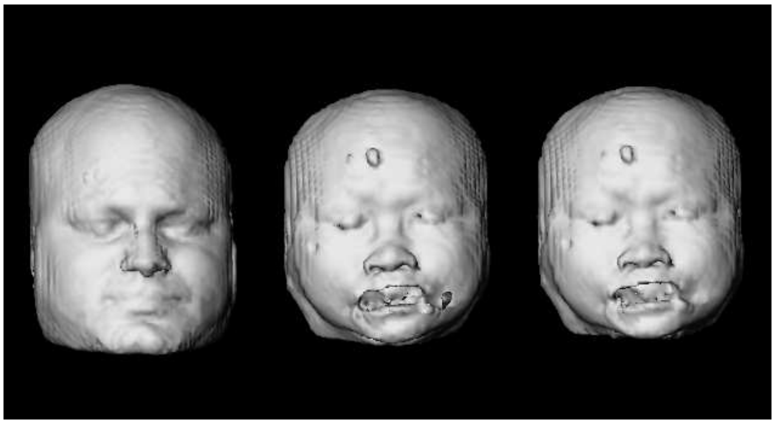

Figure 10.36: P textbook (left), FJ study (middle), and 3D fluid deformation of P to FJ (right).

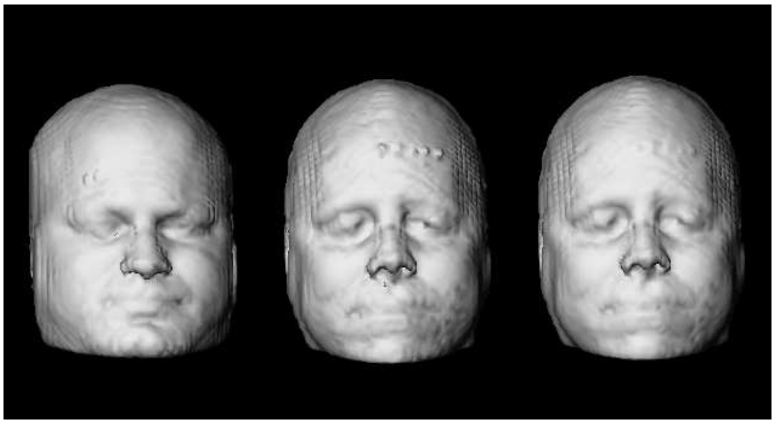

Figure 10.37: P textbook (left), S study (middle), and 3D fluid deformation of P to S (right).

pretty good. More importantly these segmentations show that gray matter is getting mapped to gray matter and white matter to white matter. This fact along with the fact that the transformation is homeomorphic support the hypothesis that the topology of the textbook is being mapped into the topology of the study.

The purpose of presenting the automatic segmentations in Figs. 10.38-10.40 was to demonstrate qualitatively how the deformable textbook will be used in the future to generate 3D automatic segmentations. Due to the memory limitations on the current computer hardware for the current algorithm, the data sets had to be downsampled to $2mm^3$ voxels. At this resolution the cortical

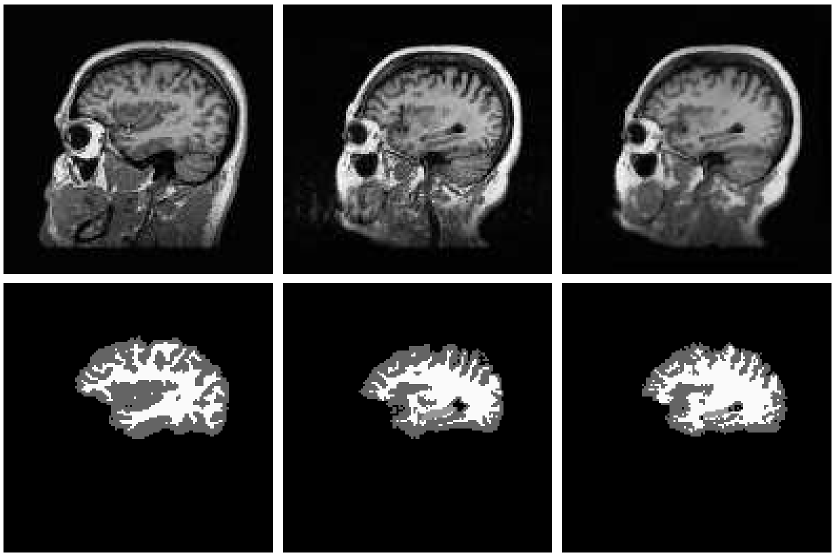

Figure 10.38: MR and segmentation of slice 32 of the P textbook (left), FJ study (middle), and deformed textbook (right).

mantel is only 2 to 4 pixels wide. This means that a one or two pixel error in the original segmentations and in the automatic segmentations shows up as a large difference. Because of the uncertainty of the original segmentations at this coarse resolution, no quantitative measure such as the relative overlap or SCHCP statistic would be meaningful measure of the performance. Therefore, no quantitative measure of the segmentation was performed. In addition, the previous 3D fluid monkey visual cortex experiment demonstrated that with enough resolution to represent the cortical mantle that the deformable textbook did a good job of segmentation.

Another reason why a quantitative measure was not performed on this automatic segmentation was that it was assumed that the topology of the textbook and the study were the same. However, the topology of the human brain is so complex that we do not know if this is a good assumption. A future area of research will be to use the deformable textbook to help classify the topology of the brain, i.e., the deformable textbook will be used as a tool that will allow researchers to compare anatomies to see whether or not they are topologically equivalent. Currently no quantitative tool exists that can be used to answer these questions.

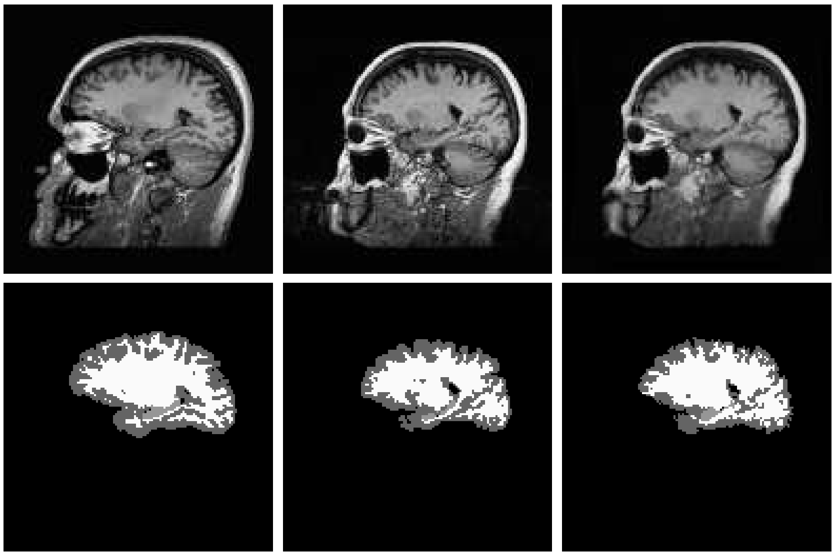

Figure 10.39: MR and segmentation of slice 38 of the P textbook (left), FJ study (middle), and deformed textbook (right).

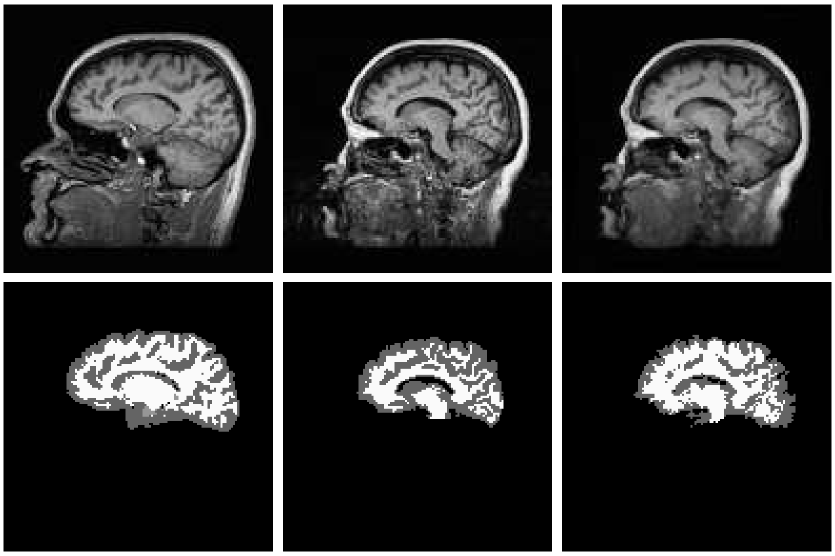

Figure 10.40: MR and segmentation of slice 43 of the P textbook (left), FJ study (middle), and deformed textbook (right).

10.3.4 3D PET-to-PET Registration Using A Deformable Textbook

The next experiment was designed to show the feasibility of using the deformable textbook to remove anatomical variation from functional modalities such as PET. The MR and PET data were provided courtesy of Marcus E. Raichle, et al. of the Department of Neurology at Washington University and Michael W. Vannier, et al. of the Mallinckrodt Institute of Radiology at Washington University. For both data sets, the PET data were rigidly registered to the MR data sets by Jing Wang and John Haller.

Anatomical shape variability is removed from functional images by finding the transformation that maps the anatomical modality of the study to the coordinate system of the textbook. This transformation is then used to map the functional modality to the coordinate system of the textbook. Thus, both an anatomical and a functional modality of the study anatomy are required in order to remove the anatomical shape variability from the functional modality.

In this experiment the anatomical and functional modalities corresponded to MR and PET, respectively. Fig. 10.41 shows the 3D outer surface of the MR modality of study D (left) and textbook S (middle) used in this experiment. After proper scaling of the PET data, the PET

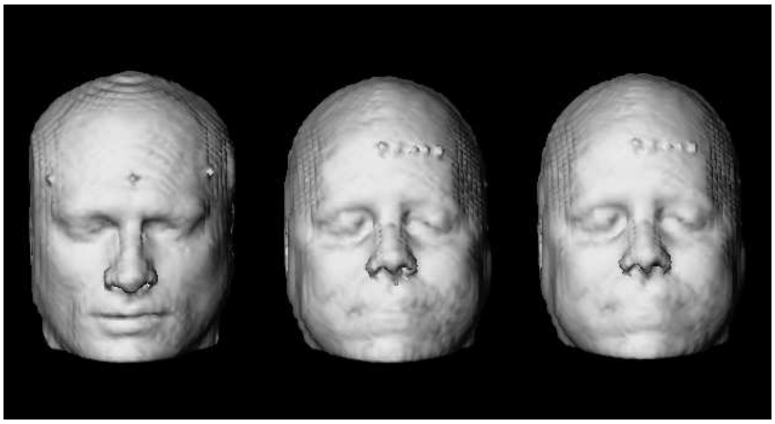

Figure 10.41: The D study (left), S textbook (middle), and the deformed study (right) for the 3D PET-to-PET fluid deformation experiment.

volume was rigidly registered to the MR volume by translations and rotations in 3-space. Next, the MR modality of the study was deformed into the shape of the the textbook modality as shown in the right panel. This transformation was estimated as a 3D fluid refinement of a 3D linear-elastic-basis transformation. The PET data were then projected into the coordinate system of the textbook using this transformation.

Note that if the anatomic modality of the study is noisy, the textbook volume would have to be deformed into the study volume. The reason for this is that the textbook modality would have to be noiseless in order to take proper derivatives. The inverse transformation would then be used to map the study back to the coordinate system of the textbook. The inverse transformation exists when the transformation is homeomorphic. Through out this chapter, we have assumed that the noise in the MR data sets is negligible and used the raw MR data as the template.

Figs. 10.42-10.44 show the 3D fluid deformation of D into S for sagittal slices 46, 56, and 68, respectively. The top row of each of these figures shows the undeformed textbook slice (left), the study slice (middle), and the deformed textbook slice (right). The bottom row shows the PET images which correspond to the MR image directly above them.

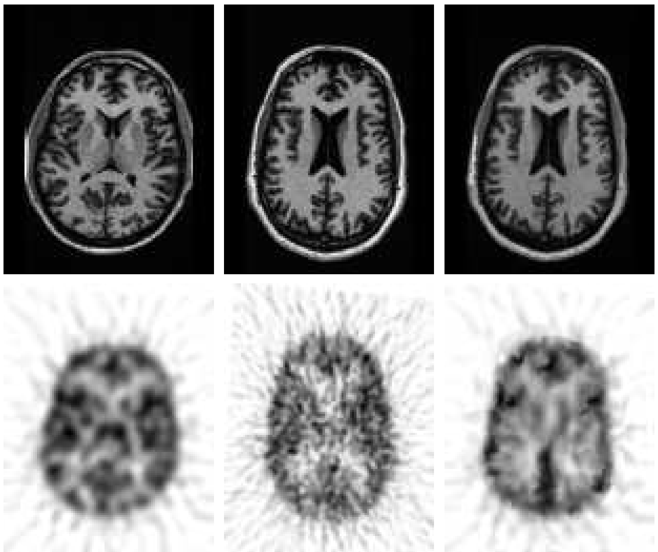

Figure 10.42: MR and PET modalities of slice 46 of the textbook (left), study (middle), and deformed textbook (right). For the PET images, higher tracer intensities are darker.

The top row of Figs. 10.42-10.44 shows that the shape differences between the D and the S anatomy where accommodated by the transformation. This implies that the deformed PET modality associated with the D anatomy was also corrected for the anatomical shape differences

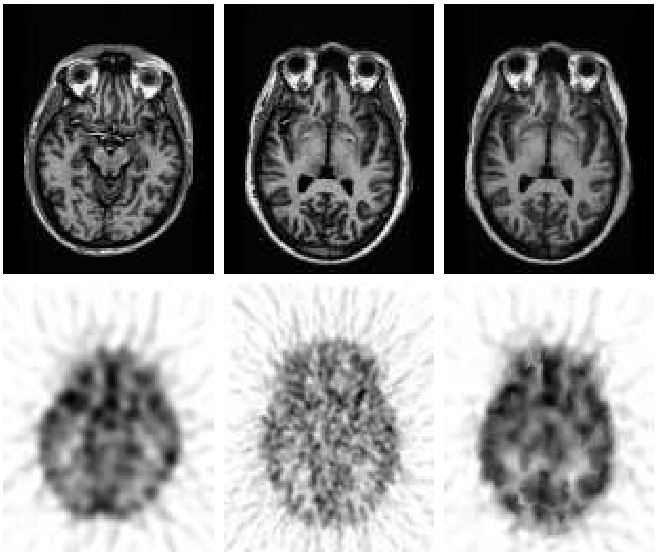

Figure 10.43: MR and PET modalities of slice 56 of the textbook (left), study (middle), and deformed textbook (right).

between D and S. It is now possible to add these two PET images together to generate a standard composite PET image that has been corrected for anatomical shape differences.

The approach described above differs greatly from the standard approach which only corrects for rigid translation, rigid rotation, and nonrigid scale. By accounting for the local shape differences between the anatomy, researchers will be able to determine if normal anatomies process data in the same place or slightly different places. For instance, suppose that ten individuals process data in exactly the same anatomical location but their anatomy differs significantly from one individual to the next. By removing the anatomical shape differences between the ten individuals as described above, the composite PET image will give a large localized area of activity. In contrast, if these same ten anatomies are corrected only for global anatomical differences, the composite PET image will give a broad area of activity because the areas of activity will not align due to the local anatomical differences. Now suppose the following situation, assume that each of the ten individuals process data in sightly different locations and their anatomies again differ significantly. Processing the PET images with the conventional method will again give a broad area of activity; this is because the areas of activity will not be aligned because of both the local anatomical differences and the local activity differences. The composite PET image generated in

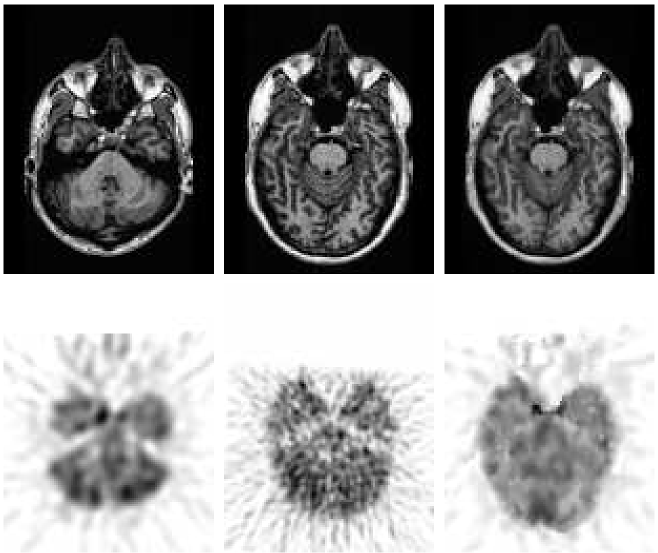

Figure 10.44: MR and PET modalities of slice 68 of the textbook (left), study (middle), and deformed textbook (right).

the coordinates of the textbook anatomy will also be broad, but it might have structure that is dependent on the textbook anatomy. In either case, the composite PET images generated with the deformable textbook should give researchers better information about how the brain works than do the conventional composite PET images that correct only for global shape differences.

11. Discussion

11.1 Viscous Fluid vs. Linear Elasticity

In previous work [41], linear-elasticity regularization was effectively used to enforce smoothness of the spatially sampled transformation. However, linear-elasticity penalizes large distance deformations thereby preventing the textbook from completely deforming into the study. As Figure (11.1) demonstrates, the viscous fluid regularization does not have this property. In this example a small wedge shaped patch is the template (top-left) and the study is a "C" (bottom-left). Shown in the top-right panel is the result of deforming the template into the study using the linear-elastic regularization. Notice that the elastic-solid regularization method prevented the template from deforming into the study. The bottom right panel shows the result of the viscous-fluid regularization which allows the template to more completely match the study. This shortcoming in the elastic map is due to the quadratic growth of the elasticity penalty and the inability of the linear-elastic theory to address large-magnitude deformations. The viscous-fluid regularization suffers from no such difficulties because the restoring forces relax over time, while at the same time accounting for the large-distance kinematic nonlinearities.

The advantage of the fluid model for deforming one anatomy into another is demonstrated in Fig.(11.2). The rows from top to bottom correspond to a fluid (top), a strong elastic (middle), and a weak elastic (bottom) transformation of the monkey visual cortex. Each row shows the deformed template (left), the difference image between the deformed template and the study (middle), and the singular Jacobian points of the deformation (right). Notice that the fluid transformation (top row) allowed the template to deform into the study as shown by the near black difference image while maintaining a positive Jacobian. Conversely, the difference image for the transformation with the large elastic penalty (middle row) shows that the template was prevented from fully deforming into the study. The all-black Jacobian image for this transformation confirms that the

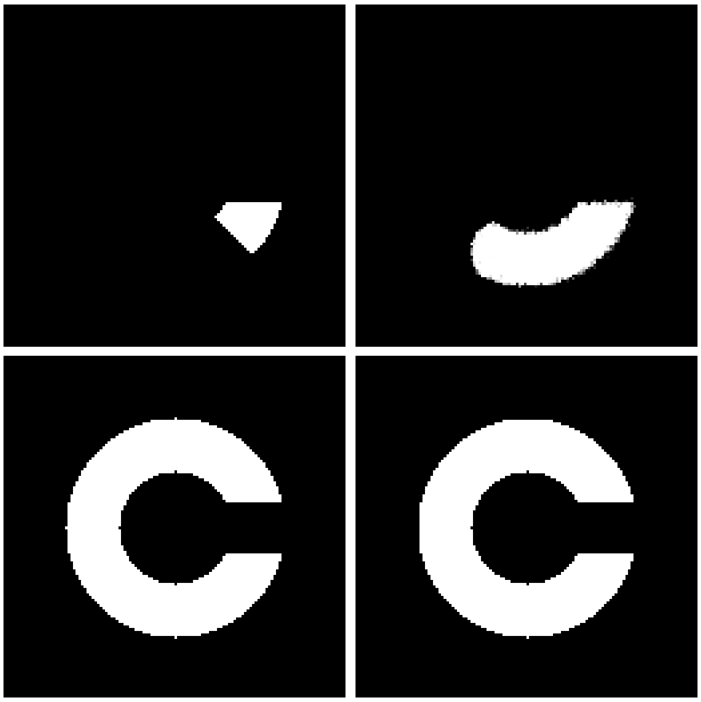

Figure 11.1: The left column shows the template (top) and the study (bottom). The template was deformed into the study using an elastic-solid regularization (top-right) and a viscous-fluid regularization (bottom-right).

fluid transformation also has a positive Jacobian. The difference image for the weak elastic penalty (bottom row) shows that the template was not held back as much as in the weak elastic case, but it was still held back more than for the fluid case. In addition, the white points in its Jacobian image show that the transformation went singular and does not reflect a 1-to-1 mapping of volumes, areas, or lines.

The linear-elasticity prior was derived assuming small angles of rotation and small linear deformations. Because it is assumed that the angles of rotation are small and the deformations are small and linear, large deformations cannot be accommodated with the linear PDE. This short coming of the linear-elastic theory could be removed by using a hyper-elastic energy form [61]. Accounting for the nonlinear kinematics would enforce a homeomorphism, but would continue to penalize large-magnitude deformations due to the elastic-strain energy.

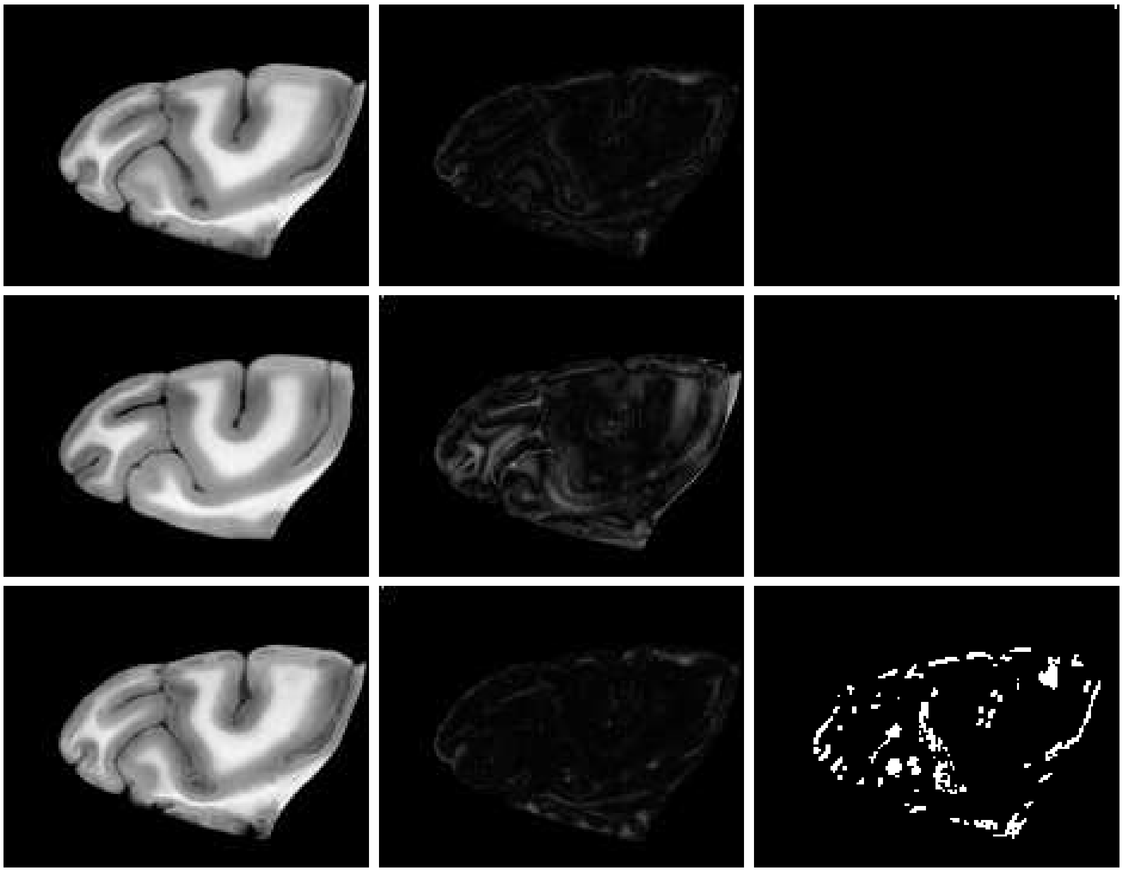

Figure 11.2: The left column shows a fluid (top), a strong-elastic (middle), and a weak-elastic (bottom) transformation. The middle column shows the difference images between the deformed template and the study. The right column shows the negative Jacobian points of the deformation (right). respectively. These data were provided courtesy of Van Essen et al.

11.2 Comparison with other Methods

Comparing our method to others, examine the choice of cost functions. We have chosen a quadratic cost $C(\vec{u})$ corresponding to a Gaussian sensor model of the form

$$C(\vec{u}) = \frac{1}{2}\int_\Omega T^2(\vec{x}-\vec{u}(\vec{x}))d\vec{x} - \int_\Omega T(\vec{x}-\vec{u}(\vec{x}))S(\vec{x})d\vec{x} + \frac{1}{2}\int_\Omega S^2(\vec{x})d\vec{x}. \tag{11.1}$$

Others [30] have investigated using the correlation which corresponds to the term $\int_\Omega T(\vec{x} - \vec{u}(\vec{x}))S(\vec{x})d\vec{x}$ above.

It is obvious that the quadratic and correlation functionals are different for non-constant-volume deformations because they differ[1] by the term $\int_\Omega T^2(\vec{x} - \vec{u}(\vec{x}))d\vec{x}$. However, it is not so obvious for constant-volume deformations in which the term $\int_\Omega T^2(\vec{x} - \vec{u}(\vec{x}))d\vec{x}$ is constant. To see the difference, compare the Fréchet derivative of the quadratic and correlation functionals which are given by

$$\vec{b}_q(\vec{x}, \vec{u}(\vec{x}, t)) = -\Big(T(\vec{x} - \vec{u}(\vec{x}, t)) - S(\vec{x})\Big)\vec{\nabla}T\Big|_{\vec{x}-\vec{u}(\vec{x},t)} \tag{11.2}$$

and

$$\vec{b}_c(\vec{x}, \vec{u}(\vec{x}, t)) = S(\vec{x})\vec{\nabla}T\Big|_{\vec{x}-\vec{u}(\vec{x},t)}, \tag{11.3}$$

respectively. The fact that these derivatives are different implies that these functionals do not have the same minimizing condition. In fact, the quadratic functional is minimized when $T(\vec{x} - \vec{u}(\vec{x})) = S(\vec{x})$ almost everywhere in Ω which can be seen by setting $\vec{b}_q$ to zero in Eq. (11.2). This condition is desirable because it means that the quadratic cost functional is minimized when the deformed template equals the data everywhere except on a set of zero measure. However, the correlation functional does not have an unique minimizer in general because neither $S(\vec{x})$ nor $\nabla T\Big|_{\vec{x}-\vec{u}(\vec{x},t)}$ is necessarily zero. For this reason, the correlation defines an ill-posed cost functional. The fact that the quadratic cost functional is well-posed allows us to have a Bayesian statistical interpretation whereas the correlation has no such interpretation.

Another difference between our work and that of others is that we evaluate the nonlinear deformation force given in Eq. (11.2) directly and do not approximate it. In contrast, others assume small deformations [62, 30] and linearize the body force using a Taylor-series approximation. But because the motivation of the work presented herein is to allow large-distance, nonlinear deformations, the small deformation assumption is not valid and linearization is not applicable.

[1]The term $\int_\Omega S^2(\vec{x})d\vec{x}$ is ignored because it is a constant for all transformations.

11.3 Timing Information

11.3.1 Introduction

The results shown in Chapter 10 were generated using a DECmpp 12000Sx/Model 200 (MasPar) massively parallel, mesh-connected, single-instruction-multiple-data (SIMD) computer. The particular computer used was configured in a 128×64 mesh. Specific details of the DECmpp implementation are given in Appendix C.

The special characteristics of the DECmpp that affected the performance of the deformable textbook algorithms were the 32-bit processors, the x-net, the global router, and the ability to perform parallel indirect addressing of memory. The fact that the DECmpp has 32-bit processors means that operations like multiplication and division are executed in $O(N)$ clock cycles where N refers to the number of bits in the operands. The previous generation of massively parallel machines like the GAPP [63] and the AMT DAP [63] were comprised of single-bit processors. For these machines, multiplications and divisions were very slow operations because they took $O(N^2)$ clock cycles. As a result of the 32-bit processors, the 128×64 DECmpp is rated at 3.2 GFLOPs, i.e., on the average, 3.2 billion floating-point operations are computed per second when only floating-point operations are performed.

The x-net was used to efficiently solve computations that required nearest-neighbor communications. These computations included solving the gradient of the posterior potential in the spatial linear-elastic model and computing the fluid PDE in the viscous fluid model.

The global router and the parallel indirect addressing were very important for the implementation of the deformation procedure used in the algorithms. These features allowed for the efficient parallel, nonlocal, random memory access of the textbook data which was needed to compute the deformation of the textbook at each iteration. Without these hardware features, their function would have had to be coded into the deformation procedure. This would have resulted in a very slow and cumbersome procedure because the procedure would need to perform a worst-case nonlocal, parallel search of the memory (constrained by the minimum and maximum displacements) to get the required textbook data values.

11.3.2 2D Execution Times

The 128×64 DECmpp12000Sx/Model 200 execution times for the 2D results shown in Chapter 10 are shown in Tables 11.1—11.3.

Table 11.1: Program-execution times for 2D linear-elastic-basis transformations.

Experiment	data size	# of param. estimated	# of itrs	# of itrs before increasing basis	total time (sec)
W to A	64×64	50	160	40	23
W to B	64×64	50	160	40	26
K to L	128×128	242	200	20	33

Table 11.2: Program-execution times for 2D spatial-linear-elastic transformations.

Experiment	data size	# of param. estimated	# of itrs	total time (sec)
W to A	256×256	1.3×10^5	250	139
W to B	256×256	1.3×10^5	250	148
K to L (strong)	320×256	1.6×10^5	250	186
K to L (weak)	320×256	1.6×10^5	250	347

Table 11.3: Program execution times for 2D viscous-fluid transformations.

Experiment	data size	# of param estimated	# of iters	# PDE itrs	# of regrids	total time (sec)
patch to C	128×128	6.6×10^4	2000	5000	60	3298
patch to C	128×128	6.6×10^4	600	5000	30	937
patch to C	128×128	6.6×10^4	600	250	33	423
circle to C	128×128	6.6×10^4	250	250	11	91
square to rect	128×128	6.6×10^4	180	250	7	45
W to A	256×256	2.6×10^5	150	250	10	107
W to B	256×256	2.6×10^5	150	250	10	110
K to L	320×256	3.3×10^5	200	250	12	756

The first column of Table 11.1 gives the name of the experiment. Experiment W to A and W to B corresponds to the elastic basis transformation of the Duke human MR data used in

Sections 10.1.1 and 10.2.2. Experiment K to L corresponds to the basis transformation used in Sections 10.1.2 and 10.2.3.

Column two shows the sizes of the data sets used in these experiments. The data sets were downsampled from their original size by factors of 4 and 6, respectively. This corresponded downsampling the MR data from 256×256 to 64×64 and the cryosectioned data from 640×480 to 107×80 and then zero padding to 128×128. The downsampled data were used in the linear-elastic-basis algorithm not because of the computation speed, but rather, because the basis model accommodated only global shape changes and the high-resolution data were not required.

The third column of Table 11.1 shows the number of parameters that were estimated for each experiment. The number of parameters for the MR experiment is determined from the fact that there are two components of the displacement field and there are two summations that were indexed over the zeroth through 4^{th} harmonics of the basis. This gives $2 \times 5^2 = 50$ parameters. Likewise, $2 \times 11^2 = 242$ parameters were estimated in the K to L experiment.

The entries in the fourth column of Table 11.1 are for the multiresolution implementation of the basis. On iteration zero, the basis consists of only the zeroth and first harmonics of the basis. After the prescribed number of iterations, the magnitude of the highest harmonic is increase by one. Thus on the 40^{th} iteration of the MR experiment, the number of basis harmonics are increased from two to three.

The fifth and sixth columns of Table 11.1 show the number of SDE iterations, and the total elapsed time of the program. The reason that the MR experiments seem to take as long as the larger visual-cortex experiment is because the size of the DECmpp mesh was reduced from 128×64 to 64×64 for computation of the later.

Table 11.2 shows the timing information for the 2D spatial-linear-elastic model corresponding to the results shown in Sections 10.1.1 and 10.1.2. The number of parameters corresponds to two components of the displacement field at each pixel in the study. Thus, the number of parameters estimated in the W to A experiment was $2 \times 256 \times 256 = 1.3 \times 10^5$.

Table 11.3 shows the timing information for the 2D viscous fluid model corresponding to the experiments in Section 10.2. The number of parameters estimated corresponds to two components of the displacement field and velocity field at each pixel in the study. Thus, there were $2 \times 2 \times 128 \times 128$ parameters estimated in the patch to the "C" experiment. Columns four and five show the number of real-time steps and the number of PDE iterations, respectively. The sixth column shows the total number of propagated templates that were used to generate the final transformation.

No significant difference was seen between the transformations of the patch to the "C" for the experiments shown in the top three rows of Table 11.3. This is because on real-time step zero, the fluid PDE requires thousands of PDE iterations to converge because the velocity field is initialized with zero. On subsequent real-time steps, the fluid PDE is initialized with the velocity field at the previous time step. Because the displacements are small from one real-time step to the next, the fluid PDE converges in only a few hundred PDE iterations. Thus in order to decrease the execution times, it was decided to used 250 PDE iterations for the other experiments.

11.3.3 3D Execution Times

The 128×64 DECmpp12000Sx/Model 200 execution times for the 3D results shown in Chapter 10 are given in Tables 11.4—11.6. These tables contain essentially the same information as their 2D counterparts in the previous section.

Table 11.4: 3D linear-elastic-basis program execution times.

Experiment	data size	# of param. estimated	# of itrs	# of itrs before increasing basis	total time (min)
S to D	$64 \times 64 \times 74$	192	100	40	11
D to S	$64 \times 64 \times 50$	192	100	40	7.8
FJ to P	$64 \times 64 \times 50$	192	100	40	7.9
P to D	$64 \times 64 \times 50$	192	100	40	7.1
P to FJ	$64 \times 64 \times 50$	192	100	40	7.4
P to S	$64 \times 64 \times 50$	192	100	40	7.7
K to L	$128 \times 128 \times 42$	192	100	40	17

Table 11.5: 3D spatial-linear-elastic program execution times.

Experiment	data size	# of param. estimated	# of itrs	total time
S to D	$128 \times 128 \times 148$	7.3×10^6	500	9.5 hours
K to L	$128 \times 128 \times 42$	2.1×10^6	100	31 min

Table 11.6: 3D viscous fluid program execution times.

Experiment	data size	# of param estimated	# of iters	# PDE itrs	# of regrids	total time (hours)
S to D	$128 \times 128 \times 148$	1.5×10^7	500	200	116	13
D to S	$128 \times 128 \times 100$	9.8×10^6	500	200	118	9.4
FJ to P	$128 \times 128 \times 100$	9.8×10^6	500	200	113	9.1
P to D	$128 \times 128 \times 100$	9.8×10^6	500	200	100	8.9
P to FJ	$128 \times 128 \times 100$	9.8×10^6	500	200	115	9.0
P to S	$128 \times 128 \times 100$	9.8×10^6	500	200	99	9.0
K to L	$128 \times 128 \times 42$	4.1×10^6	300	200	70	2.5

The number of parameters estimated for the 3D basis (see third row of Table 11.4) in all of the experiments was $3 \times 4 \times 4 \times 4 = 192$. There are three different eigenfunctions and for each eigenfunction there is a triple sum over the zeroth through third harmonics.

The 500 real-time iterations used in the six MR experiments shown in Table 11.6 was arrived at using the plot in Fig. 11.3. This plot shows the squared error or the negative of the Gaussian likelihood vs. iteration number for the P to FJ and FJ to P 3D fluid experiments. Notice that both of these curves have the same shape and seem to be leveling out as they approach the 500^{th} iteration. On the other hand, examination of the result at 300 iterations vs. that at 500 shows very little difference in the result. This result shows that more work will have to be done to determine the exact number of iterations required and how to increase the rate of convergence of the algorithms.

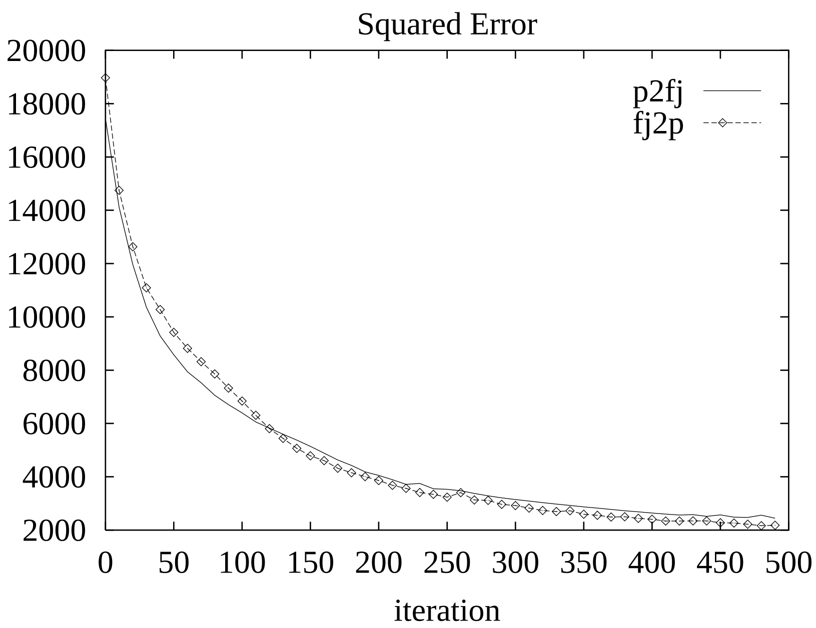

Figure 11.3: Plot of squared error of deformed textbook and study vs. iteration for 3D fluid deformation of P to FJ and FJ to P.

12. Conclusions

This dissertation presents a general approach based on Grenander's global shape models for representing the shape and variation of anatomical structures. The atlas registration methods developed in this work differ from existing methods in that shape variation is accommodated at a much finer detail than with the existing methods.

An individualized textbook is generated from the deformable textbook by deforming the textbook into the shape of the individual's anatomy. Information in the individualized textbook such as structure names, functions, volumes, etc. correspond to the particular subject's anatomy. The transformation that individualizes the textbook is useful for mapping functional information from one anatomy to another quantitatively.

Throughout this work it was assumed that the textbook and study anatomies were normal and topologically equivalent. However, it is clear that in the future these techniques will have to be extended to include abnormality due to disease. Eventually the deformable textbook may even be able to automatically identify abnormalities.

Three algorithms were developed for accommodating shape differences between topologically equivalent anatomies: the linear-elastic-basis method, the spatially parameterized linear-elastic method, and the spatially parameterized viscous-fluid model. The linear-elastic-basis method was shown to be useful for global nonrigid alignment, and the spatially parameterized algorithms were shown to be effective for local nonrigid alignment. A limitation of the linear-elasticity model is that it only ensures that the textbook transformations are homeomorphic for small deformations. In contrast, the viscous-fluid model does not assume small deformations. In fact, the viscous-fluid model ensures that the textbook transformations are homeomorphic even for long-distance, nonlinear deformations.

The two continuum mechanical models used in this work were chosen to impose desired topological properties on the textbook transformations. These models are good starting points, but in the future, new models will be needed that closely approximate the true physical laws of body tissues and fluids.

Additional work will be required to construct more complete deformable textbooks. Initiatives such as the *Visible Human* project sponsored by the National Medical Library will provide standard anatomies that can be used as the deformable textbook. Once collected, these data sets will have to be annotated by expert anatomists.

Also work needs to be done is to validate the individualized textbooks. This dissertation has presented results that show the feasibility of the deformable textbook idea. However, more studies are required to determine the repeatability and robustness of the algorithms. One specific test will be to determine how well the deformable textbook generates automatic segmentations. Another test will be to validate the registration of PET-to-PET when MR images registered with each PET image are registered to remove individual anatomical differences.

Another area of future work is to increase the convergence rate of the algorithms. The current implementations are very efficient, but they could be sped up by optimizing the step sizes and parameters used. The current implementation was not interested in the speed of the deformation other than to insure that the transformations were estimated in a reasonable time on the DECmpp. Before these techniques become widespread, it will be important to increase the rate of convergence of the transformation estimation. Faster convergence rates will allow the algorithms to be mapped onto work stations instead of supercomputers.

Finally, the deformable textbook offers the advantage that it is stored digitally and can be transmitted from one location to another using computer networks. This allows computation and visualization to be performed remotely. For example, a researcher does not necessarily have to have his/her own super-computer in order to compute the transformation that individualizes the textbook but rather can send their data over the network to a centralized computing resource which will generate the individualized textbook.

APPENDICES

A. Derivation of the Linear Elastic Model

The derivation of the linear-elastic model presented in this section has been modified for an Eulerian reference frame from the derivations presented in [53, 64]. It is presented here to expose the assumptions and approximations used in the derivation. The derivation begins by developing the general nonlinear form of the strain of a deformed differential element of material. The formulation of strain is then simplified by assuming small deformation and then linear deformation. Next, the formulation of stress of a differential element of material is discussed for linear strain. The generalized Hooke's law for elasticity is then presented which relates the strain to the stress of a deformed element of material. Then the assumption of an isotropic material is imposed to get the final stress-strain laws for an isotropic, linear-elastic material.

The final section of this appendix derives the 2D and 3D eigenvectors and eigenvalues for the linear-elasticity model.

A.1 Analysis of Strain in an Eulerian Reference Frame

The task in the analysis of strain is to chart the course of every material particle from its original to its final position [64]. Each particle can be identified on the basis of either its initial or final coordinates. When a particle is tracked with respect to its original coordinates, the reference frame is called *Lagrangian*; in contrast, when a particle is tracked with respect to its final coordinates, the reference frame is called *Eulerian*.

In this work, an Eulerian reference frame is used to track the deformation of the textbook. This reference frame was chosen because the deformation of the textbook is tracked as a volume.

In an Eulerian reference frame, the state of the deformed textbook is tracked with respect to a fixed voxel lattice. Therefore, there is no extra processing required to display the deformed textbook or to compare it to a study image. However, in a Lagrangian coordinate system, the trajectories of individual material particles of the textbook are tracked over time. As the textbook deforms the particles being tracked move to locations that do not necessarily correspond to pixel locations. Therefore an extra step is required to project the deformed continuum onto a voxel lattice for comparison with a study and for display.

Another reason that an Eulerian reference frame is used is that it is more natural for tracking large-distance, nonlinear deformations. This is because the hard problem of tracking the nonlinear trajectories of individual particles does not have to be performed. Instead, the easier task of observing the continuum as it deforms through fixed observation points is performed.

The quantity in continuum mechanics that describes how a body is deformed is called strain. The strain at P in the deformed body describes [64]:

1. the differences in length of all linear elements, dR, originating at a point P in the deformed structure and the corresponding elements, dr, originating at the point p in the undeformed structure, and

2. the differences between the angles, $\widehat{dR_1 dR_2}$, formed by all pairs of linear elements originating at point P and the angles, $\widehat{dr_1 dr_2}$ formed by the corresponding linear elements dr_1 and dr_2 originating at point p in the undeformed structure.

The quantities described in (1) are called the *extensional strains*, while those of (2) are called the *shear strains.*

Fig. A.1 shows a material body before and after deformation with respect to a single rectangular Cartesian coordinate frame x_1, x_2, and x_3. The convention of using lower case letters to describe quantities in the undeformed system and upper case letters for the deformed system will be used. It assume that a material point $p = [x_1, x_2, x_3]$ in the undeformed structure moves continuously to the point $P = [X_1, X_2, X_3]$ in the deformed structure, and that the final coordinates $\vec{X}$ must be continuous and differentiable functions of the initial coordinates $\vec{x}$. In addition, it is assumed that the applied forces deforming the continuum do not rupture the structure. Using these assumptions,

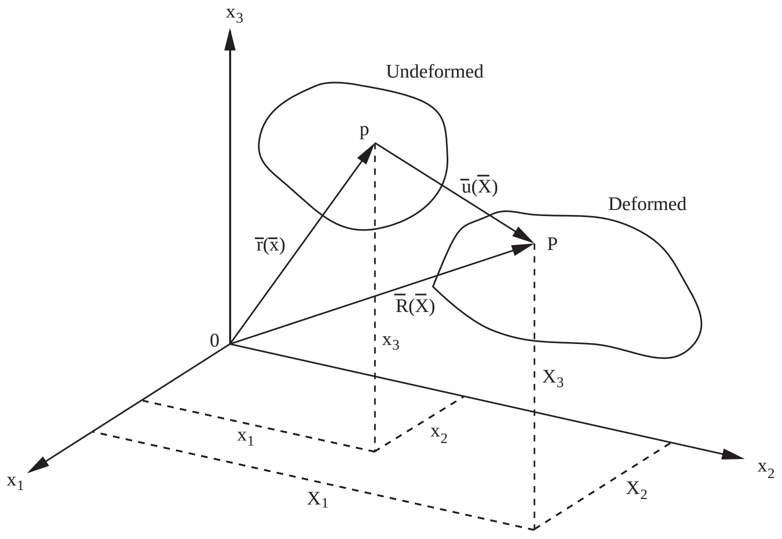

Figure A.1: Position vectors to undeformed and deformed body.

the Eulerian description of the deformed continuum is given by

$$\vec{x} = \vec{x}(\vec{X}) \ . \tag{A.1}$$

The position vector to p is given by

$$\vec{r} = \sum_{i=1}^{3} x_i \vec{e}_i = \sum_{i=1}^{3} x_i(\vec{X}) \vec{e}_i \tag{A.2}$$

and the position vector to P is given by

$$\vec{R} = \sum_{i=1}^{3} X_i \vec{e}_i \tag{A.3}$$

where $\vec{e}_i$ denotes the i^{th} unit vector.

The relationship between the material points p and q in the undeformed body to their final positions P and Q in the deformed body will now be examined with a Eulerian frame of reference. Fig. A.2 shows how the vector, $d\vec{R}$, which joins the two neighboring points P and Q in the deformed body can be enclosed by a rectangular parallelepiped with edges represented by

$$d\vec{R} = \sum_{i=1}^{3} dX_i \frac{\partial \vec{R}}{\partial X_i} = \sum_{i=1}^{3} dX_i \vec{e}_i \tag{A.4}$$

which is just the total derivative of Eq. (A.3). Note the change in perspective between the un-

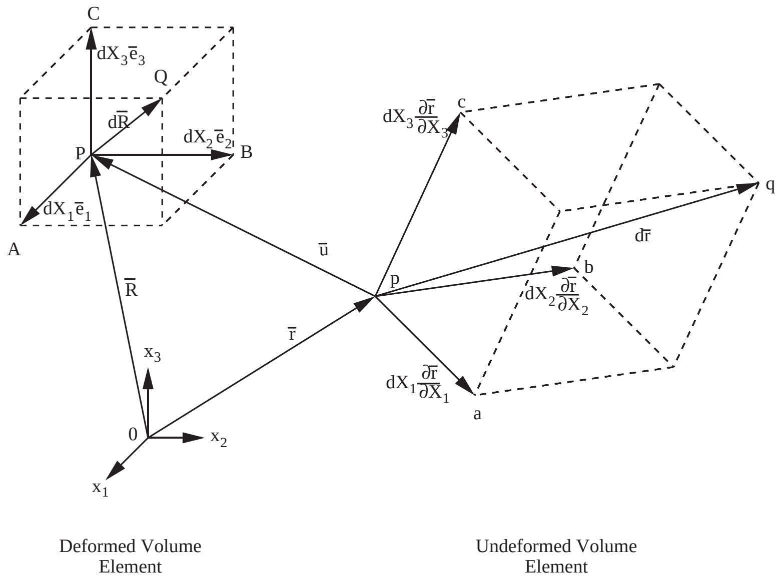

Figure A.2: Vector representations of edges of undeformed and deformed volume elements

deformed and deformed volumes in Figs. A.1 and A.2. Likewise, the vector $d\vec{r}$, which joins the points p and q in the undeformed body, is enclosed in a parallelepiped which is not rectangular. The base vectors of the parallelepiped *pabc* are given by

$$d\vec{r} = \sum_{i=1}^{3} dX_i \frac{\partial \vec{r}}{\partial X_i} \ . \tag{A.5}$$

Although it may seem backwards that the parallelepiped $PABC$ is rectangular while *pabc* is not, remember that a Eulerian reference gives the coordinates of the original material positions in terms of the final coordinates of the material particles. Thus, Fig. A.2 shows that a rectangular parallelepiped in the deformed continuum is the result of deforming a nonrectangular parallelepiped.

A.2 The Eulerian Strain Tensor

The Eulerian *strain tensor* η_{nm}, $n, m = 1, 2, 3$ is defined to be

$$\sum_{m,n=1}^{3} 2\eta_{mn} dX_m dX_n = (d\vec{R} \cdot d\vec{R}) - (d\vec{r} \cdot d\vec{r}) \tag{A.6}$$

where $d\vec{R} \cdot d\vec{R}$ and $d\vec{r} \cdot d\vec{r}$ are the square of the deformed and undeformed line elements, respectively. Solving Eq. (A.6) for η_{nm} and using Eqs. (A.4) and (A.5) gives

$$\eta_{mn} = \frac{1}{2}(H_{mn} - h_{mn}) \tag{A.7}$$

where

$$H_{mn} = \frac{\partial \vec{R}}{\partial X_m} \cdot \frac{\partial \vec{R}}{\partial X_n} = \delta_{mn} \qquad \text{and} \qquad h_{mn} = \frac{\partial \vec{r}}{\partial X_m} \cdot \frac{\partial \vec{r}}{\partial X_n} \; . \tag{A.8}$$

and δ_{mn} is the Kronecker delta function. Note that because of the symmetry of H_{mn} and h_{mn}, the Eulerian strain tensor is symmetric, i.e., $\eta_{mn} = \eta_{nm}$.

It will now be shown that the Eulerian strain tensor defined by Eqs. (A.6) and (A.7) contains all of the information required to describe the state of strain at point P in the deformed body. The relative elongation of the edge pa into edge PA using a Eulerian reference is defined by

$$E_1 = \frac{\text{change in length of } pa}{\text{final length of } pa} = \frac{|PA| - |pa|}{|PA|} = \frac{dR_1 - dr_1}{dR_1} = 1 - \sqrt{1 - 2\eta_{11}}. \tag{A.9}$$

Similarly, the relative elongations of the other two edges are given by

$$E_2 = 1 - \frac{dr_2}{dR_2} = 1 - \sqrt{1 - 2\eta_{22}} \tag{A.10}$$

$$E_3 = 1 - \frac{dr_3}{dR_3} = 1 - \sqrt{1 - 2\eta_{33}} \; . \tag{A.11}$$

Thus, the relative changes in length E_1, E_2, and E_3 are related to the strain components η_{11}, η_{22}, and η_{33} and are referred to as the *extensional components* of strain.

Next, it will be shown that the *shearing components* of strain η_{12}, η_{13}, and η_{23} as well as the extensional components of strain are related to the angular changes between adjacent edges of the rectangular parallelepiped. The angle between edges pa and pb is defined to be $(\pi/2 - \gamma_{12})$. In addition γ_{12} is also the change in the angle between apb and APB because the angle between edges PA and PB is a right angle. Using the definition of the dot product gives

$$[\frac{\partial \vec{r}}{\partial X_1} dX_1] \cdot [\frac{\partial \vec{r}}{\partial X_2} dX_2] = dr_1 dr_2 \cos(\frac{\pi}{2} - \gamma_{12}) \tag{A.12}$$

which can be rewritten using Eqs. (A.7)-(A.10) as

$$\sin \gamma_{12} = \frac{2\eta_{12}}{(1-E_1)(1-E_2)} \ . \tag{A.13}$$

Similarly, the angles bpc and cpa in the undeformed body that are deformed into the right angles BPC and CPA, respectively, are given by

$$\sin \gamma_{13} = \frac{2\eta_{13}}{(1-E_1)(1-E_3)} \tag{A.14}$$

$$\sin \gamma_{23} = \frac{2\eta_{23}}{(1-E_2)(1-E_3)} \ . \tag{A.15}$$

Thus, Eqs. (A.13), (A.14), and (A.15) illustrate how the shearing and extensional components of strain are related to the changes in angles between adjacent edges in the undeformed body parallelepiped.

A.3 Eulerian Strain-displacement Relations

The strain-displacement relations will now be developed. Fig. A.1 shows that the displacement vector $\vec{u}$ connects the original position p of a material particle in the undeformed body to its final position in the deformed body P. Taking the Eulerian view, $\vec{u}$ is the function that gives the original position p of a material particle as a function of the final position P, i.e.,

$$\vec{r}(\vec{X}) = \vec{R}(\vec{X}) - \vec{u}(\vec{X}) \ . \tag{A.16}$$

Substituting Eq. (A.16) into (A.7) gives the desired strain-displacement relations

$$\eta_{mn} = \frac{1}{2}\left(\frac{\partial u_m}{\partial X_n} + \frac{\partial u_n}{\partial X_m} - \sum_{r=1}^{3} \frac{\partial u_r}{\partial X_m}\frac{\partial u_r}{\partial X_n}\right) \quad . \tag{A.17}$$

Notice that Eulerian strain tensor is composed of a part which is linear in the derivatives of the displacements and a part due to the squares and products of the derivatives which is nonlinear.

Eq. (A.17) is linearized by assuming that the squares and products of the derivatives are negligible with respect to the first derivative terms. Because the linear approximation of strain requires modification of the nonlinear strain-displacement Eq. (A.17), it is customary to represent the linear strain tensor by the new symbol ϵ_{mn}, i.e.,

$$\epsilon_{mn} = \frac{1}{2}\left(\frac{\partial u_m}{\partial X_n} + \frac{\partial u_n}{\partial X_m}\right) \quad . \tag{A.18}$$

The geometric interpretation of linear strain is that there is a small deformation and the angles of rotation shown in Fig. A.3 are small. The vector $\frac{\partial \vec{r}}{\partial X_1}$ in this figure corresponds to the edge pa in

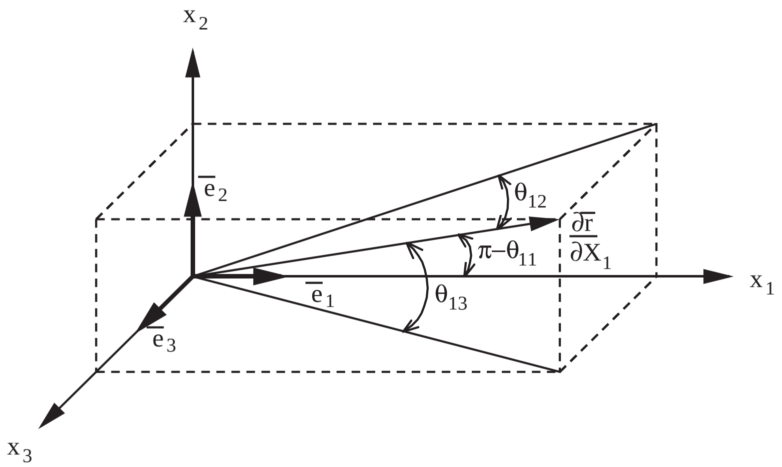

Figure A.3: Angles of rotation of $\partial\vec{r}/\partial X_1$.

Fig. A.2 translated to the origin. The angles θ_{11}, θ_{12}, and θ_{13} represent the rotation of vector $\frac{\partial \vec{r}}{\partial X_1}$ out of the planes located by $\vec{e}_1$, $\vec{e}_2$, and $\vec{e}_3$, respectively. Using the definition of the dot product, gives

$$\frac{\partial \vec{r}}{\partial X_1} \cdot \vec{e}_1 = \left|\frac{\partial \vec{r}}{\partial X_1}\right| \cos\left(\frac{\pi}{2} - \theta_{11}\right) \tag{A.19}$$

$$\frac{\partial \vec{r}}{\partial X_1} \cdot \vec{e}_2 = \left|\frac{\partial \vec{r}}{\partial X_1}\right| \cos\left(\frac{\pi}{2} - \theta_{12}\right) \tag{A.20}$$

$$\frac{\partial \vec{r}}{\partial X_1} \cdot \vec{e}_3 = \left|\frac{\partial \vec{r}}{\partial X_1}\right| \cos\left(\frac{\pi}{2} - \theta_{13}\right) \tag{A.21}$$

which can be rewritten with the aid of Eqs. (A.9) and (A.16) to give

$$\sin\theta_{11} = \frac{1}{1-E_1}\left(1 - \frac{\partial u_1}{\partial X_1}\right) \tag{A.22}$$

$$\sin\theta_{12} = -\frac{1}{1-E_1}\frac{\partial u_2}{\partial X_1} \tag{A.23}$$

$$\sin\theta_{13} = -\frac{1}{1-E_1}\frac{\partial u_3}{\partial X_1}\ . \tag{A.24}$$

Under the assumption of a small deformation, Eqs. (A.22)-(A.24) imply that

$$\left|\frac{\partial u_1}{\partial X_1}\right| << 1, \quad \theta_{12} \approx \frac{\partial u_2}{\partial X_1}, \quad \text{and} \quad \theta_{13} \approx \frac{\partial u_3}{\partial X_1}. \tag{A.25}$$

The small-deformation approximation assumes that the magnitude of the elongation E_1 is small, θ_{11} is close to $\frac{\pi}{2}$, and the angles θ_{12} and θ_{13} are small. Likewise, we can derive

$$\left|\frac{\partial u_m}{\partial X_m}\right| << 1 \quad \text{and} \quad \theta_{mn} \approx \frac{\partial u_n}{\partial X_m}\ for\ m \neq n\ , \tag{A.26}$$

for a small deformation. When the rotation angles θ_{mn} are sufficiently small then the squares and products of $\frac{\partial u_n}{\partial X_m}$ are negligible with respect to $\frac{\partial u_n}{\partial X_m}$ itself. Thus using Eq. (A.26), we see that Eq. (A.18) is approximately equal to Eq. (A.17) under the assumptions of small strain and small rotation angles.

A.4 Analysis of Stress in an Eulerian Reference Frame

The analysis of stress is concerned with a study of the cohesive forces, those which hold the structure together as an integral unit [64]. The left panel of Fig. A.4 shows a body that has been cut by an imaginary plane perpendicular to $\vec{n}$. The internal forces that hold the left half of the body together with the right half along this imaginary cut have equal magnitudes but opposite orientations. The internal force of the left half of the body acting on the differential element of area ΔA is denoted by $\Delta\vec{F}$. Stress at point p with respect to the differential element of area ΔA containing p is defined as the limit of the internal force acting on ΔA divided by ΔA as $\Delta A \to 0$, i.e.,

$$\vec{\sigma} = \lim_{\Delta A \to 0} \left(\frac{\Delta\vec{F}}{\Delta A} \right) = \frac{d\vec{F}}{dA} \ . \tag{A.27}$$

The vector denoted by $\vec{\sigma}$ is called the stress vector at p with respect to the an infinitesimal element of area dA which has orientation specified by $\vec{n}$ and is illustrated in the right panel of Fig. A.4. Note that $\vec{\sigma}$ is not in general perpendicular to the cut face and that it has dimensions of force per unit area.

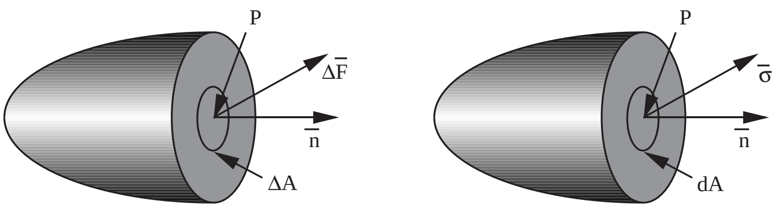

Figure A.4: Derivation of stress vector acting on dA at point P.

To simplify the analysis of strain, we make the assumption of linear strain. This assumption avoids the complicated analysis of stress of deformed bodies which entails a complicated interaction between stress and strain. The assumption of linear strain implies that the edges of the deformed element of volume under went negligible rotations from their corresponding positions in the undeformed body. Additionally, the length of the edges of the deformed element of volume only differ infinitesimally from their original length in the undeformed body. Thus, for the analysis of stress, the deformed element of volume is indistinguishable from the undeformed element of volume because the surface areas of the deformed element of volume differ infinitesimally from the surface areas in the undeformed volume.

Under the assumption of linear strain, Fig. A.5 shows the definitions of the stress vectors acting on an infinitesimally small element of volume which is a rectangular parallelepiped. The stress vectors are written in their component form as

$$\vec{\sigma}_1 = \sigma_{11}\vec{e}_1 + \sigma_{12}\vec{e}_2 + \sigma_{13}\vec{e}_3$$

$$\vec{\sigma}_2 = \sigma_{21}\vec{e}_1 + \sigma_{22}\vec{e}_2 + \sigma_{23}\vec{e}_3$$

$$\vec{\sigma}_3 = \sigma_{31}\vec{e}_1 + \sigma_{32}\vec{e}_2 + \sigma_{33}\vec{e}_3 \ . \tag{A.28}$$

The components of the stress vectors that act normal to the surfaces of the parallelepiped are called the *tensile components of stress.* And, the components of the stress vector which lie in the surfaces of the parallelepiped are called the *shear components of stress.* We also note that the components of stress vectors are symmetric, i.e., $\sigma_{mn} = \sigma_{nm}$.

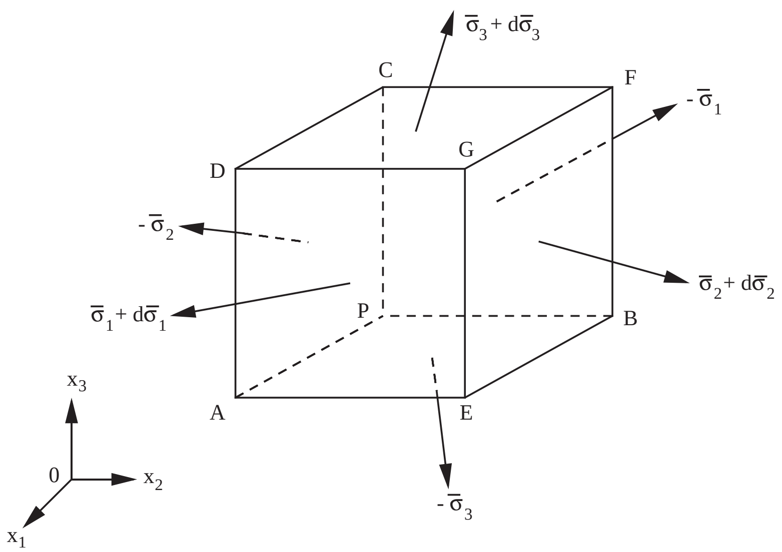

Figure A.5: Stress vectors acting on a differential volume element.

The partial differential equations assuming linear strain will now be derived by balancing the forces in Fig. A.5 with a distributed body force. The force equilibrium equation for the stresses in

Fig. A.5 and a body force $\vec{b}$ is

$$dA_1 d\vec{\sigma}_1 + dA_2 d\vec{\sigma}_2 + dA_3 d\vec{\sigma}_3 + \vec{b} dV = 0 \ . \tag{A.29}$$

The assumption of linear strain implies that

$$d\vec{\sigma}_n = \frac{\partial \vec{\sigma}_n}{\partial X_n} dX_n \tag{A.30}$$

and the force equilibrium equation becomes

$$\left(\frac{\partial \vec{\sigma}_1}{\partial X_1} + \frac{\partial \vec{\sigma}_2}{\partial X_2} + \frac{\partial \vec{\sigma}_3}{\partial X_3} + \vec{b} \right) dX_1 \, dX_2 \, dX_3 = \vec{0} \ . \tag{A.31}$$

Canceling the differential volume element $dX_1 \, dX_2 \, dX_3$ and writing each component as a separate equation gives the force balance partial differential equations assuming linear strain

$$\begin{aligned} \frac{\partial \sigma_{11}}{\partial X_1} + \frac{\partial \sigma_{21}}{\partial X_2} + \frac{\partial \sigma_{31}}{\partial X_3} + b_1 = 0, \\ \frac{\partial \sigma_{12}}{\partial X_1} + \frac{\partial \sigma_{22}}{\partial X_2} + \frac{\partial \sigma_{32}}{\partial X_3} + b_2 = 0, \\ \frac{\partial \sigma_{13}}{\partial X_1} + \frac{\partial \sigma_{23}}{\partial X_2} + \frac{\partial \sigma_{33}}{\partial X_3} + b_3 = 0. \end{aligned} \tag{A.32}$$

A.5 Linear Elasticity

The most general linear relationship which connects stress to strain is known as the generalized Hooke's Law [64] and is given mathematically by

$$\sigma_{mn} = \sum_{r,s=1}^{3} E_{mnrs} \epsilon_{rs} \ . \tag{A.33}$$

In words, this equation states that each component of the Eulerian strain tensor is linearly related to each component of the Eulerian stress tensor. Making the assumption that the material is isotropic—it possesses elastic properties independent of the orientation of the axis—Eq. (A.33) can be simplified from 81 unknown coefficients to 2. Assuming the material to be elastically isotropic, the elastic constants must be the same at a given point for all possible choices of rectangular

Cartesian coordinates used to evaluate σ_{mn} and ϵ_{rs}. Thus from [53], we can write the isotropic assumption mathematically as as

$$E_{mnrs} = \lambda\delta_{mn}\delta_{rs} + \mu(\delta_{mr}\delta_{ns} + \delta_{ms}\delta_{nr}) + \beta(\delta_{mr}\delta_{ns} - \delta_{ms}\delta_{nr}) \; . \tag{A.34}$$

Next, the observation that σ_{mn} and ϵ_{rs} are both symmetric which implies that $E_{mnrs} = E_{nmrs}$ and $E_{mnrs} = E_{mnsr}$, respectively. Combining the symmetry of E_{mnrs} with the isotropic assumption implies that $\beta = 0$ and

$$E_{mnrs} = \lambda\delta_{mn}\delta_{rs} + \mu(\delta_{mr}\delta_{ns} + \delta_{ms}\delta_{nr}) \; . \tag{A.35}$$

The pair of constants λ and μ are called *Lamé's constants*, and μ is often referred to as the *shear modulus*. Substituting Eq. (A.35) into Eq. (A.33) gives the isotropic material, stress-strain relations

$$\sigma_{mn} = 2\mu\epsilon_{mn} + \lambda\delta_{mn}\sum_{r=1}^{3}\epsilon_{rr} \; . \tag{A.36}$$

Alternatively, the above relations can be written as

$$\epsilon_{mn} = \frac{1}{E}\left[(1+\nu)\sigma_{mn} - \nu\delta_{mn}\sum_{r=1}^{3}\sigma_{rr}\right] \; , \tag{A.37}$$

where

$$\nu = \frac{\lambda}{2(\lambda+\mu)} \quad \text{and} \quad E = \frac{\mu(3\lambda+2\mu)}{\lambda+\mu} \; . \tag{A.38}$$

The constant ν is called *Poisson's ratio* and E is called *Young's modulus* or the *modulus of elasticity*. Poisson's ratio is the ratio between the lateral contraction and the axial elongation under a uniaxial stress condition.

The PDE for linear elasticity is derived by substituting the linear strain-displacement Eq. (A.17) and the stress-strain Eq. (A.36) into the force balance Eq. (A.32). The substitution of the the linear strain-displacement Eq. (A.17) into the force balance Eq. (A.32) gives

$$
\begin{aligned}
&\frac{\partial}{\partial X_1}(2\mu\epsilon_{11} + \lambda(\epsilon_{11} + \epsilon_{22} + \epsilon_{33})) + 2\mu\frac{\partial \epsilon_{21}}{\partial X_2} + 2\mu\frac{\partial \epsilon_{31}}{\partial X_3} + b_1 = 0,\\
&2\mu\frac{\partial \epsilon_{12}}{\partial X_1} + \frac{\partial}{\partial X_2}(2\mu\epsilon_{22} + \lambda(\epsilon_{11} + \epsilon_{22} + \epsilon_{33})) + 2\mu\frac{\partial \epsilon_{32}}{\partial X_3} + b_2 = 0,\\
&2\mu\frac{\partial \epsilon_{13}}{\partial X_1} + 2\mu\frac{\partial \epsilon_{23}}{\partial X_2} + \frac{\partial}{\partial X_3}(2\mu\epsilon_{33} + \lambda(\epsilon_{11} + \epsilon_{22} + \epsilon_{33})) + b_3 = 0.
\end{aligned}
\tag{A.39}
$$

This last equation can be rewritten using Eq. (A.36) to give the linear-elastic PDE

$$
\mu\nabla^2\vec{u}(\vec{X}) + (\mu + \lambda)\vec{\nabla}(\vec{\nabla}^T\vec{u}(\vec{X}) + \vec{b} = 0 \tag{A.40}
$$

where it is assumed that the Laplacian ∇^2 and the gradient $\vec{\nabla}$ are with respect to the Eulerian coordinate system $\vec{X}$.

A.6 Strain Energy

The forces that are applied to a body to deform it perform work on the structure. Assuming that the change in temperature of the deformation can be ignored, the First Law of Thermodynamics states that the work done by the external forces results in an increase in the internal energy of the structure [64]. In the theory of elasticity, the increase in the internal energy is called the *strain energy* and is given by

$$
J = \int_\Omega J^* dX_1 dX_2 dX_3 = \int_\Omega \frac{1}{2}\sum_{m,n=1}^{3} \sigma_{mn}\epsilon_{mn} dX_1 dX_2 dX_3. \tag{A.41}
$$

Using Eqs. (A.36) and (A.37), the strain-energy density can be written as

$$
J^* = \frac{1}{2}\lambda\left(\sum_{r=1}^{3}\epsilon_{rr}\right)^2 + \mu\sum_{m,n=1}^{3}(\epsilon_{mn})^2 \tag{A.42}
$$

and

$$
J^* = \frac{1}{2E}\left[(1+\nu)\sum_{m,n=1}^{3}(\sigma_{mn})^2 - \nu\left(\sum_{r=1}^{3}\sigma_{rr}\right)^2\right], \tag{A.43}
$$

respectively. In addition, the strain-energy density can be rewritten in terms of the displacement field using Eq. (A.18) to give

$$J^* = \frac{\lambda}{2}\left(\sum_{r=1}^{3}\frac{\partial u_r}{\partial X_r}\right)^2 + \frac{\mu}{4}\sum_{m,n=1}^{3}\left(\frac{\partial u_m}{\partial X_n}+\frac{\partial u_n}{\partial X_m}\right)^2. \tag{A.44}$$

At this point it is convenient to point out that the functional $\vec{u}$ that minimizes Eq. (A.44) is the linear-elastic PDE given in Eq. (A.40). To see this we use the theory of calculus of variation which states that the minimizer of

$$J = \int J^*(X_1, X_2, X_3, u_{1,1}, u_{1,2}, u_{1,3}, u_{2,1}, u_{2,2}, u_{2,3}, u_{3,1}, u_{3,2}, u_{3,3})dX_1 dX_2 dX_3 \tag{A.45}$$

where $u_{i,j} = \frac{\partial u_i}{\partial X_j}$ satisfies the Euler-Lagrange equations

$$J^*_{u_1} - \left(\frac{\partial}{\partial X_1}J^*_{u_{1,1}} + \frac{\partial}{\partial X_2}J^*_{u_{1,2}} + \frac{\partial}{\partial X_3}J^*_{u_{1,3}}\right) = 0 \tag{A.46}$$

$$J^*_{u_2} - \left(\frac{\partial}{\partial X_1}J^*_{u_{2,1}} + \frac{\partial}{\partial X_2}J^*_{u_{2,2}} + \frac{\partial}{\partial X_3}J^*_{u_{2,3}}\right) = 0 \tag{A.47}$$

$$J^*_{u_3} - \left(\frac{\partial}{\partial X_1}J^*_{u_{3,1}} + \frac{\partial}{\partial X_2}J^*_{u_{3,2}} + \frac{\partial}{\partial X_3}J^*_{u_{3,3}}\right) = 0. \tag{A.48}$$

where $J^*_{u_i} = \frac{\partial J^*}{\partial u_i}$ and $J^*_{u_{i,j}} = \frac{\partial J^*}{\partial u_{i,j}}$. Associating J^* in Eq. (A.45) with J^* in Eq. (A.44) we get the equations

$$\begin{aligned}
J^*_{u_1} &= 0\\
J^*_{u_{1,1}} &= (2\mu+\lambda)\left(\frac{\partial u_1}{\partial X_1}\right) + \lambda\left(\frac{\partial u_2}{\partial X_2}\right) + \lambda\left(\frac{\partial u_3}{\partial X_3}\right)\\
J^*_{u_{1,2}} &= \mu\left(\frac{\partial u_1}{\partial X_2}+\frac{\partial u_2}{\partial X_1}\right)\\
J^*_{u_{1,3}} &= \mu\left(\frac{\partial u_1}{\partial X_3}+\frac{\partial u_3}{\partial X_1}\right).
\end{aligned} \tag{A.49}$$

Substituting these equations into Eq. (A.46) gives

$$\mu\left(\frac{\partial^2 u_1}{\partial X_1^2}+\frac{\partial^2 u_1}{\partial X_2^2}+\frac{\partial^2 u_1}{\partial X_3^2}\right) + (\mu+\lambda)\left(\frac{\partial^2 u_1}{\partial X_1^2}+\frac{\partial^2 u_2}{\partial X_1\partial X_2}+\frac{\partial^2 u_3}{\partial X_1\partial X_3}\right) = 0. \tag{A.50}$$

Similarly the other two Euler-Lagrange equations are given by

(A.51)
$$\mu\left(\frac{\partial^2 u_2}{\partial X_1^2}+\frac{\partial^2 u_2}{\partial X_2^2}+\frac{\partial^2 u_2}{\partial X_3^2}\right)+(\mu+\lambda)\left(\frac{\partial^2 u_1}{\partial X_1\partial X_2}+\frac{\partial^2 u_2}{\partial X_2^2}+\frac{\partial^2 u_3}{\partial X_2\partial X_3}\right)=0$$

(A.52)
$$\mu\left(\frac{\partial^2 u_3}{\partial X_1^2}+\frac{\partial^2 u_3}{\partial X_2^2}+\frac{\partial^2 u_3}{\partial X_3^2}\right)+(\mu+\lambda)\left(\frac{\partial^2 u_1}{\partial X_1\partial X_3}+\frac{\partial^2 u_2}{\partial X_2\partial X_3}+\frac{\partial^2 u_3}{\partial X_3^2}\right)=0.$$

These Euler-Lagrange equations are identical to those in Eq. (A.40) and thus the minimizer of the strain energy is the linear-elastic PDE.

A.7 Converting the Linear Elasticity Model from 3D to 2D

The assumptions that were made to convert the 3D linear-elastic model from 3D to 2D are:

1. the extensional stress σ_{33} is zero, and
2. the shearing stress σ_{13} and σ_{23} are zero.

Substituting assumptions 1 and 2 into the force balance Eq. (A.32) and using the linear strain-displacement Eq. (A.17) gives

$$\begin{aligned}&\frac{\partial}{\partial X_1}(2\mu\epsilon_{11}+\lambda(\epsilon_{11}+\epsilon_{22}))+2\mu\frac{\partial\epsilon_{21}}{\partial X_2}+b_1=0,\\&2\mu\frac{\partial\epsilon_{12}}{\partial X_1}+\frac{\partial}{\partial X_2}(2\mu\epsilon_{22}+\lambda(\epsilon_{11}+\epsilon_{22}))+b_2=0.\end{aligned}\tag{A.53}$$

This last equation can be rewritten using Eq. (A.36) to give the 2D linear-elastic PDE

$$\mu\nabla^2\vec{u}(\vec{X})+(\mu+\lambda)\vec{\nabla}(\vec{\nabla}^T\vec{u}(\vec{X}))+\vec{b}=0\tag{A.54}$$

where it is assumed that the Laplacian ∇^2 and the gradient $\vec{\nabla}$ are with respect to the 2D Eulerian coordinate system $\vec{X}=(X_1,X_2)$.

Because of the similarity of the 2D and 3D linear-elastic PDEs, it is tempting to assume that the 3D strain energy density Eq. (A.44) can be reduced from 2D to 3D by reducing the number of summands from three to two. But this is not the case as will now be shown. Using assumptions 1 and 2 in Eqs. (A.17) and (A.36) gives

$$\sigma_{13} = 2\mu\epsilon_{13} = \frac{1}{2}\left(\frac{\partial u_1}{\partial X_3} + \frac{\partial u_3}{\partial X_1}\right) = 0$$

$$\sigma_{23} = 2\mu\epsilon_{23} = \frac{1}{2}\left(\frac{\partial u_2}{\partial X_3} + \frac{\partial u_3}{\partial X_2}\right) = 0$$

$$\sigma_{33} = \lambda\epsilon_{11} + \lambda\epsilon_{22} + (2\mu+\lambda)\epsilon_{33} = \lambda\frac{\partial u_1}{\partial X_1} + \lambda\frac{\partial u_2}{\partial X_2} + (2\mu+\lambda)\frac{\partial u_3}{\partial X_3} = 0 \tag{A.55}$$

which can be rewritten as

$$\frac{\partial u_1}{\partial X_3} = -\frac{\partial u_3}{\partial X_1} \tag{A.56}$$

$$\frac{\partial u_2}{\partial X_3} = -\frac{\partial u_3}{\partial X_2} \tag{A.57}$$

$$\frac{\partial u_3}{\partial X_3} = -\frac{\lambda}{\lambda+2\mu}\left(\frac{\partial u_1}{\partial X_1} + \frac{\partial u_2}{\partial X_2}\right). \tag{A.58}$$

Substitution of these equations into Eq. (A.44) gives the 2D linear-elastic potential

$$J^* = \frac{1}{2}\left[\frac{4\mu(\lambda+\mu)}{\lambda+2\mu}\left\{\left(\frac{\partial u_1}{\partial X_1}\right)^2 + \left(\frac{\partial u_2}{\partial X_2}\right)^2 + \frac{\lambda}{\lambda+\mu}\left(\frac{\partial u_1}{\partial X_1}\right)\left(\frac{\partial u_2}{\partial X_2}\right)\right\} + \mu\left(\frac{\partial u_1}{\partial X_2} + \frac{\partial u_1}{\partial X_2}\right)^2\right] \tag{A.59}$$

in terms of the Lamé constants or

$$J^* = \frac{1}{2}\left[\frac{E}{1-\nu^2}\left\{\left(\frac{\partial u_1}{\partial X_1}\right)^2 + \left(\frac{\partial u_2}{\partial X_2}\right)^2 + 2\nu\left(\frac{\partial u_1}{\partial X_1}\right)\left(\frac{\partial u_2}{\partial X_2}\right)\right\} + \frac{E}{2(1+\nu)}\left(\frac{\partial u_1}{\partial X_2} + \frac{\partial u_1}{\partial X_2}\right)^2\right]. \tag{A.60}$$

Notice that the 2D linear-elastic potential given by Eq. (A.59) is different from the potential that would be generated by reducing each of the sums in Eq. (A.44) from three summands to two. As an aside, it can be shown using the calculus of variation that the minimizer of the 2D potential in Eq. (A.59) is given by the 2D linear-elastic PDE Eq. (A.54).

A.8 Eigenvectors of the Linear Elasticity PDE

The eigenvectors and eigenvalues of the linear-elastic PDE given in Eq. (A.40) will be derived in this section in 3D and 2D.

The eigenfunctions $\phi(\vec{X})$ of a linear operator L are the functions which satisfy

$$L\phi(\vec{X}) = \kappa\phi(\vec{X}). \tag{A.61}$$

The scalar κ is called the eigenvalue associated with the eigenfunction $\phi(\vec{X})$.

The linear operator L associated with linear elasticity is given by

$$Lu = \mu\nabla^2\vec{u}(\vec{X}) + (\mu+\lambda)\vec{\nabla}(\vec{\nabla}^T\vec{u}(\vec{X})). \tag{A.62}$$

We assume that the vector $\vec{u}$ have domain $\Omega = [0,1]^3$ and have the Dirichlet boundary conditions

$$\begin{aligned} u_1(0,X_2,X_3) &= u_1(1,X_2,X_3) = 0 \\ u_2(X_1,0,X_3) &= u_2(X_1,1,X_3) = 0 \\ u_3(X_1,X_2,0) &= u_3(X_1,X_2,1) = 0 \end{aligned} \tag{A.63}$$

and the Neumann boundary conditions

$$\begin{aligned} \frac{\partial u_1}{\partial X_2}\Big|_{(X_1,0,X_3)} &= \frac{\partial u_1}{\partial X_2}\Big|_{(X_1,1,X_3)} = \frac{\partial u_1}{\partial X_3}\Big|_{(X_1,X_2,0)} = \frac{\partial u_1}{\partial X_3}\Big|_{(X_1,X_2,1)} = 0 \\ \frac{\partial u_2}{\partial X_1}\Big|_{(0,X_2,X_3)} &= \frac{\partial u_2}{\partial X_1}\Big|_{(1,X_2,X_3)} = \frac{\partial u_2}{\partial X_3}\Big|_{(X_1,X_2,0)} = \frac{\partial u_2}{\partial X_3}\Big|_{(X_1,X_2,1)} = 0 \\ \frac{\partial u_3}{\partial X_1}\Big|_{(0,X_2,X_3)} &= \frac{\partial u_3}{\partial X_1}\Big|_{(1,X_2,X_3)} = \frac{\partial u_3}{\partial X_2}\Big|_{(X_1,0,X_3)} = \frac{\partial u_3}{\partial X_2}\Big|_{(X_1,1,X_3)} = 0. \end{aligned} \tag{A.64}$$

Notice that for this choice of boundary conditions, the transformation $\vec{h}(\vec{X}) = \vec{X} - \vec{u}(\vec{x})$ maps the boundary of Ω back onto itself. The boundary conditions in Eqs. (A.63) and (A.64) are called *sliding* boundary conditions because points located on the boundary of the cube can slide along the side of the cube.

The eigenvectors for the sliding boundary conditions must be of the form

$$\vec{\phi}(\vec{X}) = \begin{bmatrix} A \sin i\pi X_1 \cos j\pi X_2 \cos k\pi X_3 \\ B \cos i\pi X_1 \sin j\pi X_2 \cos k\pi X_3 \\ C \cos i\pi X_1 \cos j\pi X_2 \sin k\pi X_3 \end{bmatrix}. \tag{A.65}$$

Due to the fact that L operates vectors in 3-space, there are three orthogonal vectors (A_r, B_r, C_r) for $r = 1, 2, 3$ which span the space of functions that satisfy Eq. (A.61). The eigenvectors $\vec{\phi}_r(\vec{X})$ generated from the coefficients (A_r, B_r, C_r) are orthogonal because $(A_r, B_r, C_r) \cdot (A_s, B_s, C_s) = 0$ for $s \neq r$.

Inserting Eq. (A.65) into (A.61) and canceling similar terms gives the equations

$$\begin{aligned} \mu(i^2A + j^2A + k^2A) + (\lambda + \mu)(i^2A + ijB + ikC) &= \kappa \frac{1}{\pi^2} A \\ \mu(i^2B + j^2B + k^2B) + (\lambda + \mu)(ijA + j^2B + jkC) &= \kappa \frac{1}{\pi^2} B \\ \mu(i^2C + j^2C + k^2C) + (\lambda + \mu)(ikA + jkB + k^2C) &= \kappa \frac{1}{\pi^2} C. \end{aligned} \tag{A.66}$$

Solving these equations for κ gives

$$\begin{aligned} \kappa &= \pi^2\mu(i^2 + j^2 + k^2) + \frac{\pi^2(\lambda + \mu)}{A}(i^2A + ijB + ikC) \\ \kappa &= \pi^2\mu(i^2 + j^2 + k^2) + \frac{\pi^2(\lambda + \mu)}{B}(ijA + j^2B + jkC) \\ \kappa &= \pi^2\mu(i^2 + j^2 + k^2) + \frac{\pi^2(\lambda + \mu)}{C}(ikA + jkB + k^2C). \end{aligned} \tag{A.67}$$

From inspection, we see that the coefficients

$$A = i, \quad B = j, \quad \text{and} \quad C = k \tag{A.68}$$

satisfy the above equations with eigenvalue

$$\kappa = \pi^2(2\mu + \lambda)(i^2 + j^2 + k^2). \tag{A.69}$$

Also notice that Eq. (A.67) is satisfied when the cross terms are zero which happens if

$$iA + jB + kC = 0. \tag{A.70}$$

This last equation implies that any vector (A, B, C) that is orthogonal to (i, j, k) will also satisfy Eq. (A.67). The coefficients

$$A = -j, \quad B = i, \quad \text{and} \quad C = 0 \tag{A.71}$$

are a solution of Eq. (A.67) because they are orthogonal to the vector (i, j, k). These coefficients were chosen to simplify the form of the resulting eigenfunction. The eigenvalue associated with these coefficients is

$$\kappa = \pi^2 \mu (i^2 + j^2 + k^2). \tag{A.72}$$

The third set of coefficients (A, B, C) are orthogonal to the solutions in Eqs. (A.68) and (A.71), i.e., they must satisfy

$$Ai + Bj + Ck = 0, \quad \text{and} \quad -Aj + Bi = 0. \tag{A.73}$$

Thus, the third set of coefficients are

$$A = ik, \quad B = jk, \quad \text{and} \quad C = -(i^2 + j^2) \tag{A.74}$$

with eigenvalues

$$\kappa = \pi^2 \mu (i^2 + j^2 + k^2). \tag{A.75}$$

Substituting the three sets of coefficients back into Eq. (A.65) gives the eigenvectors

$$\vec{\phi}_1 = \alpha_1 \begin{bmatrix} i \sin i\pi X_1 \cos j\pi X_2 \cos k\pi X_3 \\ j \cos i\pi X_1 \sin j\pi X_2 \cos k\pi X_3 \\ k \cos i\pi X_1 \cos j\pi X_2 \sin k\pi X_3 \end{bmatrix}, \quad \vec{\phi}_2 = \alpha_2 \begin{bmatrix} -j \sin i\pi X_1 \cos j\pi X_2 \cos k\pi X_3 \\ i \cos i\pi X_1 \sin j\pi X_2 \cos k\pi X_3 \\ 0 \end{bmatrix},$$

$$\vec{\phi}_3 = \alpha_3 \begin{bmatrix} ik \sin i\pi X_1 \cos j\pi X_2 \cos k\pi X_3 \\ jk \cos i\pi X_1 \sin j\pi X_2 \cos k\pi X_3 \\ -(i^2+j^2) \cos i\pi X_1 \cos j\pi X_2 \sin k\pi X_3 \end{bmatrix} \tag{A.76}$$

with associated eigenvalues

$$\begin{aligned} \kappa_{1,i,j,k} &= -\pi^2(2\mu+\lambda)(i^2+j^2+k^2) \\ \kappa_{2,i,j,k} = \kappa_{3,i,j,k} &= -\pi^2\mu(i^2+j^2+k^2) \end{aligned} \tag{A.77}$$

where α_1, α_2, and α_3 are chosen such that each eigenvector has unit energy, i.e., the eigenvectors are orthonormal. Thus, the scale factor α_1 is determined from

$$\int_\Omega \vec{\phi}_1(\vec{X}) \cdot \vec{\phi}_1(\vec{X}) d\vec{X} = \alpha_1^2 \frac{1}{8}(i^2+j^2+k^2) = 1 \tag{A.78}$$

which can be rearranged to give

$$\alpha_1 = \sqrt{\frac{8}{i^2+j^2+k^2}}. \tag{A.79}$$

The other two coefficients can be determined similarly, and are

$$\begin{aligned} \alpha_2 &= \sqrt{\frac{8}{i^2+j^2}} \\ \alpha_3 &= \sqrt{\frac{8}{(i^2+j^2)(i^2+j^2+k^2)}}. \end{aligned} \tag{A.80}$$

A second set of eigenvectors can be generated for the linear-elastic operator in Eq. (A.62) by using the *bending* boundary conditions

$$\begin{aligned}
\left.\frac{\partial u_1}{\partial X_1}\right|_{(0,X_2,X_3)} &= \left.\frac{\partial u_1}{\partial X_1}\right|_{(1,X_2,X_3)} = 0 \\
\left.\frac{\partial u_2}{\partial X_2}\right|_{(X_1,0,X_3)} &= \left.\frac{\partial u_2}{\partial X_2}\right|_{(X_1,1,X_3)} = 0 \\
\left.\frac{\partial u_3}{\partial X_3}\right|_{(X_1,X_2,0)} &= \left.\frac{\partial u_3}{\partial X_3}\right|_{(X_1,X_2,1)} = 0 \\
u_1(X_1,0,X_3) = u_1(X_1,1,X_3) &= u_1(X_1,X_2,0) = u_1(X_1,X_2,1) = 0 \\
u_2(0,X_2,X_3) = u_2(1,X_2,X_3) &= u_2(X_1,X_2,0) = u_2(X_1,X_2,1) = 0 \\
u_3(0,X_2,X_3) = u_3(1,X_2,X_3) &= u_3(X_1,0,X_3) = u_3(X_1,1,X_3) = 0.
\end{aligned} \tag{A.81}$$

These boundary conditions differ from the sliding boundary conditions in that they do not map the unit cube back to the unit cube. In fact, these boundary conditions allow the sides of the unit cube to budge in or out while keeping the six corners points anchored to their original positions. The advantage of allowing the sides of the unit cube to be flexible is that the study does not require as much zero padding to match to the template. The eigenvectors for these boundary conditions are found is a similar manner as those previously stated and are given by

$$\vec{\phi}_1 = \alpha_1 \begin{bmatrix} i\cos i\pi X_1 \sin j\pi X_2 \sin k\pi X_3 \\ j\sin i\pi X_1 \cos j\pi X_2 \sin k\pi X_3 \\ k\sin i\pi X_1 \sin j\pi X_2 \cos k\pi X_3 \end{bmatrix}, \quad \vec{\phi}_2 = \alpha_2 \begin{bmatrix} -j\cos i\pi X_1 \sin j\pi X_2 \sin k\pi X_3 \\ i\sin i\pi X_1 \cos j\pi X_2 \sin k\pi X_3 \\ 0 \end{bmatrix},$$

$$\vec{\phi}_3 = \alpha_3 \begin{bmatrix} ik\cos i\pi X_1 \sin j\pi X_2 \sin k\pi X_3 \\ jk\sin i\pi X_1 \cos j\pi X_2 \sin k\pi X_3 \\ -(i^2+j^2)\sin i\pi X_1 \sin j\pi X_2 \cos k\pi X_3 \end{bmatrix} \tag{A.82}$$

with associated eigenvalues

$$\begin{aligned}
\kappa_{1,i,j,k} &= -\pi^2(2\mu+\lambda)(i^2+j^2+k^2) \\
\kappa_{2,i,j,k} = \kappa_{3,i,j,k} &= -\pi^2\mu(i^2+j^2+k^2)
\end{aligned} \tag{A.83}$$

where the normalizing constants are again

$$\alpha_1 = \sqrt{\frac{8}{i^2+j^2+k^2}}, \tag{A.84}$$

$$\alpha_2 = \sqrt{\frac{8}{i^2+j^2}}, \tag{A.85}$$

$$\alpha_3 = \sqrt{\frac{8}{(i^2+j^2)(i^2+j^2+k^2)}}. \tag{A.86}$$

Notice that the the eigenvectors above can be converted into those in Eq. (A.76) by interchanging the sines with cosines and the cosines with sines.

Now the eigenvectors for the 2D linear-elastic PDE given Eq. (A.54) will be stated. The 2D analog of the sliding boundary conditions given in Eq. (A.63) and (A.64) are

$$\begin{gathered} u_1(0,X_2) = u_1(1,X_2) = u_2(X_1,0) = u_2(X_1,1) = 0 \\ \left.\frac{\partial u_1}{\partial X_2}\right|_{(X_1,0)} = \left.\frac{\partial u_1}{\partial X_2}\right|_{(X_1,1)} = \left.\frac{\partial u_2}{\partial X_1}\right|_{(0,X_2)} = \left.\frac{\partial u_2}{\partial X_1}\right|_{(1,X_2)} = 0. \end{gathered} \tag{A.87}$$

The eigenvectors for these boundary conditions are

$$\vec{\phi}_1 = \alpha_1 \begin{bmatrix} i \sin i\pi X_1 \cos j\pi X_2 \\ j \cos i\pi X_1 \sin j\pi X_2 \end{bmatrix} \quad \text{and} \quad \vec{\phi}_2 = \alpha_2 \begin{bmatrix} -j \sin i\pi X_1 \cos j\pi X_2 \\ i \cos i\pi X_1 \sin j\pi X_2 \end{bmatrix} \tag{A.88}$$

with associated eigenvalues

$$\kappa_{1,i,j} = -\pi^2(2\mu+\lambda)(i^2+j^2) \quad \kappa_{2,i,j} = -\pi^2\mu(i^2+j^2) \tag{A.89}$$

where

$$\alpha_1 = \alpha_2 = \frac{2}{\sqrt{i^2+j^2}}. \tag{A.90}$$

Likewise, the 2D analog of the bending boundary conditions given in Eq. (A.81 are

$$\left.\frac{\partial u_1}{\partial X_1}\right|_{(0,X_2)} = \left.\frac{\partial u_1}{\partial X_1}\right|_{(1,X_2)} = \left.\frac{\partial u_2}{\partial X_2}\right|_{(X_1,0)} = \left.\frac{\partial u_2}{\partial X_2}\right|_{(X_1,1)} = 0$$

$$u_1(X_1,0) = u_1(X_1,1) = u_2(0,X_2) = u_2(1,X_2) = 0 \tag{A.91}$$

and have eigenvectors

$$\vec{\phi}_1 = \alpha_1 \begin{bmatrix} i\cos i\pi X_1 \sin j\pi X_2 \\ j\sin i\pi X_1 \cos j\pi X_2 \end{bmatrix} \quad \text{and} \quad \vec{\phi}_2 = \alpha_2 \begin{bmatrix} -j\cos i\pi X_1 \sin j\pi X_2 \\ i\sin i\pi X_1 \cos j\pi X_2 \end{bmatrix} \tag{A.92}$$

with associated eigenvalues

$$\kappa_{1,i,j} = -\pi^2(2\mu+\lambda)(i^2+j^2) \quad \kappa_{2,i,j} = -\pi^2\mu(i^2+j^2) \tag{A.93}$$

where

$$\alpha_1 = \alpha_2 = \frac{2}{\sqrt{i^2+j^2}}. \tag{A.94}$$

It is interesting to make the connection that the 2D eigenvectors in Eq. (A.88) are equal to the 3D eigenvectors in Eq. (A.76) when $k = 0$. However, the same cannot be said for the 2D eigenvectors in Eq. (A.92) and the eigenvectors in Eq. (A.82).

B. Summaries of Deformable Atlas Literature

Fox, Perlmutter, and Raichle. A stereotactic method of anatomical localization for positron emission tomography. *Journal of Computer Assisted Tomography*, 9(1):141–153, 1985.

This paper presents a method for aligning PET data with an atlas. The PET coordinate system is estimated from a lateral x-ray taken of the patient's head once immobilized in a headholder. This coordinate system is then rigidly transformed to match the coordinate system of the atlas.

Evans, Beil, Marret, Thompson, and Hakim. Anatomical-functional correlation using an adjustable MRI-based region of interest atlas with positron emission tomography. *Journal of Cerebral Blood Flow and Metabolism*, 8:513–530, 1988.

This paper describes a method for registering 3D PET volumes with at 3D MR atlas. Extreme care was taken to acquire PET and MR data sets in the same plane orientation by using a customized headholder for each subject. This approach requires the investigator to match the VOI atlas to the MRI image and to manually adjust individual regions of interest as necessary to obtain a satisfactory match.

Bajcsy, Reivich, and Kovacic. Evaluation of Registration of PET Images and CT Images. Technical report MS-CIS-88-89, Grasp lab 163, Department of Computer and Information Science, School of Engineering and Applied Science, University of Pennsylvania, Philadelphia, PA 19104-6389, 1988.

This paper describes an algorithm for rigidly registering PET and CT images from the same individual. The CT and PET images were registered by aligning the centroid and principle axes of both images.

Pelizzari, Chen, Spelbring, Weichselbaum, and Chen. Accurate three-dimensional registration of CT, PET, and/or MR images of the brain. *Journal of Computer Assisted Tomography*, 13(1):20–26, 1989.

This paper presents a 3-D algorithm for registering CT with CT, CT with MR, CT with PET, and MR with PET. The surface of an anatomical structure visible in all scans of interest (most commonly the external surface of the head, including tissue) is produced by outlining contours on the serial slices of each scan. The surface of the lowest resolution data being matched is reduced to approximately 250 points which are match to the surface of the high-resolution data. The allowed class of transformations includes rigid body rotation and translation (three angles, three orthogonal displacements) and linear scaling along three orthogonal axes.

Greitz, Bohm, Holte, and Eriksson. A computerized brain atlas: Construction, anatomical content, and some applications. *Journal of Computer Assisted Tomography*, 15(1):26–38, January/February 1991.

This group constructed a brain atlas from digitized images of one cryosectioned cadaver brain forming an image volume of 280x512x512 voxels. Their atlas is deformed interactively using rigid and nonrigid transformations. The rigid transformations include three translations, three rotations, and three scaling factors. The nonrigid group include three skews, three scoliosis, and nine variable scalings parameters. The transformations are chosen so that reference structures of the transformed atlas (usually the brain surface and the ventricular system) coincide with corresponding structures in the CT or MR images. After interactively aligning the PET and/or SPECT images with the MR or CT images of the patient with rigid transformations, the PET and/or SPECT images are mapped back to the atlas coordinate system.

Undrill, Cameron, Cookson, Davies, Robinson, Hill, Cootes, Taylor, Thornhan, Wysocki, Liddell, and Parkinson. Integrated presentation of 3d data derived from multisensor imagery and anatomical atlases using a parallel processing system. In *Image Capture, Formatting and Display*, SPIE 1653:2–16, 1992.

An atlas was constructed from six different standard atlases including the Talairach brain atlas. Seven landmark points in both the atlas and the data volume are used for registration in an initial rigid transformation step. Structures in the atlas are represented by spline representations of their outer contour. The parameters of the spline boundary are estimated so that the spline boundary has strong edge strengths perpendicular to it in the data and it final shape is close to its initial shape.

F.L. Bookstein and W.D.K. Green. Edge information at landmarks in medical images. In *Visualization in Biomedical Computing 1992*, Richard A. Robb, editor, pages 242–258. SPIE 1808, 1992.

This method uses a landmark driven, thin-plate spline process to deform 2D MR images. Nine landmarks and the orientation of the edges through four of the landmarks were used to unwarp and average 14 midsagittal MRI images.

Evans, Dai, Collins, Neelin, and Marret. Warping of a computerized 3-d atlas to match brain image volumes for quantitative neuroanatomical and functional analysis. *Image Processing*, SPIE 1445:236–246, 1991.

A volume-of-interest (VOI) atlas was produced by manual contouring on 64 adjacent 2mm-thick MR slices to yield 60 brain structures in each hemisphere. This paper uses the landmark driven, thin-plate spline transformations of Bookstein to deform the 3D atlas. A set of 26 points in the MRI volume and the Talairach stereotactic atlas were used to deform MR brain image volumes from 16 young normal volunteers to the Talairach coordinate system.

Bajcsy and Kovacic. Multiresolution Elastic Matching. *Computer Vision, Graphics, and Image Processing*, 46:1–21, 1989.

This paper appears to be the published version of [65]. This paper describes a 3-D, multiresolutional elastic matching technique for deforming a brain atlas to match an X-ray computed tomography scan of the brain. The atlas was created by manually tracing various structures on the original digitized images. These structures were filled with a gray scale value that corresponds to the color that structure would be on a CT scan. The algorithm consists of a global rigid matching step followed by an elastic matching step. The elastic matching step is a multiresolution approach where correlation or similarity of features such as average gray level and edges drive the elastic PDE. Both the similarity function and the elastic PDE are approximated by second order Taylor's series. The forcing function of the PDE is the gradient of the similarity function. The elastic PDE is solved at each resolution by using a finite difference approximation and Jacobi relaxation.

Dann, Hoford, Kovacic, Reivich, and Bajcsy. Evaluation of Elastic Matching Systems for Anatomic (CT, MR) and Functional (PET) Cerebral Images. *Journal of Computer Assisted Tomography*, 13(4):603–611, July/August 1989.

This paper evaluates the performance of the elastic matching algorithm presented in [30]. A wide variation was reported between segmentations generated by four neurologists and in three segmentations generated by the same neuroanatomist. Even though the experts disagreed, they agreed with each other better than they agreed with the computer segmentation. They also registered PET scans from an individual with the atlas by rigidly registering the PET images with CT images of the same individual and then elastically registering the CT images with the atlas.

Collins, Peters, Dai, and Evans. Model-based segmentation of individual brain structures from MRI data. In *Visualization in Biomedical Computing 1992*, Richard A. Robb, editor, pages 10–23. SPIE 1808, 1992.

A volume-of-interest (VOI) atlas is matched to MR images using a two step approach. The first step in the procedure is global matching, consisting only of rigid transformations (translation, rotation, and scaling). The second step is performed at five different resolutions by blurring the data set with a Gaussian kernels with $\sigma = 2, 4, 8, 16,$ and 32 mm. For each node in a regular lattice with spacing σ, the optimum x,y,and z coefficients of a displacement vector is estimated that maximizes the correlation over a local 3D neighborhood of radius σ. The voxel intensity and gradient magnitude are used as features in the matching process.

Amit, Grenander, and Piccioni. Structural image restoration through deformable templates. *J. American Statistical Association*, 86(414):376–387, June 1991.

This paper presents an algorithm for representing shape variability in terms of a probabilistic model. The displacement field is modeled as a 2-dimensional Gaussian quadratic mean process with covariance structure specified by the impulse response of the Laplacian operator. The displacement field is represented by its Karhunen-Loève expansion where the basis functions diagonalize the

covariance structure. The coefficients of the basis functions are estimated using the stochastic differential equation.

Miller, Christensen, Amit, and Grenander. A mathematical textbook of deformable neuro-anatomies. *Proceedings of the National Academy of Science*, accepted January 1993.

A method for accommodating very small shape variability associated with normal human brains is presented. The displacement field is modeled as a 2-dimensional Gaussian quadratic mean process with covariance structure specified by the elasticity potential. The displacement field is parameterized by two different methods. The first parameterization uses the Karhunen-Loève expansion and the second represents the displacement field by its samples on a fine grid. The parameters for both representations of the displacement field are estimated using the stochastic differential equation.

Christensen, Rabbitt, and Miller. A deformable neuroanatomy textbook based on viscous fluid mechanics. In *1993 Conference on Information Sciences and Systems*, Prince and Runolfsson, editors, Johns Hopkins University, Baltimore, Maryland, 1993.

This paper describes a method for accommodating very small shape variability associated with normal human brains using a viscous fluid transformations. The displacement field is estimated by solving a fluid PDE which was first presented in [67].

Hanson. Bayesian reconstruction based on flexible prior models. *submitted to the Journal of the Optical Society of America A*, 1992. Los Alamos National Laboratory monograph LA-UR-92-2334.

This paper proposes a Bayesian reconstruction method that incorporates geometric information in a prior distribution. The data is modeled as a linear projection of an image with additive zero mean Gaussian noise. The image is modeled as a two-dimensional Gaussian random field with mean given by a transformed transformed image. The set of all possible objects were assumed to have the same shape and a known constant amplitude. This paper describes 2-D elastic transformations and uses polynomial coordinate transformations. warping.

Kamber, Collins, Shighal, Francis, and Evans. Model-based 3d segmentation of multiple sclerosis lesions in dual-echo MRI data. In *Visualization in Biomedical Computing 1992*, Richard A. Robb, editor, pages 590–600. SPIE 1808, 1992.

This paper describes a pixel-based approach for segmenting multiple sclerosis lesions in MR images. The segmentations of 12 individuals were registered in Talairach space to provide probabilities of gray matter, white matter, ventricular cerebrospinal fluid (CSF), external CSF, and background. The segmentation of multiple sclerosis lesions was performed in the regions with 0.5 or higher probability of white matter.

C. Massively Parallel Implementation

An Efficient Parallel Computational Method for 3D Virtualized Data

The need for virtualization arises on a massively parallel computer when the the dimensions of the data is larger than the dimensions of the array of processors. Figure C.1 illustrates the cut-and-stack method for storing a 128×128 size data set on a 64×64 processor array. The 128×128 data set is broken up into four 64×64 blocks and stored consecutively in the memory. Likewise, Fig. C.2 illustrates the cut-and-stack method in 3D. The $128 \times 128 \times 128$ data set is virtualized onto a 64×64 array of processors by first dividing it into 2D planes and then virtualizing each 2D plane separately.

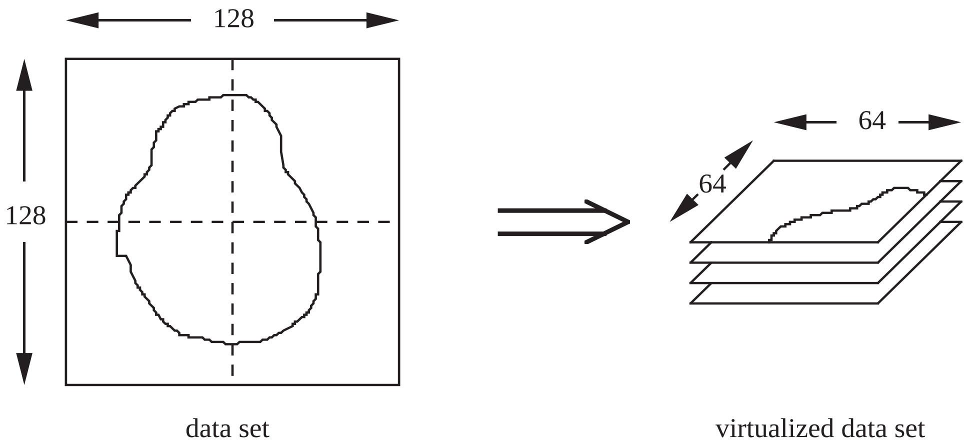

Figure C.1: Cut-and-stack virtualization of a 128×128 data set for a 64×64 processor array.

The following code fragment efficiently computes parallel commands that require nearest-neighbor communication when a 3D data set is virtualized on the DECmpp 12000Sx(MasPar).

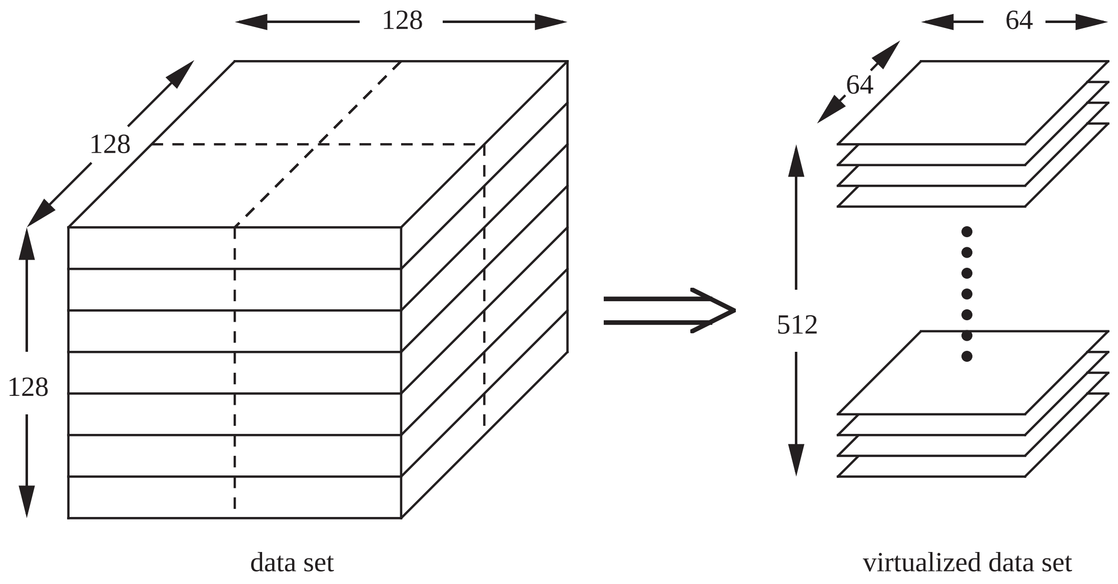

Figure C.2: Cut-and-stack virtualization of a $128 \times 128 \times 128$ data set for a 64×64 processor array.

The computations are done in place so very little additional memory is required, and the computations are fast because they are done registers.

Global variables are denoted in uppercase letters and are normally declared in a header file. The variables NXBLK, NYBLK, and NZBLK define the number of blocks that that the 3D volume is broken up into in the x, y, and z dimensions, respectively. The variables MAXNXYBLK and MAXNXYZBLK are used for memory allocation and correspond to the largest possible values of NXBLK*NYBLK and NXBLK*NYBLK*NZBLK, respectively.

The plural variable u1_orig is the 3D input data for the calculation. Computations are performed on u1_orig beginning at the top plane of u1_orig and ending at the bottom plane. The plural variables u1_xy and u1_xy_old store the original contents of u1_orig at the current plane of computation and the previous plane of computation, respectively. The variables u1_xy and u1_xy_old are needed because the data in u1_orig gets over written during the computation.

The lines of code between the comments "Initialize u1_xy to first plane of data" and "COMMANDS START HERE" are used to transfer the data from u1_orig, u1_xy, and u1_xy_old to the registered variables u1, u1_old, and u1_nxt, respectively. The commands in the subroutine are placed between the comments "COMMANDS START HERE" and "COMMANDS END HERE" and consist of arithmetic operations and nearest-neighbor communication of the variables u1, u1_old, and u1_nxt. Finally, the lines of code from the comment "COMMANDS END HERE" to the end of the code fragment are used to transfer the data from the registered variable tmp1 back to u1_orig.

The ingenuity of this code fragment is the method in which data is transfered from RAM to register store and back again. This code fragment insures sure that data located on the boundary of the data blocks are computed correctly by overlapping the fetching of data from the RAM to the register store. Similarly, only valid data from the computations are transfered from the register store back to the RAM (u1_orig).

This code fragment can easily be modified to include multiple 3D data variables by adding additional variables. For example, if two more data variables are needed for the computation, the variables u2_orig, u3_orig, u2_xy, u3_xy, u2_xy_old, u3_xy_old, u2, u3, u2_old, u3_old, u2_nxt, and u3_nxt would have to be defined and added to the code.

Finally, this code can easily be modified for data communication between processor elements separated by a distance of two or more processors by changing the ranges of the for loops.

```
/*********************************************************************
*
*  Programmer: Gary Christensen
*  Organization: Electrical Engineering
*      Washington University, St. Louis, MO 63130
*  Function: To efficiently compute nearest-neighbor computations in
*      place using the register store.
*
**********************************************************************/

#define MAXNXYBLK 16
#define MAXNXYZBLK 1024
int NXBLK=4,NYBLK=4,NZBLK=128;
int NXYBLK=NXBLK*NYBLK, NXYZBLK=NXBLK*NYBLK*NZBLK;

plural float u1_orig[MAXNXYZBLK];
plural float u1_xy[MAXNXYBLK], u1_xy_old[MAXNXYBLK];
register plural float u1, u1_old, u1_nxt, tmp1;
int i,j,k,m,n,o,r,s,t;

  /********** Initialize u1_xy to first plane of data *********/
  for(i=0;i<NXYBLK;i++)
        u1_xy[i]=u1_orig[i];
```

```
for(k=1,n=NXYBLK;n<NXYZBLK-NXYBLK;k++,n+=NXYBLK){
  /********** Save a copy of current plane and previous plane *********/
  for(t=0;t<NXYBLK;t++){
      u1_xy_old[t]=u1_xy[t];
      u1_xy[t]=u1_orig[t+n];
  }
  for(i=0,m=n,r=0;m<n+NXYBLK;i+=2,m+=NXBLK,r+=NXBLK){
    for(j=0,o=m,s=r;o<m+NXBLK;j+=2,o++,s++){
      /*************** Load variables into registers **************/
      u1=u1_xy[s];
      u1_old=u1_xy_old[s];
      u1_nxt=u1_orig[o+NXYBLK];

      if((ixproc<(nxproc-j))&(iyproc>=(nyproc-i))){
          u1=u1_xy[s-NXBLK];
          u1_old=u1_xy_old[s-NXBLK];
          u1_nxt=u1_orig[o+NXYBLK-NXBLK];
      }
      if((ixproc>=(nxproc-j))&(iyproc<(nyproc-i))){
          u1=u1_xy[s-1];
          u1_old=u1_xy_old[s-1];
          u1_nxt=u1_orig[o+NXYBLK-1];
      }
      if((ixproc>=(nxproc-j))&(iyproc>=(nyproc-i))){
          u1=u1_xy[s-1-NXBLK];
          u1_old=u1_xy_old[s-1-NXBLK];
          u1_nxt=u1_orig[o+NXYBLK-1-NXBLK];
      }

      /****************COMMANDS START HERE*********************/
      tmp1=u1+xnetW[1].u1+xnetE[1].u1+xnetN[1].u1+xnetS[1].u1+u1_old+u1_nxt;
      /****************COMMANDS END HERE***********************/

      /******* Save updated values in u1_orig ************/
      if((iyproc>-i)&(iyproc<(nyproc-i-1))&(ixproc>-j)&(ixproc<(nxproc-j-1)))
          u1_orig[o]=tmp1;
      if((iyproc>-i)&(iyproc<(nyproc-i-1))&(ixproc>(nxproc-j)))
          u1_orig[o-1]=tmp1;
      if((iyproc>(nyproc-i))&(ixproc>-j)&(ixproc<(nxproc-j-1)))
          u1_orig[o-NXBLK]=tmp1;
      if((iyproc>(nyproc-i))&(ixproc>(nxproc-j)))
          u1_orig[o-NXBLK-1]=tmp1;
    }
  }
}
```

Parallel Textbook Deformation

The code fragment below is a simplification of the routine that is used to calculate the deformation of the textbook on the DECmpp 12000Sx (MasPar) computer. This routine uses the global router and the parallel indirect memory addressing capabilities of the DECmpp. Trilinear interpolation is used to evaluate the deformed textbook and its derivatives.

The derivatives of the textbook are computed "on the fly" to save memory as opposed to precomputing them. The additional computation time required is negligible with respect to the time needed to fetch the data for the computation. Fortunately, the time required to fetch the data is equivalent to the time that would be required to fetch the precomputed derivatives.

The global variables used in this code fragment have the same definition as in the previous code fragment example. In addition, the global variables xIS, yIS, and zIS are initialized to the number of voxels in the x, y, and z directions, respectively.

The overview of the code fragment is the following. First, the addresses of the deformed coordinate system are calculated from the plural variables u1, u2, and u3 which contain the displacement field. Next, the 32 data values closest to each address are fetched from the plural variable template and are stored in the plural array node_value (this step requires the global router and parallel indirect memory addressing). The diagram in the comments below depicts the fetching order of the data values and gives the index for the data values in node_value. Finally, the template and its derivatives are computed with respect to the deformed coordinate system using trilinear interpolation. The derivatives are computed with the symmetric difference formula.

```
/*******************************************************************
*
*  Programmer: Gary Christensen
*  Organization: Electrical Engineering
*      Washington University, St. Louis, MO 63130
*  Function: Calculate the deformed template and its x-derivative
*      using trilinear interpolation.
*  Note:  The fetching ordering of the nodes for router is as follows:
*
* y=-1                0--1            2--3                 z in
*                     |  |            |  |                 ----> x
* y= 0   4--5       6 --7--8-- 9    10--11--12--13    14--15     |
*        |  |       |   |  |   |     |   |   |   |     |  |      |
* y= 1  16--17     18--19--20--21   22--23--24--25    26--27     \/
*                     |  |            |  |                       y
* y= 2               28--29          30--31
*
*       z=-1          z=0             z=1           z=2
*
*        x  x       x   x  x   x     x   x  x   x      x  x
*        "  "       "   "  "   "     "   "  "   "      "  "
*        0  1      -1   0  1   2    -1   0  1   2      0  1
*
*******************************************************************/

plural unsigned int template[MAXNXYZBLK];
plural float u1[MAXNXYZBLK], u2[MAXNXYZBLK], u3[MAXNXYZBLK];
float xnumdiv ynumdiv, znumdiv;

register plural float xtmp,ytmp,ztmp;
register plural int address, depth, PEdepth, notdone;
register plural int xloc, yloc, zloc, ixprocIS, iyprocIS;
register plural float xdelta,ydelta,zdelta;
register plural int xlocation, ylocation,zlocation;
register plural unsigned int itmp;
register plural float template_value, dtdx;
register plural float *tp;

int yoff, xoff;
register int i,j,k,m,n,r,s,t,u,v,w,inode;
plural float node_value[32];

static int xnodeoffset[32]={0,1,0,1,0,1,-1,0,1,2,-1,0,1,2,0,1,0,1,-1,0,1,2,-1,
0,1,2,0,1,0,1,0,1};
static int ynodeoffset[32]={-1,-1,-1,-1,0,0,0,0,0,0,0,0,0,0,0,0,1,1,1,1,1,1,1,
1,1,1,1,1,2,2,2,2};
static int znodeoffset[32]={0,0,1,1,-1,-1,0,0,0,0,1,1,1,1,2,2,-1,-1,0,0,0,0,1,
1,1,1,2,2,0,0,1,1};
```

```
tp=node_value;

for(i=0,u=0,xoff=0;i<NXBLK;i++,u++,xoff+=nxproc){
  for(j=0,v=u,yoff=0;j<NYBLK;j++,v+=NXBLK,yoff+=nyproc){
    for(k=0,w=v;k<NZBLK;k++,w+=NXYBLK){

        ixprocIS = ixproc + xoff;       /* virtual x-coordinate of processor */
        iyprocIS = iyproc + yoff;       /* virtual y-coordinate of processor */

        xtmp=fp_floor(u1[w]*xnumdiv);
        xdelta=u1[w]*xnumdiv - xtmp;
        xlocation=ixproc+(plural int)xtmp+xoff;

        ytmp=fp_floor(u2[w]*ynumdiv);
        ydelta=u2[w]*ynumdiv - ytmp;
        ylocation=iyproc+(plural int)ytmp+yoff;

        ztmp=fp_floor(u3[w]*znumdiv);
        zdelta=u3[w]*znumdiv - ztmp;
        zlocation=k+(plural int)ztmp;

        /* Collect the points required to calculate the interpolated
           values of the template, and derivatives of the template */

        for(inode=0;inode<32;inode++){
            xloc=xlocation+xnodeoffset[inode];
            yloc=ylocation+ynodeoffset[inode];
            zloc=zlocation+znodeoffset[inode];

            address=(xloc%nxproc)+((yloc%nyproc)<<lnxproc);
            depth=(xloc>>lnxproc)+(yloc>>lnyproc)*NXBLK+zloc*NXYBLK;

            notdone=1;
            itmp=0.0;

            if ((xloc>=0)&(xloc<xIS)&(yloc>=0)&(yloc<yIS)&(zloc>=0)&(zloc<zIS)){
                while(notdone){
                    if (connected(address)){
                        router[address].PEdepth = depth;
                        itmp = router[address].template[PEdepth];
                        notdone=0;
                    }
                }
            }
            node_value[inode]=(plural float)itmp;
        }

        /* Calculate interpolated template intensity and derivates */
        xtmp=1.0-xdelta;
```

```
        ytmp=1.0-ydelta;
        ztmp=1.0-zdelta;

        template_value = ztmp*(ytmp*(xtmp*tp[ 7]+xdelta*tp[ 8])
                            +ydelta*(xtmp*tp[19]+xdelta*tp[20]))
                      +zdelta*(ytmp*(xtmp*tp[11]+xdelta*tp[12])
                            +ydelta*(xtmp*tp[23]+xdelta*tp[24]));

        dtdx =   ztmp*(ytmp*(xtmp*(tp[ 8]-tp[ 6])
                          +xdelta*(tp[ 9]-tp[ 7]))
                    +ydelta*(xtmp*(tp[20]-tp[18])
                          +xdelta*(tp[21]-tp[19])))
              +zdelta*(ytmp*(xtmp*(tp[12]-tp[10])
                          +xdelta*(tp[13]-tp[11]))
                    +ydelta*(xtmp*(tp[24]-tp[22])
                          +xdelta*(tp[25]-tp[23])));
        dtdx = dtdx*0.5*((float)xIS-1.0)/255.0;

        dtdy =   ztmp*(ytmp*(xtmp*(tp[19]-tp[ 0])
                          +xdelta*(tp[20]-tp[ 1]))
                    +ydelta*(xtmp*(tp[28]-tp[ 7])
                          +xdelta*(tp[29]-tp[ 8])))
              +zdelta*(ytmp*(xtmp*(tp[23]-tp[ 2])
                          +xdelta*(tp[24]-tp[ 3]))
                    +ydelta*(xtmp*(tp[30]-tp[11])
                          +xdelta*(tp[31]-tp[12])));
        dtdy = dtdy*0.5*((float)yIS-1.0)/255.0;

        dtdz =   ztmp*(ytmp*(xtmp*(tp[11]-tp[ 4])
                          +xdelta*(tp[12]-tp[ 5]))
                    +ydelta*(xtmp*(tp[23]-tp[16])
                          +xdelta*(tp[24]-tp[17])))
              +zdelta*(ytmp*(xtmp*(tp[14]-tp[ 7])
                          +xdelta*(tp[15]-tp[ 8]))
                    +ydelta*(xtmp*(tp[26]-tp[19])
                          +xdelta*(tp[27]-tp[20])));
        dtdz = dtdz*0.5*((float)zIS-1.0)/255.0;
    }
  }
}
```

Acknowledgments

I first want to thank Jesus Christ for blessing me with the excellent family, teachers, friends, and opportunities in my life.

My thanks go to the faculty and students of the Electronic Systems and Signal Laboratory (ESSRL) and Biomedical Computing Laboratory (BCL) for their help and comments during my years at Washington University. I am especially grateful to my mentor and academic advisor, Michael Miller, for his guidance, encouragement, and for his undying enthusiasm.

Special thanks go to Ulf Grenander and Yali Amit who played a crucial role in my understanding of Global Shape Models.

I also am very thankful to Richard Rabbitt who derived the original dilatational viscous fluid model and who helped me in my understanding of continuum mechanics.

Thanks to Donald Snyder, Director of the Electronic Systems and Signals Research Laboratory, and Lewis Thomas Jr., Director of the Biomedical Computer Laboratory, for their guidance, encouragement, and allowing me to be apart of their laboratories.

I would like to thank Russ Hermes, Martin Arthur, Lloyd Brown, Dan Fuhrmann, Bob Morley, Bixio Rimoldi, Fred Rosenberger, John Ollinger, and Jody O'Sullivan for their guidance and help throughout the years.

I am thankful to Michael Vannier, Marcus Raichle, Tom Videen, John Haller, David Van Essen, Tom Coogan, and Scott Nadel for providing me with excellent data and for their invaluable collaborations.

Additional thanks go to my fellow graduate students and friends for their help. Specifically, I would like to thank Tim Barton, Chris Butler, Ken Clark, Heather Drury, George Engel, Rob Grothe, Shirley Gonzalez-Rubio, Walter Gude, Abed M. Hammoud, Sarang Joshi, Seh Wah Kwa, Aaron Lanterman, Quinn Li, Dave Maffitt, Kevin Mark, Leslie Nelson, Dave Politte, Don Porter, Jim Prater, Chrysanthè Preza, Tim Schaewe, Rob Schmich, Tim Schulz, Steven Schwartz, Debbie Schwab, Kurt Smith, Anuj Srivastava, Tom Sullivan, Rob Teichman, and Jing Wang.

I would especially like to thank my wife Shannon for all of her love, support, understanding, and patience during this long process. I thank my family for giving me constant support and encouragement in pursuing my studies.

Finally, I am thankful for the helpful suggestions and comments of Michael Miller, Richard Rabbitt, Donald Snyder, Lewis Thomas Jr., Michael Vannier, Bixio Rimoldi, and David Van Essen, on the final version of this dissertation as members of my defense committee.

This research was supported by the NIH-NCRR-RR01380 and a grant from the Digital Equipment Corporation.

Bibliography

[1] U. Grenander and M. I. Miller. Representations of knowledge in complex systems. *Journal of the Royal Statistical Society B*, 56(3), 1994.

[2] K.H. Höhne, A. Pommert, M. Rimer, T. Schiemann, R. Schubert, and U. Tiede. Framework for the generation of 3d anatomical atlases. In Richard A. Robb, editor, *Visualization in Biomedical Computing 1992*, pages 510–520. SPIE 1808, 1992.

[3] K.H. Höhne, A. Pommert, M. Rimer, T. Schiemann, R. Schubert, and U. Tiede. Representing spatial knowledge using the intelligent volumes concept. In *Proceedings of the American Association for Artificial Intelligence, Spring Symposium: Applications of Computer Vision in Medical Image Processing*, Stanford, California, March 21-23 1994. Stanford University.

[4] Institute of Medicine. Mapping the brain and its functions: Integrating enabling technologies in neuroscience research. Report of the institute of medicine, 1991.

[5] A.C. Evans, C. Beil, S. Marret, C.J. Thompson, and A. Hakim. Anatomical-functional correlation using an adjustable mri-based region of interest atlas with positron emission tomography. *Journal of Cerebral Blood Flow and Metabolism*, 8:513–530, 1988.

[6] M. Vannier, T. Pilgram, G. Bhatia, J. Nemecek, and V. Young. Quantitative three-dimensional assessment of face-lift with an optical facial surface scanner. *Annals of Plastic Surgery*, 30:204–211, 1993.

[7] S. Joshi, M. I. Miller, J. McNally, Y. Amit, and U. Grenander. Global shape models for optical sectioning microscopy. *Journal of the Optical Society*, October 1992.

[8] M.I. Miller, S. Joshi, D. R. Maffitt, J. G. McNally, and U. Grenander. Mitochondria, membranes and amoebae: 1,2 and 3 dimensional shape models. In Kanti Mardia, editor, *Statistics and Imaging*, volume II. Carfax Publishing C., Abingdon, Oxfordshire - England, 1994.

[9] D. Metaxas and D. Terzopoulos. Sahpe and nonrigid motion estimation through phsics-based synthesis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(6):580–591, 1993.

[10] M.I. Miller, D. Maffitt, J. Shrauner, B. Roysam, and U. Grenander. Automated segmentation of biological shapes in electron microscopic autoradiography. In *Proceedings of the Twenty-Fifth Annual Conference on Information Sciences and Systems*, pages 637–642, Baltimore, MD, 1991. Johns Hopkins University.

[11] M.I. Miller, B. Roysam, K. Smith, and J.A. O'Sullivan. Representing and computing regular languages on massively parallel networks. *I.E.E.E. Transactions on Neural Networks*, 2,No.1:56–72, January 1991.

[12] M.I. Miller and B. Roysam. Bayesian image reconstruction for emission tomography: Implementation of the EM algorithm and Good's roughness prior on massively parallel processors. *Proc. of the Natl. Acad. of Sci.*, 88:3223–3227, April 1991.

[13] A. Schweikard, R. Tombropoulos, T.R. Adler, and J.C. Latombe. Planning for image-guided radiosurgery. In *Proceedings of the American Association for Artificial Intelligence, Spring Symposium: Applications of Computer Vision in Medical Image Processing*, Stanford, California, March 21-23 1994. Stanford University.

[14] E. Grimson, T. Lozano-Perez, W. Wells, G. Ettinger, S. White, and R. Kikinis. Automated registration for enhanced reality visualization in surgery. In *Proceedings of the American Association for Artificial Intelligence, Spring Symposium: Applications of Computer Vision in Medical Image Processing*, Stanford, California, March 21-23 1994. Stanford University.

[15] K.R. Smith and R.D. Bucholz. Computer methods for improved diagnostic image display applied to stereotactic neurosurgery. *Automedica*, 14:371–382, 1992.

[16] R.D. Bucholz, K.R. Smith, J.M. McDurmont, and W. Schulz. Intraoperative localization using a three dimensional optical digitizer. *Proceedings of SPIE*, 1894, 1993.

[17] P. Sautot, S. Lavalle, J. Troccaz, P. Cinquin, and P. Merloz. Computer assisted spine surgery. In *Proceedings of the American Association for Artificial Intelligence, Spring Symposium: Applications of Computer Vision in Medical Image Processing*, Stanford, California, March 21-23 1994. Stanford University.

[18] E. Bainville, P. Chaffanjon, and P. Cinquin. Computer generated visual assistance to a surgical operation: the retroperitoneoscopy. In *Proceedings of the American Association for Artificial Intelligence, Spring Symposium: Applications of Computer Vision in Medical Image Processing*, Stanford, California, March 21-23 1994. Stanford University.

[19] A.C.F. Colchester, R.L. Evans, J. Zhao, D.J. Hawkes, C. Studholme, D.L.G. Hill, P. Roberts, M. Stephens, N. Maitland, A.J. Strong, M.J. Gleeson, D.G. Thomas, and T.C.S. Cox. Vislan: Combining intra-operative video and pre-operative images for surgical guidance. In *Proceedings of the American Association for Artificial Intelligence, Spring Symposium: Applications of Computer Vision in Medical Image Processing*, Stanford, California, March 21-23 1994. Stanford University.

[20] P.L. Gleason, R. Kikinis, W. Wells, W. Lorensen, H. Cline, G. Ettinger, P.M. Black, E. Alexander III, and F. Jolesz. Enhanced reality for neurosurgical guidance. In *Proceedings of the American Association for Artificial Intelligence, Spring Symposium: Applications of Computer Vision in Medical Image Processing*, Stanford, California, March 21-23 1994. Stanford University.

[21] Richard A. Robb, editor. *Visualization in Biomedical Computing 1992*. SPIE 1808, 1992.

[22] P.T. Fox, J.S. Perlmutter, and M.E. Raichle. A stereotactic method of anatomical localization for positron emission tomography. *Journal of Computer Assisted Tomography*, 9, No.1:141–153, 1985.

[23] Greitz, Bohm, Holte, and Eriksson. A computerized brain atlas: Construction, anatomical content, and some applications. *Journal of Computer Assisted Tomography*, 15(1):26–38, Janurary/February 1991.

[24] R. Bajcsy, , M. Reivich, and S. Kovacic. Evaluation of Registration of PET Images and CT Images. Technical report ms-cis-88-89, grasp lab 163, Department of Computer and Information Science, School of Engineering and Applied Science, University of Pennsylvania, Philadelphia, PA 19104-6389, 1988.

[25] C.A. Pelizzari, G.T.Y. Chen, D.R. Spelbring, R.R. Weichselbaum, and C.T. Chen. Accurate three-dimensional registration of ct, pet, and/or mr images of the brain. *Journal of Computer Assisted Tomography*, 13(1):20–26, 1989.

[26] P.E. Undrill, G.G. Cameron, M.J. Cookson, C. Davies, N.L. Robinson, A. Hill, T.F. Cootes, C.J. Taylor, A. Thornhan, J. Wysocki, H. M. Liddell, and D. Parkinson. Integrated presentation of 3d data derived from multi-sensor imagery and anatomical atlases using a parallel processing system. In *SPIE*, volume 1653 Image Capture, Formatting and Display, San Diego, CA., 1992.

[27] F.L. Bookstein and W.D.K. Green. Edge information at landmarks in medical images. In Richard A. Robb, editor, *Visualization in Biomedical Computing 1992*, pages 242–258. SPIE 1808, 1992.

[28] A.C. Evans, W. Dai, L. Collins, P. Neelin, and S. Marret. Warping of a computerized 3-d atlas to match brain image volumes for quantitative neuroanatomical and functional analysis. *Image Processing*, 1445:236–246, 1991.

[29] Y. Amit, U. Grenander, and M. Piccioni. Structural image restoration through deformable templates. *J. American Statistical Association*, 86(414):376–387, June 1991.

[30] R. Bajcsy and S. Kovacic. Multiresolution Elastic Matching. *Computer Vision, Graphics, and Image Processing*, 46:1–21, 1989.

[31] D.L. Collins, T.M. Peters, W. Dai, and A.C. Evans. Model-based segmentaion of individual brain structures from mri data. In Richard A. Robb, editor, *Visualization in Biomedical Computing 1992*, pages 10–23. SPIE 1808, 1992.

[32] D. Terzopoulos and K. Waters. Physically-based facial modelling, analysis, and animation. *The Journal of Visualization and Computer Animation*, 1:73–80, 1990.

[33] Y. Amit. A non-linear variational problem for image matching. *to appear in SIAM Journal on Scientific Computation*, 1993.

[34] R. C. Buck. *Advanced Calculus.* McGraw-Hill Book Company, St. Louis, 3 edition, 1978.

[35] H. L. Van Trees. *Detection, Estimation and Modulation Theory, Part I.* John Wiley and Sons, New York, 1968.

[36] D.L. Snyder, T. J. Schulz, and J.A. O'Sullivan. Deblurring subject to nonnegativity constraints. *J. Acoustics, Speech and Signal Processing*, pages 1143–1150, 1992.

[37] D. L. Snyder and M. I. Miller. *Random Point Processes in Time and Space.* Springer-Verlag, ii edition, 1991.

[38] R. Szeliski. *Bayesian Modeling of Uncertainy in Low-Level Vision.* Kluwer Academic Publisher, Boston, 1989.

[39] Grace Wahba. *Spline Models for Observational Data.* Society for Industrial and Applied Mathematics, Philadelphia, 1990.

[40] F.L. Bookstein. *Morphometric Tools for Landmark Data.* Cambridge University Press, New York, 1991.

[41] M.I. Miller, G.E. Christensen, Y. Amit, and U. Grenander. Mathematical textbook of deformable neuroanatomies. *Proceedings of the National Academy of Science*, 90(24), December 1993.

[42] G. E. Christensen, R. D. Rabbitt, and M.I. Miller. 3D brain mapping using a deformable neuroanatomy. *Physics in Medicine and Biology*, 39, 1994.

[43] Y. Amit and M. Piccioni. A nonhomogenous markov process for the estimation of gaussian random fields with non-linear observations. *The Annals of Probability*, 19:1664–1678, 1991.

[44] I. Karatzas and S.E. Shreve. *Brownian Motion and Stochastic Calculus.* Springer-Verlag, 1987.

[45] E. Wong and B Hajek. *Stochastic Processes in Engineering Systems.* Springer-Verlag, New York, 1985.

[46] U. Grenander and M. I. Miller. Jump-diffusion processes for abduction and recognition of biological shapes. *Monograph of the Electronic Signals and Systems Research Laboratory*, 1991.

[47] R.D. Rabbitt. private communication. *BCL monograph*, 1994.

[48] L. Meirovitch. *Analytical Methods in Vibrations.* The McMillan Co., 1967.

[49] J.N. Reddy. *Energy and Variational Methods in Applied Mechanics.* Wiley-Interscience, J. Wiley and Sons, 1984.

[50] U. Grenander, M.I. Miller, and R. D. Rabbitt. Space time measures for fluid deformable models. *BCL monograph*, 1994.

[51] L. Mittag, M.J. Stephen, and W. Yourgrau. Variational principles in hydrodynamics. In W. Yourgrau and S. Mandlestam, editors, *Variational Principles in Dynamics and Quantum Theory.* Dover Publications, New York, 1979.

[52] U. Grenander, M.I. Miller, and G. Christensen. Deformable anatomical data bases using global shape models: A position paper for the 1992 electronic imaging of the human body workshop. In *Proceedings of the Cooperative Working Group on Whole Body 3-D Electronic Imaging of the Human Body*, Administered by the Crew System Ergonomics Information Analysis Center -Wright Patterson Air Force Base, March 9-11, 1992.

[53] L.E. Malvern. *Introduction to the Mechanics of a Continuous Medium.* Prentice-Hall, 1969.

[54] E.R. McVeigh, R.M. Henkelman, and M.J. Bronskill. Noise and filtration in magnetic resonance imaging. *Medical Physics*, 12 (5):586–591, 1985.

[55] T.J. Schaewe and M.I. Miller. A maximum likelihood reconstruction algorithm for magnetic resonance imaging. In R.H. Schneider, editor, *Medical Imaging IV : Image Formation*, pages 188–194. SPIE 1231, 1990.

[56] M.I. Miller, T.J. Schaewe, C.S. Bosch, J.J.H. Ackerman, R.G. Jost, A.S. Greene, and S.N. Nadel. A mathematical signal model for phase and frequency encoded magnetic resonance imaging data. *Monograph of the Electronics Signals and Systems Research Laboratory*, (ESSRL-91-47), 1991.

[57] D. L. Snyder, Jr. L. J. Thomas, and M. M. Ter-Pogossian. A mathematical model for positron emission tomography systems having time-of-flight measurements. *IEEE Transactions on Nuclear Science*, NS-28:3575–3583, 1981.

[58] J.C. Strikwerda. *Finite Difference Schemes and Partial Differential Equations.* Wadsworth and Brooks/Cole, 1989.

[59] S.J. DeArmond, M.M. Fusco, and M.M. Dewey. *Structure of the Human Brain: A Photographic Atlas.* Oxford University Press, New York, third edition, 1989.

[60] R. Dann, J. Hoford, S. Kovacic, M. Reivich, and R. Bajcsy. Evaluation of Elastic Matching Systems for Anatomic (CT, MR) and Functional (PET) Cerebral Images. *Journal of Computer Assisted Tomography*, 13(4):603–611, July/August 1989.

[61] J.C. Simo. A framework for finite strain elastoplasticity based on maximum dissipation and the multiplicative decomposition: Part i. continuum formulation. *Computational Methods in Applied Mechanics and Engineering*, 66:199–219, 1988.

[62] Berthold Horn. *Robot Vision.* MIT Press, Cambridge, Mass., 1986.

[63] R.E. Morley and T.J. Sullivan. A massively parallel systolic array processor for image processing applications. In *Monograph of Electronics Systems and Signals Research Laboratory*, St. Louis, MO., July, 1987. Washington University at St. Louis.

[64] R.L. Bisplinghoff, J.W. Marr, and T.H.H. Pian. *Statics of Deformable Solids.* Dover Publications, Inc., 1965.

[65] R. Bajcsy and S. Kovacic. Multiresolution Elastic Matching. Technical report ms-cis-87-94, grasp lab 123, Department of Computer and Information Science, School of Engineering and Applied Science, University of Pennsylvania, Philadelphia, PA 19104-6389, 1987.

[66] M.I. Miller, G.E. Christensen, Y. Amit, and U. Grenander. A mathematical textbook of deformable neuro-anatomies. *Biomedical Computer Lab Monograph 548*, 1992.

[67] G. E. Christensen, R. D. Rabbitt, and M.I. Miller. A deformable neuroanatomy textbook based on viscous fluid mechanics. In Jerry Prince and Thordur Runolfsson, editors, *Proceedings of the Twenty-Seventh Annual Conference on Information Sciences and Systems*, pages 211–216, Baltimore, Maryland, March 24-26 1993. Department of Electrical Engineering, The Johns Hopkins University.

[68] K.M. Hanson. Bayesian reconstruction based on flexible prior models. *submitted to the Journal of the Optical Society of America A*, 1992. Los Alamos National Laboratory monograph LA-UR-92-2334.

[69] M. Kamber, D.L. Collins, R. Shighal, G.S. Francis, and A.C. Evans. Model-based 3d segmentation of multiple sclerosis lesions in dual-echo mri data. In Richard A. Robb, editor, *Visualization in Biomedical Computing 1992*, pages 590–600. SPIE 1808, 1992.

Vita

Gary Edward Christensen

Date of Birth:	December 24, 1965
Place of Birth:	St. Louis, Missouri
Undergraduate Study:	Washington University, St. Louis, Missouri B. S. Electrical Engineering, May, 1988 B. S. Computer Science, May, 1988
Graduate Study:	Washington University, St. Louis, Missouri M. S. Electrical Engineering, December, 1989 Washington University, St. Louis, Missouri D. Sc. Electrical Engineering, August, 1994
Professional Societies:	IEEE Tau Beta Pi Eta Kappa Nu, president
Honors/Awards:	Doctoral Fellowship, Dept. of Electrical Engineering, W.U. (1988 – 1992) Outstanding Senior, Dept. of Electrical Engineering, W.U. (1988) Academic Excellence Award, Dept. of Computer Science, W.U. (1988) Robert W. Otto Scholarship, Washington University (1984 – 1988) Scholarship, Monsanto (1984 – 1988) Third Place in Physics, International Science and Engineering Fair (1984) First Place Overall, St. Louis Science Fair (1984) Eagle Scout (1980)

Publications:

Christensen, Rabbitt, and Miller. 3D brain mapping using a deformable neuroanatomy. *Physics in Medicine and Biology*, 39, pp. 609–618, 1994.

Christensen, Miller, and Vannier. A 3D deformable magnetic resonance textbook based on elasticity. *Proceedings of the American Association for Artificial Intelligence, 1994 Spring Symposium: Applications of Computer Vision in Medical Image Processing*, Stanford University, March 21-23, 1994.

Miller, Christensen, Amit, and Grenander. Mathematical textbook of deformable neuroanatomies. *Proceedings of the National Academy of Science*, 90(24) pp. 11944–11948, December 15, 1993.

Christensen, Rabbitt, and Miller. A deformable neuroanatomy textbook based on viscous fluid mechanics. Invited paper. In Prince and Runolfsson, editors, *Proceedings of the Twenty-Seventh Annual Conference on Information Sciences and Systems*, pp. 211–216, Baltimore, Maryland, March 24–26, 1993. Department of Electrical Engineering, The Johns Hopkins University.

Christensen, Miller, Amit, and Grenander. Global shape models for anatomical structures. In Poor and Schwartz, editors, *Proceedings of the 26th Conference on Information Sciences and Systems*, p. 356, Princeton University, March 18-20, 1992.

Grenander, Miller, and Christensen. Deformable anatomical data bases using global shape models: A position paper for the 1992 electronic imaging of the human body workshop. In *Proceedings of the Cooperative Working Group on Whole Body 3-D Electronic Imaging of the Human Body*, March 9-11, 1992.

Christensen. The design of a bit-serial coprocessor to perform multiplication and division on a massively parallel architecture, *Electrical Engineering Masters Thesis*, Washington University, St. Louis, December 1989.

Morely, Christensen, and Sullivan. The design of a bit-serial coprocessor to perform multiplication and division on a massively parallel architecture, *Systolic Array Processors*, Prentice Hall, New York, pp. 497-503, 1989.

Morely, Christensen, Sullivan, and Kamin. The design of a bit-serial coprocessor to perform multiplication and division on a massively parallel architecture, *Frontiers of Massively Parallel Computation*, IEEE Computer Society Press, Washington, pp. 419-422, October 1988.

August, 1994

Short title: **Deformable Shape Models for Anatomy** Christensen, D. Sc. 1994