3-D Deformable Registration Using a Statistical Atlas with Applications in Medicine

Mei Chen

CMU-RI-TR-99-20

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Robotics

October, 1999

To grandma, who teaches me hardship can be endured with grace

To grandpa, who nurtured my passion for art

To mom, who proves women can do rocket science

To dad, who gives me love unconditional

To auntie, who shows me beauty of compassion

To brother and sister, who keep reminding me how long I have been in grad school

Abstract

Registering medical images of different individuals is difficult due to inherent anatomical variabilities and possible pathologies. This thesis focuses on characterizing non-pathological variations in human brain anatomy, and applying such knowledge to achieve accurate 3-D deformable registration.

Inherent anatomical variations are automatically extracted by deformably registering training data with an expert-segmented 3-D image, a digital brain atlas. Statistical properties of the density and geometric variations in brain anatomy are measured and encoded into the atlas to build a statistical atlas. These statistics can function as prior knowledge to guide the automatic registration process. Compared to an algorithm with no knowledge guidance, registration using the statistical atlas reduces the overall error on 40 test cases by 34%.

Automatic registration between the atlas and a subject's data adapts the expert segmentation for the subject, thus reduces the months-long manual segmentation process to minutes. Accurate and efficient segmentation of medical images enable quantitative study of anatomical differences between populations, as well as detection of abnormal variations indicative of pathologies.

Acknowledgments

I wish to thank my advisors, Professor Takeo Kanade and Dr. Dean Pomerleau, for their support and advice during my graduate studies. I also wish to thank other members of my committee, Professor W. Eric L. Grimson and Dr. Steven Seitz, for their valuable comments concerning my work.

I owe much to the Brigham and Women's Hospital of the Harvard Medical School for the brain atlas. I am also grateful to those who provide me MRI data: Kate Fissell in the Carnegie Mellon Psychology Department, Dr. Daniel Rio and Mike Kerich in the National Institute of Health, Dr. Matcheri Keshavan, Dr. Michael Debellis, Elizabeth Dick, Stephen Spencer, and Amy Boring in the Western Psychiatric Institute and Clinicl, Dr. Keith Thulborn and Steve Uttecht in the Magnetic Resonance Imaging Center of the University of Pittsburgh Medical School, and Greg Hood of the Pittsburgh Supercomputing Center. I would like to thank Dr. Williams Rothfus at Allegheny General Hospital for helpful discussions in the early stages of this work.

Interaction with members of the Vision and Autonomous Systems Center at Carnegie Mellon has been indispensable. From them I have not only learned about different aspects of computer vision and image understanding, but also how to communicate my thoughts and ideas. Special thanks to Shumeet Baluja, Parag Batavia, John Hancock, David LaRose, Daniel Morris, Jim Rehg, Jeff Schneider, Henry Schneiderman, David Simon, and Todd Williamson for the discussions, for proofreading my paper drafts, and for attending my practice talks.

My time at Pittsburgh has been enriched by too many people to name them all. However, I want to thank my friends who offered generous help during my recovery from two accidents: Atsushi Yoshida, Mike (Chuang) He, Kate Fissell, Maria Mosely, Mei Han and Wei Hua; Tech Khim Ng, Belinda Thom, Taku Osada, Stephanie Land, Kimmaree Murday, Cecilia Wandiga, John Hancock, Dirk Langer, Rahul and Gita Sukthankar, Parag Batavia, Daniel Morris, Yokiko Wada, and Joyoni Dey. I also wish to express my gratitude to friends who offered consistent support from afar: Richard (Xiaosong) Zhang, Jane (Jingxi) Chu, Jennifer Kay, Xuemei Wang, Yolanda Wang, Takayuki Yoshigahara, Mark Maimone, James Lien, and John Long.

Contents

1. Introduction

2. 3-D Hierarchical Deformable Registration

3. Building a Statistical Atlas

4. Registration Using the Statistical Atlas

5. Model Neighborhood Context

6. Implementation Details

7. Quantitative Study of the Anatomy

8. ADORE: Anomaly Detection thrOugh REgistration

9. Conclusion

Appendix

A. Levenberg-Marquart Algorithm

B. Principal Component Analysis

C. Segmentation Comparison

Bibliography

CHAPTER 1 Introduction

Rapid development in medical imaging devices such as magnetic resonance imaging (MRI), computer assisted tomography (CT), and positron emission tomography (PET), has brought a revolution in the medical domain. For the first time, physicians can view people's internal structures in three dimensions, in a non-invasive way.

Among current medical imaging modalities, MRI reveals soft tissues the best. An MRI is an image volume consisting of a series of parallel cross-sections along one of three principal axes: coronal, axial, or sagittal. Example slices of a human brain MRI taken along each axis are shown in Figure 1. The number of cross-sections taken depends on the purpose of the MRI. Within the realm of MRI there are three common sub-modalities: T1-weighted, T2-weighted, and proton density. While tissue density is reflected in image intensity, the same tissue can appear of different intensities in different image modalities, or due to gain artifacts in the scanner.

It is of great interest in medicine to study anatomical structures and to detect possible anomalies. While qualitative analysis may be sufficient for disease diagnosis, quantitative analysis of specific anatomical regions is required for longitudinal monitoring of disease

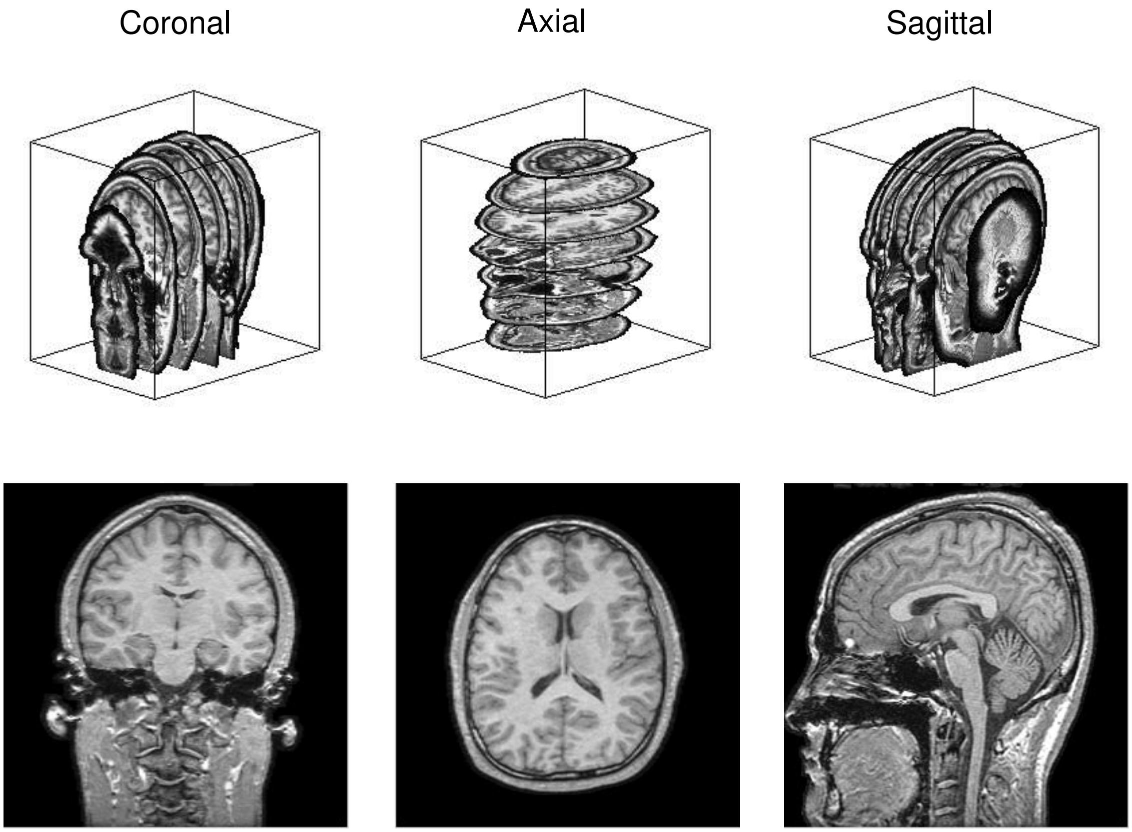

Figure 1: Example cross-sections of brain MRI taken along each of the three principal axes.

progression or remission, pre-operative evaluation and surgical planning, radiosurgery and radiotherapy treatment planning, mapping of functional neuroanatomy of sensorimotor and cognitive processes, as well as the analysis of neuroanatomical variability among normal brains.

Quantitative regional analysis is possible only with explicit segmentation to separate and identify anatomical structures. Traditionally, the segmentation of anatomical structures is done manually. The process is painstakingly slow and involves a cadre of workers who laboriously outline and edit each region of interest [53]. The left image in Figure 2 is one slice of a T1-weighted coronal MRI of a non-pathological brain. The MRI volume contains 123 slices, and each slice is a 256x256 pixel matrix. Each pixel is a 0.9375x0.9375 mm^2 square, and all slices are of thickness 1.5 mm. The right image in Figure 2 is a color coding of the classification done by 7 operators (courtesy of Brigham and Women's hospital of Harvard Medical School). This is valuable data, but the process is too time-consuming to study anatomical properties of a population. In a recent study, "one man-month" was required to outline the thalami of 200 subjects [25]. Moreover, errors occur because of human subjectivity in slice selection, structure interpretation, poor software interface design, poor hand-eye coordination, low tissue contrast, and unclear tissue boundaries caused by *partial volumes* (individual pixels contain more than one tissue type). Human performance is also inherently inconsistent. It is reported that segmentations of the same brain structure given by the same neuroanatomist three months apart can differ by 4.9% in volume, and with only 87.8% overlap [14]. As for segmentations given by different human operators, the inter-operator reliability can range from 64% to 87%.

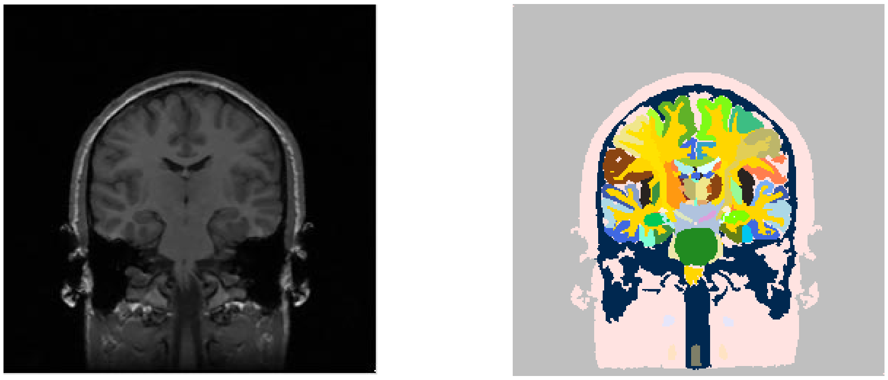

Figure 2: One cross-section of a T1-weighted MRI of a normal brain (left), and the corresponding color-coded classification of brain tissues (right).

Although MRI data is digital by nature, it was not until mid-1980s for computers to be powerful enough to process and analyze them, especially in 3-D. However, from the dawn of its emergence, research in biomedical image analysis and computer assisted intervention has been vigorous, and has raised much enthusiasm among medical researchers.

1.1 Problem Definition

This thesis tackles the problem of structural segmentation and classification. The technique is 3-D deformable registration between an atlas and the input of a particular subject's data. An atlas is reference data with anatomical localization or functional interpretation, such as the aforementioned non-pathological brain MRI volume with structural segmentation and classifications given by human experts. Figure 3 illustrates the approach to the problem: given a subject's brain MRI volume, the atlas is deformed in 3-D to match the input data. Under the same deformation, classification labels of anatomical structures in the atlas can also be warped in 3-D to register with corresponding structures of the subject. In this way, anatomical information in the atlas is *customized* for the subject. Therefore, the problem of structural segmentation and classification becomes one of finding the optimal 3-D deformation that best registers the atlas and the subject image volumes.

1.2 Importance of registration

As discussed above, 3-D deformable registration between an atlas and a particular subject's image data customizes the segmentation and labelling of anatomical structures for the subject. This facilitates pathology detection and quantitative study of individual structures. Given data from different populations, such as normal control subjects versus schizophrenics, this also allows statistical study of possible anatomical differences between populations. Further, while deformable registration between different subjects' data enables cross-subject diagnosis, it is also essential for similar case retrieval in medical image databases. Applications of much interest also include registration of the same subject's data over time. This supports post-treatment analysis and longitudinal study of anatomical changes, such as the relation between brain loss and aging.

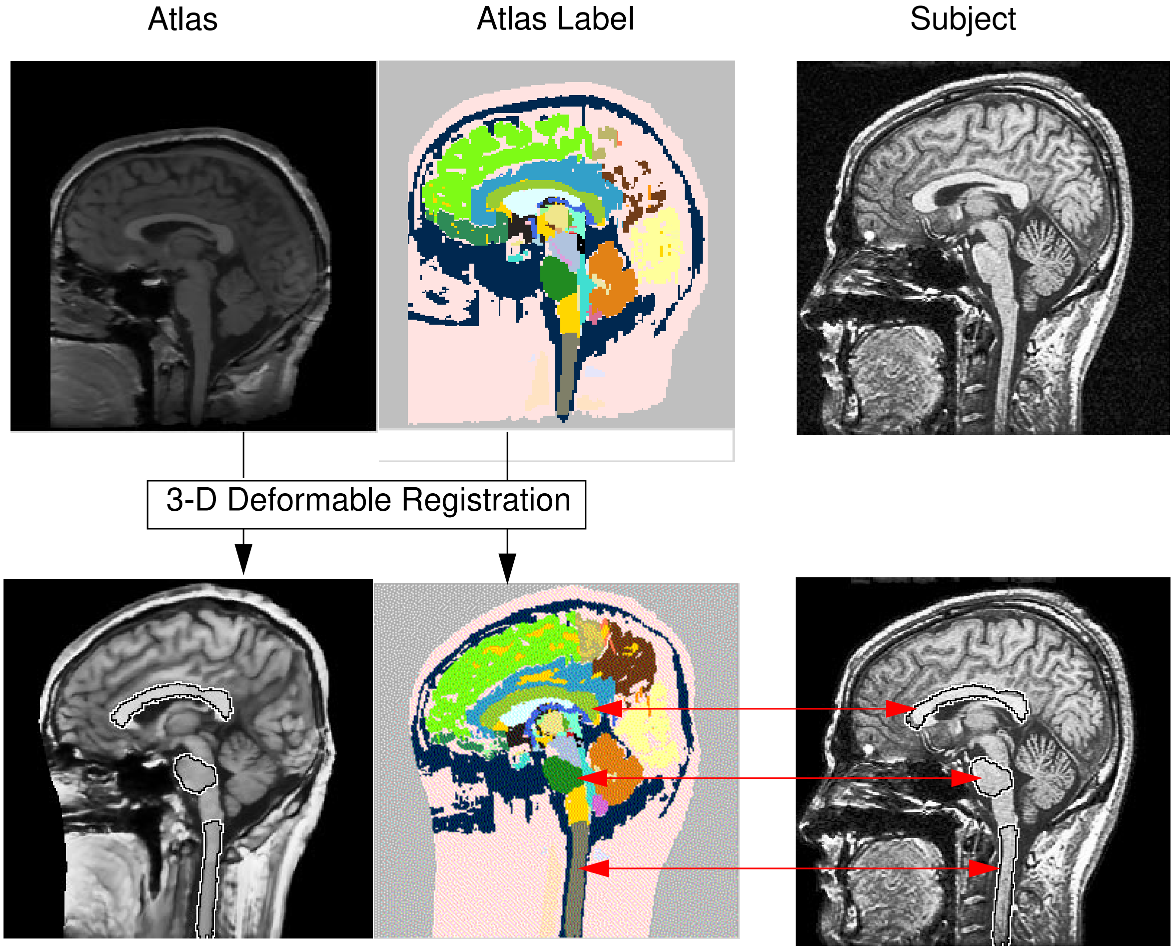

Figure 3: Illustration of classifying a subject's anatomical structures through registration with the atlas. The atlas is deformed in 3-D to match the input data. Under the same deformation, classification labels of anatomical structures in the atlas can also be warped in 3-D to register with corresponding structures of the subject. (the atlas volume doesnot contain partions the nose, mouth, etc.)

1.3 Research in registration

Registration of medical images has been an active research area--the comprehensive survey article by van den Elsen et al. [101] lists 161 citations. Two popular schools of registration using image properties are the *feature-based* approach and *the voxel-based* approach.

Feature-based methods attempt to extract the contours or surfaces, i.e. features, of anatomical structures in the images to be registered, and find the correspondence between them, [2], [11], [14], [28], [33], [80], [94]. They are efficient in representation and independent of imaging modality. However, feature-based registration is critically dependent on the quality of feature extraction, which is not trivial since anatomical structures tend to have complex shapes and ill-defined boundaries. Human interaction is generally necessary to help select and extract features or to guide the registration procedure. Consequently, it is prone to user subjectivity, is inconvenient, and can be time-consuming.

The landmark work in feature-based medical image registration was done by Bajcsy et al. They assume anatomical variations between individuals can be modeled by elastic deformation, and model the atlas as a physical object with elastic properties. The registration procedure first extracts contours of anatomical structures in the subject's image, then elastically deforms contours in the atlas to match with those in the subject [2], [33]. However, without user interaction, their atlas can have difficulty matching complicated object boundaries. In addition, this method is computationally expensive, and requires time-consuming interactive preprocessing that involves removing the skull in the image volume. Davatzikos proposed a method that uses outer cortical surface mapping to drive a 3-D deformation with linear elasticity [20]. This method also requires human interaction in outlining features such as sulci and fissures. Sandor and Leahy automated the feature extraction by using boundary-finding and a morphological procedure, whose result is highly dependent on the accuracy of detected boundaries and may still require manual corrections [80]. Szeliski and Coughlan used tensor product splines to represent transformations between two 3-D anatomical surfaces, and introduced octree splines for fast computation of the distance between surfaces [90]. Thompson and Toga employed 3-D active surfaces initialized by a hybrid surface model based on superquadrics and spherical harmonics, and use surface deformation maps to drive the volumetric warp [96]. Collins et al. presented an automatic 3-D registration method that matches features using a hierarchical local affine transformation [14]. Feldmar and Ayache developed a surface to surface non-rigid registration scheme using locally affine transformation [30]. There has also been research on feature selection for medical image registration, such as 2-D and 3-D ridge seeking operators suggested by Maintz et al. [63], and near-optimal intra data selection developed by Simon et al. for fast and accurate 3-D rigid-body registration [84].

As an alternative, *voxel-based* algorithms obviate the need for explicit feature extraction or segmentation. Christensen et al. [11], [12] represented small deformations between brain volumes using a linear-elastic model, and large magnitude deformations with a viscous fluid model. The fluid dynamic model allowed large deformations by relaxing the motion-restraining stress over time. However, the computation took hours on a massively parallel computer, and the relevance of the physical model to registration is still questionable. Thirion assumed image volumes to be registered had similar intensity distributions and were already grossly aligned, and modeled the deformation as a 3-D voxel flow. He used the gradient of the non-deforming volume instead of the deforming one in determining the deformation, because computation of the latter requires tri-linear interpolation of each voxel's gradient. However, this quicker method will be erroneous when the atlas does not resemble the subject closely. Also, it is prone to failure when there are large intensity differences between image volumes [92]. Vemuri et al. formulated registration as a motion estimation problem, and proposed a hierarchical optical flow motion model which represented the flow field using a B-spline basis [102]. Instead of relying on similar intensity distributions, Woods et al. assumed that voxel intensities in accurately aligned image volumes could be related by a single multiplicative factor. They computed this ratio for each pair of corresponding voxels, and minimized its standard deviation across all pairs to determine the optimal affine transformation between the image volumes [111].

Registration algorithms using intensity correspondence have shown encouraging results. However, they are problematic when there are significant intensity differences. Moreover, they can only register multi-modal data if there exists a linear intensity mapping between images from different modalities. Viola et al. [103] and Maes et al. [62], independently investigated the mutual information criterion (MI). MI measures the statistical dependence between two random variables, or the amount of information that one variable contains about the other. Modeling image intensities as random variables, registration using the MI criterion assumes that the statistical dependence between corresponding voxel intensities is maximal if the image volumes are geometrically aligned. Because no assumptions are made regarding the nature of this dependence, the MI criterion is highly data independent. Both groups presented accurate results on intra-subject (the same person) intermodality registration. Pokrandt et al. studied multi modality registration using Gaussian entropy as an evaluation function [75]. Rangarajan and Duncan applied MI criterion to feature-based registration [79]. However, so far registration using an MI criterion has been limited to affine transformations, due to the difficulty of defining analytic expressions of image gradient during deformation. Meyer et al. explored MI-based registration allowing affine transformation and thin-plate spline warp, but with prohibitive computational expense [69]. The potential of applying MI criterion to inter-subject deformable registration remains to be further studied.

1.4 Difficulties in registration

It is clear from the above discussion that medical image registration is far from a solved problem. Figure 4 shows the result of registering the atlas with a particular subject's data, using Thirion's algorithm [92]. Segmentations of brain structures in the atlas are adapted for the subject, and contours of several structures are overlaid on the subject's image volume. Note that even Thirion's method has difficulty matching these two image volumes: there is significant mis-match between the adapted segmentation and the subject's brain structures.

Many factors hamper accurate registration, such as image degradation caused by artifacts and noise, low tissue contrast, and blurring due to partial volume effects (tissue-mixing within a single voxel). One major factor comes from image variations caused by different acquisition processes, another major factor relates to image variations because of inherent anatomical differences between individuals. In the example in Figure 4, the sub-

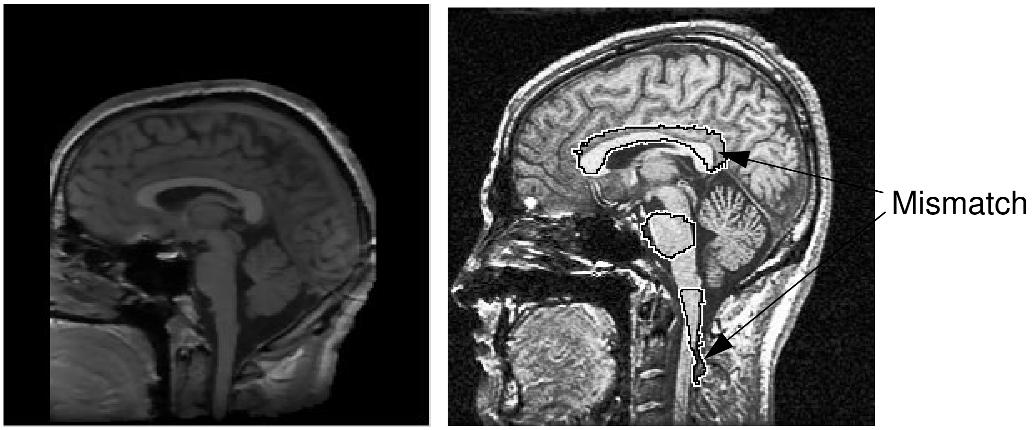

Figure 4: The result of matching the atlas (left), to a particular subject (right), using Thirion's method. Segmentations of brain structures in the atlas are adapted for the subject, and outlines of several examples are overlaid on the subject's image data.

ject's brain structures are of different shape and size from those in the atlas; moreover, the subject's image data is considerably brighter than the atlas, which poses more difficulty for registration methods based on intensity correspondence.

1.4.1 Image acquisition process

Currently there is no enforced standard in the image acquisition process of MRI. The interface between the MRI device and the subject is a horizontal tube, in which the subject should lie still during the whole imaging process. There is no fine calibration of the subject's 3-D position and orientation. As a result, the axis along which the images are taken is generally at an angle to the principal axis. Further, each cross-section in the image volume represents the average response of a partial volume, the thickness of which may vary. An MRI can focus on a sub-section of the brain if so desired. Moreover, the numerous parameter settings and the drifting of the magnetic field can cause inhomogeneities in image intensities.

These factors cause variations in the orientation, position, scale, resolution, and intensity distributions between different image volumes, as illustrated in Figure 5. These variations affect the whole image volumes, i.e. their effects are global. They are regarded as extrinsic image variations because they are caused by factors external to the subject. A registration algorithm needs to cope with these variations in order to perform well.

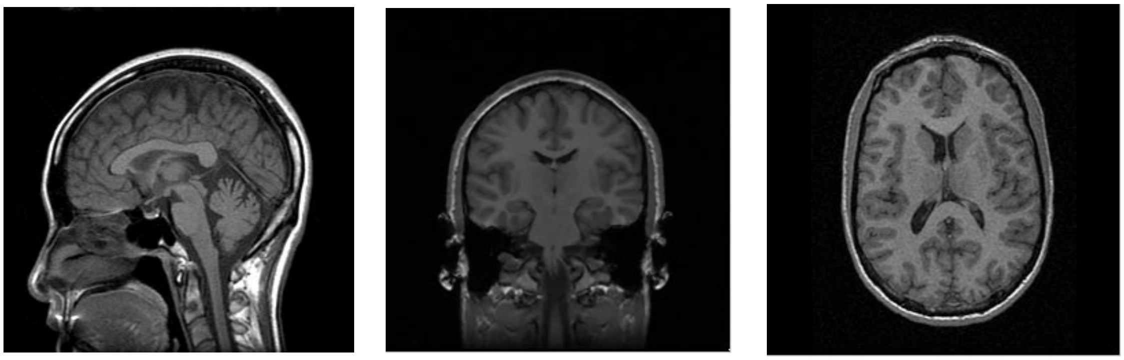

Figure 5: Example cross-sections from different MRI volumes. There exist differences in orientation, scale, and intensity distribution.

1.4.2 Non-pathological anatomical differences

Due to genetic and life-style factors, there are also inherent non-pathological differences in the appearance and location of anatomical structures between individuals. Figure 6 displays cross-sections from two non-pathological brains' MRI volumes. The example structure, corpus callosum, has different shape, size, and location in these two brains. These variations are characteristic for the particular structure of the individuals, i.e. they are local and intrinsic. For registration algorithms that assume the same structure should have the same appearance or location in different individuals, these inherent variations make accurate inter-subject registration difficult.

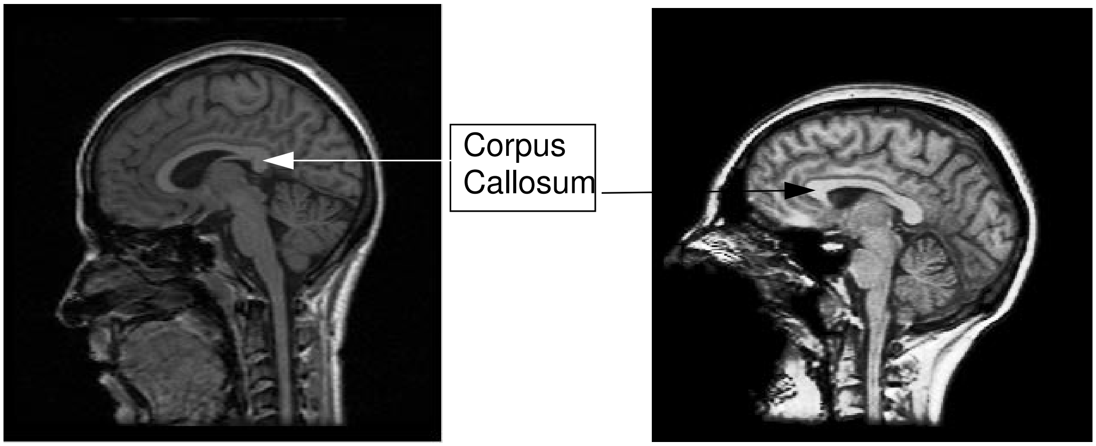

Figure 6: Innate variations between individuals. Note the differences in the size, shape, and location of corpus callosum

1.5 Bootstrap Strategy

The goal of this thesis is to achieve accurate automatic segmentation via atlas-subject 3-D deformable registration, so as to facilitate applications in medicine. The core approach is a closed-loop bootstrap framework for characterizing the appearance of anatomical structures and their non-pathological variations between individuals, then applying such knowledge to improve registration performance, and further using the improved registration to refine the anatomical characterization, which helps obtaining more accurate registration. This closed-loop bootstrap process can keep going as more image data becomes available, Figure 7 illustrates the concept. Knowledge of anatomical variations not only allows the registration algorithm to tolerate non-pathological differences that exist between individuals, but also facilitates anomaly detection and quantitative study of anatomical differences between populations, such as *normal* control subjects versus schizophrenics. Algorithms developed in this thesis apply to any imaging modalities, however, due to limited data source, most experiments were conducted on T1-weighted magnetic resonance imaging (MRI) of human brain.

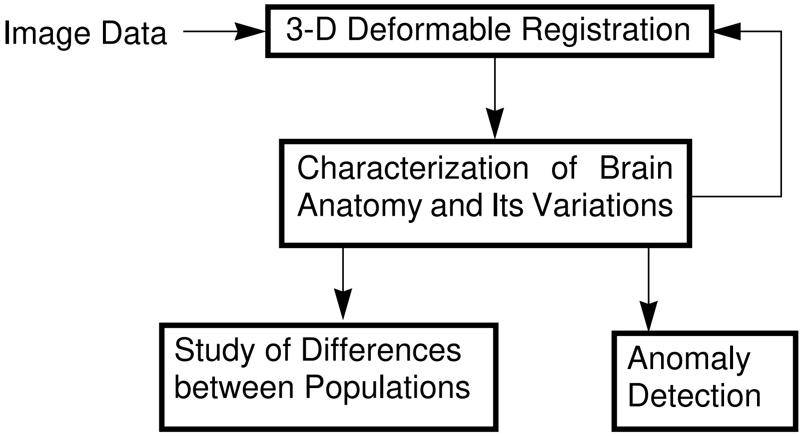

Figure 7: Bootstrap Strategy

1.6 Dissertation overview

Figure 7 presents a seemingly cyclic problem: knowledge of the anatomy and its variations will provide guidance to deformable registration, however, a deformable registration algorithm is necessary to extract the knowledge of variations. Therefore, Chapter 2 will first introduce a 3-D deformable registration method without knowledge guidance, and give quantitative evaluations of its performance. This algorithm will bootstrap the closed-loop in Figure 7. Chapter 3 will focus on the extraction and characterization of anatomical variations between individuals. Then Chapter 4 closes the loop with a registration algorithm guided by this knowledge of anatomical variations, and compares its performance to that of the method in Chapter 2. Further improvement on knowledge representation and application will be discussed in Chapter 5. Chapter 6 covers details on current implementation of the algorithms. The following two chapters are devoted to explorations of medical applications, with Chapter 7 showing supportive results on quantitative study of anatomical differences between populations, and Chapter 8 presenting approaches to anomaly detection. In the end, Chapter 9 concludes this thesis by highlighting the contributions and discussions on future research directions.

CHAPTER 2 3-D Hierarchical Deformable Registration

This chapter introduces a 3-D deformable registration method that does not utilize knowledge of anatomical variations. This algorithm will bootstrap the closed-loop of extracting knowledge for guiding registration, and function as a baseline for performance comparisons.

The 3-D registration algorithm matches image volumes based on intensity correspondence. The task of registration is to find a 3-D deformation function $\mathcal{D}$ that maps any particular atlas voxel (x_A, y_A, z_A) to the corresponding voxel (x_S, y_S, z_S) in the subject's volume, illustrated in Figure 8. As mentioned in section 1.4, there are not only intrinsic anatomical differences between people's MRIs, but also extrinsic differences resulted from different image acquisition processes. Therefore, a hierarchical scheme is adopted to first address the extrinsic variations, and then deformably match corresponding anatomical structures between individuals. Note that this is an approach employed in this thesis, other methods also apply to the framework presented in Section 1.5.

$$(x_A, y_A, z_A) \xrightarrow{\mathcal{D}} (x_S, y_S, z_S)$$

$$\mathcal{D}(x_A, y_A, z_A) = (x_S, y_S, z_S)$$

Figure 8: The concept of registration.

Because different image acquisition processes can result in different intensity distributions, intensity equalization procedures are necessary to ensure that corresponding anatomical structures have comparable intensities. However, this is a chicken-and-egg problem because anatomical correspondence is the goal of registration, and therefore unknown. One solution is a hierarchical approach that starts with a crude intensity normalization, and then refines the equalization scheme as better anatomical correspondence is achieved in the registration process. This is a novel approach compared to existing related work.

This 3-D hierarchical deformable registration algorithm involves two interleaving hierarchies, one is a hierarchy of intensity equalization schemes, the other is a hierarchy of geometrical deformable models, as illustrated in Figure 9. It first grossly equalizes the intensities between the image volumes, and applies transformations to globally align them (Section 2.1). Based on this initial alignment, it then uses a more localized intensity normalization, and employs a smooth deformation to roughly match the anatomical structures in the image volumes (Section 2.2). This correspondence allows a more informed intensity transformation between the two image volumes, and a fine-tuning deformation adjusts the correspondence of anatomical structures more precisely (Section 2.3). This algorithm is automated by a random initialization method (section 6.3). Iterative optimization algorithms are used to determine the deformation parameters at all levels.

2.1 Global alignment with whole volume intensity equalization

The first level deformable model adjusts the extrinsic variations between the atlas and the subject's image volume. Because the extrinsic variations correspond differences in orientation, position, and scale of the image volumes, a similarity transformation composed of 3-D rotation, translation, and uniform scaling is sufficient to compensate for them.

As mentioned earlier, different imaging processes may result in different intensity distributions in the atlas and the subject volume. This difference can make a method using intensity correspondence unreliable, as shown in the example in Figure 4. Before global alignment, the image volumes can be of significantly different orientation, position, and scale. With unknown geometrical correspondence, the first level intensity equalization is a whole volume intensity normalization. It equalizes the mean and standard deviation of the intensities in the head volumes to roughly correct the intensity discrepancy. The head volumes are separated from the background in preprocessing (Section 6.1).

2.1.1 Representing similarity transformation

Figure 10 shows the coordinate systems used in the similarity transformation, T. The origins of the coordinate systems in the atlas and the subject volume are placed at their centroids (center of mass of the head volumes). The Z axis coincide with the axis along which the cross-sections were scanned. Note that the Z axis do not necessarily coincide with any principal axes, as discussed in section 1.4.1. The atlas is first rotated about its origin to the same orientation of the subject volume, then uniformly scaled about its origin to be of the

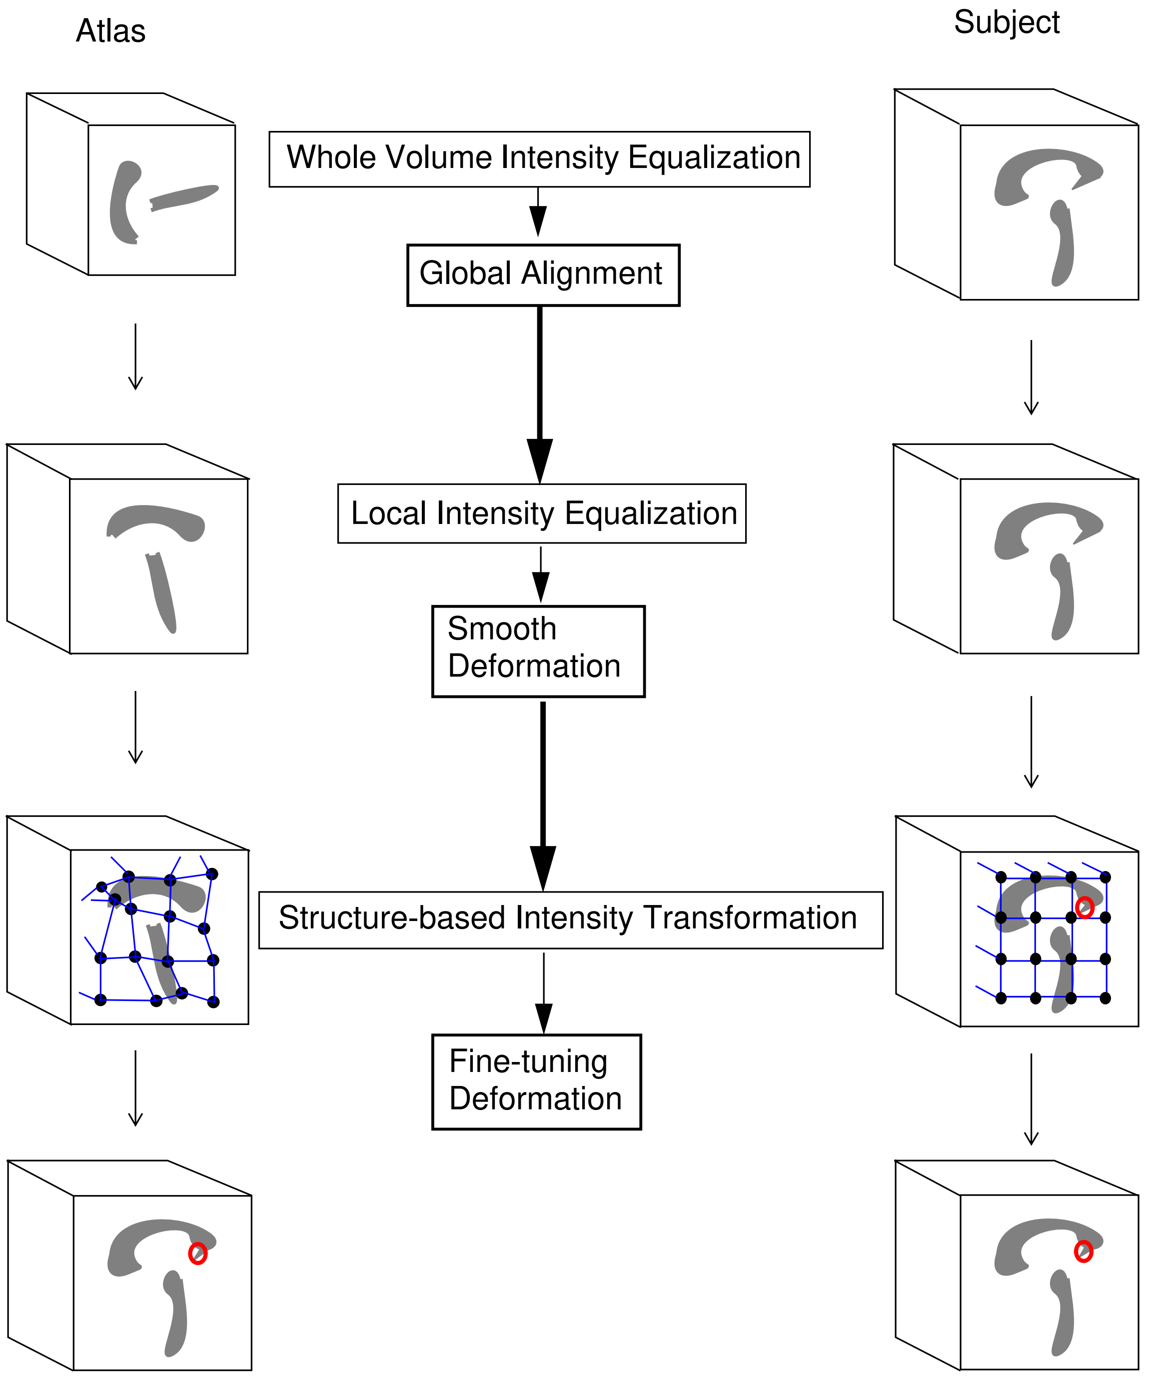

Figure 9: 3-D Hierarchical Deformable Registration with two interleaving hierarchies: a deformation hierarchy and an intensity equalization hierarchy.

same overall size, and then translated to align with the subject volume. The similarity transformation T has 7 degrees of freedom.

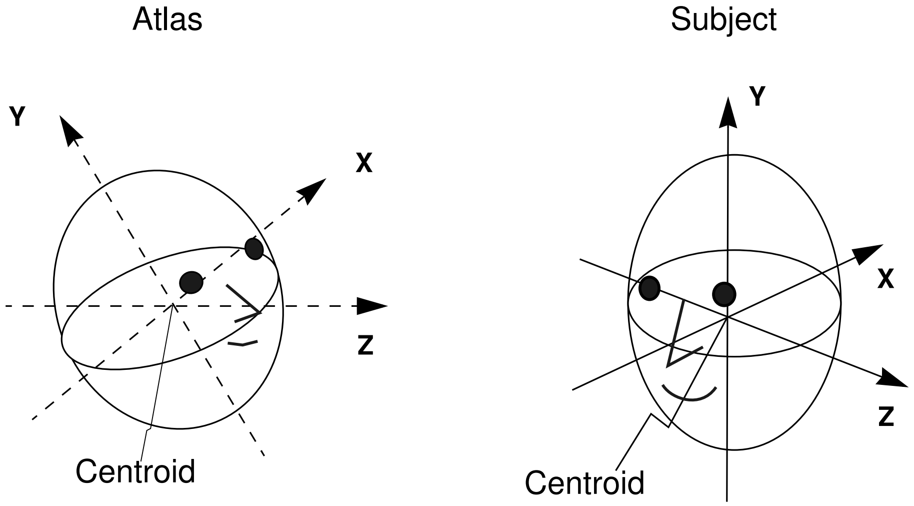

Figure 10: Coordinate systems used in the similarity transformation.

2.1.2 Determining similarity transformation

Because the atlas and the subject are inherently different, there is no similarity transformation that exactly matches them. The best transformation only minimizes the differences. The quality of a transformation is measured by the sum of squared differences (SSD) between the intensities of geometrically corresponding voxels in the image volumes [111].

Suppose $I_s(x, y, z)$ is the intensity of voxel $[x, y, z]_s$ in the subject's volume. The similarity transformation T matches $[x, y, z]_s$ to $T(x, y, z)_a$ in the atlas, and $I_a(T(x, y, z))$ is the intensity at $T(x, y, z)_a$ in the atlas. The squared difference between $I_s(x, y, z)$ and $I_a(T(x, y, z))$ is summed over the whole volume to compute SSD:

$$SSD = \sum_{(x, y, z)} (I_s(x, y, z) - I_a(T(x, y, z)))^2 \qquad (1)$$

SSD is a function of similarity transformation *T*. To find the optimum transformation, Levenberg-Marquardt non-linear optimization algorithm [76] is used to iteratively adjust *T* to reduce SSD. Figure 11 is a block diagram illustrating the iterative process for determining the best transformation. More discussion on the Levenberg-Marquardt non-linear optimization algorithm can be found in Appendix A.

The iteration continues until changes in *T* are below a preset criteria. This criteria is not particularly critical since the similarity transformation is meant only to bring the image volumes into approximate alignment. The registration will be refined by further steps in the deformation hierarchy. The criteria used in this thesis requires that the change in 3-D rotation between iterations be less than 2 degrees about each axis, that the change in uniform scaling between iterations be below 1%, and that the change in 3-D translation be smaller than 1 pixel along each direction. Multi-resolution processing and stochastic sampling are employed for efficiency and to help prevent the optimization from becoming trapped in local minima (section 6.2).

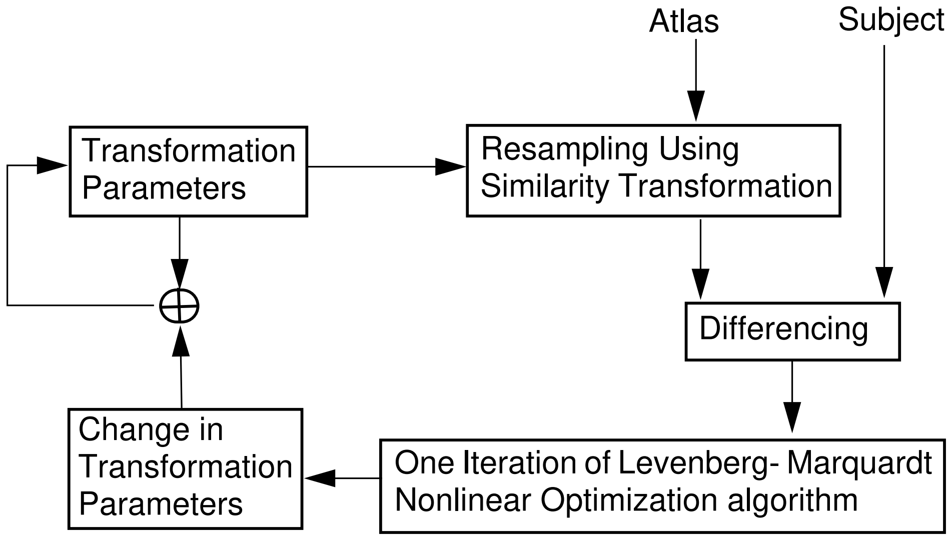

Figure 11: Block diagram of determining global alignment

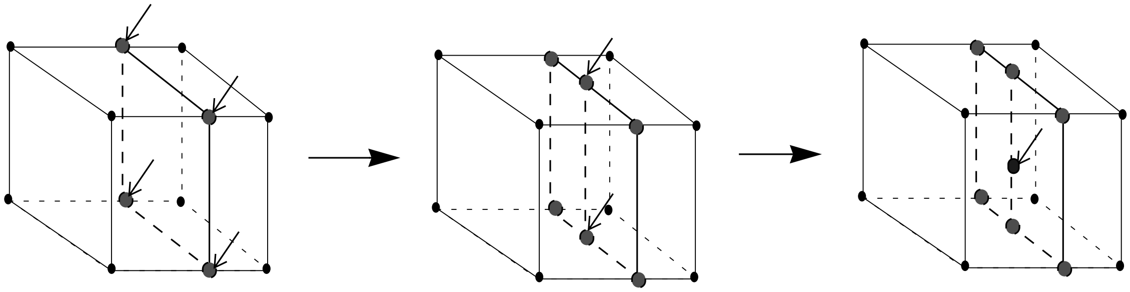

Linear interpolations among 8 neighbors along one axis gives 4 intermediate values

Linear interpolation among the 4 intermediate values along another axis then yields 2 intermediate values

Linear interpolation between the 2 intermediate values along the third axis gives the voxel intensity

Figure 12: Trilinear Interpolation gives the intensity of a voxel with non-integral coordinates by doing linear interpolations among its 8 bounding neighbors along each of the three axes.

Note that the transformed atlas coordinates, $T(x, y, z)_a$, may not be integral, and $I_a(T(x, y, z))$ will not be given by the original data. In this case, tri-linear interpolation is used to determine the voxel's intensity from its eight bounding neighbors, see Figure 12. If $T(x, y, z)_a$ falls outside of the volume, that voxel is ignored for the SSD computation.

Figure 13 shows an example of globally aligning the atlas to a particular subject. The first row are example slices of the atlas and the subject's volume before the global alignment. The atlas is then resampled to match the subject's volume via a similarity transformation, and corresponding slices are shown on the second row. Note that the atlas is now of the same orientation, scale, and position as the subject's volume. The segmentation of one brain structure, the corpus callosum, is adapted for the subject under the same transformation. The outline of the corpus callosum is shown in the resampled atlas, and directly applied to the subject's volume. Note that although the two volumes are grossly aligned, the example structure does not match well with its counterpart. This is because there exist intrinsic anatomical differences between individuals, as discussed in Section 1.4.2, and global alignment only addresses extrinsic differences caused by separate image acquisition processes.

2.2 Smooth deformation with local intensity equalization

The global alignment adjusts extrinsic variations between image volumes, but cannot address intrinsic variations between individual structures, as shown in Figure 13. Transformations that allow local deformations are necessary. An intuitive solution is to allow each voxel to shift freely in 3-D space to align with its counterpart. But this will require the voxels' initial positions to be close to their desired positions so as not to be trapped in local minima. Since the global alignment cannot guarantee a precise enough initial correspondence for individual structures, the second level deformable model takes an intermediate step which allows 3-D deformation at a local neighborhood level.

The intensity equalization scheme can be refined now that there is more information on the correspondence between the two volumes. The second level intensity equalization evens the intensity mean and standard deviation between the overlapping portions of the atlas and the subject volume after they are globally aligned. Similar to the case in global alignment, intensity differences between corresponding voxels in the subject volume and the atlas act as the deforming force, which shift voxel neighborhoods in 3-D space to align with their counterparts.

2.2.1 Representing smooth deformations

To represent local deformations at a neighborhood level, the smooth deformation procedure uses a 3-D control grid which is a coarser grid than the voxel grid. Control grids used in this thesis are generated by regularly sub-sampling the voxel grid, so each *cell* in the control grid is a parallelepiped or a cube. Vertices of the control grid are control points that can shift independently in 3-D space. 3-D displacements of the control points are deformation

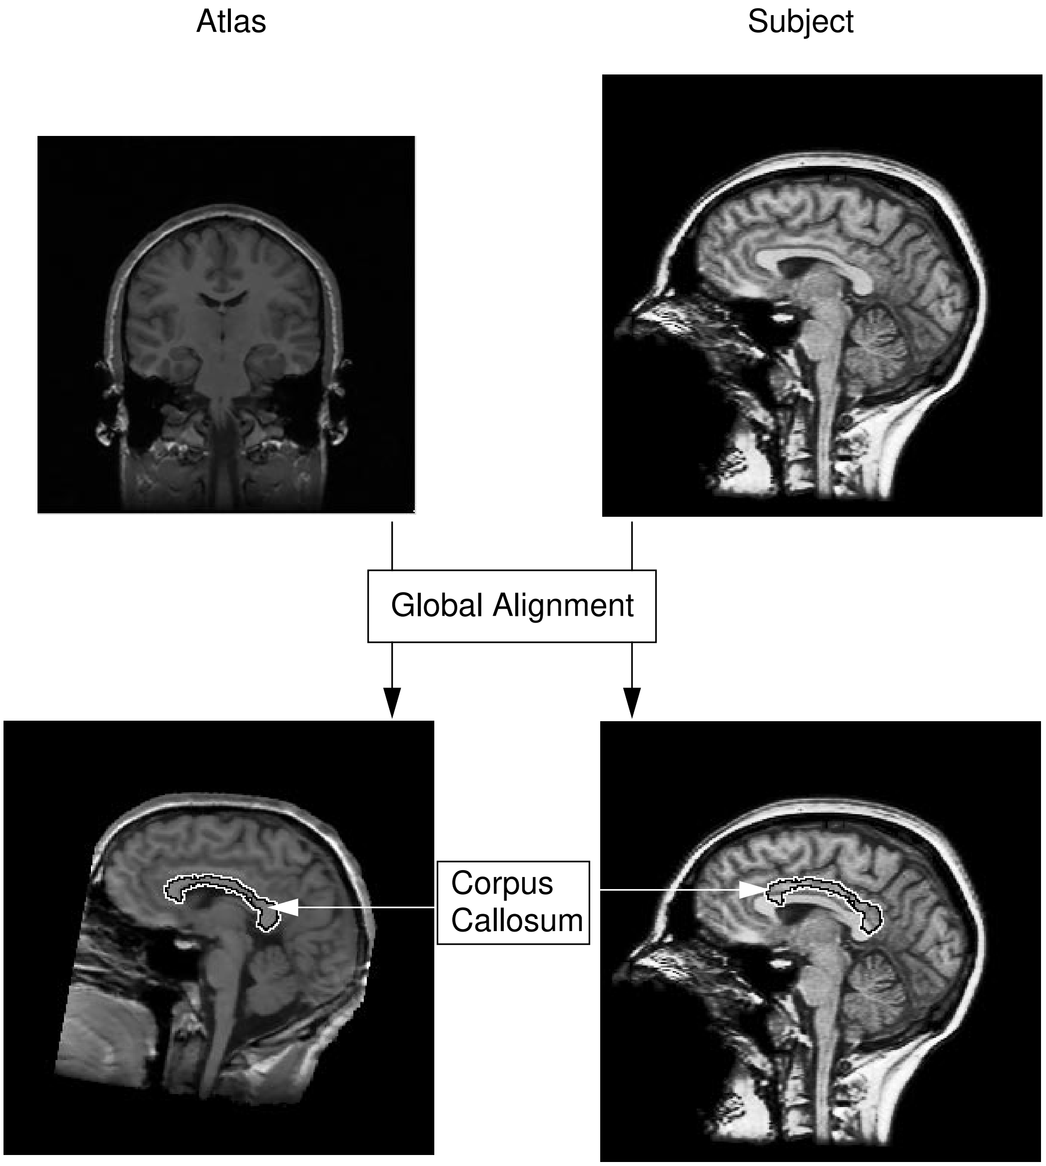

Figure 13: Global alignment matches the atlas to the subject's volume, and adapts segmentations of brain structures for the subject. However, individual structures do not align well.

parameters which determine the displacements of the voxels they enclose. This imposes an implicit smoothness on the displacement field. Since the number of control points is orders of magnitude lower than the number of voxels, this representation makes the deformation optimization process more stable. Figure 14 illustrates the 3-D smooth deformation. Adaptable multi-resolution processing and stochastic sampling are again employed for efficiency and to help prevent the optimization from becoming trapped in local minima (section 6.2). The highest resolution control grid employed is 7x7x7. Szeliski used similar approaches in 2-D image registration and 3-D surface registration [89], [90]. Vemuri et al. proposed an analogues scheme for motion analysis [102]. Collins et al. employed 3-D grids for feature-based registration, though deformation of the grid is constrained to a local affine model [14].

2.2.2 Estimating smooth deformation

Similar to the case for global alignment, the goodness of the smooth deformation is measured by the SSD between intensities of corresponding voxels in the atlas and the subject volume. For a voxel at $[x, y, z]_s$ in the subject volume, suppose it belongs to the *i*th control cell $Cell_s[i]$. The control cell of the same index in the atlas is $Cell_a[i]$. From the relative position of voxel $[x, y, z]_s$ with respect to the eight vertices of $Cell_s[i]$ in the subject's volume, the location of its corresponding voxel in the atlas, $S(x, y, z)_a$, can be determined: assume the same relative position holds between voxel $S(x, y, z)_a$ and control cell $Cell_a[i]$ in the atlas, tri-linearly interpolate the position $S(x, y, z)_a$ from the eight vertices of $Cell_a[i]$. Here *S* denotes smooth deformation. In the case that $S(x, y, z)_a$ does not fall on the voxel grid, the intensity of voxel $S(x, y, z)_a$ is tri-linearly interpolated from its eight neighboring voxels in the atlas. Note that the intensity SSD is computed over all voxels in the atlas and the subject volume, not just for the control points.

If the atlas completely aligns with the subject volume, intensities between corresponding voxels should be equal. However, in practice they may differ. The best deformation minimizes the intensity SSD. Similar to section 2.1.2, a Levenberg-Marquardt iterative non-linear optimization method is used to determine the best smooth deformation parame-

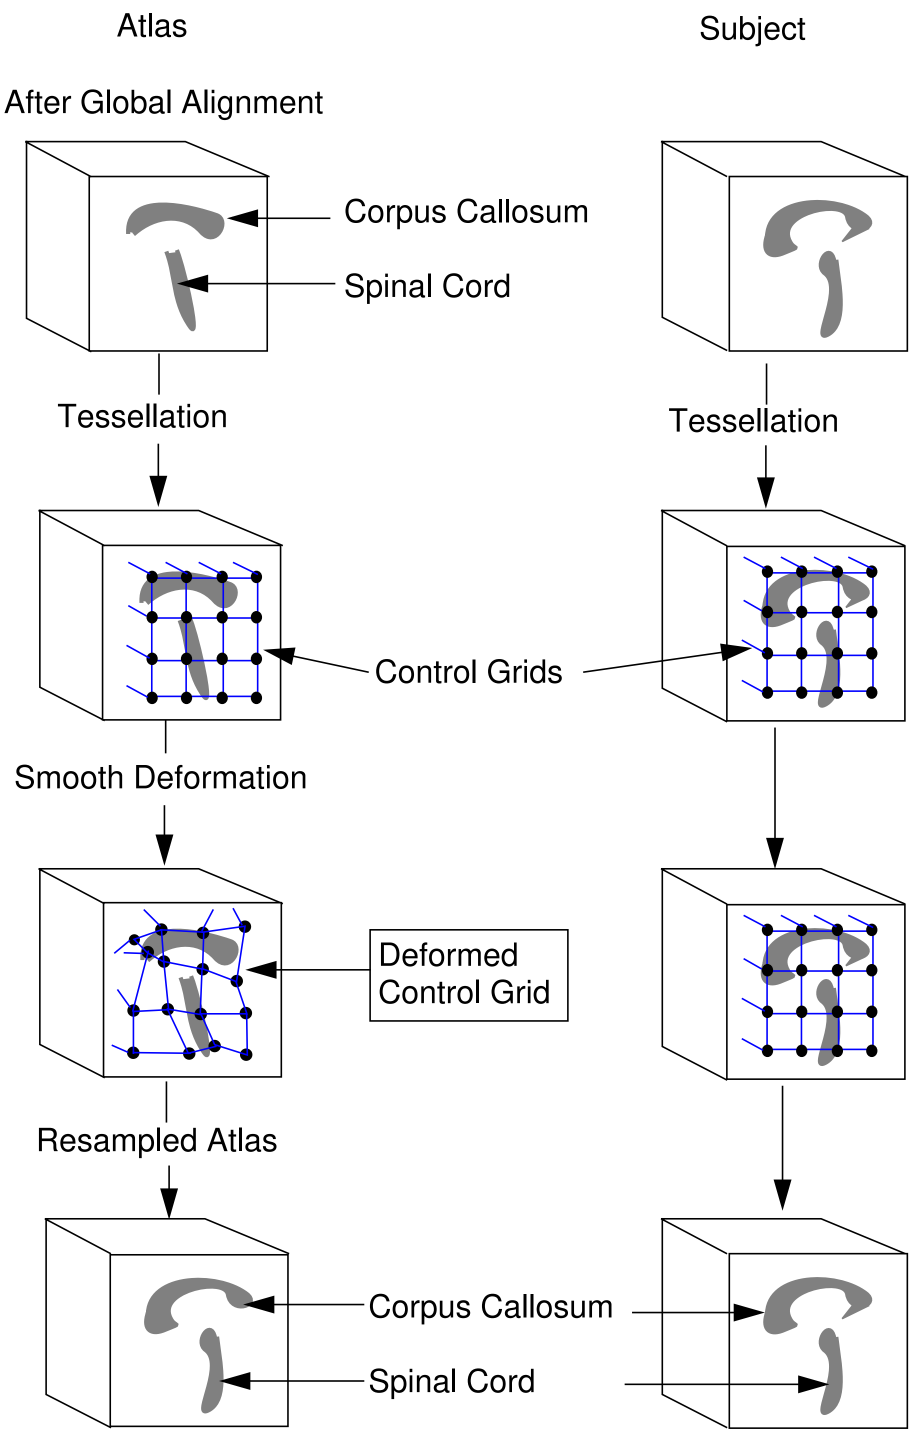

Figure 14: An illustration of the 3-D smooth deformation. The control points in the atlas shift to match their counterparts in the subject.

ters. Figure 15 is a block diagram illustrating the iterative optimization process of smooth deformation.

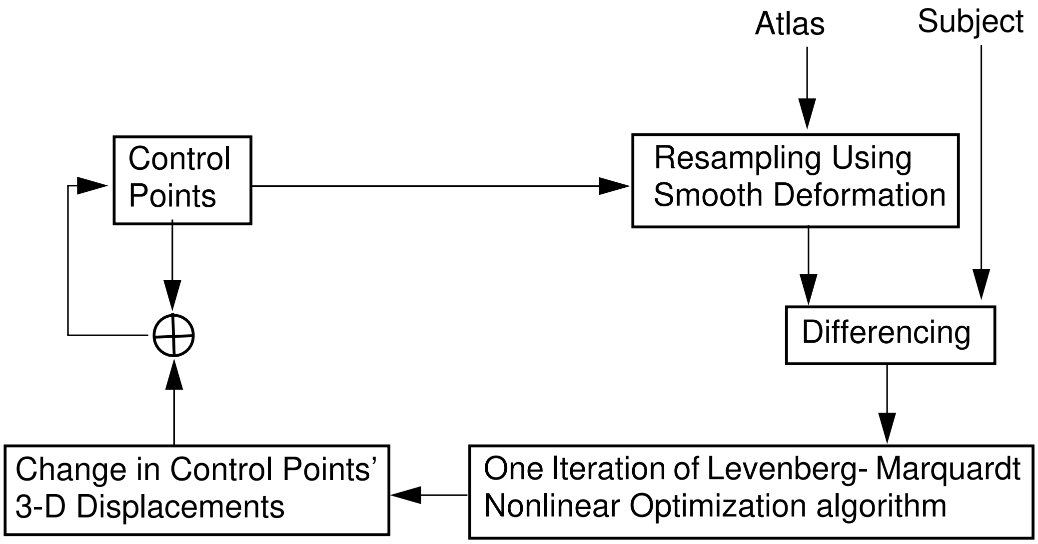

Figure 15: Block diagram of smooth deformation

The best smooth deformation is considered to be determined when the change of each parameter drops below a preset threshold. The threshold used in this thesis is that the change of each control point's position along each of the 3-D directions between iterations is below 1 voxel unit. Control grids and image volumes in multi-resolution are used to improve efficiency and avoid local minima, which will be discussed in 6.2.

Figure 16 shows the effect of applying smooth deformation to the intermediate result after global alignment. The atlas is warped in 3-D to match with the subject volume. Segmentation of corpus callosum is further adapted for the subject. When directly applied to the subject volume, its outline roughly aligns with the subject's corpus callosum. Compared to the result after global alignment, the registration for individual brain structures improved significantly, but there still exists misalignment at fine details. This is because smooth deformation implicitly enforces local neighborhoods to deform coherently.

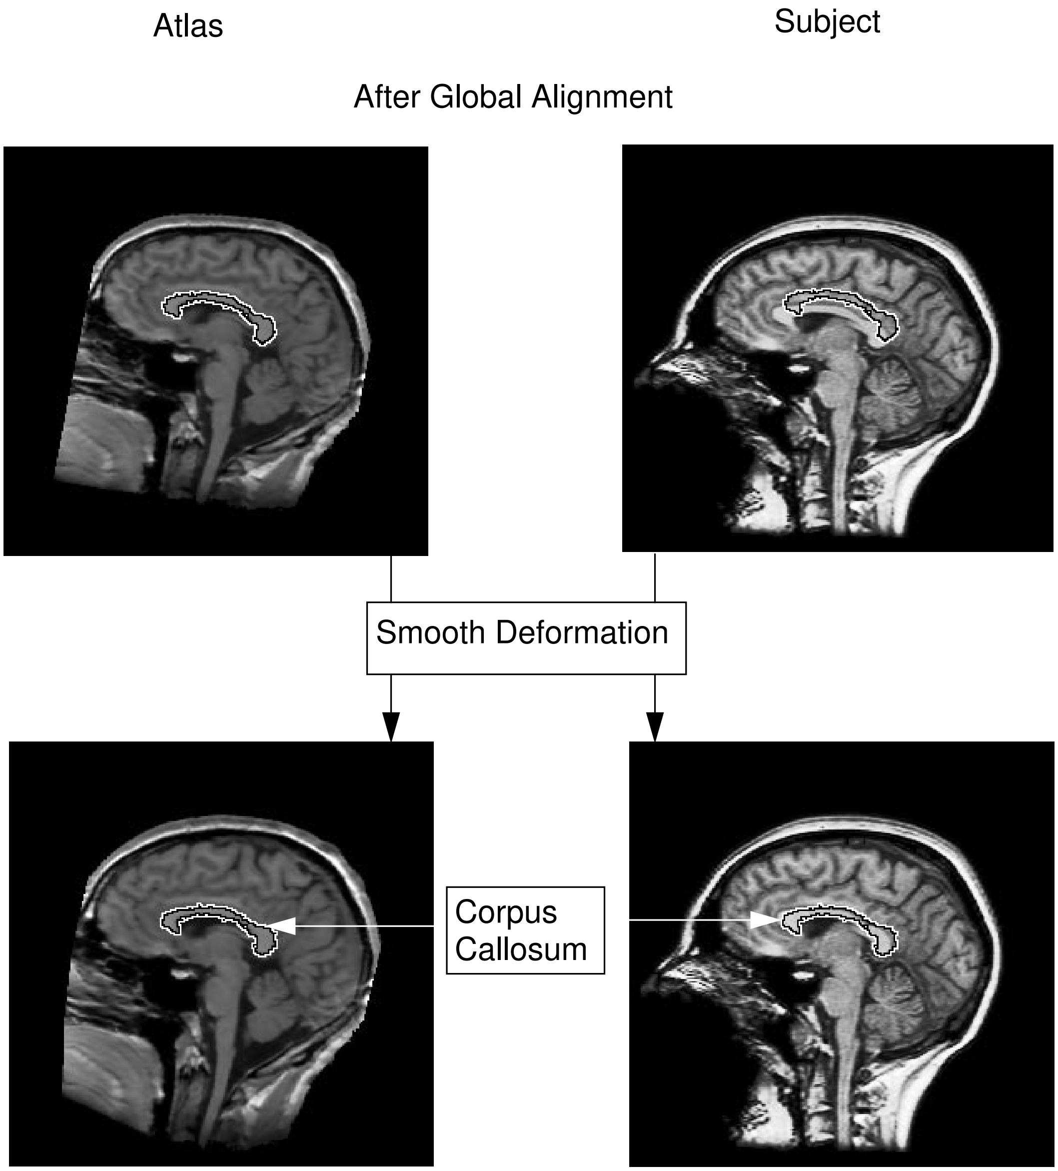

Figure 16: Smooth deformation improved the alignment of individual structures between the atlas and the subject's volume

2.3 Fine-tuning deformation with linear intensity transformation

Smooth deformation only allows control points to shift freely in 3-D space, displacements of voxels inside control grid cells are determined by the bounding control points. This means geometrical differences between the atlas and the subject that are smaller than the size of a control grid cell cannot be adjusted. This necessitates the last level of deformation. It fine-tunes the 3-D alignment by permitting each voxel to shift independently in 3-D space to match with its counterpart. Note that this is a special case of the smooth deformation in which the control grid is the voxel grid, and each voxel is a control point.

The improved alignment of individual anatomical structures after smooth deformation enables a more precise intensity equalization between the atlas and the subject volume. From observation, the match is generally more reliable for structures with relatively simpler shape and distinct intensity, such as corpus callosum and skull. Corpus callosum displays a high signal intensity, whereas skull has a low signal intensity, as shown in Figure 17. Their representative intensities can jointly determine a linear transformation that equalizes the two volumes' intensity distributions. Figure 18 displays intensity histograms of these two structures in the atlas, and histograms of what was automatically segmented as corpus callosum and skull in the subject's volume. The representative intensity of corpus callosum is defined as the highest peak in its intensity histogram after Gaussian smoothing, and the representative intensity of skull is denoted as the lowest peak in its smoothed intensity histogram. These representative intensities from the atlas and the subject form a linear intensity transformation between them. This transformation further equalizes the intensities in the atlas and the subject's volume.

2.3.1 Representing fine-tuning deformation

Similar to global alignment and smooth deformation, the intensity difference between spatially corresponding voxels in the atlas and the subject's volume serves as the deforming force. What is different is that now each voxel can shift independently in 3-D space. The

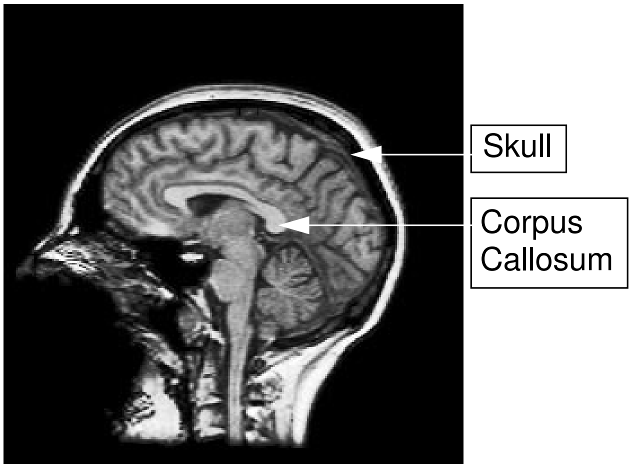

Figure 17: Corpus Callosum has a high signal intensity (bright), and skull has a low signal intensity (dark).

deformation parameters are 3-D displacements of all voxels, which is 3 times the amount of intensity data. Note that this is an under-constrained problem.

2.3.2 Estimating fine-tuning deformation

Suppose D denotes the current fine-tuning deformation. For a voxel $[x, y, z]_s$ in the subject's volume with intensity $I_s(x, y, z)$, its corresponding voxel $D(x, y, z)_a$ in the atlas has intensity $I_a(D(x, y, z))$. Consider the case that corresponding voxels have the same intensity, and the optimum deformation is achieved one step from D. Use δD to denote the difference between D and the optima, we have:

$$I_a(\mathcal{D}(x, y, z) + \delta\mathcal{D}(x, y, z)) = I_s(x, y, z) \tag{2}$$

The first order Taylor expansion of the left-hand-side in (2) gives:

$$I_a(\mathcal{D}(x, y, z)) + [\nabla I_a(\mathcal{D}(x, y, z))]^{\mathcal{T}} \delta\mathcal{D} \tag{3}$$

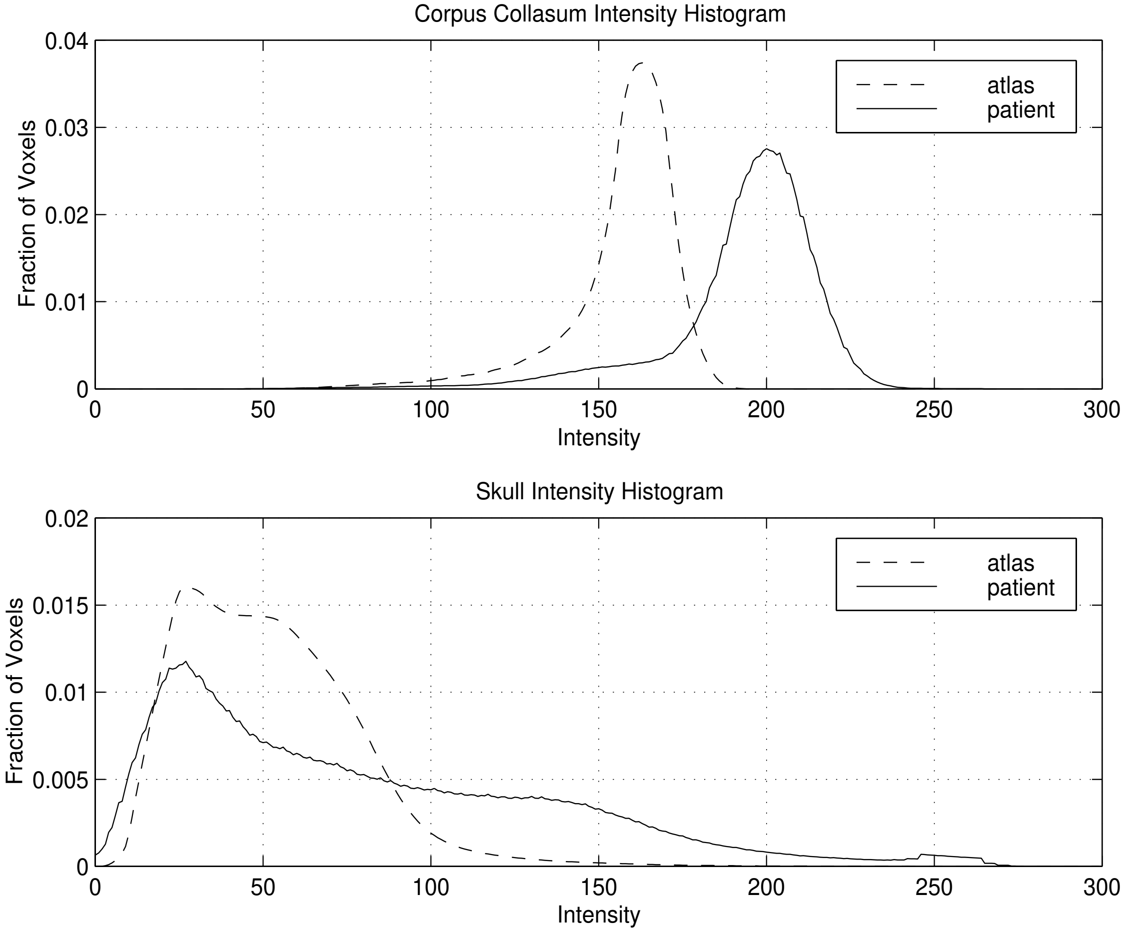
Corpus Collasum Intensity Histogram

Skull Intensity Histogram

Figure 18: Intensity histograms of corpus callosum (top) in the atlas (dotted line) and the subject's volume (solid line), as well as the corresponding ones of skull (bottom).

$\nabla I_a(D(x, y, z))$ is the first order derivative, i.e. image gradient, at voxel $D(x, y, z)_a$ in the atlas. Substitute (3) into (2) gives the *image brightness constraint* [77]*:*

$$I_a(\mathcal{D}(x, y, z)) + [\nabla I_a(\mathcal{D}(x, y, z))]^{\mathcal{T}} \delta \mathcal{D} - I_s(x, y, z) = 0$$

The above equation yields one solution for δD:

$$\delta \mathcal{D} = \frac{I_s(x, y, z) - I_a(\mathcal{D}(x, y, z))}{\|\nabla I_a(\mathcal{D}(x, y, z))\|^2} \nabla I_a(\mathcal{D}(x, y, z)) \tag{4}$$

To stabilize the deformation parameters when the intensity gradient in the atlas, $\nabla I_a(D(x, y, z))$, is close to zero, a stabilizing factor α is added:

$$\delta \mathcal{D} = \frac{I_s(x, y, z) - I_a(\mathcal{D}(x, y, z))}{\|\nabla I_a(\mathcal{D}(x, y, z))\|^2 + \alpha} \nabla I_a(\mathcal{D}(x, y, z)) \tag{5}$$

The deformation D is recovered by computing δD, adding it to D, and iterating until δD is smaller than a preset threshold. In the current implementation, the iteration stops when the root-mean-square (RMS) between the intensities of spatially corresponding voxels between two iterations decreases by less than 0.5%. 3-D isotropic Gaussian smoothing is applied to the volume's 3-D displacement flow after each iteration to regularize the under-constrained problem. Thirion used a similar approach [92]. Figure 19 shows a block diagram that illustrates the fine-tuning deformation procedure.

Figure 20 shows the effect of applying fine-tuning deformation to the intermediate result after smooth deformation. The atlas is further warped to match with the subject's volume. The segmentation of corpus callosum is adapted to match better with the structure in the subject's data.

Another example of registering the atlas with a subject's data using the hierarchical deformable registration is displayed in Figure 21. Segmentation of anatomical structures is also adapted for the subject, and outlines of several structures, e.g. the pair of lateral ventricles, are projected onto the subject's data to illustrate the improvement in matching individual structures. Figure 22 is a close-up of the subject's lateral ventricles in Figure 21. The alignment between the adapted segmentation and the subject's lateral ventricles improves

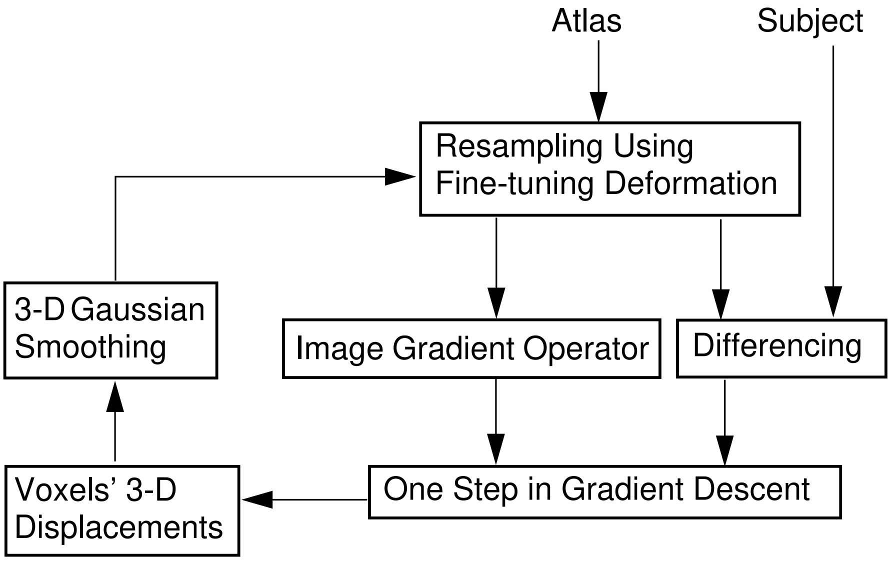

Figure 19: Block diagram of fine-tuning deformation

significantly along the deformation hierarchy. It is note-worthy that the shape of the adapted segmentation seems to be changing along the deformation hierarchy. That is because the atlas is deformed in 3-D to match with the subject's volume, whereas what is being displayed is a fixed 2-D cross-section of the subject's volume.

The 3-D hierarchical deformable registration algorithm reduced the typical months-long manual segmentation time of a whole human brain to minutes. Using an SGI computer with four 194 MHz R10K processors, it takes 18 minutes to adapt the atlas' segmentation for a subject's MRI with 256x256x124 voxels (124 slices, and each slice is a 256x256 pixel matrix). Parameters can be tuned to further improve efficiency. The algorithm is fully automated using a random initialization of the global alignment (Section 6.3), which makes its application convenient.

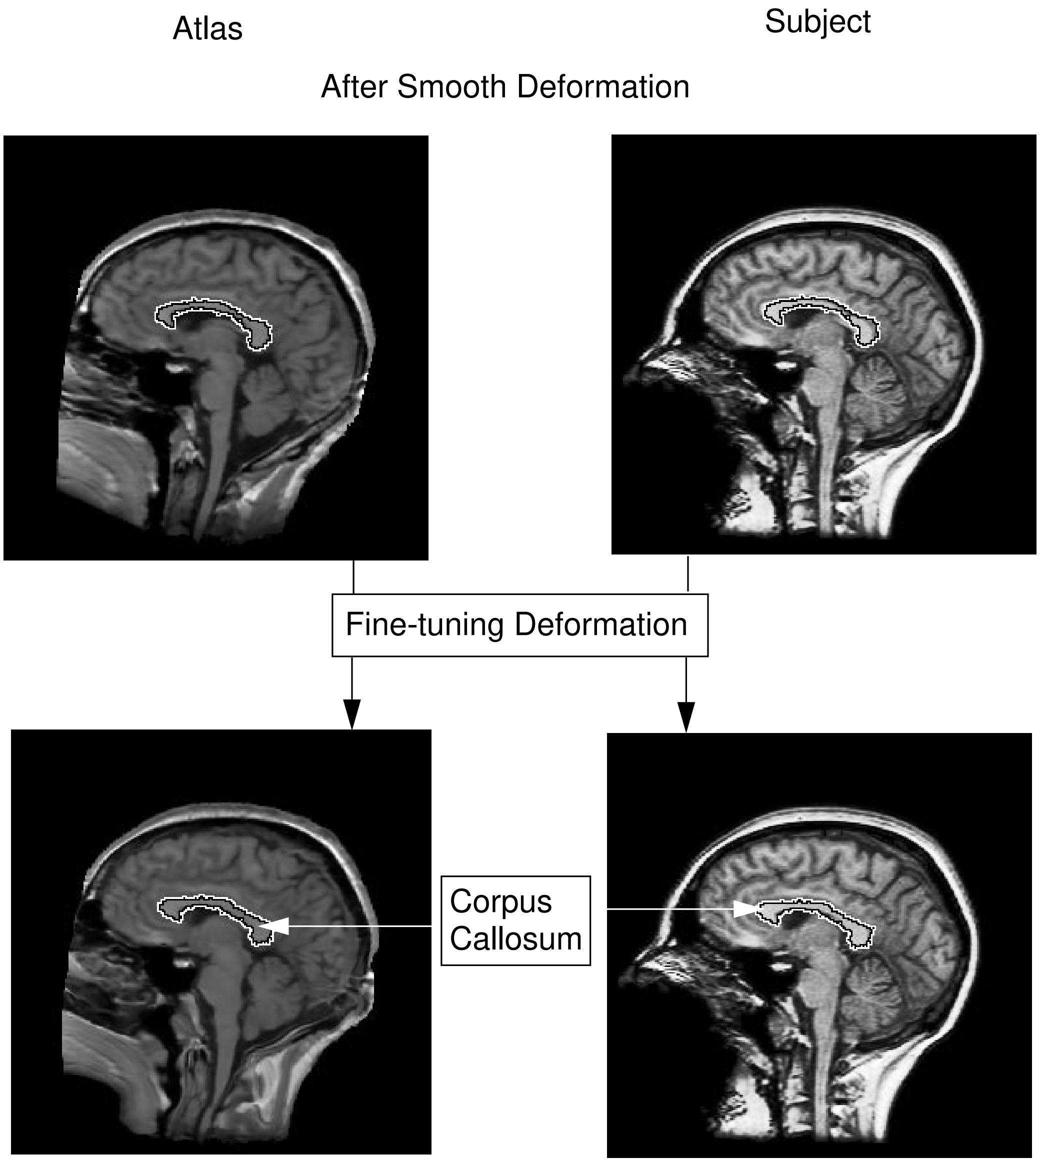

Figure 20: Fine-tuning deformation further improves registration accuracy of anatomical structures.

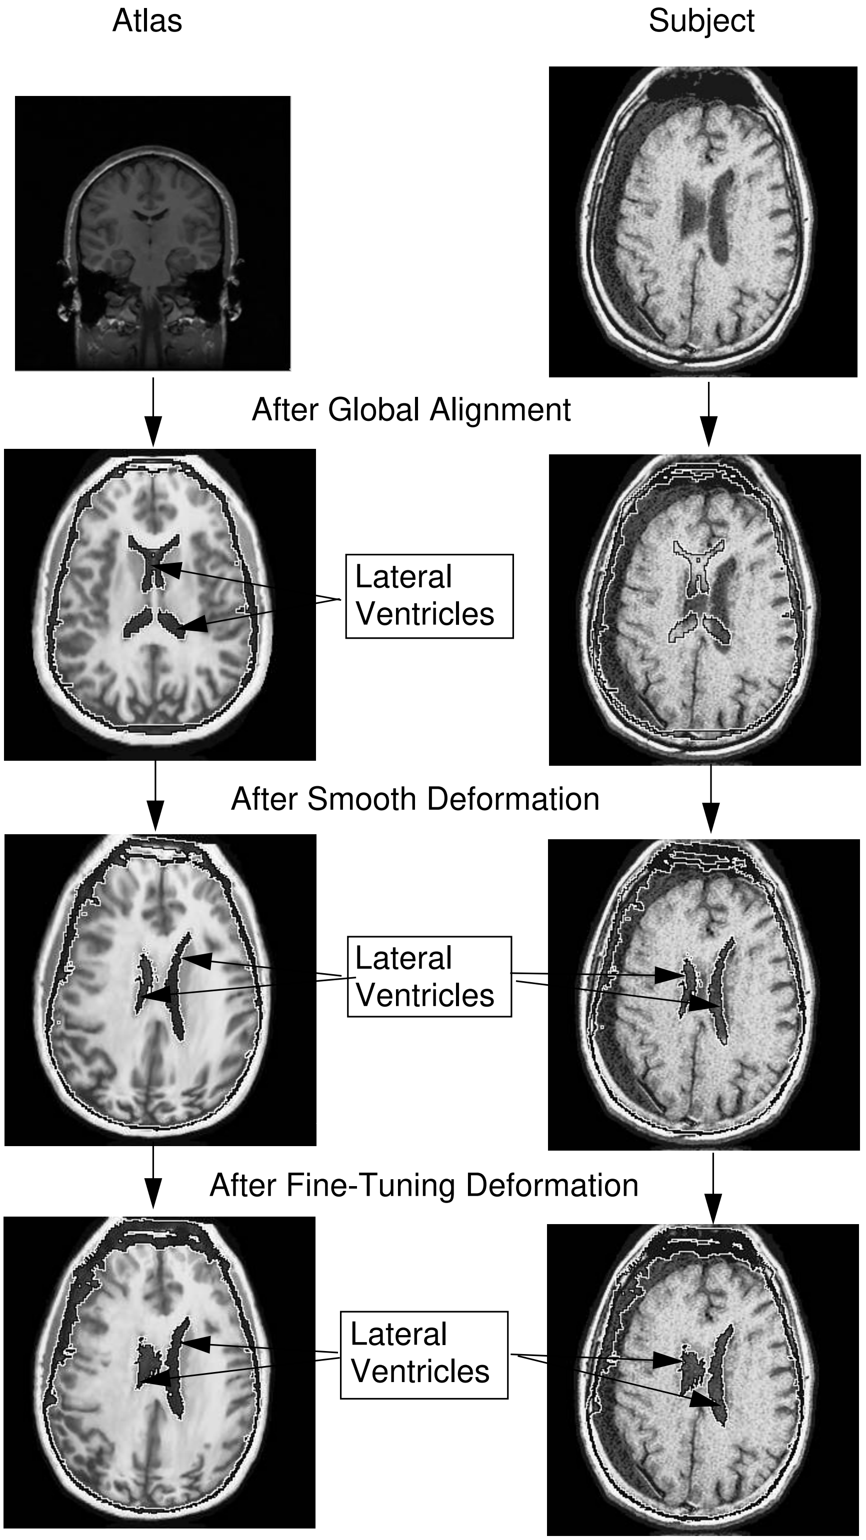

Figure 21: The progressive results of hierarchical deformation.

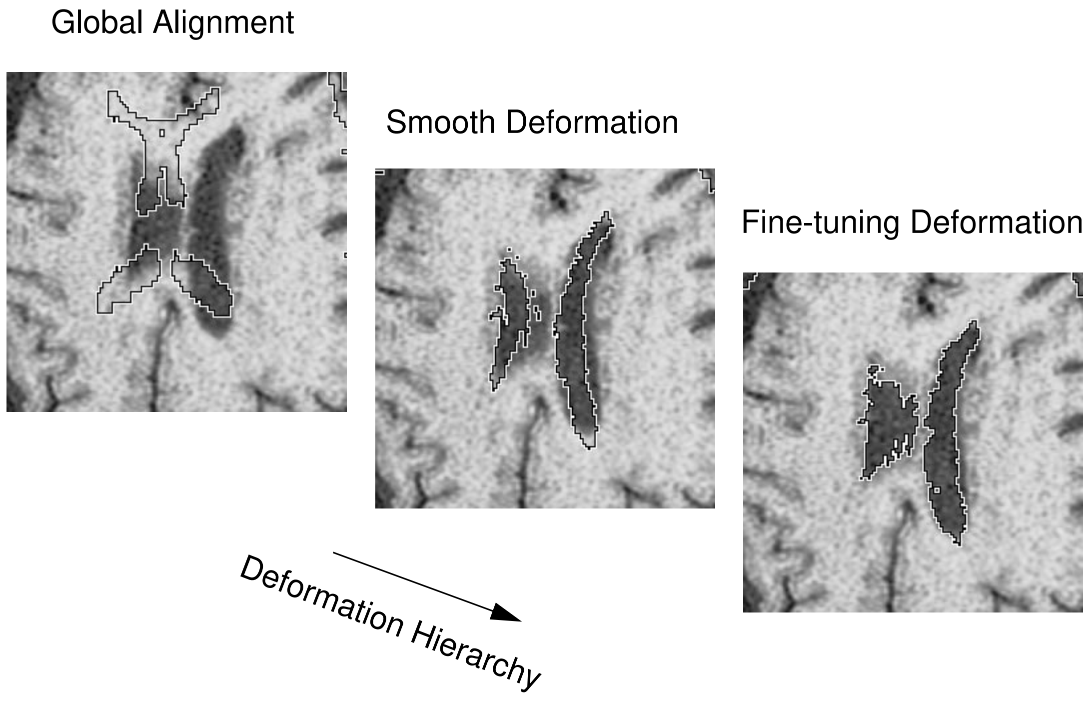

Figure 22: A close-up on the subject's lateral ventricles in Figure 21. The alignment between the adapted segmentation from the atlas and the subject's structures improve significantly.

2.4 Quantitative evaluation

Qualitatively, registration results shown in Figure 20 and Figure 22 are encouraging, however, objective and quantitative evaluation is necessary. For rigid registration, transformation parameters can be compared to those derived from stereotaxic fiducial markers rigidly fixed to a subject's skull [110]. Although fiducial-based registration itself has inherent measurement errors, it is generally adopted as *a ground-truth* transformation. Unfortunately, no such *ground-truth* can be acquired for deformable registration. One common solution is to assess how well the adapted segmentation from an atlas matches the corresponding anatomical structure in the subject's image volume. In this way, the demand for *ground-truth deformation* becomes the requirement for *ground-truth segmentation* of the subject's anatomical structures.

2.4.1 Ground-truth segmentation

In vivo data, the kind of data used in this thesis, presents another difficulty to quantitative validation: *ground-truth segmentation* of anatomical structures in living people is not available. As a solution, expert segmentation and classification of a subject's anatomical structures are regarded as *ground-truth*, or the *gold-standard*. A similar approach was employed in [2], [14], [19], and [33].

The *ground-truth* used in this thesis is comprised of 40 subjects' MRI volumes that have expert segmentation of one structure, corpus callosum, in one plane, the mid-sagittal plane. Figure 23 shows an example. The dimensionality of the test set is listed in Table 1. The experts are trained operators, and these image volumes are used as the test set. It is arguable that the corpus callosum is a distinct structure that is easy to segment, and therefore may not be appropriate for performance validation. On the other hand, this is also a structure that experts are confident to segment as ground-truth, and a structure of much research interest among collaborators. The goal of automatic segmentation is eventually being able to accurately segment subtler structures so as to facilitate research in the medical domain.

Number of Volumes	Volume Orientation	Slices/ Volume	Slice Thickness (mm)	In Plane Pixel Matrix	Pixel Size (mm^2)	Bits/ Pixel
1	Sagittal	124	1.5	256 x 256	0.9375 x 0.9375	16
1	Coronal	124	1.5	256 x 256	0.9375 x 0.9375	16
1	Axial	187	1.2	256 x 256	0.98 x 0.98	16
13	Axial	124	1.3	256 x 256	0.9375 x 0.9375	16
24	Sagittal	256	1	256 x 256	1 x 1	8

Table 1. Dimensionality of the test set.

Some researchers validate their methods by registering one image volume with a transformed version of itself, and comparing the computed transformation to the known transformation. Note that when the known transformation is of the same formulation as the transformation used in the registration process, this scheme is testing the algorithm's consistency, but not the accuracy.

Figure 23: An example of expert segmentation of corpus callosum in the mid-sagittal plane.

2.4.2 Measurement

Currently, there is no standard metric for evaluating segmentation accuracy. Dann et al. introduced a relative overlap measure for comparing two segmentations when neither is necessarily correct. It is defined as the ratio between the area of intersection and the area of their union [19]. Collins et al. used three measures to evaluate segmentation accuracy [14]. One is the ratio of absolute volume difference between ground-truth and the computed segmentation w.r.t. ground-truth. Since a small volume difference does not indicate accurate segmentation, another measure is defined as the ratio between the overlapping volume and ground-truth. However, this measure gives 100% for an accurate segmentation or any segmentation that is a superset of ground-truth. This necessitates the third measure, which is the ratio between the overlapping volume and the computed segmentation. The smaller of the second and third measure is used for validation. Bajcsy et al. employed the number of correctly matched voxels, false positives, and false negatives to assess the performance of their global matching procedure [2]. Gee et al. examined segmentation of 32 cortical and subcortical structures in a test set of 6 subjects, using the relative overlap measure in [19] and the second measure in [14]. In addition, they employed the distance between the centroids of a segmentation and its ground-truth to indicate the localization accuracy [33]. For feature-based registration, Davatzikos defined the registration error at each feature point or landmark to be the distance between its computed location and its ground-truth [20].

In this thesis, segmentation error is measured by the ratio between the number of mislabelled voxels and ground-truth, as illustrated in Figure 24. The number of mislabelled voxels includes both false positives, i.e. voxels classified as corpus callosum by the algorithm but not in the ground-truth, and false negatives, i.e. voxels classified as corpus callosum in the ground-truth but not by the algorithm. Note that this error can be bigger than 100%. If the three areas in Figure 24 are of equal size, then this error will be 100%, whereas [19]'s relative overlap measure will be 33%, and measures defined in [14] will give zero volume difference and 50% overlap. The error measured defined here is therefore the most stringent.

2.4.3 Performance

The performance evaluation process involves applying the 3-D hierarchical deformable registration algorithm to match the atlas to each image volume in the test set, and thus adapting the atlas' anatomical segmentation respectively. Overall error ratio of the whole test set is computed at each level of the deformation hierarchy, and compared in Figure 25. For the final result after fine-tuning deformation, the algorithm reached an overall error ratio of 4.4%. This is a significant reduction over the 22.8% error yielded by smooth deformation, and a drastic reduction over the 55.5% error after global alignment. This quantitative assessment of registration accuracy validates the observation in Figure 22.

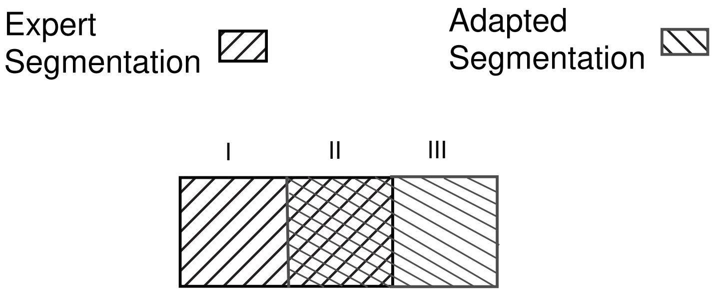

Area II: Correct Segmentation
Area I: False Negative
Area III: False Positive

$$\mathcal{E}rror = \frac{\mathcal{A}reaI + \mathcal{A}reaIII}{\mathcal{A}reaI + \mathcal{A}reaII}$$

Figure 24: Illustration of the error metric

It would be valuable to compare this algorithm's relative performance with methods developed by other researchers. However, such comparison has been difficult due to variations in the type of data being aligned, the type of machine used for computation, and the type of validation methodology. A retrospective registration evaluation project is being conducted at Vandelbilt University which provides clinical evaluation of the accuracy of retrospective techniques. Currently, evaluation is only available for PET-to-MR and CT-to-MR affine registration [107], and is therefore not applicable to work described in this thesis.

The algorithm is fully automatic, where global alignment is automated by random initialization. It takes 18 minutes to register 3-D images of 256x256x124 voxels on a SGI workstation with four 194 MHz processors; whereas typical manual segmentation of 3-D brain images takes months.

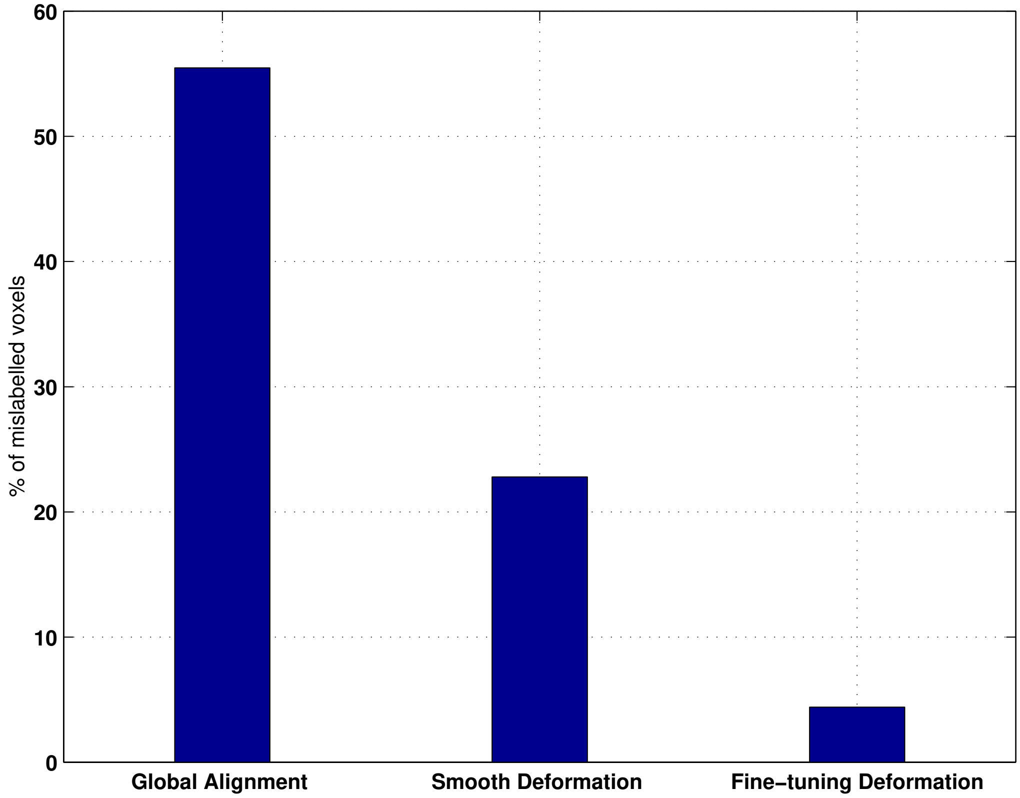

Figure 25: Overall error ratio of the test set at each level of the deformation hierarchy.

2.5 Error analysis

Errors in segmentation can be caused by several sources. One is bias incurred by using manually segmented structures from a single individual as the model atlas. Any errors in the atlas are carried through the deformable registration, yielding errors in the segmentation of subjects' data. This will be true in general for any atlas defined on the basis of a single brain, and can only be compensated when the atlas is extended to represent anatomical variations between individuals. Another source is inter-subject inter-observer variability. Anatomical differences between individuals complicate the manual segmentation processes involved in creating the *ground truth* and the atlas. Definition of the corpus callosum varies from expert to expert, from subject to subject. Segmentations from multiple experts can be integrated to improve consistency. Finally, errors from deformable registration. Since geometric correspondence is established based on intensity correspondence, any discrepancy in intensity can lead to errors in registration. Improvement of the registration algorithm will be further discussed in the following chapters.

2.6 Algorithm analysis

While it is validated that registration accuracy improves along the deformation hierarchy, it is also important to assess the effectiveness of the intensity equalization scheme. Moreover, while the overall error ratio reveals the algorithm's performance, it is still of interest to examine the distribution of error over the test set.

2.6.1 Effectiveness of intensity equalization

Three experiments are designed to assess the efficacy of the intensity equalization hierarchy, as shown in Table 2. These experiments are conducted over the whole test set. Overall error ratios show that the first and second levels of intensity equalization generally help the registration to reduce error by 6%. The third level, structure-based equalization, proves to be remarkably effective. It brought the error rate from 26.2% to 4.4%, which is a 83% error reduction. This demonstrates that smooth deformation is able to align the atlas and the subject's volume well enough for the structure-based equalization to be reliable. Note that registration results deteriorated by 15% when fine-tuning deformation was performed without the structure-based intensity equalization. This validates the necessity for the structure-based equalization, and proves the effectiveness of interleaving an intensity normalization hierarchy with a 3-D deformation hierarchy.

Experiment	Geometric Transformation	Intensity Equalization	Error Ratio
1	Global Alignment	None	61.3%
		Whole-volume equalization	55.5%
2	Global Alignment Smooth Deformation	Whole-volume equalization	25.6%
		Whole-volume equalization & Local equalization	22.8%
3	Global Alignment Smooth Deformation Fine-tuning Deformation	Whole-volume equalization & Local equalization	26.2%
		Whole-volume equalization & Local equalization & Structure-based equalization	4.4%

Table 2. Evaluating the effectiveness of the intensity equalization hierarchy

2.6.2 Error distribution

Histogram distribution of error ratios over the test set is computed at each level of the deformation hierarchy, and displayed in Figure 26. After global alignment only 14.6% of all samples have an error ratio below 20%. After smooth deformation 43.9% of all samples have less than 20% error, and 14.6% of all samples have an error ratio below 10%. Fine-tuning deformation eliminated cases with more than 20% error, and brought 92.7% of all samples to an error ratio below 5%. The increase in registration accuracy is evident.

2.7 Discussion

So far the discussion has been focused on the design and evaluation of the hierarchical deformable registration algorithm, and its interleaving intensity equalization scheme. It is also interesting to remark on certain considerations and alternative approaches investigated during the design.

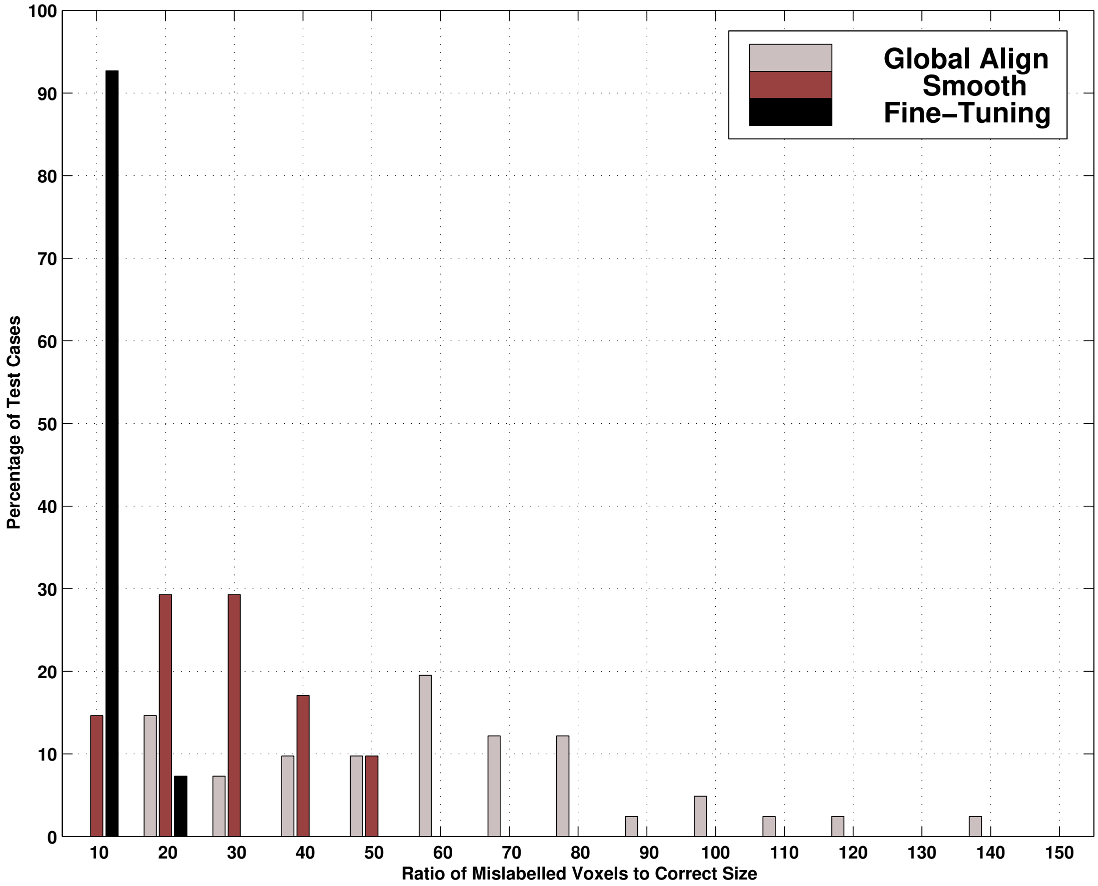

Figure 26: Histogram distribution of error ratio over the test set, at each level of the deformation hierarchy.

2.7.1 Transformation and resampling

To register the atlas with a particular subject's volume, the transformations can be either from the atlas' coordinate frame to the subject's, or vice versa. However, for applications that need the segmentation and classification of the subject's anatomical structures, it is necessary to resample the atlas into the subject's coordinate frame so as to adapt its expert segmentation and classification for the subject. This resampling will be straight forward if transformations involved in the registration have 1-on-1 mapping. Unfortunately this is not true for deformable registration, due to its highly nonlinear nature. Figure 27 gives a 2-D illustration of the nonlinear mapping: it can be a multiple-to-one correspondence, or no correspondence.

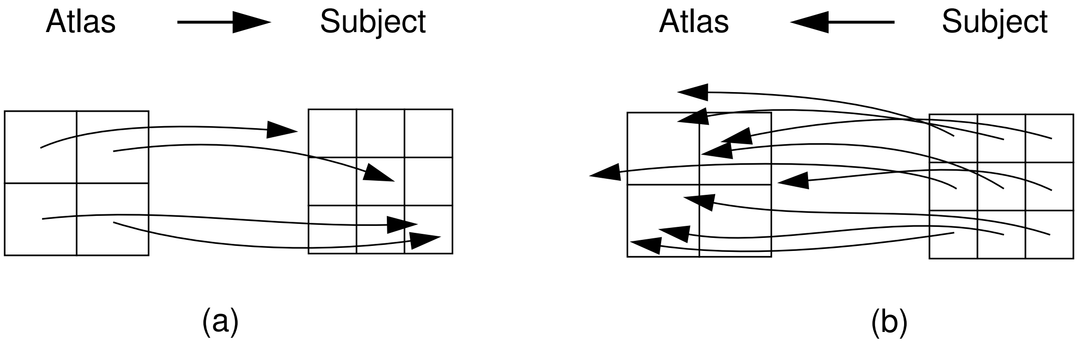

Figure 27: 2-D illustration of the transformation's nonlinear nature.

As illustrated in Figure 27 (a), an atlas-to-subject transformation will not guarantee each voxel in the subject's volume a correspondence in the atlas. This means some voxels in the subject's volume will have multiple anatomical classifications adapted from the atlas, whereas some will have undetermined classification. A typical solution to this *resampling problem* is to use the subject-to-atlas transformation, as illustrated in Figure 27 (b). In this way, each voxel in the subject's volume will either have an anatomical classification adapted from its corresponding atlas voxel, or be labelled as *background*. This thesis presents registration transformations as from the atlas to a subject's volume for the ease of understanding. However, all algorithms apply to both directions. To adapt the anatomical classification in the atlas for a particular subject, a subject-to-atlas transformation is employed.

2.7.2 Other intensity equalization methods

A number of intensity normalization schemes were explored before the final design, such as using histogram equalization to remove bright and dark outliers, and using lateral ventricles instead of skull in determining linear intensity transformation. Histogram equalization did not prove superior to normalization of intensity mean and variance, whereas intermediate segmentation of lateral ventricles were not reliable enough to facilitate computation of the correct intensity transformation. A more rigorous intensity normalization method should account for signal distortions that are unique to the MRI process. Wells et al. developed an EM-segmentation algorithm that used an imaging model to account for that, and Kapur et al. extended the theme by adding a regularizer to combat salt-and-pepper noise [49], [106].

2.7.3 Smoothness of deformation

The assumption of smoothness in deformation guarantees that neighboring voxels in the atlas be mapped to neighboring points in the target. Similar approaches were used by Black and Anandan [6]. However, there exist neighboring points in unconnected structures (such as on opposite sides of a sulcus, or on either side of the longitudinal fissure) that do not need to be mapped to neighboring voxels in the subject. Therefore, it may be desirable to allow a discontinuity in the transformation at internal brain structures, e.g. surfaces that separate the cerebellum from the occipital lobe or that separate the temporal lobe from the inferior frontal lobe.

2.7.4 Quantitative evaluation

Since the corpus callosum seems an easy structure to identify, quantitative evaluation based on its segmentation may not be the most convincing. However, during this thesis work, this has been the structure experts can segment confidently enough to provide as ground truth. Although performance evaluation based on the segmentation of structures of more complex shape will be more rigorous, it has been observed and reported that structures with a complex boundary shape are more difficult for people to segment than those with a simple boundary [52]. People have a tendency to over- or under-estimate the boundary. Consequently, manual segmentation exhibits a consistent variability in the segmentation of voxels at the boundary of complicated shapes such as the cortical grey matter.

2.8 Chapter summary

This chapter has presented and evaluated a 3-D hierarchical deformable registration algorithm that does not use guidance of anatomical knowledge. This algorithm is a necessary starting point to achieve the next goal, the characterization of brain anatomy and its variations, which will be discussed in the next chapter.

CHAPTER 3 Building a Statistical Atlas

The 3-D hierarchical deformable registration algorithm discussed in Chapter 2 achieved encouraging performance. However, considerable inaccuracies still exist. Figure 28 compares segmentation results given by the automatic algorithm and that from an expert. Note that certain inaccuracies are caused by the algorithm's insufficient knowledge of the anatomy, and cannot be corrected by exploiting image features solely.

This observation is supported by the analysis in Section 2.6, that one major error source is bias introduced by the atlas. The current *atlas* is not an average or representative model of any population, it is a single non-pathological brain MRI whose structures were manually classified. It is possible that the atlas is an extreme of the normal distribution. The current algorithm is matching this potentially biased atlas to any particular subject's data, which is inevitably erroneous. Moreover, as discussed in Section 1.4.2, anatomical differences between individuals present a major difficulty to inter-subject registration. Therefore, building an atlas based on multiple subjects' data so that it represents the *average* brain as well as its normal anatomical variability will help improve registration performance. This leads to the focus of this chapter, which is the characterization of brain anatomy and its non-pathological variations in a population:

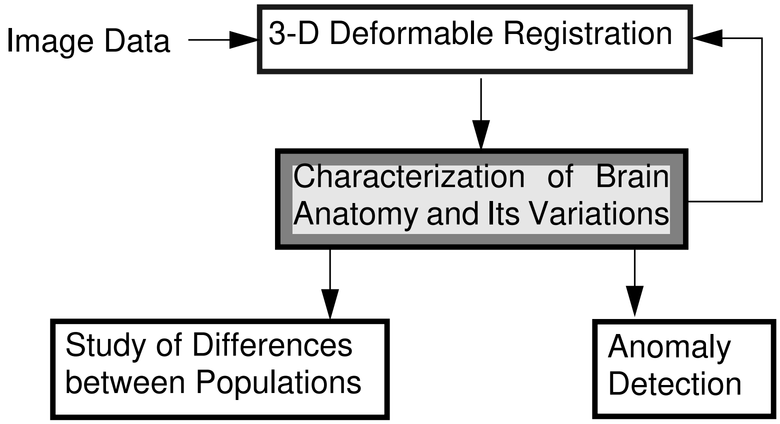

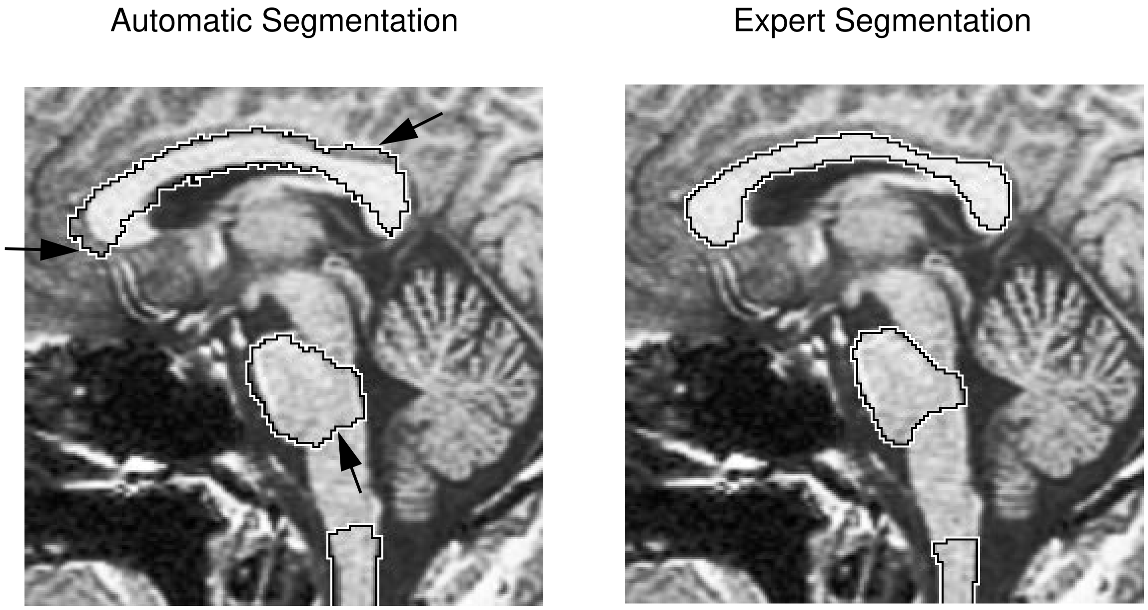

Figure 28: Inaccuracies in the automatic segmentation, compared to expert segmentation.

3.1 Related work

Techniques solely base on image features have appeared inadequate for many segmentation, registration, and measurement tasks in medical imaging. The inclusion of prior knowledge in the analysis procedure has become increasingly important. Prior knowledge may concern the physics of image formation; the anatomy, pathology or (patho)physiology of the organs or tissues under investigation; or the expertise of the medical specialist on the interpretation of the image data.

During the past decade, many researchers have worked on modeling shape variation of organs or physiological variation of tissue characteristics. One approach is landmark-based morphometrics for multivariate analysis of curving outlines in biomedical images. In this approach, shapes are defined as equivalence classes of discrete point-sets under the operation of the Euclidean similarity group. A Procrustes distance between every pair of shapes is characterized by a least-squares formula. Bookstein introduced a combination of Procrustes analysis and thin-plate splines to provide a range of complementary filters, from high pass to low pass, for effects on outline shape in grouped studies [7]. He applied this hybrid method to compare the shape of the corpus callosum from mid-sagittal sections of 25 human brain MRIs, 12 normal and 13 with schizophrenia.

However, in any structural anatomical data set there will be more information about structural differences than can be captured by landmarks. Davatzikos et al. studied callosal outlines acquired by an automatic active contour method, and relate the outlines of samples to an a priori norm using elastic relaxation. The relation of each individual to the norm is then described by an areal distortion function, which captures the anatomical variability [21]. Groups (eight male and eight female) are compared by averaging the distortion at corresponding points of the normative form or over regions, and thresholding the differences at various effect sizes. Sclaroff and Pentland represented shape in terms of an object's physical deformation modes. The representation consists of two levels, one is modal deformations that describe the overall shape of a solid, the other is displacement maps that employ a multiscale wavelet representation to provide local and fine surface detail [73]. They applied the method to compress dense 3-D point data from surfaces, using displacement maps and wavelets [82]. The limitations of methods using physical models include the questionable accuracy of the models of brain stiffness. Further investigation of the brain's material properties will be required.

Instead of physically modeling the structure under study, researchers have also sought to obtain shape descriptions directly from sample data. Cootes et al. used principal components to describe the modes of variation inherent in a training set of 2-D heart images [17]. Hill et al. extended the technique to 3-D, and performed multivariate analysis on shape measurements derived from manually collected homologous points [40]. Only objects of simple shapes and limited variations can be accounted using this method. It is not applicable to modeling the whole brain. Thompson et al. quantified the variability of gyral and subcortical surfaces as a non-stationary Gaussian random tensor field [95]. Tensor field corrections based on Christoffel fields and Winslow theory were applied to create an average image template with the average anatomical configuration and intensity for the group under study. Their earlier work studied the variability in location and geometry of five sulci in each hemisphere of six postmortem human brains placed in the Talairach stereotaxic grid. The sulci were modeled as 3-D surfaces, and heterogeneous profiles of 3-D variations were quantified locally within individual sulci [99]. Joshi et al. examined the neuro-anatomical variation of the geometry and shape of 2-D surfaces in the brain, e.g. the cortical and hippocampal surfaces. Shape was quantified via the construction of templates, while variations were represented by defining probabilistic deformations of the templates [46]. Empirically, they represent the probabilistic deformations as Gaussian random vector fields on the embedded sub-manifolds. Manifold information on easily identifiable landmarks need to be provided by an operator.

The high complexity and variability of the anatomy and related pathology make it difficult to represent using a single model. Cootes has examined both physical and statistical shape models, with the goal of smoothly transitioning from a physical to a statistical shape description as more data become available [18]. Zhu and Yuille considered physical and statistical shape models in the context of representing and recognizing objects from their 2-

D silhouettes [112]. Szekely et al. developed a model-based segmentation technique combining desirable properties of physical models, shape representation, and modeling of natural shape variability. They designed flexible Fourier parametric shape models by combining the mean contour with a set of eigenmodes of the parameters. The method was applied to characterize the biological variability of the shape of the corpus callosum outline [87]. Wang and Staib also incorporated statistical shape models into elastic model based 2-D non-rigid registration [104]. The elastic model constrains the transformations to maintain smoothness, and statistical shape information embedded in corresponding boundary points function as Bayesian priors to improve the registration.

Most of the above work concerns with statistics of individual structures, i.e. models that are local. But local models lose information on the relative position and interaction between neighboring structures. While this is promising for applications aiming at segmenting particular structures and lesion, a global model capturing variations of all brain structures is necessary in order to register and segment the whole brain simultaneously. Evans et al. constructed a probabilistic atlas by averaging 305 brains in a stereotactic space. The variability is illustrated by the sharpness of the grey level contours, but the method does not provide suitable information to derive a shape model [26]. Gee and Le Briquer aligned eight normal subjects' image data with a reference configuration, and modeled their anatomical variability using principal component analysis of the spatial mappings [31]. Due to the complexity of brain anatomy, more samples are necessary to validate their model. Similarly, Guimond et al. developed an automatic method to acquire average intensity and shape models of the human brain [38]. This thesis will attempt to build a global model that captures both the average and the variations of all brain structures between individuals, using a non-pathological population of considerable size.

3.2 Capture Anatomical Variations

In order to extract anatomical variations between individuals, image data of a population needs to be compared in a common reference frame. The atlas is chosen as the common reference frame, for the convenience of further augmenting it into a multi-subject atlas. The 3-D hierarchical deformable registration algorithm discussed in Chapter 2 is used to map image data of a population to the atlas' framework. A population of 105 T1-weighted brain MRI volumes constitute the training set. Dimensionality of the training set is listed in Table 3. All brains in the training set are non-pathological, and they are separate from the test set introduced in Section 2.4.1. Figure 29 shows example cross-sections from several individuals brain MRI volumes. Notice that there exists differences in the density (reflected in intensity), shape, size, and location of brain structures between individuals, i.e. differences that are intrinsic. In addition, these image volumes also differ in overall intensity distribution, orientation, position, and scale, due to different image acquisition processes. These variations are extrinsic to the anatomical differences, and therefore should be addressed separately from the intrinsic variations. This is similar to Martin et al.'s approach of separating important and unimportant shape variation so as to quantitatively describe pathological shape variations [66].

Number of Volumes	Volume Orientation	Slices/ Volume	Slice Thickness (mm)	In Plane Pixel Matrix	Pixel Size (mm^2)	Bits/ Pixel
58	Coronal	124	1.5	256 x 256	0.9375 x 0.9375	16
21	Axial	124	1.3	256 x 256	0.9375 x 0.9375	16
26	Sagittal	256	1	256 x 256	1 x 1	8

Table 3. Dimensionality of the Training Set.

The process of capturing anatomical variations involves eliminating extrinsic differences caused by image acquisition, and extracting intrinsic differences related to brain anatomy, as illustrated in Figure 30. Eliminating extrinsic differences in geometry entails bringing all image volumes to the same orientation, position, and overall size as the atlas. This is conveniently achieved using global alignment, the first level in the 3-D hierarchical deformable registration algorithm (Section 2.1). To extract intrinsic geometric differences between individuals, anatomical structures of each sample volume in the training set need

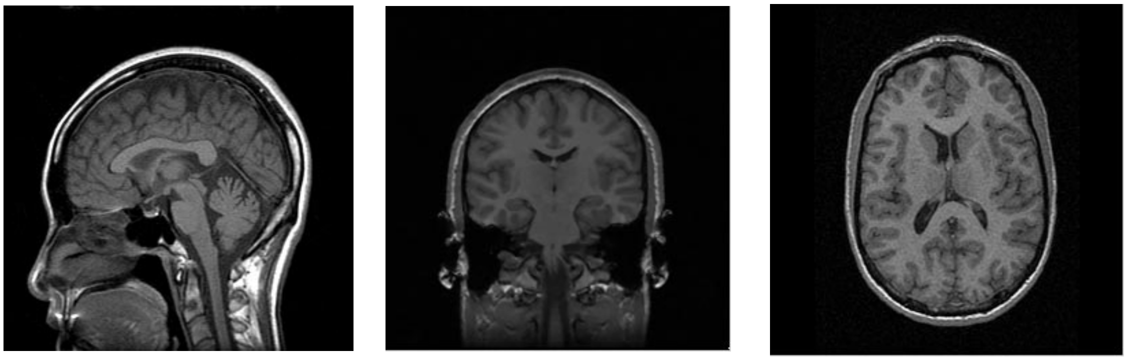

Figure 29: Example cross-sections of different individuals' brain MRI data from the training set.

to be compared to obtain information on variability. This is achieved by deformably matching each image volume in the training set to the atlas, using the second and third levels of the 3-D hierarchical deformable registration method (Section 2.2, Section 2.3). The respective deformation flows embody intrinsic geometric variations between each sample in the training set and the atlas. A similar approach is taken to capture inherent variations in structure density between individuals. Extrinsic intensity variations from separate image acquisition processes are removed via the intensity equalization hierarchy interleaved with the deformation hierarchy (Section 2.1, Section 2.2, and Section 2.3). Intrinsic variations in structure density are the residual intensity differences between corresponding voxels in each deformed sample volume and the atlas.

3.3 Model Anatomical Variations

While anatomical variations are captured in deformation flows and corresponding voxels between each deformed sample volume in the training set and the atlas, further abstraction of this knowledge is necessary. Variations in structure density and geometry are modeled separately. Because of the limited sample size and the complex nature of morphological variations in human anatomy, knowledge of anatomical variations are modeled using the first and second order statistics of the measurements.

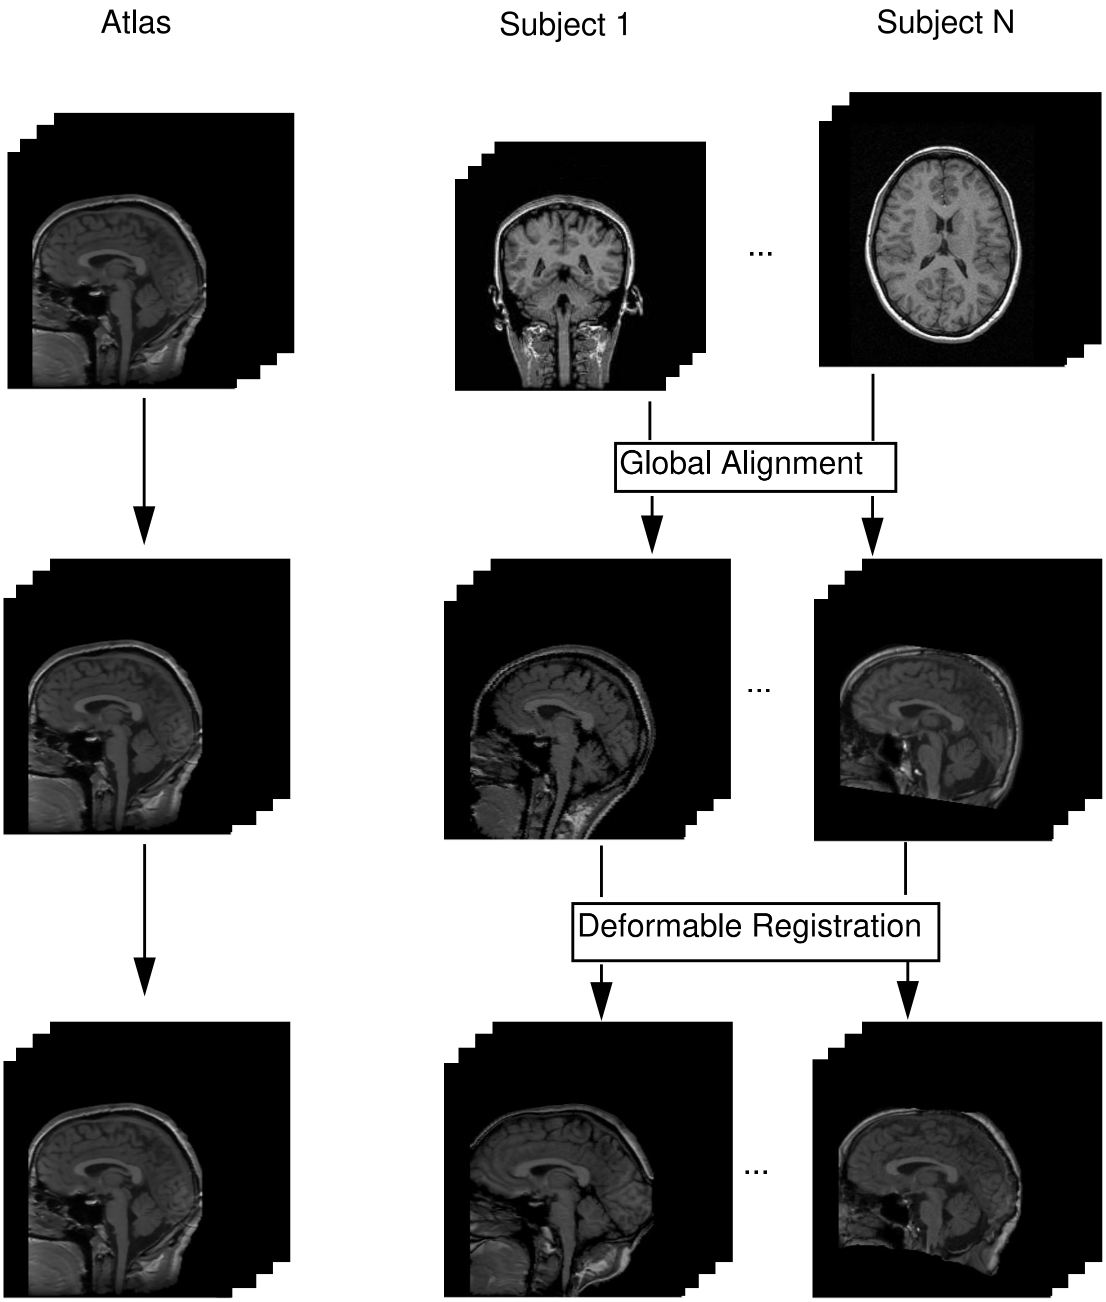

Figure 30: The process of eliminating extrinsic variations between individual image volumes, and extracting anatomical variations.

3.3.1 Modeling density variations

After each sample volume in the training set is deformed to match with the atlas, the residual intensity differences between corresponding voxels reflect variations in structure density between the sample and the atlas. Suppose there are *N* sample volumes in the training set (currently *N*=105), then each atlas voxel spatially corresponds to *N* voxels, with one voxel from each deformed sample volume. Tissue density variation in the training set population at each atlas voxel location can be examined in a histogram, as illustrated in Figure 31. Note that the atlas voxel may not have the average intensity in the histogram, because the atlas is simply one non-pathological brain MRI volume and is not necessarily representative of any population. Each atlas voxel is associated with one such histogram. The collection of such intensity histograms over the atlas volume embodies intrinsic variations of structure density in a population.

The histogram distribution can be modeled as 1-D Gaussian distribution, $P(dI|D)$, where D is the specific deformation that maps a sample volume to the reference frame of the atlas, and dI is the intensity difference between each atlas voxel under examination and the corresponding voxel in the deformed sample volume:

$$\mathcal{P}\langle dI|\mathcal{D}\rangle = \frac{1}{\sqrt{2\pi}\sigma}e^{-\frac{(dI-\mu)^2}{2\sigma^2}} \tag{6}$$

here μ is the mean density difference between all samples in the training set and the atlas at this voxel; σ^2 is the variance of the density difference distribution. Note that dI is computed after having removed extrinsic intensity differences caused by separate image acquisition processes. Therefore, this 1-D Gaussian distribution models the inherent variations in structure density between individuals.

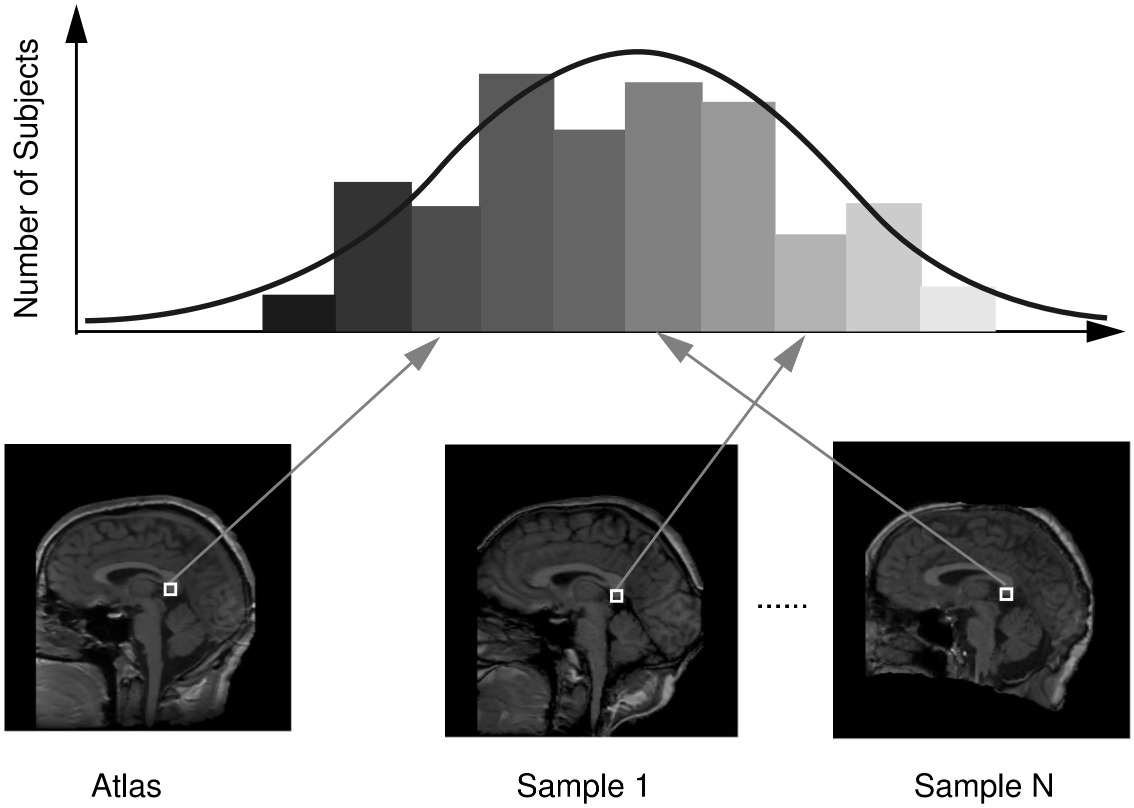

Figure 31: Modeling intrinsic variations in structure density: each atlas voxel is associated with a 1-D Gaussian distribution.

3.3.2 Modeling geometric variations

As discussed in Section 3.2, intrinsic geometric variations are captured in the 3-D deformation flow that warps each sample volume to register with the atlas. Now each atlas voxel is associated with N 3-D flow vectors, with each vector being the spatial distance between the atlas voxel and the anatomically corresponding voxel in one particular sample volume. Distribution of the N 3-D flow vectors associated with each atlas voxel can be captured in a 3-D histogram. Figure 32 shows a 2-D illustration. Each atlas voxel is associated with one such histogram. The collection of such 3-D histograms over the atlas volume embodies intrinsic geometric variations of brain anatomy in a population.

Similar to Section 3.3.1, the 3-D histogram distribution at each atlas voxel is modeled as a 3-D Gaussian distribution, $P(D)$:

$$\mathcal{P}(\mathcal{D}) = \frac{1}{\sqrt{(2\pi)^3|\Phi|}}e^{-\frac{(\overrightarrow{\Delta\vartheta} - \vec{\omega})^{T}\Phi^{-1}(\overrightarrow{\Delta\vartheta} - \vec{\omega})}{2}} \tag{7}$$

here $\overrightarrow{\Delta\vartheta}$ is a 3-D flow vector that associates the atlas voxel with its anatomically corresponding voxel in each particular sample volume. $\vec{\omega}$ is the mean 3-D flow vector that spatially relates this atlas voxel to its anatomical counterparts in the training set; and Φ is the 3x3 covariance matrix of the distribution. Note that $\overrightarrow{\Delta\vartheta}$ has been adjusted for extrinsic geometrical variations resulted from image acquisition. This 3-D Gaussian distribution models the inherent geometric variations between individuals.

3.4 A statistical atlas

The atlas was originally one single non-pathological brain MRI volume, whose voxels are related to their intensity values and classification labels of respective anatomical structures they belong to. It is valuable data, however non-representative. The above procedure of extracting and modeling anatomical variations associates each atlas voxel with a mean intensity and its variance in a population, and a mean position and its variance in the population. In this way, the atlas is enriched to represent the anatomical variability in tissue density and geometry of a population. This means the atlas is no longer an atlas built upon information of one possibly biased example, it is augmented into a statistical atlas that embodies anatomical information of a population. Figure 33 illustrates the concept.

The statistical atlas was built by examining distributions of intensity and spatial location of corresponding voxels in a training set. The correspondence was established using

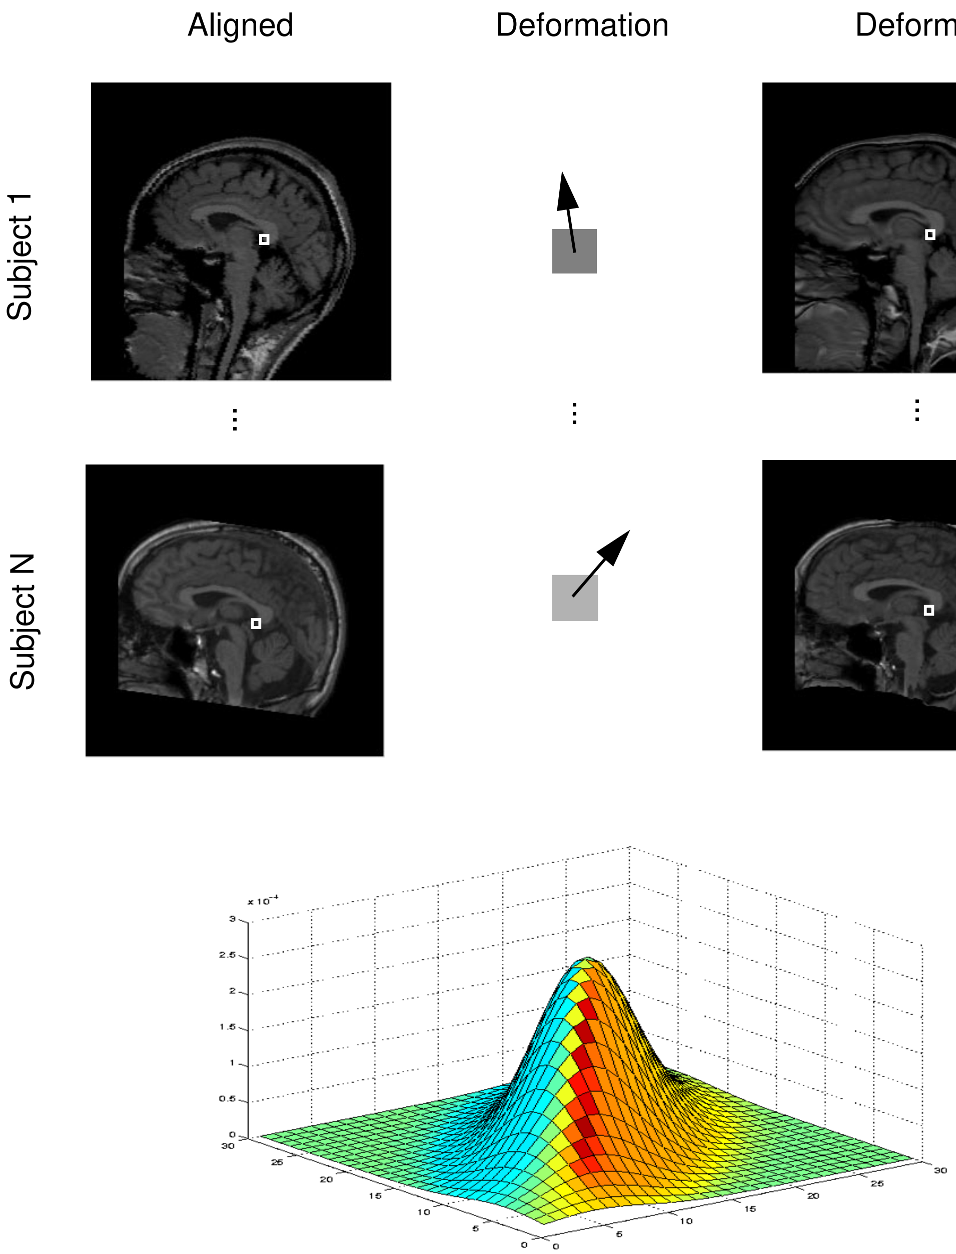

Figure 32: Modeling intrinsic variations in geometry as 3-D Gaussian.

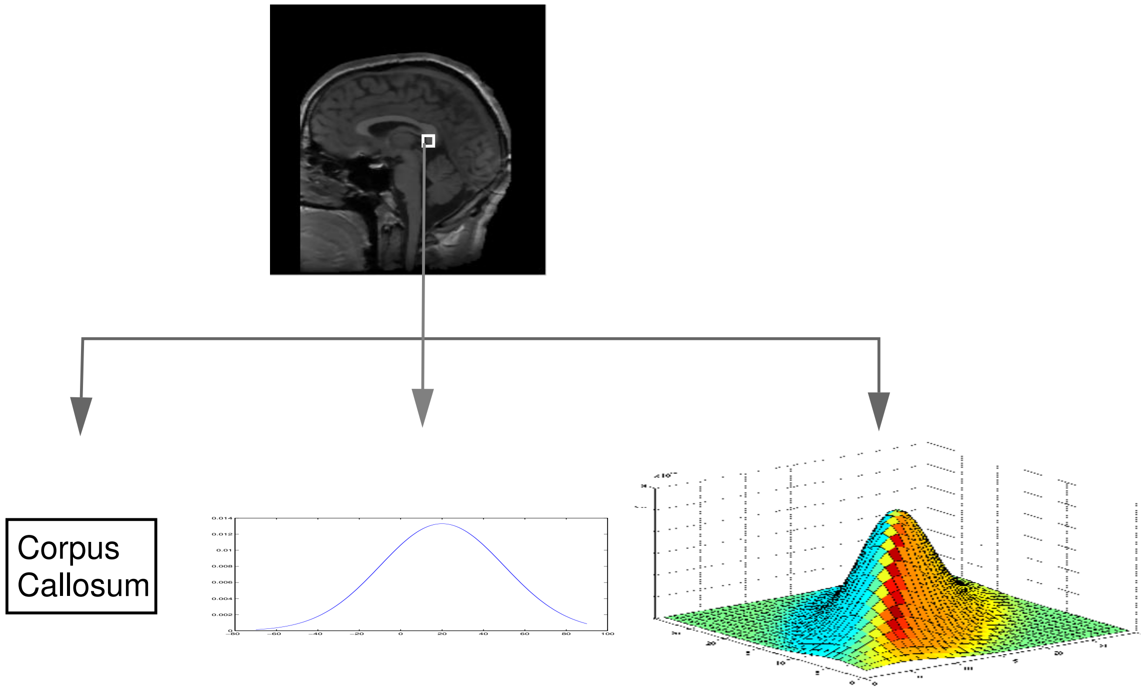

Figure 33: A statistical atlas. In addition to its intensity value and anatomical classification, atlas voxel is associated with a 1-D distribution of tissue density variations as well as a 3-D distribution of geometric variations in a population

an automatic 3-D deformable registration algorithm introduced and evaluated in Chapter 2. However, as discussed in Section 2.5, there exist inaccuracies in results given by this algorithm. Inaccuracies in registration inevitably propagate to statistical models built thereupon, and reduce the credibility of the statistical atlas. To overcome this problem, a bootstrap approach is adopted (Section 1.5): still employ the same scheme of extracting and modeling anatomical variabilities, however, instead of using the entire training set, first build a statistical atlas from a subset of training samples that were more precisely registered; as will be presented in Chapter 4, anatomical knowledge embedded in the statistical atlas helps improve registration accuracy, thus enables the building of a more reliable statistical atlas from a bigger subset of accurately registered training samples. This process continues as more training samples are available, which allows the statistical atlas to be progressively improved.

Interesting observations can be made from the variance maps of intensity and geometry in the statistical atlas, as shown in Figure 34. The brightness of the variance maps corresponds to the magnitude of the variance. The intensity variance map (on the left) reveals distinctively low variations in the intensity of certain structures, such as the corpus callosum, dorsomedial nucleus, fronix, pons, midbrain, medulla, and vermis. This indicates that these structures have rather consistent densities between individuals. On the other hand, the geometric variance map (right) showed no distinguishable structure-wise differences.

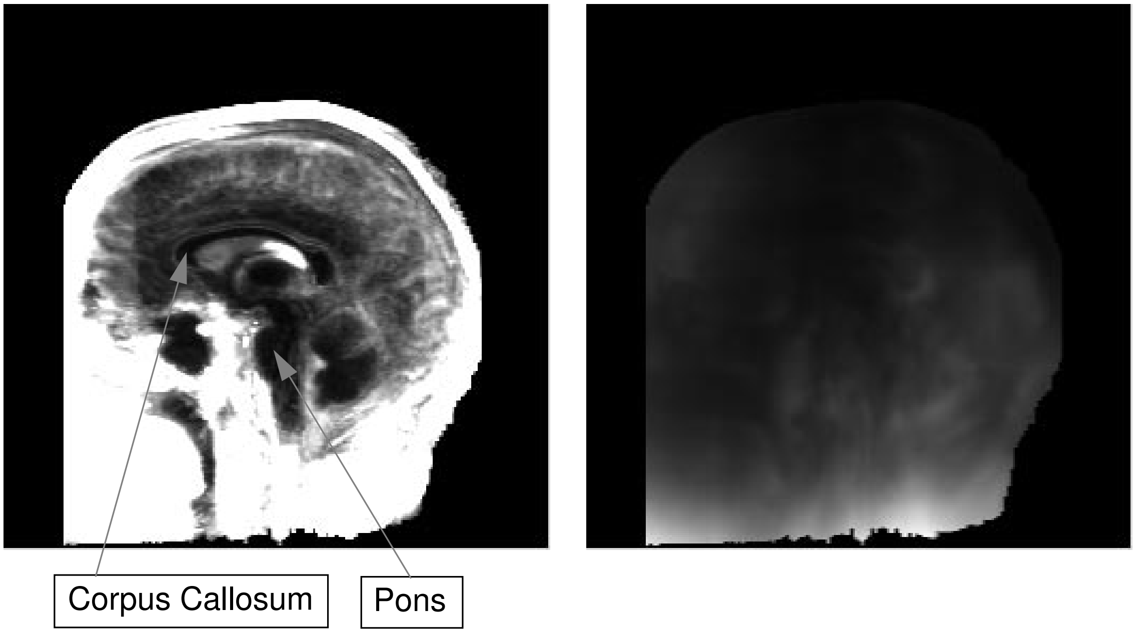

Figure 34: The intensity variance map (left) and the geometric variance map (right) of the statistical atlas. Brightness reflects the magnitude of the variance. Note that the intensity of certain structures, such as the corpus callosum and pons have low variance across the training set.

3.5 Discussion

The problem of modeling anatomical variability is approached by augmenting an atlas based on a single example into a statistical atlas built from a population. The statistics are gleaned by comparing samples from a training set in the common reference frame of the original atlas, using 3-D deformable registration. Patterns of anatomical variations are modeled using the first and second order statistics of the measurements. Other options were considered and investigated during the design process, which are presented here for future reference.

3.5.1 Population-specific training sets

The current statistical atlas was built on a training set of non-pathological brain MRI volumes. Due to the difficulty of acquiring sufficient data, subjects in this training set are not stratified into subpopulations according to age, gender, or any other demographic criteria. However, research in neurology has revealed that structural changes due to growth, aging, and certain psychiatric diseases are observable at both cellular and gross anatomic scales. In the case of the corpus callosum, it undergoes profound changes in morphology and composition during brain development, and these patterns are dramatically altered in psychiatric diseases such as multi-infarct dementia, schizophrenia, Alzheimer's Disease, etc. Therefore, a static representation is ill-suited to determining the dynamic effects of development and disease. With the availability of large image archives, population-specific atlases should be built to encode population variability.

3.5.2 Manual versus automatic classification

It is arguable that studying anatomical variations from a training set of expert classified data may produce more precise models. However, manually classifying high-resolution volumetric data is extremely time-consuming (it took an expert eight months to classify 144 structures in the atlas volume). This makes it difficult if not impossible to acquire a sizable training set. Moreover, the inherent subjectiveness and inconsistency in human performance, between people or for the same person over time, will likewise decrease the precision of models built upon it ([14], Chapter 1).

3.5.3 Multiple experts' opinions

The current statistical atlas encodes anatomical variations of a population. However, the anatomical classification is still one expert's opinion, which can be subjective and biased. Combining different experts' opinions in a probabilistic manner will provide a more impartial classification.

3.5.4 Global versus structure-based models

As discussed in Section 3.1, models of anatomical variations can be local or global. In this thesis, the primary goal of modeling anatomical variability is to achieve accurate registration and segmentation of all brain structures simultaneously. Therefore, a global model that encompasses all structures as well as their spatial relations and interactions is preferable. The current statistics are voxel-based, but structure or multi-structure based statistics can be conveniently modeled and incorporated within the framework.

3.5.5 Choice of method

In modeling geometric variations from the training set's 3-D deformation flow, other methods than voxel-based statistics were also considered, such as principal component analysis (PCA), modal analysis, and finite element surface model. The potential of using PCA was further investigated, and presented in Appendix B.

PCA is a more compact way to represent the average geometry and major patterns of variations. However, due to the high dimensionality of MRI data, it is currently impractical to conduct PCA on 3-D deformation flow of the whole volumes. While PCA can be used to model the mean geometry and major modes of variation of individual anatomical structures, critical information on their relative position and inter-structure interaction is lost.

One possible solution was investigated and presented in Appendix B, but it did not promise to improve registration accuracy over the existing algorithm.

3.6 Chapter summary

This chapter has focused on augmenting the atlas from a non-representative, single example into a statistical atlas that embodies anatomical variations of a population. Although there are many applications for a statistical atlas, the major interest of this thesis is to use it to achieve accurate registration and segmentation, which will be discussed in the following chapters.

CHAPTER 4

Registration Using the Statistical Atlas

The purpose of building the statistical atlas is to achieve more accurate registration through bootstrap. This chapter focuses on using the knowledge of anatomical variations to guide registration, so as to close the loop proposed in the framework:

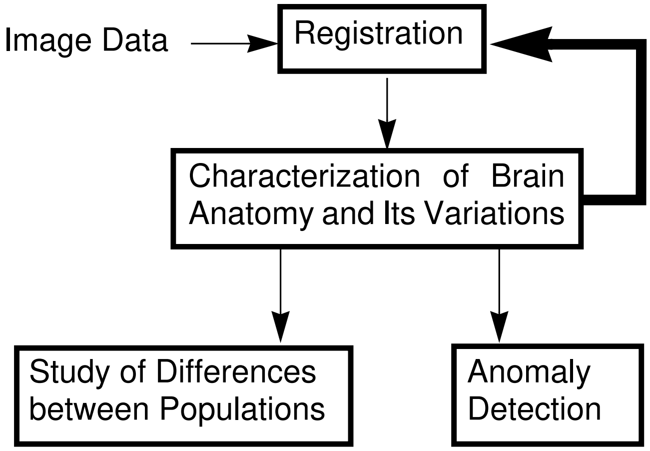

4.1 Related Work

Despite much effort on modeling individual variabilities of anatomical structures, few researchers have reported on using knowledge of anatomical variations to improve registration. Moghaddam et al. presented a probabilistic similarity measure for direct image matching based on a Bayesian analysis of image deformations [71]. They modeled variations in object appearance as probabilistic distributions, and computed a similarity measure based on the a posteriori probabilities. This approach was applied to 2-D image recognition. Gee et al. [31] proposed a Bayesian formulation of the image matching problem, and tested it on 2-D MRI data generated by applying manually defined distortions to real MRI data. Results on matching image volumes from different individuals are as yet unavailable. Wang and Staib posed 2-D deformation estimation problem in a maximum a posterori framework, in which statistical shape model based on corresponding boundary points were used as prior probabilities [104].

4.2 Registration using the statistical atlas

Similar to the approaches in related work, image registration is formulated within a Bayesian estimation-theoretic framework: the statistics derived from past observations specify the relevant prior probability model for the unknown spatial mapping. A Bayesian framework is well suited to the task of modeling variations in morphology. Because we are dealing with uncertain quantities, such as noisy data, from which information or decisions are derived (that must themselves then be uncertain), it is natural to adopt a probabilistic approach. Second, Bayesian analysis formally embodies the use of prior information that we have about the problem. In matching, the prior serves to constrain the mappings by favoring certain configurations, in effect, regularizing the problem. Statistical information about morphological variability, accumulated over past samples, can be formally introduced into the problem formulation to guide the matching or normalization of future data sets. Among the many operational advantages of Bayesian analysis, the most relevant to this work is that the result is a posterior distribution for the unknown mapping, which expresses the probability of any mapping given the observed images. The existence of this distribution makes possible a range of analyses, including the estimation of the variance or reliability of the estimated mapping.

The registration of two intensity volumes, such as a subject volume and the statistical atlas, can be formulated as finding the voxel deformation D that gives the highest a posterior probability $P\langle D|dI\rangle$. dI is the intensity difference between corresponding voxels in the two volumes. Using Bayes' Rule, $P\langle D|dI\rangle$ can be expressed as:

$$\mathcal{P}\langle \mathcal{D}|dI\rangle = \frac{\mathcal{P}\langle dI|\mathcal{D}\rangle \mathcal{P}(\mathcal{D})}{\mathcal{P}(dI)} \tag{8}$$

where

$$\mathcal{P}(dI) = \sum_{\mathcal{D}} \mathcal{P}(dI|\mathcal{D})$$

Therefore, the problem of finding the highest posterior probability, $P\langle D|dI\rangle$, is that of maximizing the right hand side of equation (8). Here $P(dI)$ is a constant C for two given image volumes, and the numerator has the same maximum as its logarithm. Substituting from equations (6) and (7) and taking logarithms yield:

$$\log \mathcal{P}\langle \mathcal{D}|dI\rangle = \log\frac{1}{\sqrt{2\pi}\sigma} - \frac{(dI-\mu)^2}{2\sigma^2} + \log\frac{1}{\sqrt{(2\pi)^3|\Phi|}} - \frac{(\overrightarrow{\Delta\vartheta} - \vec{\varpi})^{\mathcal{T}}\Phi^{-1}(\overrightarrow{\Delta\vartheta} - \vec{\varpi})}{2} - C$$

and hence maximizing the posterior probability is equivalent to minimizing the term:

$$\frac{(dI-\mu)^2}{2\sigma^2} + \frac{(\overrightarrow{\Delta\vartheta} - \vec{\varpi})^{\mathcal{T}}\Phi^{-1}(\overrightarrow{\Delta\vartheta} - \vec{\varpi})}{2} \tag{9}$$

Gradient descent is used to minimizes the term in (9). The 3-D gradient, ∇, at each step of the descent is given by the first order derivative of (9):

$$\nabla = \frac{dI - \mu}{\sigma^2}\nabla I + \Phi^{-1}(\overrightarrow{\Delta\vartheta} - \vec{\varpi}) \tag{10}$$

where ∇I is the 3-D image gradient, which is a function of the voxel's position. The first term will shift the current voxel by moving it in the direction that will reduce the difference between its current intensity and its mean intensity across all training samples. If the variance in this voxel's intensity distribution is large, the magnitude of the step along the image gradient will be reduced. The second term will shift the current voxel by moving it towards the mean geometric deformation of this voxel in the training set. If the variance of this voxel's geometric distribution is large, the magnitude of the step will be reduced. Since σ and Φ can have small values, a stabilizing factor is added to σ^2 and to the diagonal elements of Φ to regularize the gradient. The empirical value used is $\|\nabla I\| + 1$. The 3-D shift δD of the voxel is then:

$$\delta\mathcal{D} = -\lambda\nabla \tag{11}$$

here λ is a step size constant. In this way, each voxel in the statistical atlas is guided to search for a counterpart in the subject, so their match is most probable according to the statistics gathered from a population. 3-D Gaussian smoothing is applied to the voxels' 3-D displacements after each iteration to smooth the deformation. This compensates for the fact that the dependence between the deformation of neighboring voxels is not modeled in the currently voxel-based statistics.

This algorithm differs from the hierarchical deformable registration algorithm discussed in Chapter 2 in the measurement of the goodness of the voxel deformation flow. In this method, the goodness is the posterior probability of the current deformation, which is maximized using prior statistics gathered from a training set; in the previous algorithm, the *badness* is the intensity difference between spatially corresponding voxels in the atlas and the subject volume, which was minimized over the volumes. Before undergoing deformation, both algorithms globally align the two image volumes to eliminate extrinsic variations.

4.3 Performance evaluation

To validate this approach, the algorithm is applied to the same test set as described in Section 2.4. Experiments are conducted to examine the effectiveness of the intensity and geometric statistics alone as well as their combined strength.

4.3.1 Registration using the intensity statistics

The effectiveness of the intensity statistics model is tested by assuming a constant geometric prior probability. The maximization problem in equation (8) simplifies to:

$$\mathcal{P}\langle\mathcal{D}|dI\rangle \propto \mathcal{P}\langle dI|\mathcal{D}\rangle$$

and equation (10) becomes:

$$\nabla = \frac{dI - \mu}{\sigma^2}\nabla I$$

Using this method to register the atlas with the test set, the overall mislabelled voxel ratio is 3.8%. This is a 13.6% error reduction over the algorithm with no knowledge guidance, which had an error rate of 4.4% (Section 2.4).

4.3.2 Registration using the geometric statistics

In this experiment, the effectiveness of the geometric statistics model is assessed, with the intensity statistics assumed to be a constant. The maximization problem in equation (8) becomes:

$$\mathcal{P}(\mathcal{D}|dI) \propto \mathcal{P}(\mathcal{D})$$

and the 3-D gradient at each step is the second term in equation (10):

$$\nabla = \Phi^{-1}(\overrightarrow{\Delta\vartheta} - \vec{\bar{\omega}})$$

however, ∇ alone is insufficient to determine the deformation, because it ignores the images being registered, and forces the deformation toward the *average* brain. To combine the prior model prediction and the observed data, the inner product between ∇ and image gradient is used to obtain a 3-D deformation gradient:

$$\tilde{\nabla} = \left(\frac{\nabla}{\|\nabla\|} \circ \frac{\nabla I}{\|\nabla I\|}\right)\nabla \tag{12}$$

this helps to balance the influence of the prior distribution and the fidelity to the observations. When applied to register the atlas to the test set, this method yields an overall mislabelled voxel ratio of 4.05%. This is an 8% error reduction over the algorithm with no knowledge guidance.

4.3.3 Registration using the statistical atlas

When both intensity and geometric prior distributions are used, as derived in equation (10), registration between the atlas and the test set yields an overall mislabelled voxel ratio of 3.6%. This a 18.2% error reduction over the algorithm using the original atlas with no knowledge guidance. Figure 35 shows an example of improved registration performance using the statistical atlas, compared to result given by using the original atlas.

These experiments show that while the intensity statistics model is more effective than the geometric one at guiding the deformable registration, the best registration result is achieved when the intensity and geometric statistics models are combined. Figure 36 compares registration performance using the original atlas, using voxel-based intensity statistics only, using voxel-based geometric statistics only, and using the statistical atlas.

Registration using the full statistical atlas takes 35 minutes to register 3-D images of 256x256x124 voxels on a SGI workstation with four 194 MHz processors.

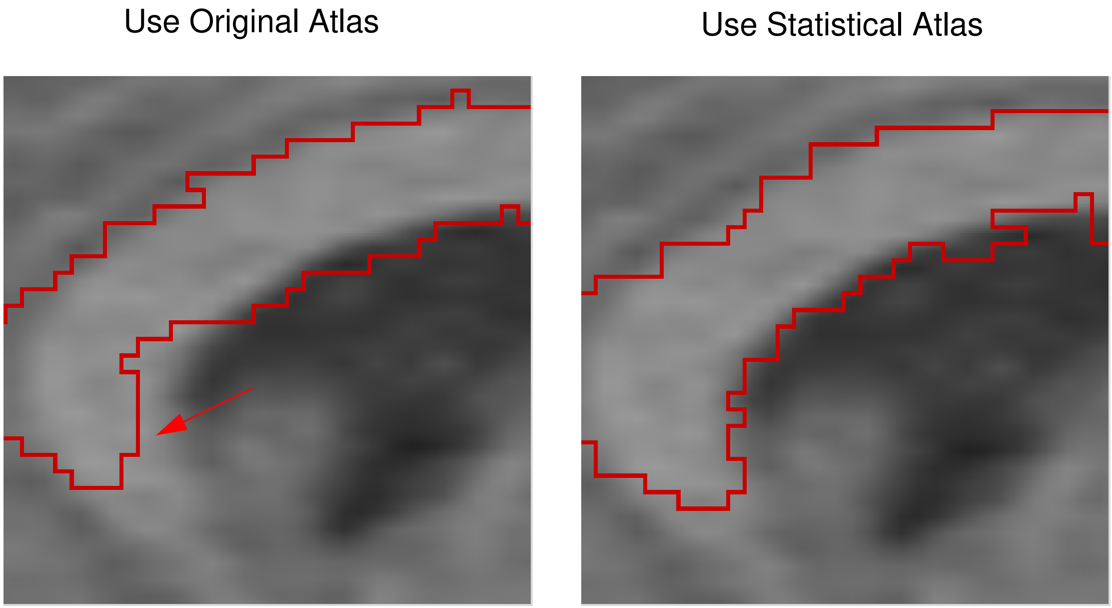

Use Original Atlas

Use Statistical Atlas

Figure 35: Registration using the statistical atlas achieves better performance than using the original atlas.

4.4 Discussion

In order to build the statistical atlas, each image volume in the training set was warped to match the atlas. Therefore, the geometric statistics represent the distribution of 3-D displacements from each atlas voxel to voxels in different subject volumes. However, for most applications, it is preferable to compute the inverse deformation i.e. to warp the atlas to match the subject volume. In this way, anatomical classifications in the atlas can be customized for the particular subject, thus facilitate interesting applications. Unfortunately, the analytical inverse of the voxel-based 3-D deformation does not exist. One practical solution is to interpolate the inverse deformation.

Figure 37 gives an 1-D illustration. Assume $\overrightarrow{\vartheta}_a$ is one dimension in the atlas, and $\overrightarrow{\vartheta}_s$ is the corresponding dimension in the subject volume. From the statistical atlas, the correspondence from each position on $\overrightarrow{\vartheta}_a$ to that on $\overrightarrow{\vartheta}_s$ can be computed. What is unknown is

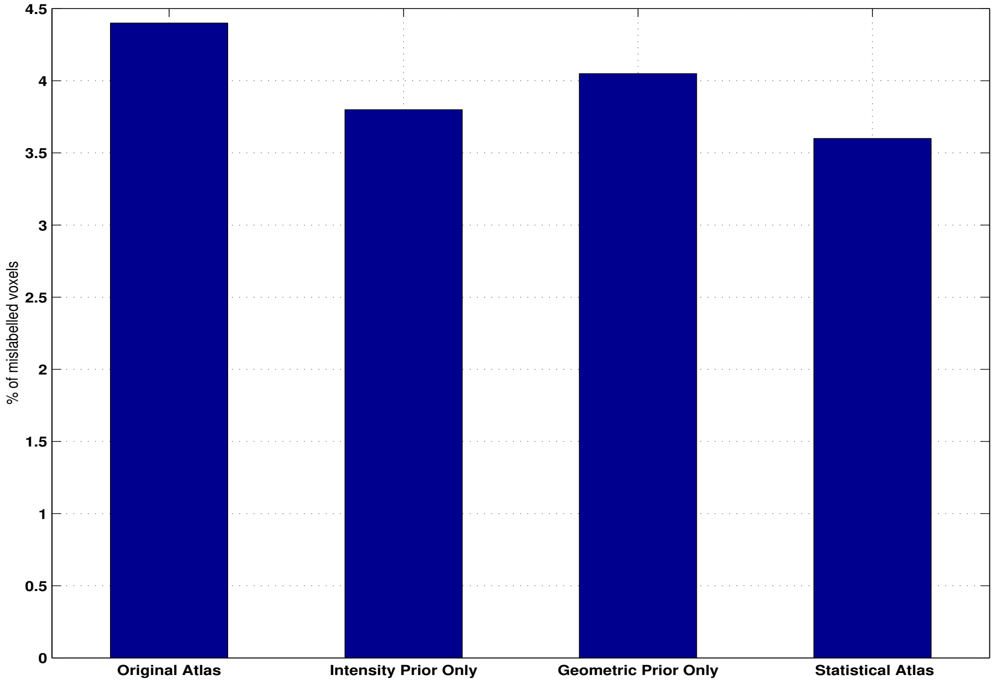

Figure 36: Error ratios of the test set given by registration using the original atlas, using voxel-based intensity statistics only, using voxel-based geometric statistics only, and using the statistical atlas. The intensity statistics prior is more effective than the geometric one at guiding registration; the best result is achieved using the statistical atlas with both intensity and geometric statistics.

the correspondence from each position on $\overrightarrow{\vartheta_s}$ to that on $\overrightarrow{\vartheta_a}$. Take a particular position $\overrightarrow{\vartheta_{s2}}$, the current hypothesis in gradient descent corresponds it to $\overrightarrow{\vartheta_{a1}}$. However, from the prior probabilities encoded in the statistical atlas, $\overrightarrow{\vartheta_{a1}}$ should correspond to $\overrightarrow{\vartheta_{s1}}$ via a displacement $\mu(\overrightarrow{\vartheta_{a1}})$. What needs to be determined is the correspondence for $\overrightarrow{\vartheta_{s2}}$ that is consistent with the prior knowledge in the statistical atlas. Assume that the correct correspondence for $\overrightarrow{\vartheta_{s2}}$ is only one small step $\overrightarrow{\delta\vartheta_a}$ away from $\overrightarrow{\vartheta_{a1}}$, and that changes between neighboring displacements are linear, we have

$$\overrightarrow{\vartheta_{s2}} = \overrightarrow{\vartheta_{a1}} + \overrightarrow{\delta\vartheta_a} + \mu(\overrightarrow{\vartheta_{a1}} + \overrightarrow{\delta\vartheta_a}) \tag{13}$$

Take Taylor expansion of $\mu(\overrightarrow{\vartheta_{a1}} + \overrightarrow{\delta\vartheta_a})$, and ignore higher order terms give:

$$\mu(\overrightarrow{\vartheta_{a1}} + \overrightarrow{\delta\vartheta_a}) = \mu(\overrightarrow{\vartheta_{a1}}) + (\nabla\mu(\overrightarrow{\vartheta_{a1}}))\overrightarrow{\delta\vartheta_a} \tag{14}$$

Combining the above two equations gives

$$\overrightarrow{\delta\vartheta_a} = (1 + \nabla\mu(\overrightarrow{\vartheta_{a1}}))^{-1}(\overrightarrow{\vartheta_{s2}} - \overrightarrow{\vartheta_{a1}} - \mu(\overrightarrow{\vartheta_{a1}}))$$

In this way, the correspondence for $\overrightarrow{\vartheta_{s2}}$, $\overrightarrow{\vartheta_{a2}}$, is determined:

$$\overrightarrow{\vartheta_{a2}} = \overrightarrow{\vartheta_{a1}} + \overrightarrow{\delta\vartheta_a}$$

Empirically, the assumptions of small step and linear change in neighboring flows hold well. More study on neighboring flows will be presented in Chapter 5.

4.5 Chapter summary

This chapter applies the statistical atlas to 3-D deformable registration. Using Bayesian formulation, statistics embedded in the atlas function as prior probabilities. The optimal deformation maximizes the posterior probability. This approach is evaluated by registering the statistical atlas to the same test set as described in Section 2.4. The overall error rate is 3.6%, which is a 18.2% reduction over the previous algorithm (Section 2.4). However, the current knowledge model is simply voxel-based statistics, considerations on neighborhood context will be discussed in the next chapter.

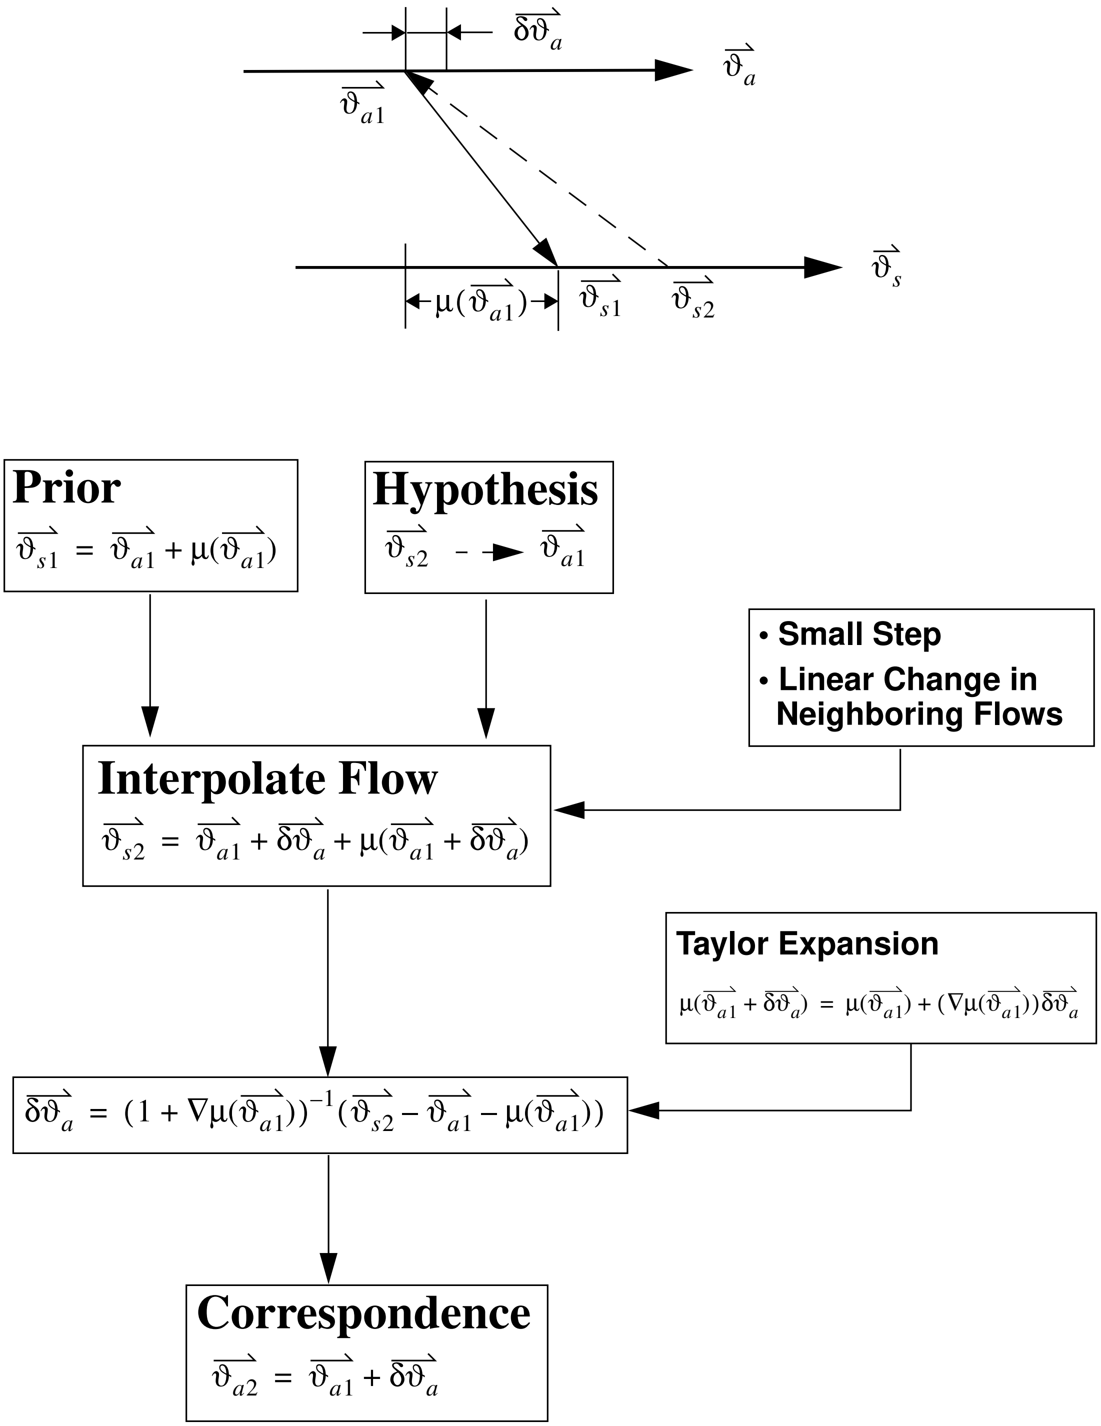

Figure 37: 1-D illustration of interpolating the inverse deformation

CHAPTER 5 Model Neighborhood Context

Currently, anatomical variations between individuals are modeled as voxel-based statistics. This representation is simple and efficient, as demonstrated in Section 4.3. However, information on voxel neighborhood context is missing. Figure 38 is a close-up on the posterior part of corpus callosum in the mid-sagittal plane. 2-D projections of the 3-D deformation flow are overlaid on top of the intensity image. Note that the deformation flow is smooth and congruous locally, which means the deformation at each voxel is not completely independent of its neighbors. Therefore, a more comprehensive model should consider this dependency between neighboring voxels, which will be the focus of this chapter.

5.1 Modeling neighborhood context

A convenient way to model neighborhood dependency is a direct higher dimensional extension of the voxel-based statistics models. Consider a 3-D neighborhood of NxMxK centered at an atlas voxel; the distribution of intensity variations of this neighborhood is modeled as an $\Re$ dimensional Gaussian distribution, where $\Re$ equals NxMxK:

$$\mathcal{P}\langle\overrightarrow{dI}|\mathcal{D}\rangle = \frac{1}{\sqrt{(2\pi)^{\Re}|\Sigma|}}e^{-\frac{(\overrightarrow{dI}-\vec{\mu})^{\mathcal{T}}\Sigma^{-1}(\overrightarrow{dI}-\vec{\mu})}{2}} \tag{15}$$

here $\overrightarrow{dI}$ is an $\Re x 1$ vector of intensity differences between the corresponding neighborhoods in the subject volume and the atlas; $\mathcal{D}$ is the current 3-D deformation. $\vec{\mu}$ is the $\Re x 1$ mean vector of the neighborhood intensity variation distribution; Σ is the $\Re x \Re$ covariance matrix of the intensity variation distribution. Here $\overrightarrow{dI}$ has been adjusted for extrinsic intensity variations induced by separate image acquisition processes.

Similarly geometric variations of a NxMxK neighborhood centered at each atlas voxel can be modeled as a $3\Re$ dimensional Gaussian distribution:

$$\mathcal{P}(\mathcal{D}) = \frac{1}{\sqrt{(2\pi)^{3\Re}|\Psi|}}e^{-\frac{(\overrightarrow{\Delta\vartheta}-\vec{\varpi})^{\mathcal{T}}\Psi^{-1}(\overrightarrow{\Delta\vartheta}-\vec{\varpi})}{2}} \tag{16}$$

where $\overrightarrow{\Delta\vartheta}$ is the $3\Re x 1$ vector of the neighborhood's 3-D displacements; $\vec{\varpi}$ is the $3\Re \times 1$ mean vector of the neighborhood's 3-D deformation flow; Ψ is the $3\Re \times 3\Re$ covariance matrix of the distribution of geometric variations. Here $\overrightarrow{\Delta\vartheta}$ has been adjusted for extrinsic geometrical variations. Note that the voxel-based statistics models are a special case of the neighborhood statistics models with a 1x1x1 neighborhood. Also note that the size of Σ and Ψ grows exponentially as $\Re$ increases, which makes computation expensive, as will be discussed in the following section.

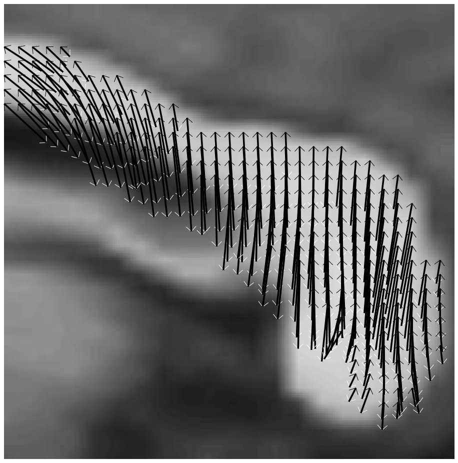

Figure 38: the posterior portion of the corpus callosum in the mid-sagittal plane, overlaid with the 2-D projection of the 3-D deformation flow. The deformation is smooth locally.

5.2 Registration using neighborhood context

The same deduction procedure as in Section 4.2 is used to determine the deformation that maximizes the posterior probability $\mathcal{P}\langle\mathcal{D}|dI\rangle$ for a voxel neighborhood; the 3-D gradient of voxels in the neighborhood is:

$$\nabla = \Sigma^{-1}(\overrightarrow{dI} - \vec{\mu})\nabla I + \Psi^{-1}(\overrightarrow{\Delta\vartheta} - \vec{\omega}) \qquad (17)$$

Theoretically this algorithm can be implemented in the same way as the voxel-based statistics models; however, the $\Re x \Re$ intensity covariance matrix Σ has $\Re(\Re + 1)/2$ distinct entries, and the $3\Re \times 3\Re$ geometric covariance matrix Ψ has $(3\Re(3\Re + 1))/2$ distinct entries. MRI volumes typically have more than 8 million voxels. Even if all entries in the covariance matrices can be stored as bytes, the covariance information for each 2x2x2 voxel neighborhood will still require 336 MByte memory. Together with other memory requirements, the total demand makes this approach impractical.

To simplify the problem, only the interactions between immediate neighbors are considered. Instead of storing interactions between immediate neighbors, they are computed on the fly. The voxel-neighbor interactions are estimated using the *goodness* of the neighbors' current match according to their respective prior distributions. Figure 39 illustrates the idea: the goal is to find a match for the black voxel in the subject; however, when no neighborhood context is considered, there is ambiguity as to which black voxel in the atlas it should match to. But when considering the goodness of how well the neighbors of the black voxel in the subject match with their counterparts in the atlas, it is clear which atlas voxel it should match to.

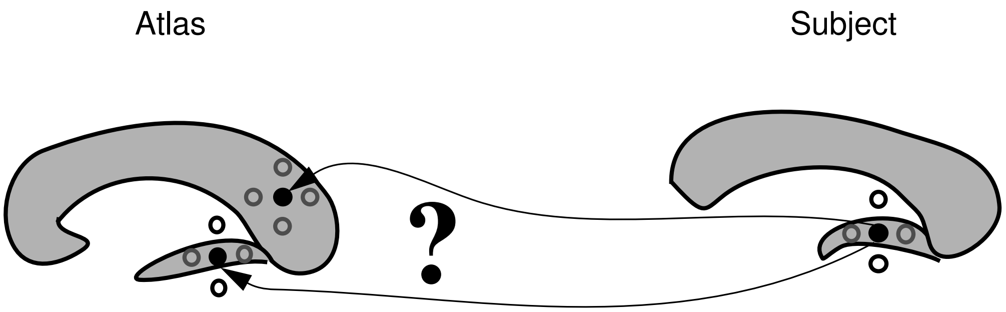

Figure 39: Registration using neighborhood context

Using a weighted-window matching approach, the *goodness* is weighted by the distance between the current voxel under investigation and the particular neighbor. Therefore, for a voxel neighborhood $\aleph$, the 3-D gradient determined using neighborhood context is a direct extension of equation (10):

$$\nabla = \sum_{i,j,k} w_{ijk}\left(\frac{dI - \mu}{\sigma^2}\nabla I\right)_{ijk} + \sum_{i,j,k} w_{ijk}[\Phi^{-1}(\overrightarrow{\Delta\vartheta} - \vec{\omega})]_{ijk} \qquad ijk \; \text{Œ} \; \aleph \tag{18}$$

5.3 Performance evaluation

The algorithm is evaluated by registering the atlas to the same test set as in Section 4.3. The size of the voxel neighborhood used is 3x3x3. Weights of neighboring voxels are inversely proportional to their distance from the voxel under examination. Empirical numbers used are 0.01 for 8 farthest neighbors, 0.014 for 12 intermediate neighbors, and 0.017 for 6 closest neighbors. Experiments showed that as long as major weight is put on the central voxel, registration results are not sensitive to the particular numbers. Registration using neighborhood context achieves an overall error ratio of 2.9%. This is a 34% error reduction over the algorithm with no knowledge guidance, and a 19.4% error reduction over registration using the statistical atlas. Figure 40 shows an example of better registration achieved by using neighborhood context, and Figure 41 compares examples of segmentation given by the algorithm, and the respective segmentation given by experts.

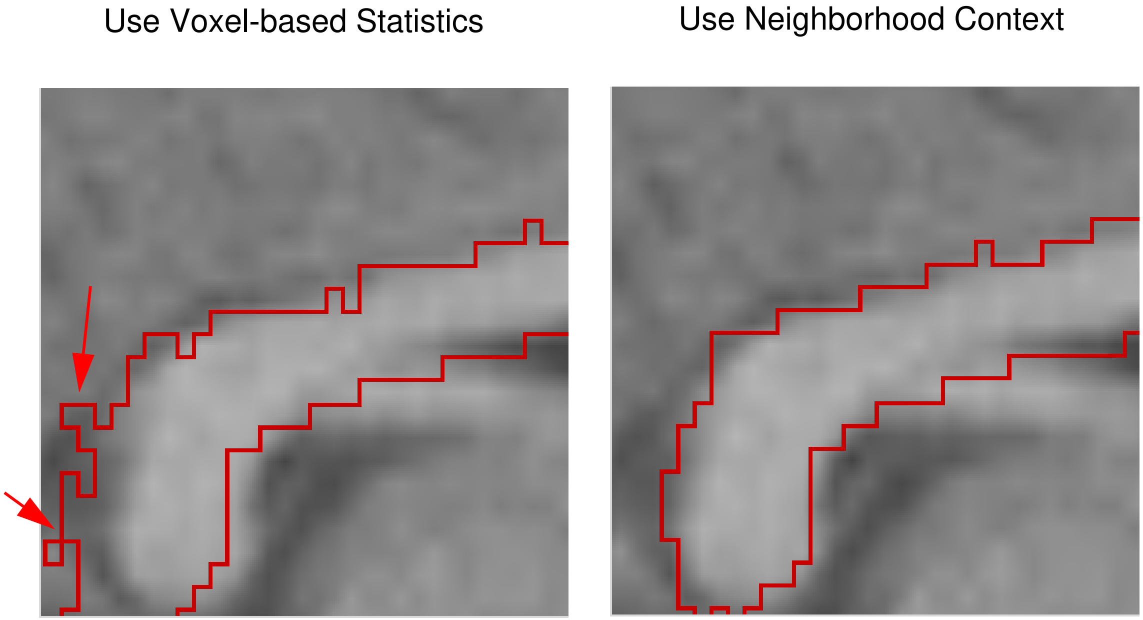

Figure 40: Better registration is achieved by considering neighborhood context.

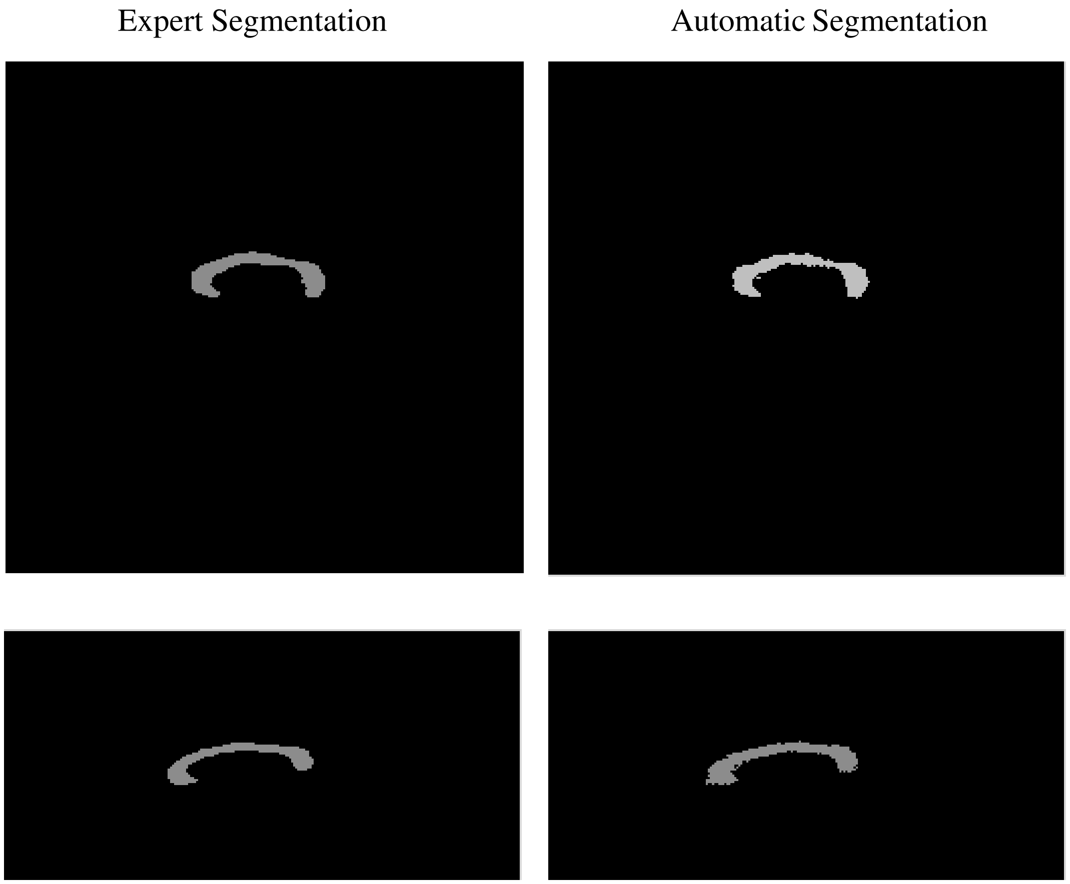

Figure 41: Examples of automatic segmentation compared to their respective expert segmentation

This weighted window matching approach yields better results than using voxel-based statistics without increasing memory demand, as compared in Figure 42. However, it is much more computationally expensive. To register 3-D images of 256x256x124 voxels on a SGI workstation with four 194 MHz processors will take 57 minutes. Appendix C compares the expert segmentation and the computed segmentation of the corpus callosum in the mid-sagittal plane of each of the 40 test cases. The images are binary, with the corpus callosum shown at intensity 140, and the rest shown as zero intensity background. Error ratios given by the weighted window matching approach are listed under respective computed segmentations.

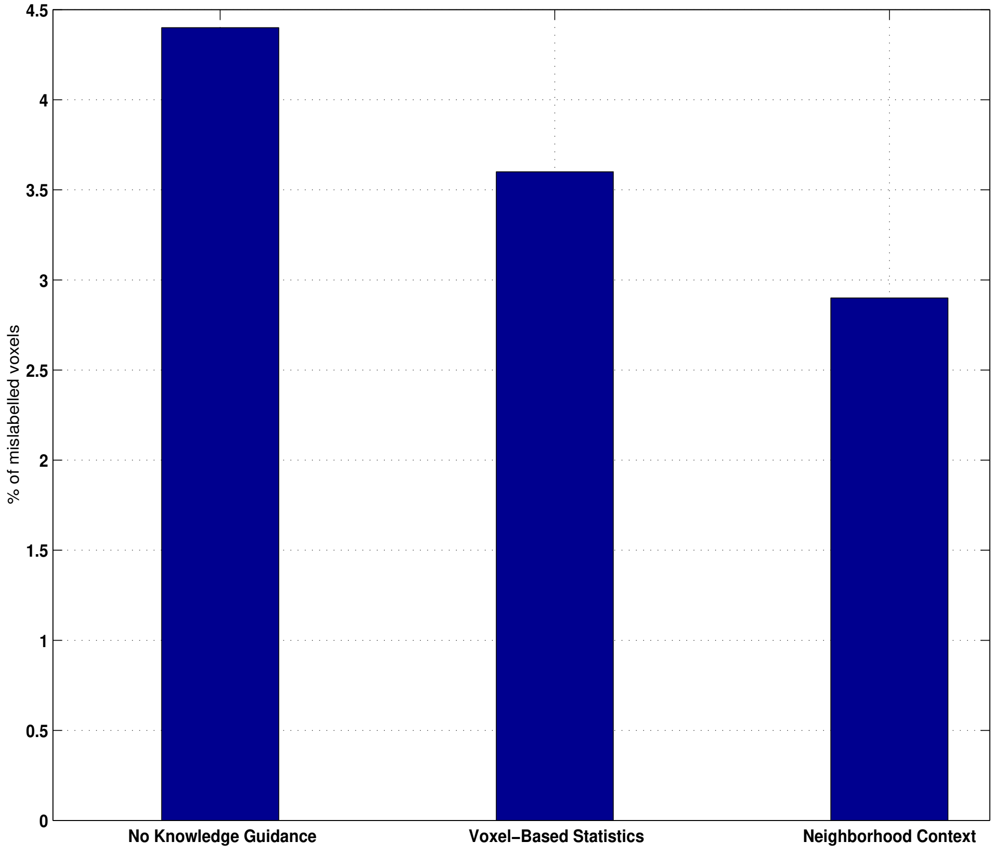

Figure 42: Error ratios of the test set given by registration without knowledge guidance, registration using voxel-based statistics, and registration using neighborhood context. The best result is achieved when neighborhood context is considered.

5.4 Distance Error Metric

So far the quantitative evaluation of registration performance has been using the ratio of mislabelled voxels as the measurement. While this metric is effective, it does not differentiate various error voxels. Figure 43 illustrates a scenario where more discriminative evaluation is necessary. The shaded region is the ground truth segmentation given by an expert, and dashed outlines are results yielded by the registration algorithm. Suppose the two examples shown in Figure 43 have the same number of mislabelled voxels, the current error metric will give them the same evaluation. However, errors in these two results are of different significance. For the example shown on the left, its error voxels occur close to the boundary of the ground truth segmentation; whereas for the result on the right, many of its error voxels stray considerably away from the ground truth segmentation. Therefore, in terms of segmentation quality, the example on the left is better than the one on the right.

Figure 43: A senario that demands more discriminative evaluation. The shaded region is the ground truth segmentation given by an expert, and dashed outlines are results yielded by the registration algorithm. Suppose the two examples have the same number of mislabelled voxels, using the ratio of mislabelled voxels will give them the same evaluation. However,in terms of segmentation quality, the example shown on the left is better than the one shown on the right.

A more discerning measurement is the distance between each boundary voxel of the computed segmentation and the closest boundary voxel of the ground truth segmentation. Under this metric, error voxels that are far away from the true boundary will be penalized more than error voxels that occur close by the ground truth., i.e. it de-emphasizes errors that are near the true boundary. This is especially advantageous considering boundary voxels are of higher uncertainty even in expert segmentation.

Using the distance error metric to evaluate the registration performance of the weighted window matching approach, the mean distance of error voxels for the test set is 1.535 mm, and the standard deviation is 0.762 mm. Figure 44 shows histogram distributions of distance measurements of boundary voxels. The histogram at the top includes the correct boundary voxels, i.e. the voxels that lie on the true boundary. Because the number of correct boundary voxels is orders of magnitudes larger than the number of erroneous boundary voxels, it is impossible to visualize the distribution of distance measurements of error voxels. Therefore, the correct boundary voxels are excluded in the histogram on the bottom, thus allowing visualization of the distribution of error voxels in terms of the distance metric. The vertical axis of this histogram is at a scale $7/9000$ of the top one. Note that most error voxels are within 2 mm from true boundaries, with few voxels being farther than 4 mm away, and no voxels are beyond 9.4 mm distance. Considering the average length of the corpus callosum is around 100 mm (4 inches), this level of accuracy is satisfactory.

5.5 Discussion

Markov Random Field (MRF) has been used in computer vision to model smoothness as well as different textures. Typical solvers for MRFs include Gibbs sampling [35], the Metropolis algorithm [68], Iterated Conditional Modes (ICM) [5], and Mean-Field (MF) methods [34]. Kapur et al. used a discrete vector valued Markov Random Field model as a regularizing prior in a Bayesian classification algorithm to minimize the effect of salt-and-pepper noise. They presented preliminary results on segmentation of white matter, gray matter, fluid, and fat in brain MRIs [49]. The possibility of using MRF to model neighborhood context was investigated. However, the dimensionality of volumetric data and the 3-D nature of deformation fields have so far made the memory demand prohibitive even for the 1-order model. With computer memory more readily available, this is an alternative worth studying.

5.6 Chapter summary

This chapter discussed the modeling of neighborhood context in the statistical priors, and an implementation of applying such knowledge to guide registration. Quantitative evaluation demonstrated better performance than using just voxel-based statistics. A new error metric based on distances between boundary voxels of the computed segmentation and the ground truth is introduced, and employed to assess the registration performance using neighborhood context.

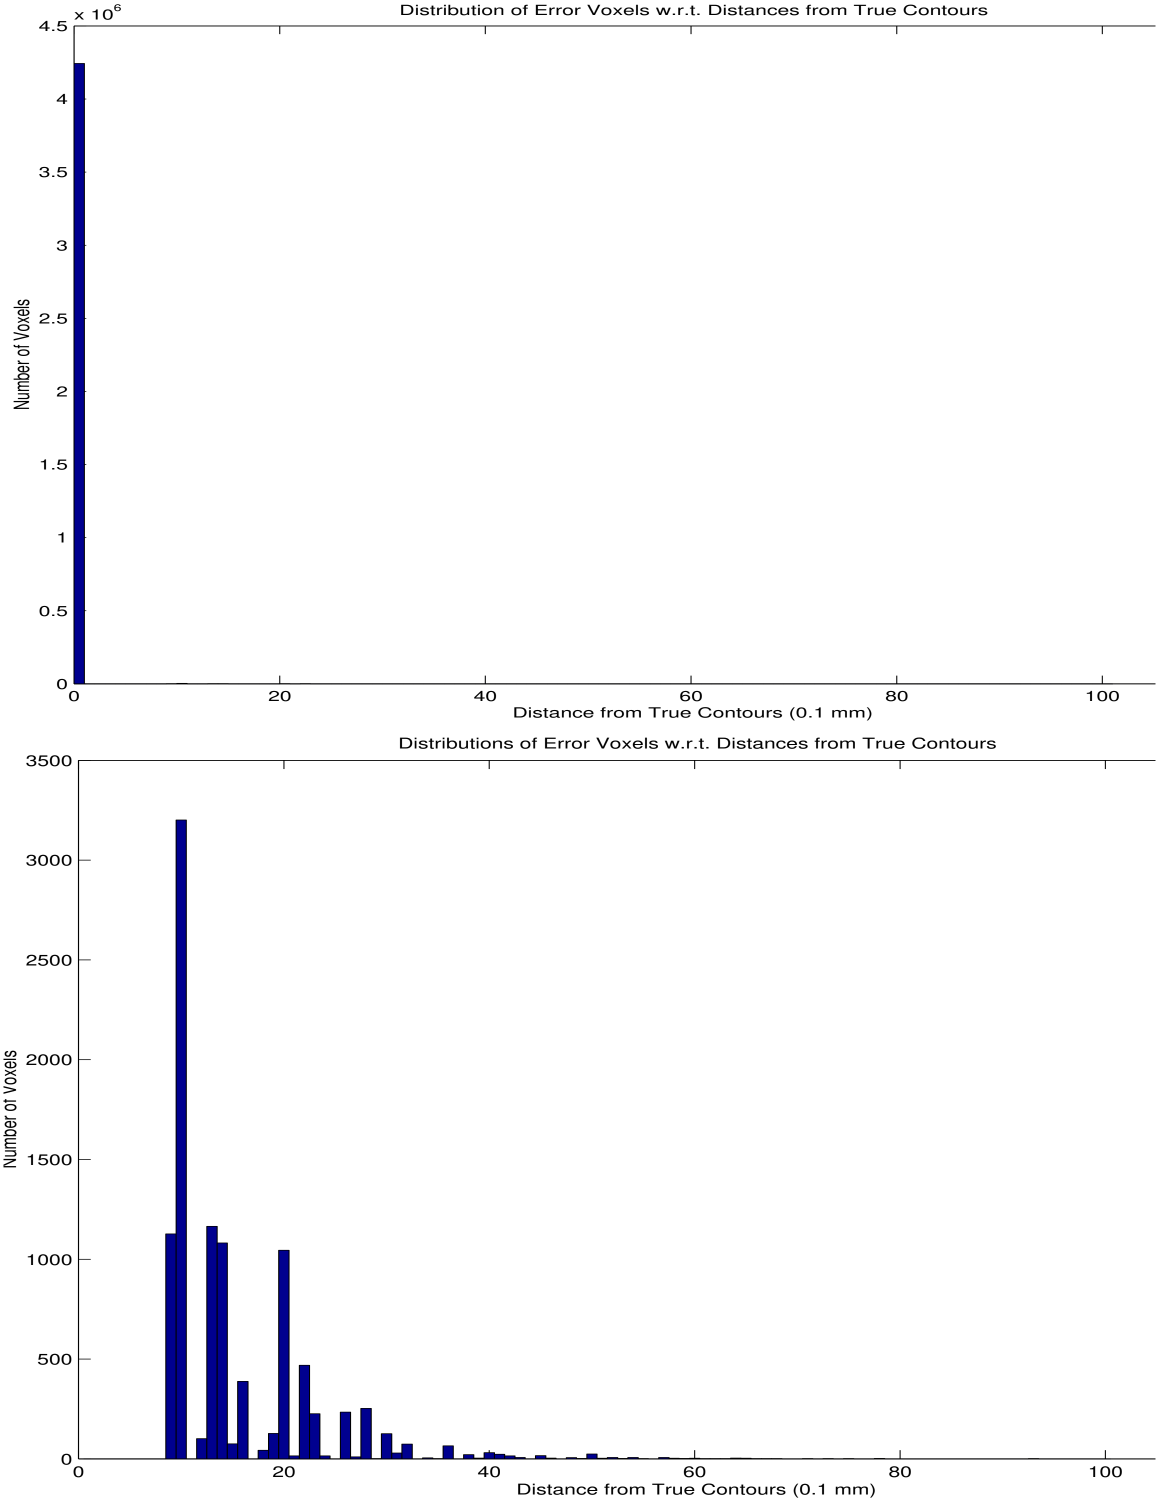

Figure 44: Histogram distributions of distance measurements of boundary voxels. The top graph includes the correct boundary voxels, whereas the bottom graph shows the distribution of error voxels only.

CHAPTER 6 Implementation Details

Several issues in the implementation are crucial to robust and efficient performance. They apply to both the algorithm using the statistical atlas, and the earlier method without knowledge guidance. This chapter will discuss them in detail.

6.1 Background Separation

Since the algorithms do not require prior segmentation or feature extraction, each voxel in the image data is given equal weight in the registration. However, there is often background noise in the data, which necessitates background separation to ensure only relevant voxels contribute to the registration process.

A thresholding method is applied to eliminate dim background noise. The threshold is automatically selected as the first valley of the intensity histogram of the 3-D image, after it is smoothed using a Gaussian filter (the Gaussian kernel used is of size 7). This produces a binary image volume with the background intensity set to zero.

Apart from dim background noise, there also exists bright outliers that have similar intensities as certain anatomical structures. Since the region containing the head volume is connected, a connected-component algorithm is applied to the thresholded binary image volume, and the largest connected component is selected. Any holes caused by dark regions inside the head volume are filled in. The result is a binary image volume with only voxels inside the connected component being above zero intensity. This binary volume functions as a mask to ensure only voxels within the head volume are considered in later processing.

6.2 Adaptive Multi-resolution Processing

Multi-resolution processing has been extensively used in computer vision algorithms, such as stereo matching and motion estimation. This approach first solves the matching problem on smaller, lower-resolution data, use the result to initialize higher-resolution estimates, and refine it using details present at the higher resolutions. Its advantage includes increased computation efficiency, as well as the ability to escape from local minima by first focusing on global patterns.

MRI volumes are high dimensional data, which makes efficient processing imperative. The current registration algorithm uses two strands of adaptable multi-resolution processing: one is an image pyramid, the other is a control grid pyramid (in smooth deformation). Sizes of MRI volumes vary depending on the number of slices. Among the over two hundred MRIs processed by the registration algorithm, the number of slices ranges from 7 to 256. To make the multi-resolution scheme adaptive to different volume sizes, the minimum number of slices in the subsampled image volumes and the control grid cells are enforced. The current empirical number is 30 slices for image volumes, and 15 slices for each control grid cell. This ensures the subsampled volumes and the control grid cells to contain sufficient image information. In the image pyramid, each level has half the resolution of the next higher level; in the control grid pyramid, the numbers of control points along x, y, and z increase from 2 at an interval of 1 at each higher lever. For image data so far examined, the highest image pyramid has 3 levels, and the highest control grid pyramid has 7x7x7 control points.

6.3 Random Initialization

For the ease of application, the algorithms are fully automated. Global Alignment, the first step of registration, is initiated by trying many random similarity transformations and picking the one which gives the minimum SSD [74].

Because brain MRI volumes are not taken in a completely random manner, it is efficient to limit the infinite choices to those that are possible. Since the registration between the atlas and different subjects is of particular interest, the coordinate frame of one volume, the atlas, is thus fixed. There are three principal axes along which brain MRIs are taken (Section 1.4.1), which gives six possible gross orientations of the volumes. For initialization, it is sufficient to limit the possible rotations from the atlas to these six orientations. As for translation and scaling, possible transformations are randomly sampled at small intervals within empirical ranges. Currently, random initialization is conducted on images volumes at the coarsest level of the image pyramid. The range for translation is [-8, 8] voxels along each X-Y-Z dimension, the range for uniform 3-D scaling is [0.9, 1,1]. The number of random translation and scaling tried for each rotation angle is 300.

6.4 Stochastic Sampling

To further improve algorithm efficiency, not every voxel is processed during the optimization in global alignment and smooth deformation. Instead, a random set of voxels is sampled at each iteration of the optimization processes. This method is valid because the optimization processes are over-constrained, i.e. the number of voxels is orders of magnitudes greater than the number of parameters to be optimized. It is not necessary to involve each voxel to solve for the parameters. Moreover, the stochastic nature of the sampling helps the optimization processes to escape from local minima [103]. In the current implementation, the percentage of voxels sampled is 1%. For fine-tuning deformation, it has 3 times as many parameters as the number of voxels, i.e. the problem is highly underconstrained, therefore stochastic sampling is not appropriate.

6.5 Parallel Processing

In the registration algorithms developed in this thesis, the computation for each voxel is identical. Therefore, the process is voxel-wise parallelizable. To reduce overhead, parallel processing is employed at a higher level than the voxel representation. During the random initialization of global alignment, each combination of transformation parameters is examined in parallel; during smooth deformation, each control grid cell is processed in parallel; whereas in fine-tuning deformation, each slice of the image volume is processed in parallel.

CHAPTER 7 Quantitative Study of the Anatomy

A quantitative study of anatomical structures brings broad prospects to medical research. It allows measurement of the range of normal variation, and the detection of abnormalities. Statistics on anatomy and pathologies can help express expert knowledge so as to enhance medical education. Furthermore, quantitative description of medical image content will facilitate efficient retrieval in the ever increasing medical databases.

Imaging techniques such as computed tomography and MRI have made it possible to make non-invasive measurements of anatomical structures. However, traditional morphometry or volumetry (i.e. the measurement of substructural volumes), has been performed manually or semi-automatically. These methods are limited by relatively high error and low repeatability attributable to the subjective assessment of substructure boundaries required for manual outlining. However, the magnitude of the volume changes in diseases such as neurological and psychiatric disorders may be small relative to the error associated with these measurement techniques. Also, manual operations are tedious and time-consuming, which hampers study of statistically significant data sets.

Major sources of error in manual measurement of anatomical volumes, in order of magnitude, come from (a) subject variability, (b) the observer, (c) the method of volume estimation from the digital image, (d) voxel representation, and (e) the imager. Automatic algorithms can help reduce the second and third sources of error listed above. For example, it was reported that using interactive thresholding methods, expert observers could learn to measure anatomical structures with a precision of about 2%; i.e. with practice a single person could learn to measure structures such as grey and white matter repeatedly on the same brain image within that degree of consistency. However, the same observers produced measurements that differed between observers by 15% [41]. Similar intra- and inter-operator reliability of 5-20% was also mentioned in [70]. Since a long-term clinical study cannot depend on the availability of the same observer throughout the study, a more objective automatic algorithm with comparable accuracy is advantageous. Besides, characterization of subtle anatomic aberrations, such as have been found in schizophrenia, requires quantifying local variability across subjects. This requirement makes high dimensionality of the transformations on the coordinate systems essential, which can be readily provided by automatic algorithms that allow local deformations. In addition, the efficiency of an automatic algorithm allows the study of large data sets, which is basic to sound conclusions. Furthermore, it has been the case in this field that different research groups collect separate data sets, and use different operators or semi-automatic tools to analyze the data. Although they have had conflicting conclusions, it has been hard to reproduce or cross-validate each other's results. The performance consistency of automatic algorithms makes cross-examination possible. This chapter will discuss the application of the aforementioned registration algorithm to quantitative study of the anatomy:

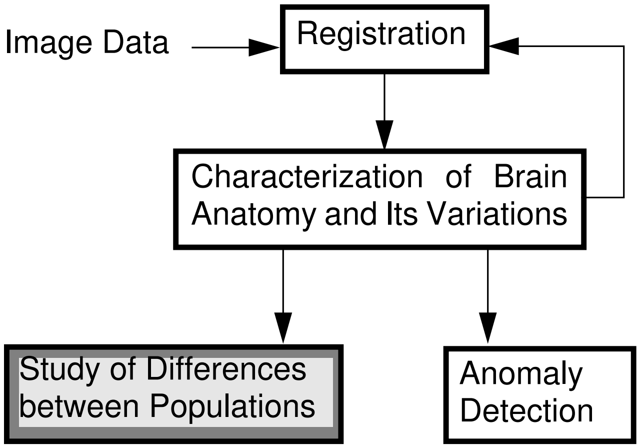

7.1 Related work

Due to its interdisciplinary nature, the application of automatic algorithms in the quantitative study of anatomy has been scarce. Haller et al. examined automated 3-D MRI morphometry of the human hippocampus by deformably matching a canonical atlas to 10 test volumes [39]. Their results showed that the automated method estimated hippocampal volumes with less variability than that of manual outlining. Davatzikos et al. proposed an approach for quantifying the shape of the corpus callosum, based on the Talairach space normalization and on the measurement of a 2-D deformation function resulting from elastically registering the Talairach atlas with subject images [21]. They compared the corpus callosum in the mid-sagittal plane of eight male and eight female subjects. Results suggested that the splenial (posterior) region in adult females' corpus callosum is larger than in males, which agreed with some previous reports.

7.2 An automatic system for the quantitative study of anatomy

The 3-D deformable registration algorithm developed in this thesis can readily apply to the quantitative study of the anatomy. As discussed in Section 1.1, by registering the atlas with a particular subject's image volume, the anatomical classification in the atlas can be customized for the subject, therefore the 3-D segmentation of the subject's anatomical structures (Figure 3). Once the segmentation is acquired, quantitative study on the size and shape of anatomical structures can be conducted.

In order to investigate and demonstrate the feasibility of such approach, a convenient interface has been developed for quantitatively examining brain anatomy, as shown in Figure 45. The left image is one slice of a subject's T1-weighted MRI volume, and on the right is the corresponding segmentation automatically customized from the atlas. The segmentation is color-coded to represent different anatomical structures. To study any structure, the user simply clicks on it, and the system instantly overlays the segmentation outline(s) on the intensity data. The name of the structure and several measurements such as its volume and average density are automatically computed for illustration. For structures having symmetric counterparts, measurements of the parallel are also displayed. The slide bar in the left bottom allows the user to examine the 3-D volume. The underlining computer languages of this interface are C/C++ and Tcl/Tk, which are commonly available on SGI Irix and SUN Solaris platforms. For applications looking at large populations, the automatic atlas-subject registration, customization of anatomical segmentation, and quantitative analysis can be readily pipelined.

7.3 Feasibility study

In order to prove the potential of automatic quantitative analysis of the anatomy, and to test its acceptability to medical professionals, several collaborative study have been conducted with research groups in medical institutions.

7.3.1 Lateral Ventricles and Schizophrenia

Ever since the evolution of the concept of the schizophrenic illness, it has been argued that at least some patients may have a deteriorative process at work. However, studies examining the progression of brain structural alterations in schizophrenia have led to con-

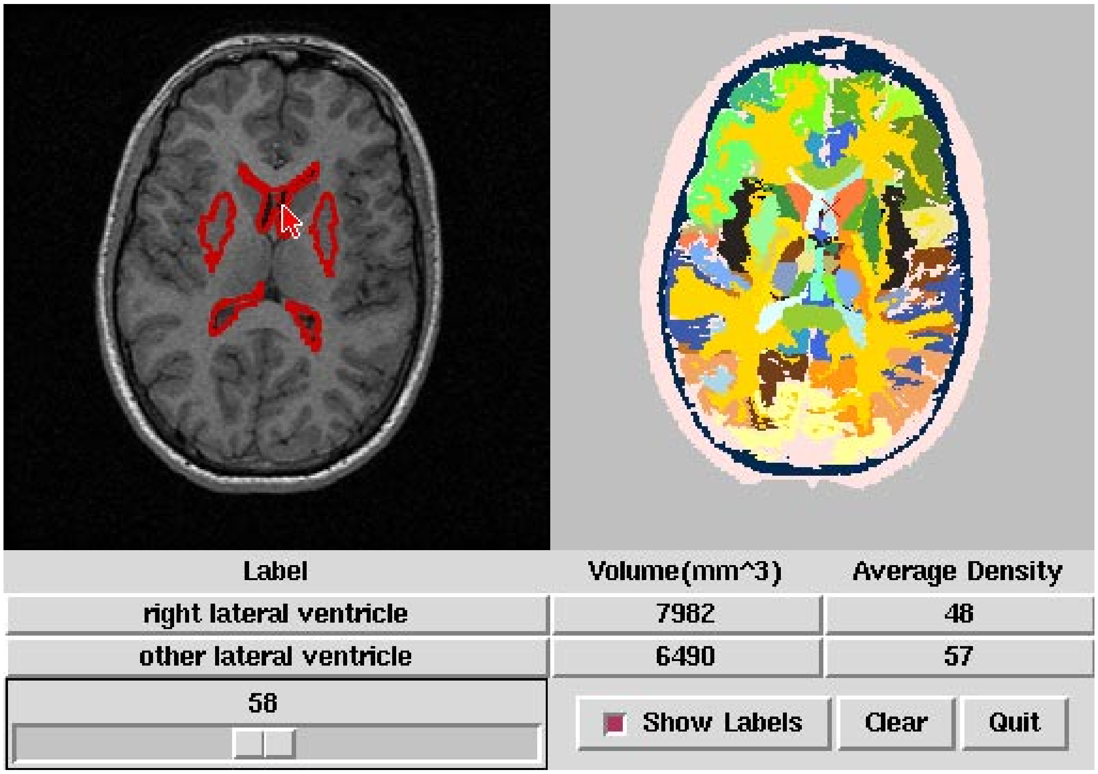

Label	Volume(mm^3)	Average Density
right lateral ventricle	7982	48
other lateral ventricle	6490	57

Figure 45: An interface showing one slice of a subject's data (left), and its automatically computed anatomical segmentation (right). When the user selects a structure, its segmentation outline is overlaid on the intensity data. The structure's name and several measurements are automatically displayed.

flicting results. Some researchers have found accumulating evidence for abnormalities in brain structure in adults with schizophrenia, such as lateral ventricular enlargement[24], [81], whereas some other researchers contend that no substantial differences are seen in their studies [3], [85]. Researchers at Western Psychiatric Institute and Clinic of the University of Pittsburgh Medical Center have been studying the possible structural and functional changes in schizophrenic patients [51]. However, so far their experiments employ human operators using semi-automatic tools to acquire anatomical segmentation.

Based on mutual interest in automatic quantitative analysis, segmentation of the right lateral ventricle in nine schizophrenics and twelve normal controls were conducted both by their operators and the automatic algorithm. Histograms of respective volumetric measurement given by both methods are compared in Figure 46. The top row shows results given by their operators, with the left histogram for normal controls, and the right one for schizophrenics; the bottom row gives corresponding results from the automatic analysis. Despite a systematic tendency of larger volumes estimated by the automatic algorithm, results from these two methods are highly correlated (Pearson's correlation coefficient = 0.95). Both seem to suggest lateral ventricular enlargement in schizophrenics.

7.3.2 Corpus Callosum and Female Alcoholics

The corpus callosum is the main fiber tract connecting the two brain hemispheres, consisting of approximately 200-350 million fibers in man. Surgical transection of the callosum in humans provides evidence that it functions to communicate perceptual, cognitive, mnemonic, learned and volitional information between the two brain hemispheres. Given the importance of sensory, motor, and cognitive callosal relay between brain hemispheres, this anatomic region has been a focus of studies examining structural and functional neuropathology [97]. In disease, effects on callosal structure are observable at both cellular and gross anatomic scales. However, intense controversy exists on the question of whether different callosal regions undergo selective changes in certain diseases.

Researchers at National Institute on Alcohol Abuse and Alcoholism, National Institute of Health, have been investigating gender differences in alcoholism-related brain atrophy. They examine mid-sagittal MRI scans from of similarly aged alcoholic women and men and normal controls, and measure the cross-sectional area of the corpus callosum to examine brain damage. Unlike other investigators who found that alcoholic women have equal [45], [64] or even smaller [56] ventricles than alcoholic men, they observed that the corpus callosum area was significantly smaller (which leads to larger ventricles) among alcoholic women than nonalcoholic women or alcoholic men. These results suggest an increased sen-

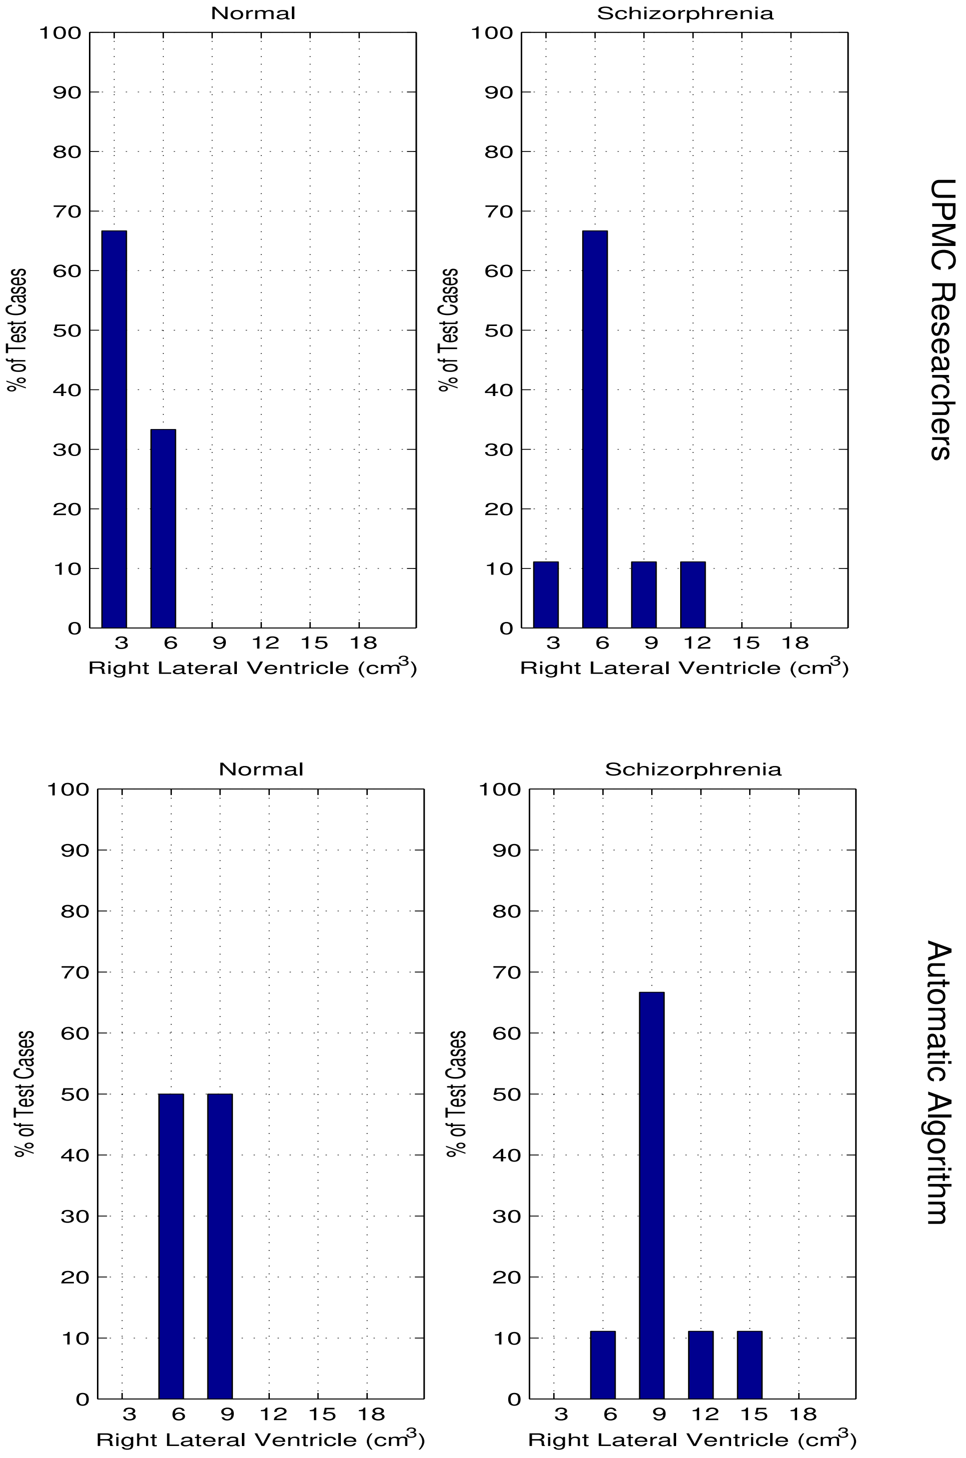

Figure 46: Histograms of right lateral ventricle volumes of normal controls and schizophrenics, resulted from manual estimation (top) and automatic analysis (bottom). High correlation is observed

sitivity to alcohol-induced brain damage among alcoholic women compared with alcoholic men [42].

In an attempt to validate their study as well as to test the automatic analysis, MRI data of 14 alcoholic women and 9 nonalcoholic women were examined by both their human operator and the automatic algorithm. Figure 47 compares histograms of the corpus callosum size in the mid-sagittal plane given by researchers at the NIH (top), and that given by automatic analysis (bottom). Histograms on the left are for nonalcoholic controls, whereas those on the right are for alcoholics. Despite an overall bias of larger areas computed by the algorithm, the correlation between the results is fairly high (Pearson's correlation coefficient = 0.89). The automatic analysis supports the observation from manual estimation.

The above studies demonstrated high correlation between results given by the automatic analysis and human operators. Limitations of the automated method include potential biases introduced by the classification in the reference atlas, as observed in both experiments (Figure 46, Figure 47). This is analogous to the differences found between two observers segmenting the same data. In theory, a representative average atlas combining multiple experts' opinions would eliminate this bias, and provide a better reference template.

7.4 Chapter summary

This chapter introduced an automatic system for conducting quantitative analysis of the anatomy, and demonstrated its feasibility in two collaborative clinical studies. Quantitative study of anatomical differences between populations is attracting more and more attention, such as the recent publication on the differences between Einstain's brain and that of ordinary people. With the assistance from automatic systems, it will be possible to study large data sets of various subpopulations to look for minute differences that are so far unknown.

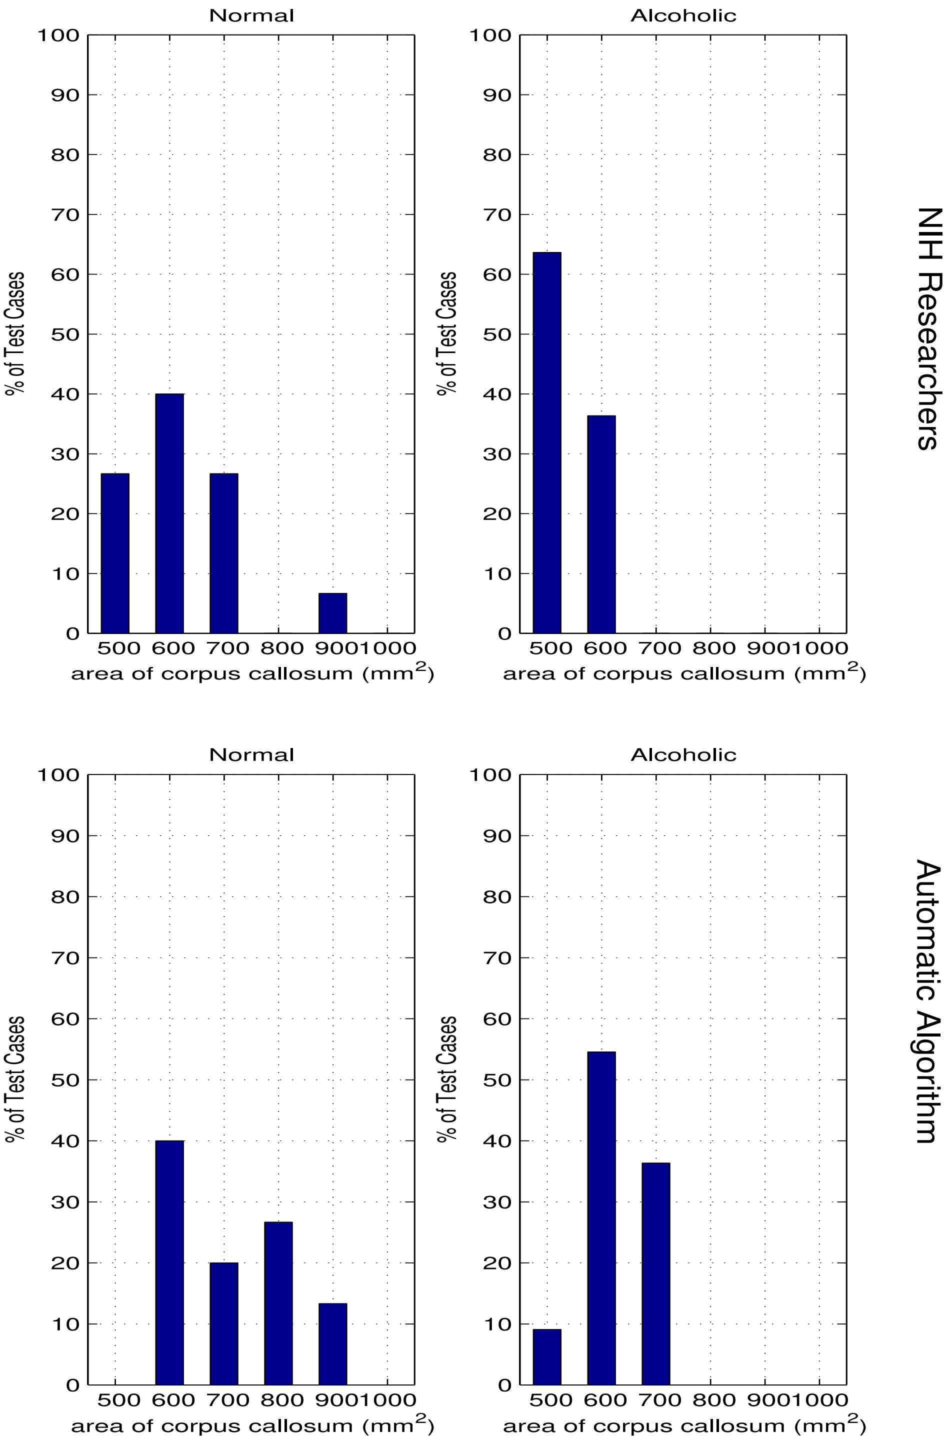

Figure 47: Histograms of corpus callosum size in the mid-sagittal plane of normal controls and alcoholics, resulted from manual estimation (top) and automatic analysis (bottom). Good correlation is observed.

CHAPTER 8 Anomaly Detection thrOugh REgistration

Knowledge of normal anatomical variations can not only enable quantitative study of anatomical differences between populations, as discussed in the previous chapter, it can also facilitate detection of abnormal variations that are indicative of possible neuropathology, which will be the focus of this chapter.

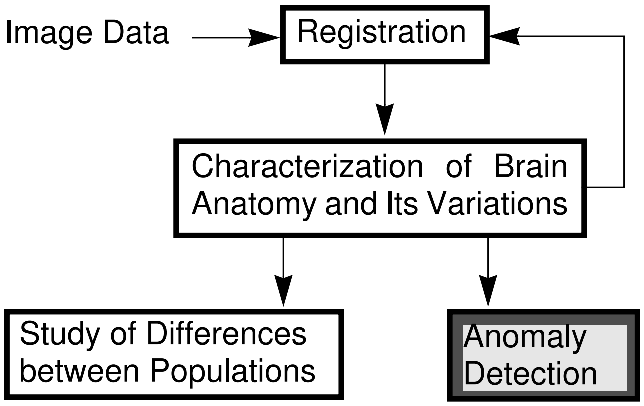

8.1 Related work

Many researchers have approached automatic pathology detection through segmentation, using methods such as neural networks [54], automatic segmentation using a 3-D geometric brain tissue probability model [48], and spatially weighted K-means histogram-based clustering [61].

A different research avenue is using registration to detect abnormalities. Warfield et al. developed an elastic matching algorithm that warps a reference data set containing information about the location of the grey matter into the approximate shape of the patient's brain. White matter and white matter lesions were then segmented without interference from grey matter, using a two class minimum distance classifier [105]. Experiments on segmenting multiple sclerosis in sixteen patients were promising. Kyriacou and Davatzikos studied finite element analysis to simulate soft tissue deformation in the brain caused by tumor growth, and tested on 2-D elastic registration of a standard brain atlas to simulated data and one tumor-bearing image [57]. Martin et al. proposed a method to separate out important from unimportant shape variation across a class of objects, so as to quantitatively describe pathological shape variations. By modeling the brain's elastic properties, they were able to compensate for some of the non-pathological modes of shape variation, and characterize modes of variation that are indicative of neuropathology [66]. Thompson et al. discussed tracking tumor growth rate by rigidly transforming subject data into Talairach stereotaxic space, and aligning a population-based probabilistic tissue map to the subject data. A Gaussian mixture distribution reflecting the intensities of specific tissue classes at each time-point in the scan series was determined, and tissue types were differentiated using a nearest neighbor algorithm [98].

8.2 Symmetry versus asymmetry

There are many types of neuropathology, of particular interest are pathologies that affect the structural symmetry. The organizational scheme of many animal species is based on bilateral symmetry: some organs appear in pairs in the body, “symmetrical” with respect to the mid-plane; other organs are placed near the mid-plane and are also approximately symmetrical. This symmetry is rather general for the human head, including the brain and its two hemispheres. One example is shown in the left of Figure 48. There is approximate symmetry about the central line.

On the other hand, symmetrical anatomical structures are sometimes also asymmetric, which means each of the two organs in a pair can present a specialization and therefore a slightly different morphology. For the human brain, some normal functional asymmetries are well known, which translate into morphological asymmetries. Furthermore, certain pathologies cause *mass effect*, which forces nearby structures to shift from their normal positions and incur abnormal asymmetry, as shown in the example on the right of Figure 48.

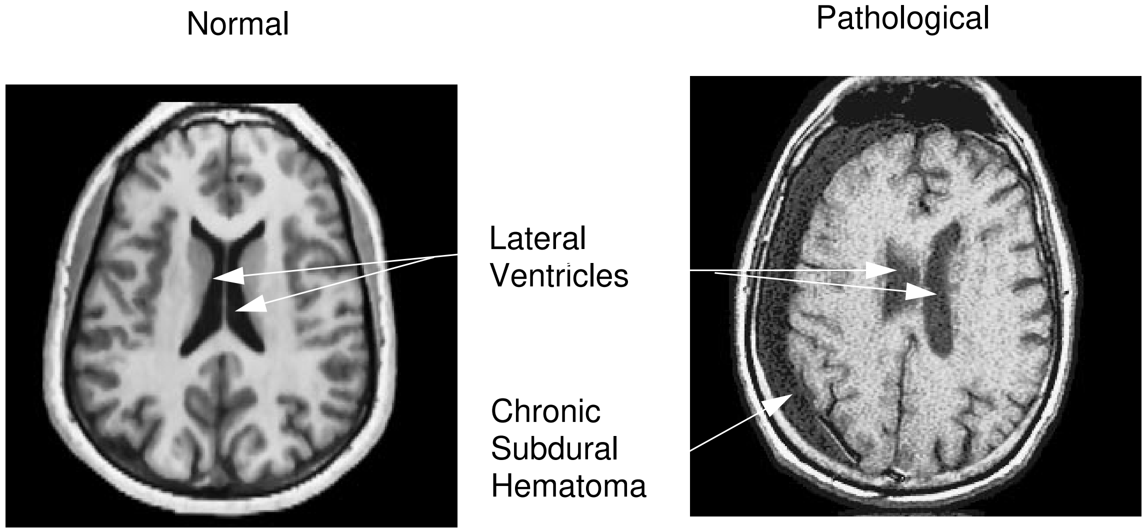

Figure 48: The corresponding cross-sections of a normal and a pathological brain. The symmetry in the normal brain is lost in the pathological one.

8.3 Anomaly detection using asymmetry

Using asymmetry to detect pathology is a familiar concept to specialists such as radiologists [22]. However, studies conducted by specialists suffer from sensitivity with respect to the operator. In addition, structures in the 3-D image are studied independently slice by slice, with the underlying assumption that the slices are exactly perpendicular to the symmetry plane. Moreover, rich 3-D information of anatomical structures are generally reduced to measurements such as lengths or widths.

Many researchers have worked on automating the analysis of symmetry and asymmetry. Liu et al. developed a technique for extracting the symmetry plane in axial images, and demonstrated its effectiveness in clinical data containing pathologies [59]. Marias et al. fit a 3-D mid-plane to a set of mid-lines detected using 2-D Snakes, and realigned the 3-D image with respect to the mid-plane. They extracted the cortical surface using a propagation of 2-D Snakes, and measured the perpendicular distance from the mid-plane to the cortical surface for both sides to quantify asymmetry. Information from several subjects was then fused by matching the cortical surfaces [65]. Thirion et al. adopted a 3-D volumetric approach along the same theme [93]. They detected the symmetry plane by a least squares fitting from features matched in both object sides, and defined the asymmetry map over the 3-D image, instead of only on the mid-plane. Moreover, they used 3-D volumetric non-rigid matching for inter-subject data fusion, which allows local analysis such as expansions or atrophies. The potential of their method was illustrated in preliminary results.

8.3.1 Anomaly detection via quantifying asymmetry

As noted above, quantification of normal asymmetry can be a powerful tool to detect abnormalities. Comparing the examples in Figure 48, brain symmetry and asymmetry are the most conspicuous in the size difference between the pair of lateral ventricles. Using the system developed in the previous chapter, the normal asymmetry of the lateral ventricles has been studied. The measurement for asymmetry is the absolute volume difference between the left and right lateral ventricles. The population studied consists of 128 non-pathological human brains. Figure 49 shows the histogram distribution of the asymmetry in this population. The horizontal axis is the absolute ventricular volume difference in cm^3, the vertical axis is the percentage of normal brains whose lateral ventricle size differences fall between certain range. Note that for 85% of the subjects, the asymmetry between the pair of lateral ventricles is within 1 cm^3; however, about 1% of the normal subjects have lateral ventricles that differ in size by more than 3 cm^3.

To examine the potential of detecting abnormal asymmetry, the asymmetry measurement of the pathological brain in Figure 48 is located in the distribution (Figure 49). Note that although the lateral ventricles in the pathological brain is strikingly asymmetric, which is consistent with the fact that it falls in the outmost bin of the histogram, it still has company from normal subjects. Therefore, such asymmetry is only indicative of a possible anomaly, but not conclusive. It needs to be combined with other diagnosis to give a definite answer.

8.3.2 Anomaly detection via mirror registration

A more direct application of registration to detect asymmetry is to match a 3-D image's *mirror image* to the original one. The mirror image is created by flipping the volume about its mid-sagittal plane. For a normal brain with approximate symmetry, little deformation is needed to align its mirror image to itself; for a brain with pathologies that cause asymmetry, a more significant deformation will be necessary. Examples in Figure 50 demonstrate this idea. Subjects studied include one normal control, one patient with chronic subdural hematoma (bleed), and one with a tumor. The top row shows axial cross-sections from each subject's volume. The bottom row shows the same cross-sections overlaid with the 2-D projection of the 3-D deformation flows necessary for matching each subject's mirror image to the original volume. For the normal brain, the magnitude of the deformation vectors is negligible because it is approximately symmetric. For the brain with a bleed, the deformation flows have a much larger magnitude, and follow a uniform direction. This direction suggests the source of the asymmetry, i.e. the bleed. Unlike most image volumes examined in this thesis, the 3-D image with a tumor is a CT scan. Its deformation flows also have a large magnitude, and demonstrate a swirl pattern. This is because the tumor is so close to the mid-sagittal plane that its "counterpart" in the mirror image partly matched with itself. Characterization of the deformation flows can help detecting, locating, and classifying anomalies that affect brain symmetry.

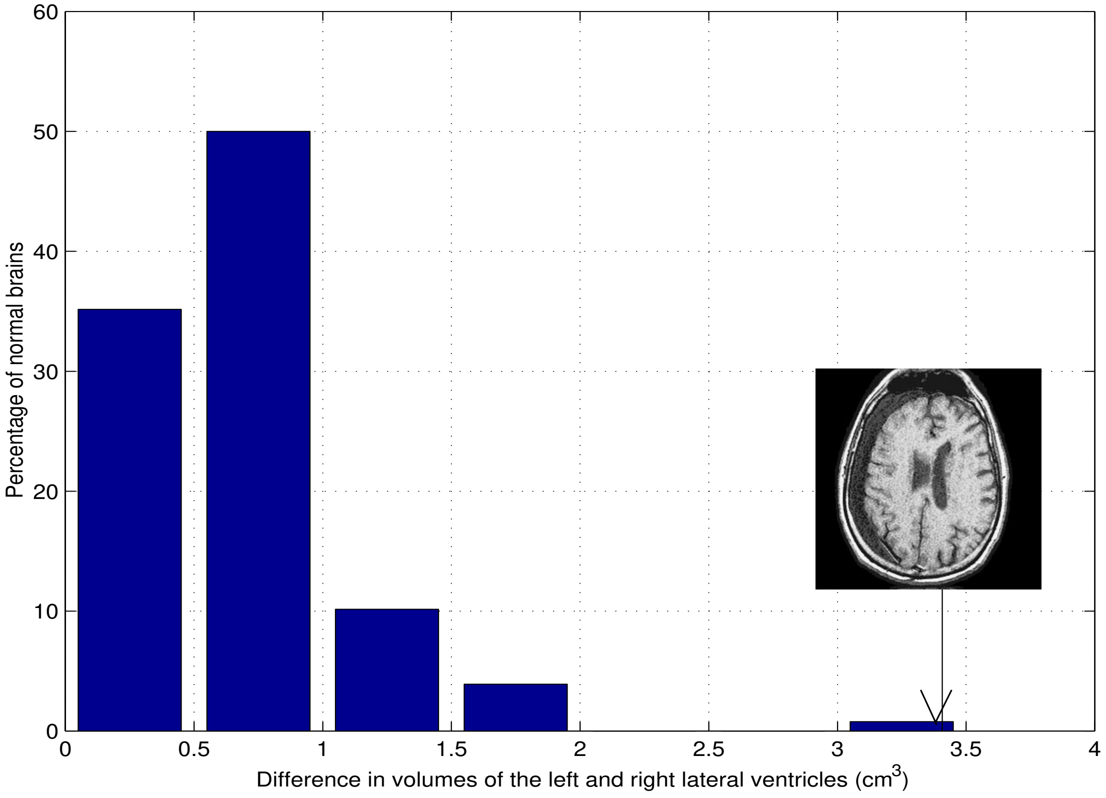

Figure 49: Histogram of the absolute volume difference between the pair of lateral ventricles of 128 normal subjects. A pathological brain falls in the outmost bin of the histogram.

8.4 Detect abnormal variations

Despite drastic individual variability in the anatomy, there still exists a distinction between the normal range of variability and pathology-afflicted alteration. Figure 51 com-

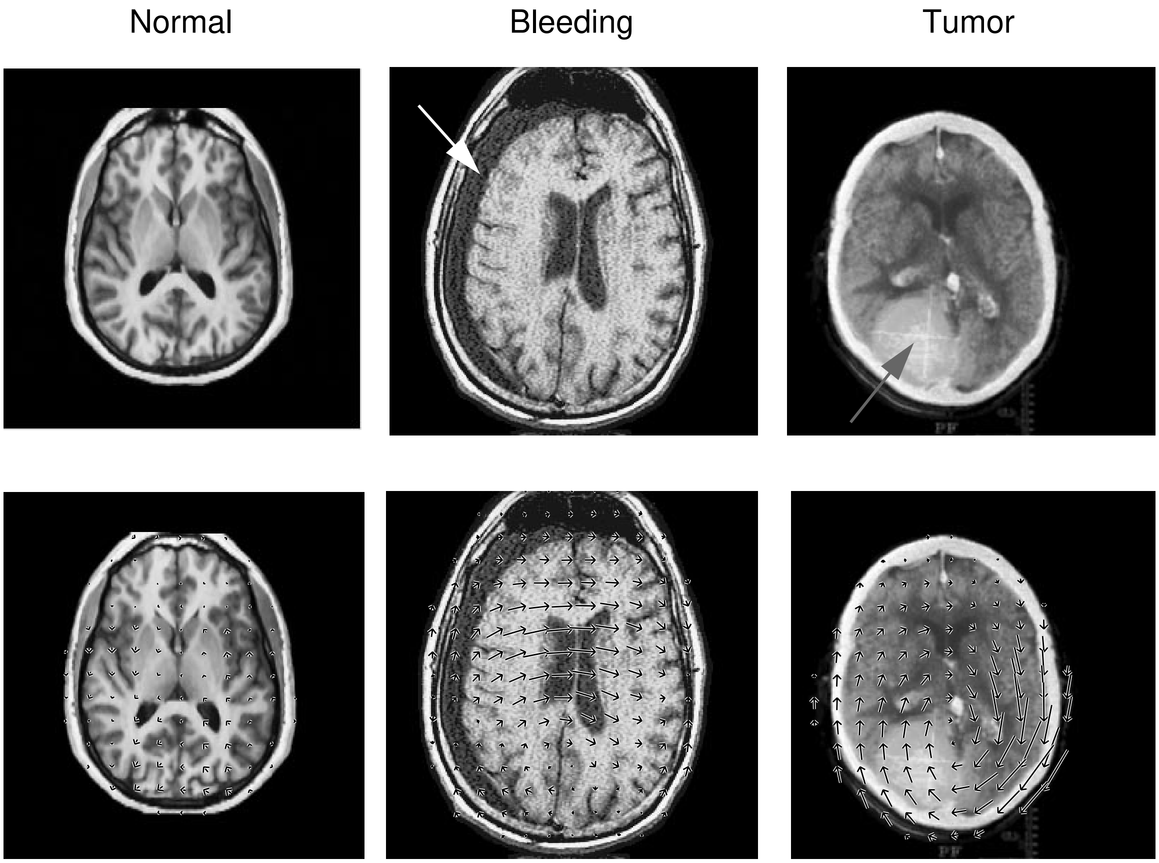

Figure 50: Several subjects (top) and the deformations incurred in matching their mirror images to the original volumes (bottom overlay). Large deformations present for pathological cases.

pares the label, *skull,* between the atlas and a brain with chronic subdural hematoma. The *Skull* in the pathological brain was segmented by registering with the atlas, and thus customizing the volumetric classification in the atlas. Note that the skull in the atlas has uniform thickness, with variations within a small range, whereas the thickness of *skull* in the pathological brain varies significantly. This is because the pathology, hematoma, is nonexistent in the atlas. During registration, it was mistakenly matched to the neighboring structure with similar intensity, the skull.

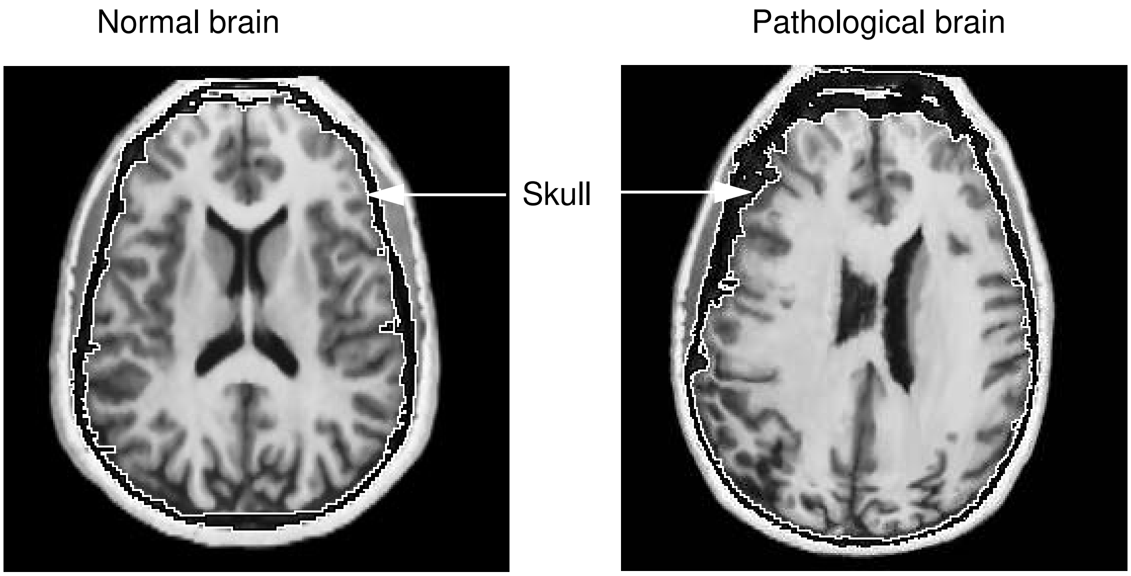

Figure 51: Comparing the label, *skull,* between a normal brain and a brain with chronic subdural hematoma.

A statistical study of the skull volume over a population of 48 normal subjects gives a histogram distribution shown in Figure 52. The horizontal axis is the volume of the skull in cm^3, the vertical axis is the percentage of subjects studied. Note the wide range of variations among normal subjects. The estimated skull volume in the pathological brain is 106.57 cm^3, which falls beyond the normal distribution. Combined with the analysis using symmetry, this can be used as another clue for anomaly detection.

8.5 Chapter summary

This chapter presented explorations in using registration for pathology detection. Combining analysis of normal asymmetry and normal range of variations in a population, pathologies that affect brain morphology can be indicated. However, white matter lesions such as multiple sclerosis do not cause significant distortion of a patient's anatomy, which makes it difficult for deformable registration to detect. Approaches combining statistical classification and deformable registration may promise to overcome the limitations, as proposed in [105].

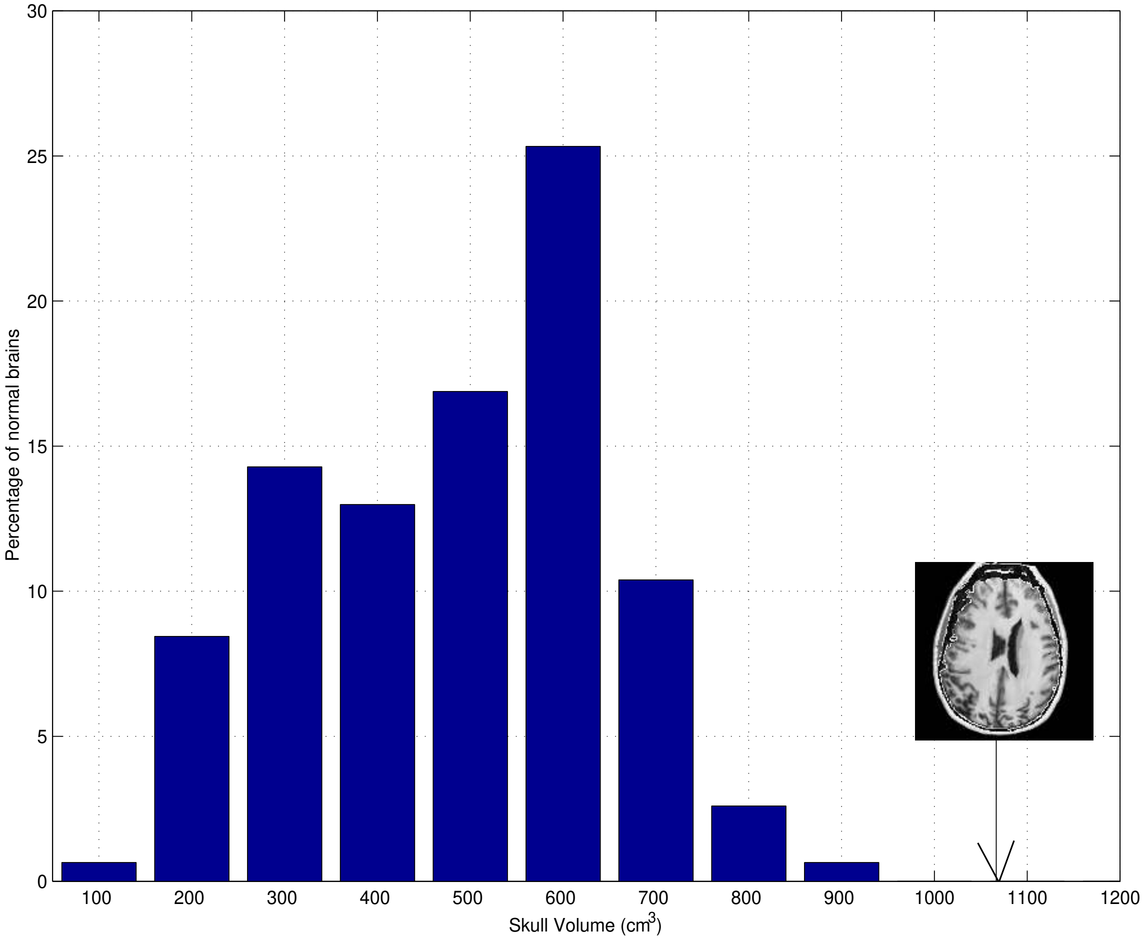

Figure 52: Histogram of the skull volume of 48 normal subjects. A pathological brain falls beyond the normal range of variations.

CHAPTER 9 Conclusion

This thesis has focused on extracting and modeling the knowledge of non-pathological anatomical variabilities between individuals, and the application of such knowledge to achieve more accurate registration than traditional image-based correspondence. Inherent anatomical variations are automatically extracted by deformably registering training data with an expert-segmented 3-D image, a digital brain atlas. Statistical properties of the density and geometric variations in brain anatomy are evaluated and encoded into the atlas to build a statistical atlas. These statistics can function as prior knowledge to guide the automatic registration process. Compared to an algorithm with no knowledge guidance, registration using the statistical atlas reduces the overall error on 40 test cases by 34%.

The designed algorithm is fully automatic, efficient, and has an accuracy comparable to human performance. When a pre-segmented image volume, referred to as an atlas, is registered to a new input image data, referred to as a subject, the anatomical segmentation in the atlas can be customized for the subject. This reduces the months-long manual segmentation process to minutes, thus enables applications such as quantitative study of anatomical differences between populations on statistically significant data archives, and objective cross-validation of different researchers' discoveries. In addition, characterization of non-pathological variations facilitates detection of abnormal variations indicative of possible pathologies.

9.1 Highlights of Contributions

Registration Algorithm

(i) 3-D hierarchical deformable registration algorithm (Chapter 2): an automatic, efficient, and reasonably accurate algorithm that can be employed to extract anatomical variabilities between individuals.

  - 3-D grid-based deformable registration (Section 2.2).
  - Hierarchical intensity equalization, specifically the structure-based intensity transformation (Section 2.3).

(ii) Registration using the statistical atlas to achieve higher accuracy (Chapter 4 and Chapter 5).

  - Probabilistic matching under Bayesian formulation (Chapter 4).
  - Weighted-window matching to incorporate neighborhood context (Chapter 5).

(iii) Efficient implementation techniques, enabling fast deformable registration of 3-D images (Chapter 6). To register 3-D images of 256x256x124 voxels on an SGI workstation with four 194 MHz processors, the 3-D hierarchical deformable registration algorithm takes 18 minutes, registration using voxel-based statistics takes 35 minutes, and registration with neighborhood context takes 57 minutes.

  - Random initialization (Section 6.3).
  - Adaptive multi-resolution processing (Section 6.2).

Extraction of Anatomical Variabilities

(i) Automatic extraction and characterization of non-pathological anatomical variations between individuals (Chapter 3).

  - Knowledge extraction from a considerable population of 105 normal brain MRI volumes. This is a larger population than what is employed in most related work.

(ii) A 3-D statistical atlas (Chapter 3): a global, volumetric representation of anatomical variabilities between individuals.

(iii) Bootstrap strategy to reduce imprecisions in knowledge models propagated from inaccurate automatic registration.

Quantitative Evaluation of Registration Performance

(i) Quantitative evaluation of performance on a sizable test set of 40 MRI volumes. This is considerably larger than what is used in most related work.

(ii) Developed and demonstrated the efficacy of two new performance metrics for quantitative evaluation of 3-D registration results (Section 2.4.2, Section 5.4).

(iii) Registration accuracy of 2.9% overall error compared to expert performance. This is arguably one of the most accurate deformable registration results published.
Figure 53 compares registration results given by each level of the 3-D hierarchical deformable registration, registration using voxel-based statistics, and registration using neighborhood context. The reduction in error ratio is significant along the hierarchical deformation (Figure 53 a), and considerable while more comprehensive knowledge model is employed (Figure 53 b).

Applications

(i) Feasibility study of applications in quantitative analysis of anatomical differences between populations, with results highly correlated to human operation (Chapter 7).

(ii) Experiments on using knowledge of non-pathological variabilities to detect abnormal variations that are indicative of pathologies (Chapter 8).

(iii) Mirror-registration for asymmetry detection (Chapter 8).

9.2 Future Directions

This thesis has investigated the problems of extracting and modeling anatomical variabilities between individuals, the application of such knowledge to achieve accurate registration, and the use of registration to quantitatively study anatomical differences between populations. The current research is an early investigation in the vast and largely unknown field of computational anatomy, much remains to be explored and refined.

9.2.1 Methodology

One important issue is to refine the statistical atlas by building population-specific atlases. To produce population specific representations of anatomy, data from large popu-

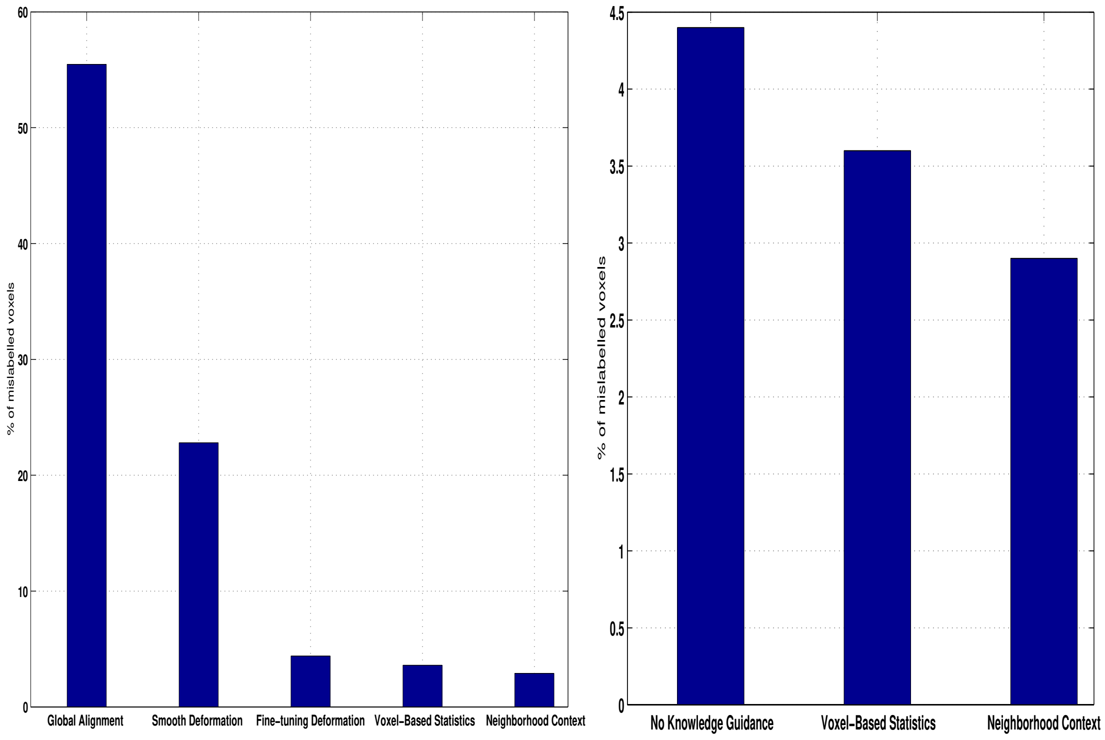

Figure 53: Comparison of registration performance using the algorithm with no knowledge guidance, using the statistical atlas , and using neighborhood context.

lations need to be stratified into subpopulations according to age, disease state, gender and other demographic criteria. Brain atlases constructed from subpopulations encode information on population variability, and therefore can facilitate study of identifying group-specific patterns of anatomic or functional alterations. Researchers at the Laboratory of Neuro Imaging in the University of California at Los Angeles, are building population-based digital brain atlases for Alzheimer's Disease and schizophrenia. Atlases for different age groups will also help studying anatomical changes in development and aging.

Under the current implementation, the training set is mapped to the atlas' reference frame using an automatic registration algorithm with inaccuracies; intensity differences caused by image acquisition are separated from inherent tissue density variations using hierarchical intensity equalization, which also contains imprecisions. Errors in registration and intensity normalization affect the accuracy of models of anatomical variabilities, and thus the rigorousness of using such models as prior knowledge to further guide registration. One practical approach to alleviate this problem is to build an initial statistical atlas from a small but accurately registered training set, then bootstrap it into a model based on a considerable training set. Modeling the effects of different parameter settings on an MRI machine will help more accurately remove extrinsic intensity discrepancies (Section 2.6.1).

Another issue warrants further study is the choice of knowledge models. The current models prove to be effective and simple, but higher-level, more compact models will be necessary to express 3-D shape variations of individual structures. Many researchers are exploring shape models based on modal analysis, finite element analysis, and morphometrics. The preliminary investigation on using PCA (Appendix B) is also one avenue to pursue.

As observed in registering more than two hundred image volumes, some structures are easily registered, whereas some others are more difficult to be matched. Therefore, a coarse-to-fine approach that first registers the *easy* structures and then use them to bootstrap the match of the more difficult ones will be helpful. Similar idea was used in the hierarchical intensity equalization scheme (Section 2.3). Kapur et al. employed such strategy in MRI segmentation using a relative geometric prior (prior distribution on the geometry of difficult structures, given the geometry of easy structures) [49].

It is noteworthy that the underlining assumption to the construction of the statistical atlas is that an individual anatomical image can model the morphometric variations observed in other individual images. Under this assumption, the mappings between different images are homologous, which means they only map corresponding structures and fail to capture the changes due to topological discrepancies. Thus a morphometric analysis only partially describes the shape differences. While the assumption that there exists a 1-to-1 correspondence between two brains is never strictly true, such as the recently reported difference in sulcal patterns between Einstain and average people, the validity of this assumption depends on the spatial scale of comparison. Topology of the cortex is not consistent for different subjects, but this is a completely different problem, complex enough to warrant a separate research project. Approaches proposed by researchers include extracting the cortical surface and mapping it onto a simple parametric space with the same topology, e.g., a sphere.

9.2.2 Applications

Quantitative analysis of the anatomy and pathology detection are the few applications investigated in this thesis, additional exciting areas remain to be explored.

9.2.2.1 Database retrieval

The introduction of large scale PACS (Picture Archival and Communication Systems) has resulted in the creation of large digital image databases. A typical radiology department generates between 100, 000 and 10, 000, 000 images per year. A filmless imaging department such as the Baltimore VA Medical Center generates approximately 1.5 terabytes of image data annually. This has substantially increased the complexity and number of images presented to radiologists and other physicians. Efficient retrieval in medical databases will benefit experts, as well as non-experts, in applications such as computer-aided diagnosis, time evolution analysis and forecasting.

Albeit a relatively new domain, medical image database retrieval has attracted much research attention. Korn et al. examined the problem of finding similar tumor shapes so as to search for nearest neighbors in large collections of tumor-like shapes [55]. Berrut et al. developed a specialized user interface which allows a radiologist to semi-automatically index images represented by a general model [4]. Liu and Dellaert employed image classification to retrieve similar cases in a well-defined database, where the distance metric defining a classifier functioned as the similarity metric [60]. Guimond and Subsol approached the retrieval of *volume of interest* (VOI) using intensity-based deformable registration between the VOI given by a user and images in the database [37]. Morphological similarity is measured using correlation. Researchers at Los Alamos National Lab have developed CANDID algorithm to retrieve similar CT slices with lung diseases using texture analysis [50]. Quantification of registration results can also be used to compare image data and to classify similar cases.

9.2.2.2 Data-mining

Accurate and efficient registration algorithms can also facilitate data mining that detect correlations among shapes, diagnosis, symptoms, and demographic data. This will help form and test hypothesis about the development and treatment of diseases, study potential relationships between subtle neuroanatomical changes and symptom severity, prognosis, and the capacity to respond to treatment.

9.2.2.3 Model anatomical abnormalities

Various neurological disorders affect the gross anatomical shape of different brain structures. These changes have been studied for several decades, using both postmortem and invasive in vivo methods. Recent advances in the contrast and resolution of MRI scanners make it possible to study these shape effects in vivo and noninvasively, with the potential of better diagnosis and treatment.

To properly study pathological deformations, the non-pathological inter-patient variations, such as that modeled in the statistical atlas in this thesis, need to be normalized. The goal is to detect and quantify distributed patterns of deviation from normal anatomy [66]. For the study of asymmetry, statistical analysis of asymmetry fields will help answer questions such as the significantly asymmetrical regions in a population, the difference in asymmetry between populations, and the detection of abnormal asymmetry [93]. Potential approach include using vector operators such as the norm of the vector field which characterizes its magnitude, the divergence of the vector field which captures its radial aspect. High divergence and high magnitude are characteristic of atrophies or expansions due to lesions or tumor growths,

9.2.2.4 Functional atlas

The current atlas is a structural atlas with anatomical classifications, however, there could also be atlases on the functional organization of the brain, or on the response of different brain regions to therapeutic procedures such as radiotherapy. Atlases from different modalities will facilitate integration of functional and anatomical data across individuals.

9.2.2.5 From volume to surface

Morphometric variance of the human brain is qualitatively observable in surface features of the cortex. Even sulci and gyri that consistently appear in all normal subjects exhibit pronounced variability in size and configuration. Sulci separate functionally distinct regions of the brain, and provide a natural topographic partition of its anatomy. Most of the junctional zones between adjacent microanatomic fields run along the beds of major or minor cortical sulci. Striking intersubject variations in sulcal geometry have been reported in primary motor, auditory cortex [78], visual cortex [86], as well as the recent compelling report on the partial absence of one of Einstain's sulci. Although the intrinsic variability in sulcal configuration across individuals is well known and much research has been conducted in the quantitative study [99], the ranges of the normal variations have not yet been determined. Statistical analysis of sulcal geometry will facilitate multi-subject atlasing and neurosurgical studies, and help studying neuro-degenerative diseases that have dramatic decrease in the fractal dimension of the cerebral cortex, such as epilepsy [16].

9.2.2.6 Beyond human brain MRI

This thesis emphasizes results on T1-weighted MRI data of human brains, however, the algorithm is applicable to other imaging modalities, as well as medical data of other body parts. Applications of the statistical atlas include detection, quantification, and mapping of local shape changes in 3-D medical images in disease, and during normal or abnormal growth and development. For non-medical applications, the algorithm could potentially be used for non-destructive diagnosis of malfunctioning mechanical objects such as engines, and quality inspection of manufactured objects to detect interior defects.

Although the current algorithm is developed for volumetric data, the general approach is suitable for a whole cadre of shape categorization and classification problems in which knowledge of shape variation can be gleaned over a particular object class. For instance, the formulation could have fruitful application in tracking and recognition of gesture, facial expression, and gait. Furthermore, such models can provide a parameterized estimate of principal deformations due to a specific process: manipulation, locomotion, growth, manufacture, disease, wind, heat, etc.

APPENDIX A Levenberg-Maquardt Algorithm

Levenberg-Marquardt algorithm is a non-linear optimization algorithm based on gradient descent. Suppose a vector function $\mathcal{R}(p)$ expresses the residual difference at each voxel, where p is the vector of transformation parameters. The derivative of $\mathcal{R}(p)$ with respect to p shows how each component of $\mathcal{R}(p)$ will change given a change in the transformation parameters Δp.

$$\Delta\mathcal{R}(p) = \frac{\partial\mathcal{R}}{\partial p}\Delta p \tag{19}$$

Ideally, the goal would be to find the transformation parameters that make $\mathcal{R}(p)$ equal to the zero vector:

$$\mathcal{R}(p_{new}) = \mathcal{R}(p_{old}) + \Delta\mathcal{R}(p) = 0 \tag{20}$$

Because of the existence of noise and the innate discrepancies in different volumes, $\mathcal{R}(p)$ can not be reduced to the zero vector in practice. Levenberg-Marquardt algorithm attempts to minimize $\mathcal{R}(p)$ by adjusting the transformation parameter vector p. From (19) and (20) we have

$$-\mathcal{R}(p_{old}) = \frac{\partial\mathcal{R}}{\partial p}\Delta p \tag{21}$$

For a given p, there are far more voxels than the number of transformation parameters, so $\frac{\partial\mathcal{R}}{\partial p}$ is not a square matrix. To solve this over-constrained system for p, Levenberg-Marquardt algorithm employs the pseudo inverse method:

$$\Delta p = \left[\left(\frac{\partial\mathcal{R}}{\partial p}\right)^{\mathcal{T}}\frac{\partial\mathcal{R}}{\partial p}\right]^{-1}\left[\left(\frac{\partial\mathcal{R}}{\partial p}\right)^{\mathcal{T}}(-\mathcal{R}(p_{old}))\right] \tag{22}$$

$$\left(\frac{\partial\mathcal{R}}{\partial p}\right)^{\mathcal{T}}\frac{\partial\mathcal{R}}{\partial p} \tag{23}$$

If matrix (23) is not of full rank, the pseudo inverse will not exist. Levenberg-Marquardt algorithm adds a stabilizing term λ to the diagonal elements of the matrix

$$\Delta p = \left[\left(\frac{\partial\mathcal{R}}{\partial p}\right)^{\mathcal{T}}\frac{\partial\mathcal{R}}{\partial p} + \lambda I\right]^{-1}\left[\left(\frac{\partial\mathcal{R}}{\partial p}\right)^{\mathcal{T}}(-\mathcal{R}(p_{old}))\right] \tag{24}$$

In each iteration of Levenberg-Marquardt optimization, Δp is computed from the current $\mathcal{R}(p_{old})$ and p_{old}. The summation of Δp and p_{old} gives p_{new}. The $\mathcal{R}(p_{new})$ is in turn used to compute the new Δp. The iteration goes on until Δp is smaller than a user defined threshold, at which point the transformation parameters that minimize the residual difference are considered to be recovered [76].

APPENDIX B Principal Component Analysis

Principal component analysis (PCA) is an elegant method for handling high dimensional data. In the context of studying anatomical variations, one potential approach is using PCA to extract eigen-variations of a population. Major modes of variations may be captured by principal eigen-variations, as well as probability distributions of the respective coefficients for reconstructing data of the population.

Of course, it is questionable whether a small number of principal eigen-variations will be sufficient in representing the immensely intricate anatomical variations of human brain, and the high dimensionality of the data further complicates the problem. This thesis conducts preliminary study on using PCA to improve registration performance, as well as using PCA for classification.

1. Registration

The registration problem is still formulated as finding the deformation that maximizes the posterior probability:

$$\mathcal{P}\langle \mathcal{D}eformation | \mathcal{D}ata \rangle = \frac{\mathcal{P}\langle \mathcal{D}ata | \mathcal{D}eformation \rangle \mathcal{P}(\mathcal{D}eformation)}{\mathcal{P}(\mathcal{D}ata)}$$

$\mathcal{P}\langle \mathcal{D}ata | \mathcal{D}eformation \rangle$ is defined the same as that in Section 3.3.1., the difference is in the representation of $\mathcal{P}(\mathcal{D}eformation)$. Instead of being the distribution of per voxel deformation flow, it is defined as the distribution of coefficients a_i of the respective eigen-deformations $\vec{e}_i$:

$$\mathcal{P}(\mathcal{D}eformation) = \mathcal{P}(a_1, \ldots, a_i, \ldots)$$

Because eigen-deformations are orthonormal, their respective coefficients are statistically independent. The above expression simplifies into

$$\mathcal{P}(\mathcal{D}eformation) = \prod_i \mathcal{P}(a_i)$$

The number of principal eigen-deformations is decided using the plot of the sum of eigenvalues with respect to the number of eigenvalues. Figure B-1 shows an example where the sum of eigenvalues plateaus after 20 eigenvalues. Therefore, the first 20 eigen-deformations are considered sufficient for representing major variations in the population.

2. Classification

It is observed that there may exist populational differences in the shape of certain anatomical structures, e.g. many researchers find the posterior section of the corpus callosum more bulbous in female subjects than their male counterparts [21]. Using principal eigenvectors to represent each data category, a new input can be classified based on how well it can be reconstructed by different categories. Mahalanobis distance is used as the distance

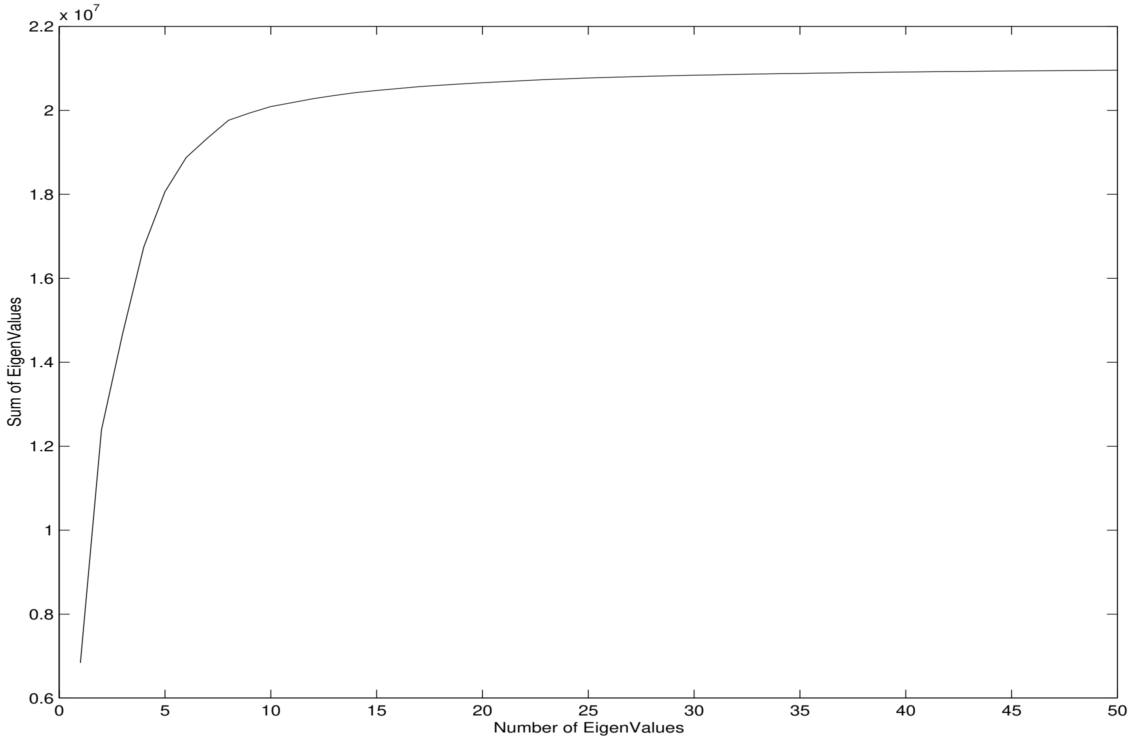

Figure B-1. The sum of eigenvalues plateaus after 20 eigenvalues. Therefore, the first 20 eigenvectors are sufficient for representing major variations.

measurement. Suppose the category under examination has a mean deformation $\vec{\mu}$, and a covariance matrix Σ. For a particular deformation $\vec{\upsilon}$, its Mahalanobis distance from this category is:

$$\mathcal{D}_{\mathcal{M}} = (\vec{\upsilon} - \vec{\mu})^{\mathcal{T}} \Sigma^{-1} (\vec{\upsilon} - \vec{\mu})$$

Using $\vec{e_1}, \vec{e_2}, \ldots \vec{e_k}$ to represent the k principal eigenvectors given by PCA, and $a_1, a_2, \ldots, a_k$ to represent the corresponding coefficients for reconstructing $\vec{\upsilon}$, $\mathcal{D}_{\mathcal{M}}$ becomes

$$\mathcal{D}_{\mathcal{M}} = \left(\sum_i a_i \vec{e_i}^{\mathcal{T}}\right) \Sigma^{-1} \left(\sum_j a_j \vec{e_j}\right)$$

Expand the above equation gives

$$\mathcal{D}_{\mathcal{M}} = \sum_{i,j} a_i a_j \vec{e_i}^{\mathcal{T}} \Sigma^{-1} \vec{e_j}$$

Because $\vec{e_1}, \vec{e_2}, \ldots \vec{e_k}$ are orthonormal vectors, $\mathcal{D}_{\mathcal{M}}$ simply becomes

$$\mathcal{D}_{\mathcal{M}} = \frac{\sum_i a_i^2}{\lambda}$$

where λ is the largest eigenvalue. The intuition of this distance measure is illustrated in Figure B-2. c_m and c_f are centers of the two categories, $\overrightarrow{e_{f1}}$ and $\overrightarrow{e_{f2}}$, $\overrightarrow{e_{m1}}$ and $\overrightarrow{e_{m2}}$ are their respective eigenvectors. Point p is classified as category f because of shorter Mahalanobis distance.

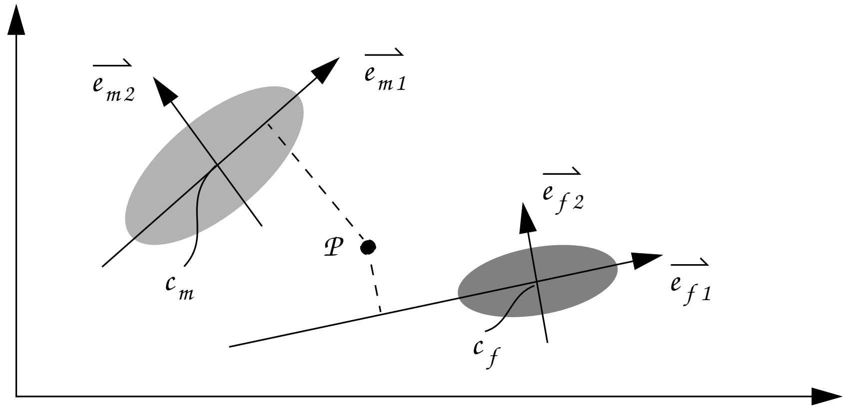

Figure B-2. The intuition of using Mahalanobis distance in classification.

Interesting directions for further investigation are conducting PCA on deformation flows of control points, initializing the deformation by the mean deformation and searching for coefficients of each eigen-deformation, de-coupling shape information from location and scale, etc.

APPENDIX C Segmentation Comparison

This appendix compares experts' segmentation of the corpus callosum in the mid-sagittal plane of 40 test cases to that from the automatic registration algorithm. All results are displayed as binary images, with the corpus callosum shown at intensity 140, and the rest of the mid-sagittal plane shown as zero intensity background. Error ratio of each case is displayed below its respective automatic segmentation result.

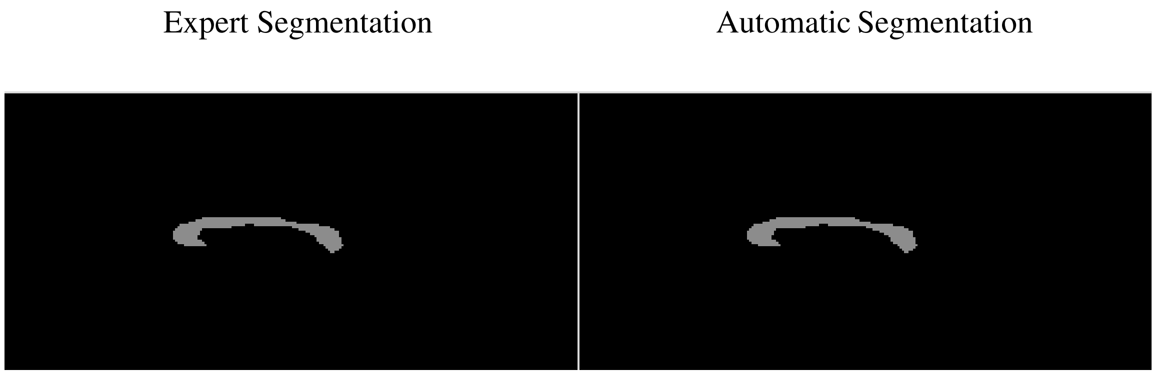

Error Ratio = 0.6%

Expert Segmentation

Automatic Segmentation

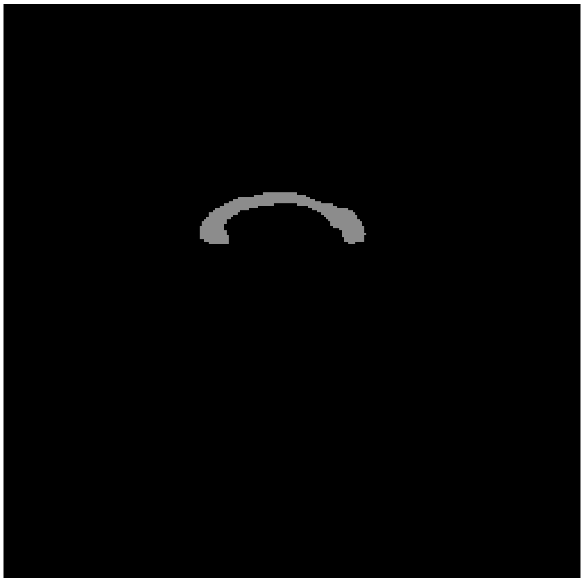

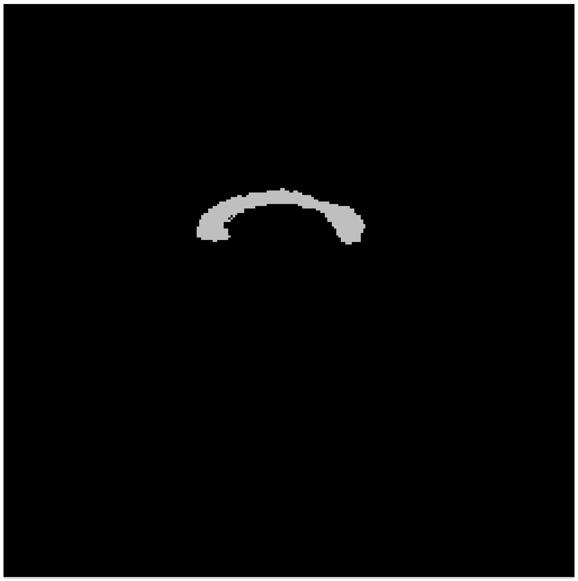

Error Ratio = 3.03%

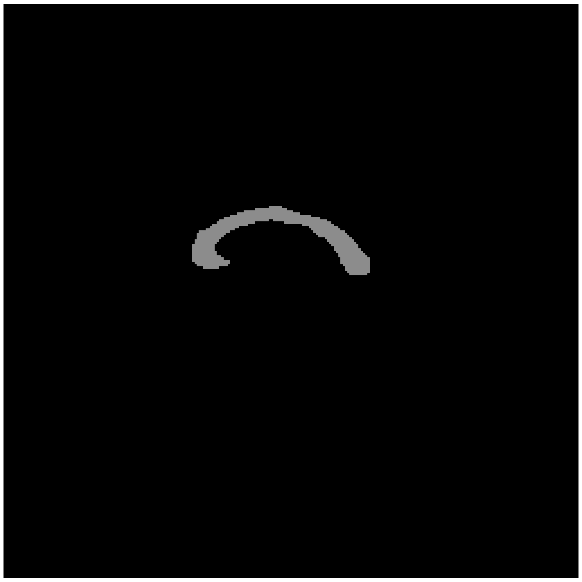

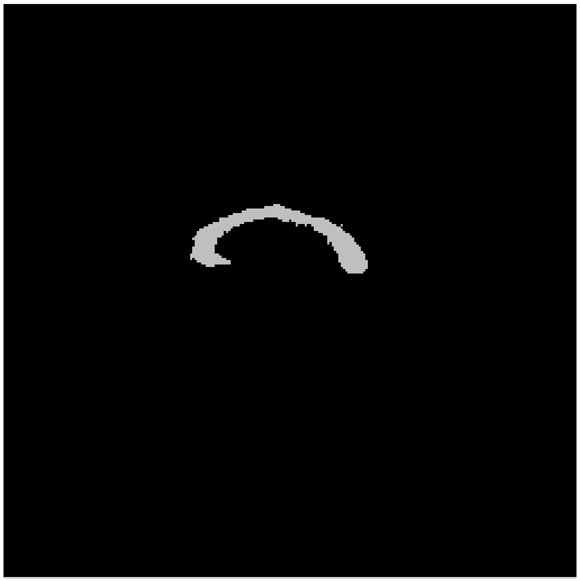

Error Ratio = 0.39%

Expert Segmentation

Automatic Segmentation

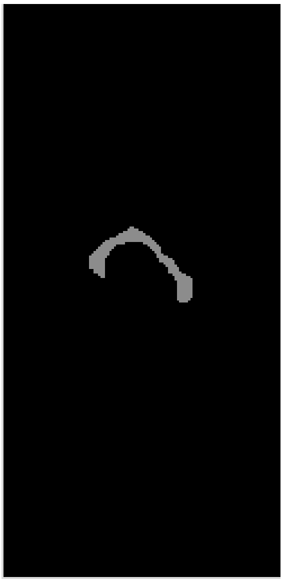

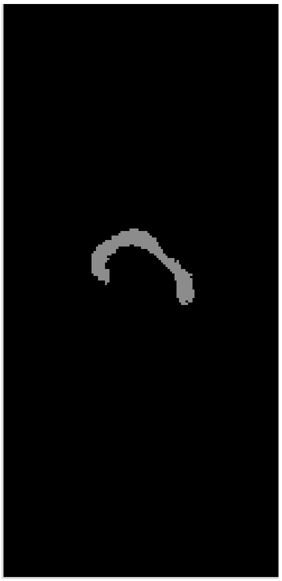

Error Ratio = 4.92%

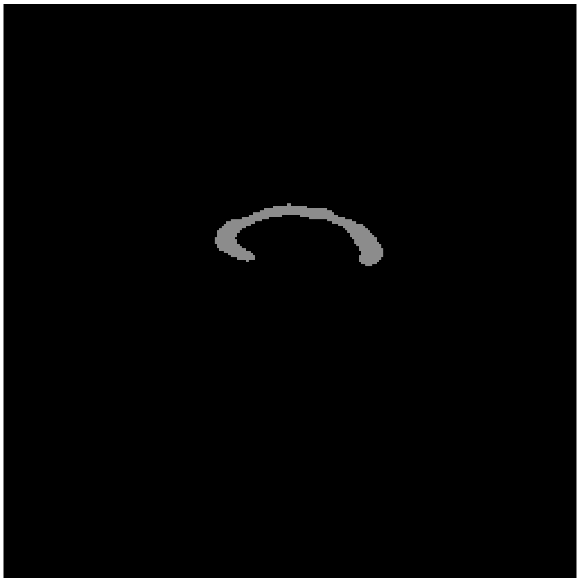

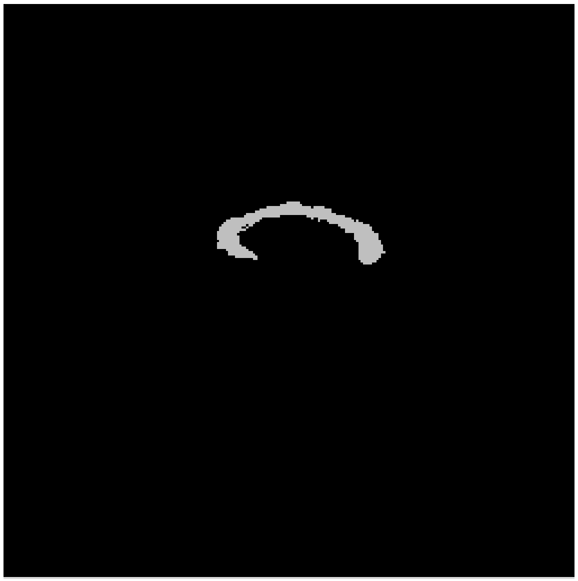

Error Ratio = 0.35%

Expert Segmentation Automatic Segmentation

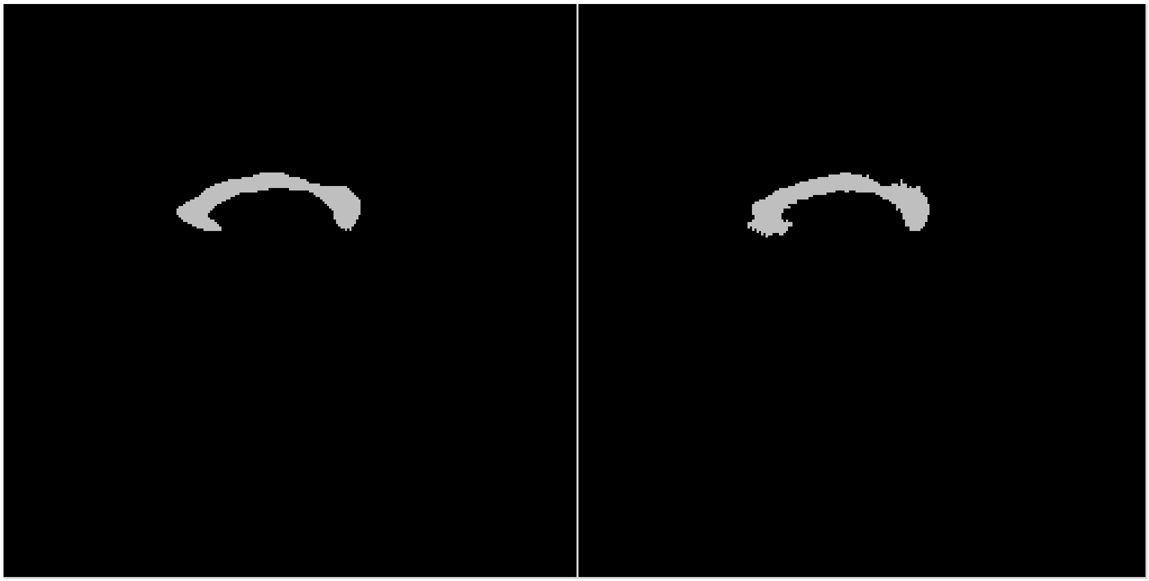

Error Ratio = 5.0%

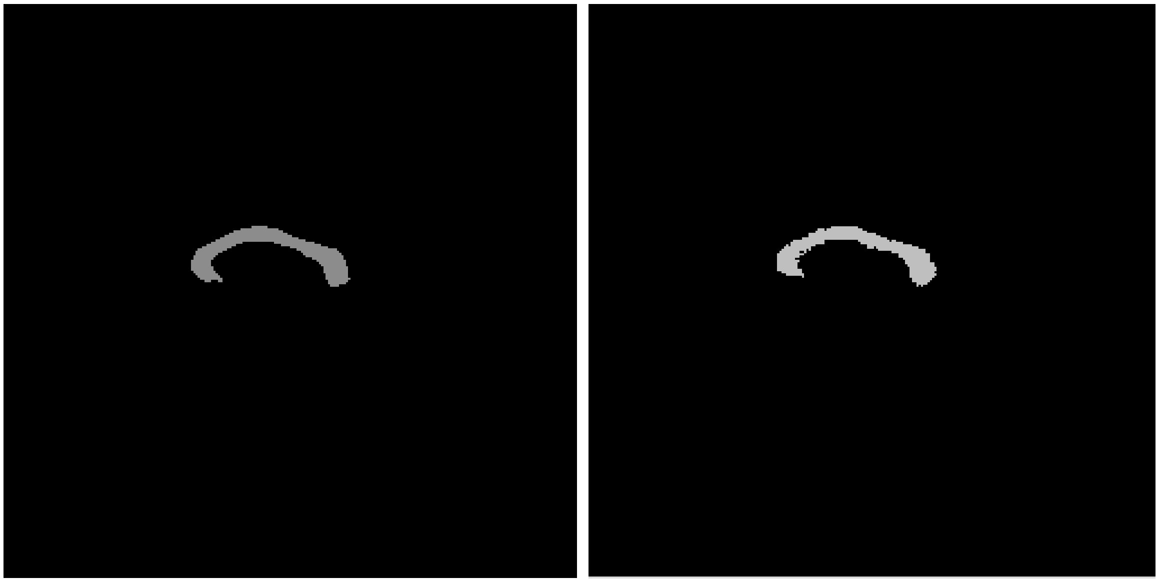

Error Ratio = 1.13%

Expert Segmentation Automatic Segmentation

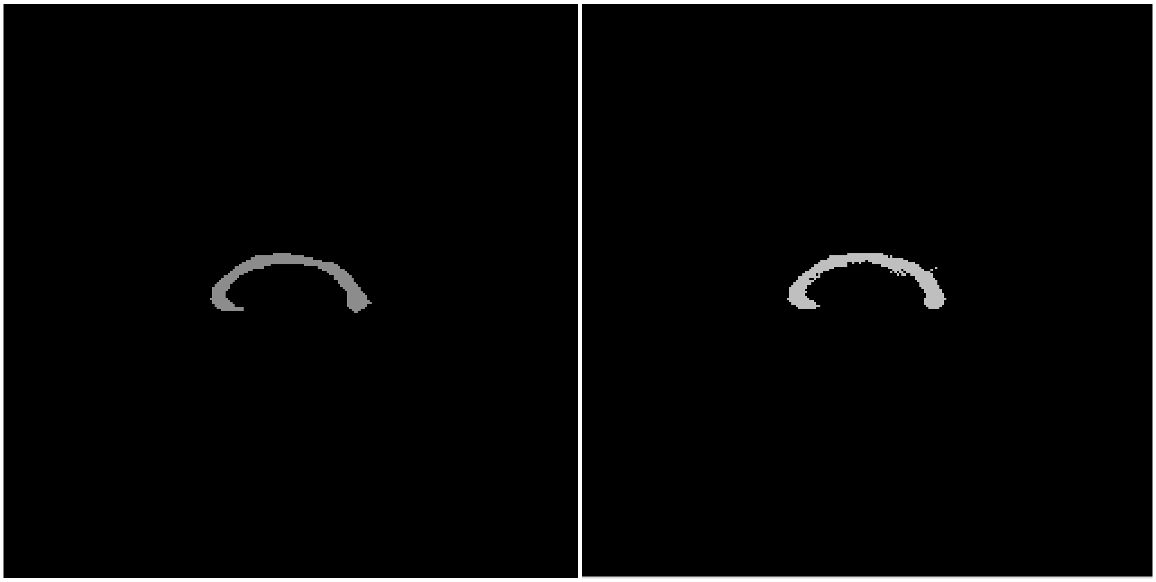

Error Ratio = 1.42%

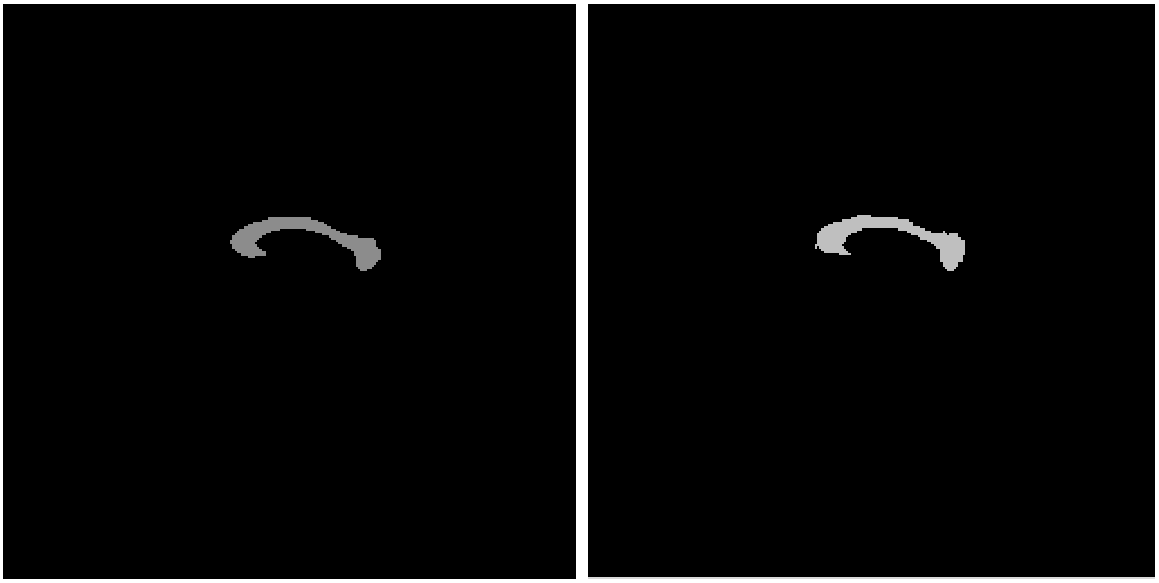

Error Ratio = 1.50%

Expert Segmentation

Automatic Segmentation

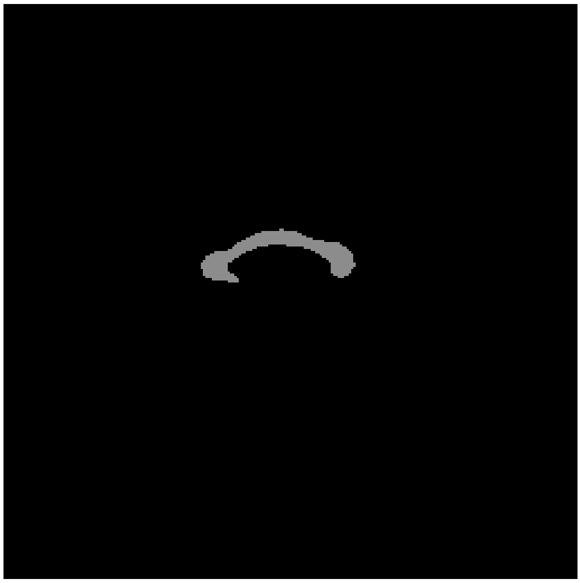

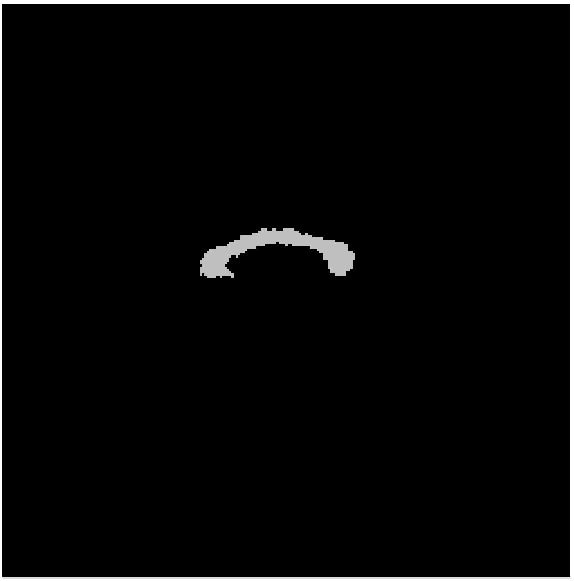

Error Ratio = 1.04%

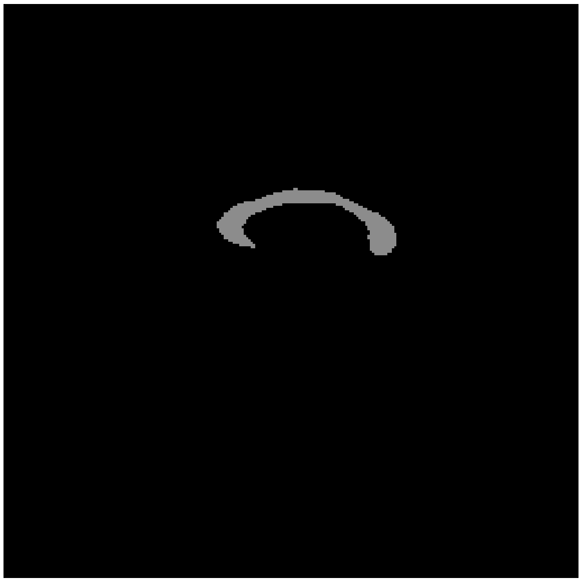

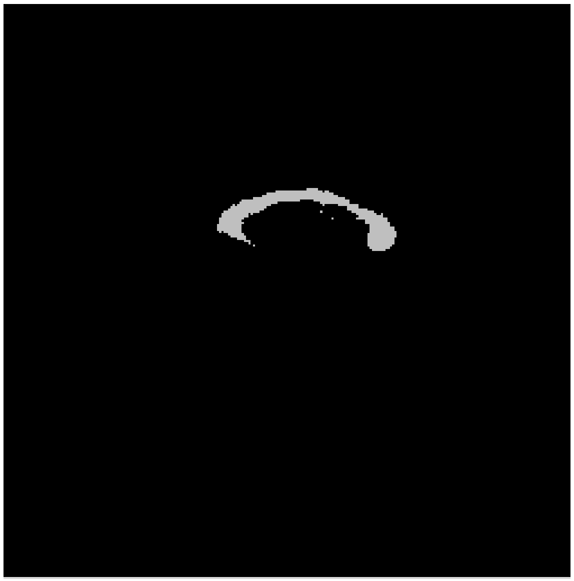

Error Ratio = 2.87%

Expert Segmentation Automatic Segmentation

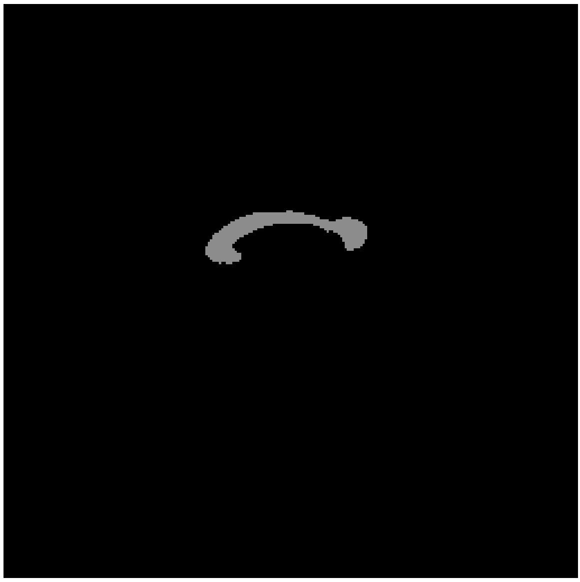

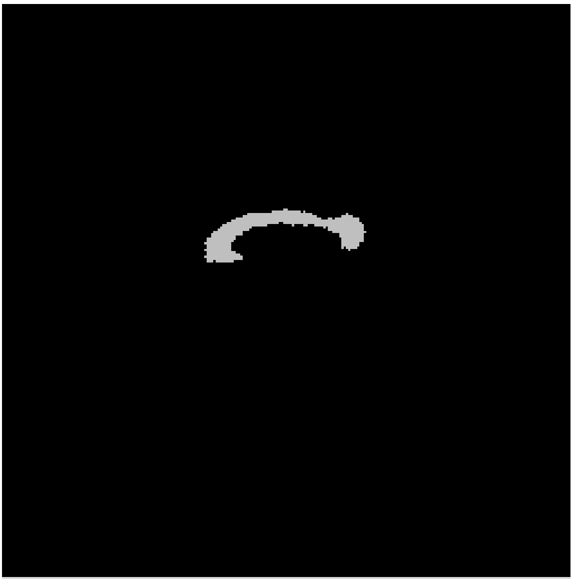

Error Ratio = 2.64%

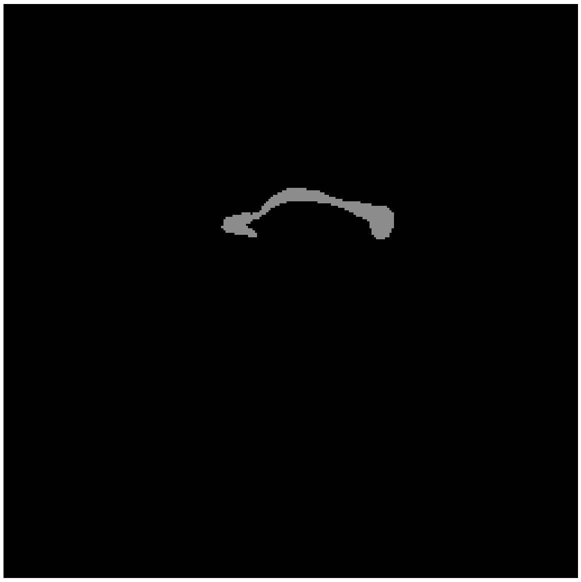

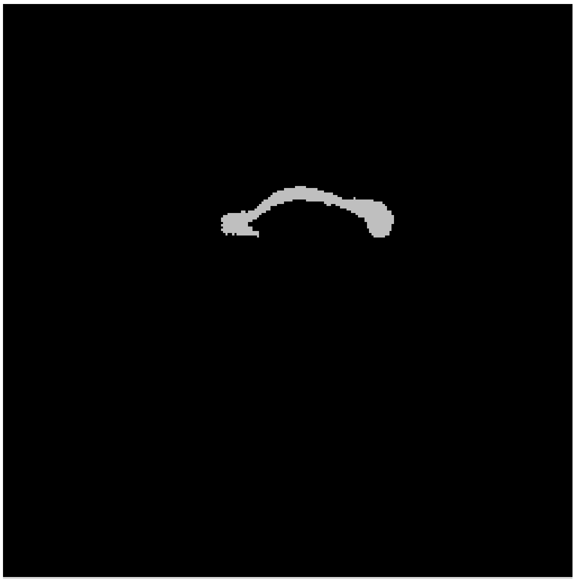

Error Ratio = 1.16%

Expert Segmentation Automatic Segmentation

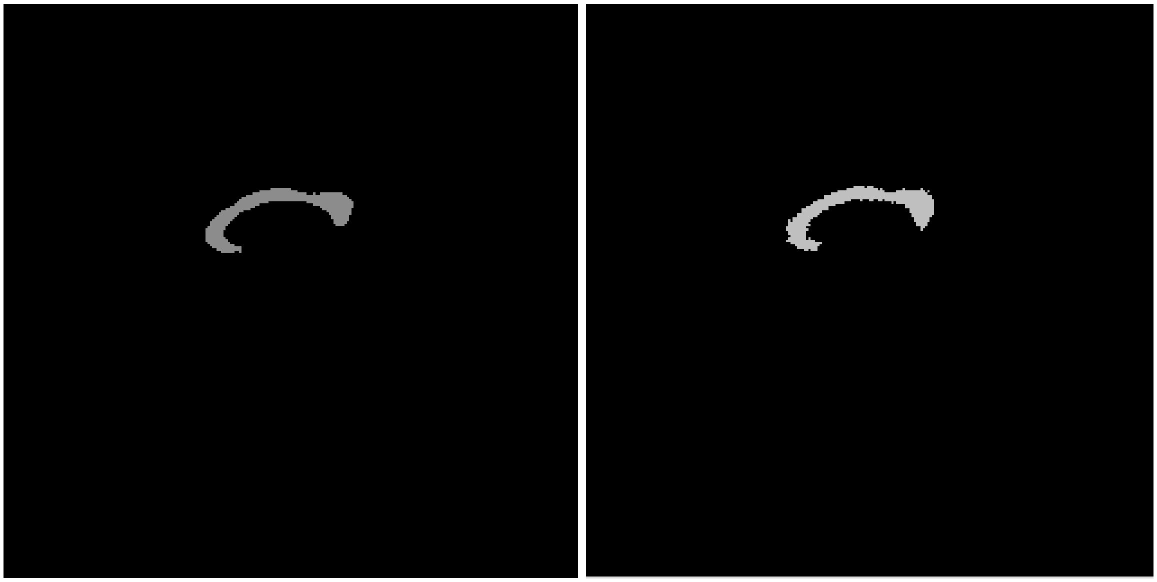

Error Ratio = 0.91%

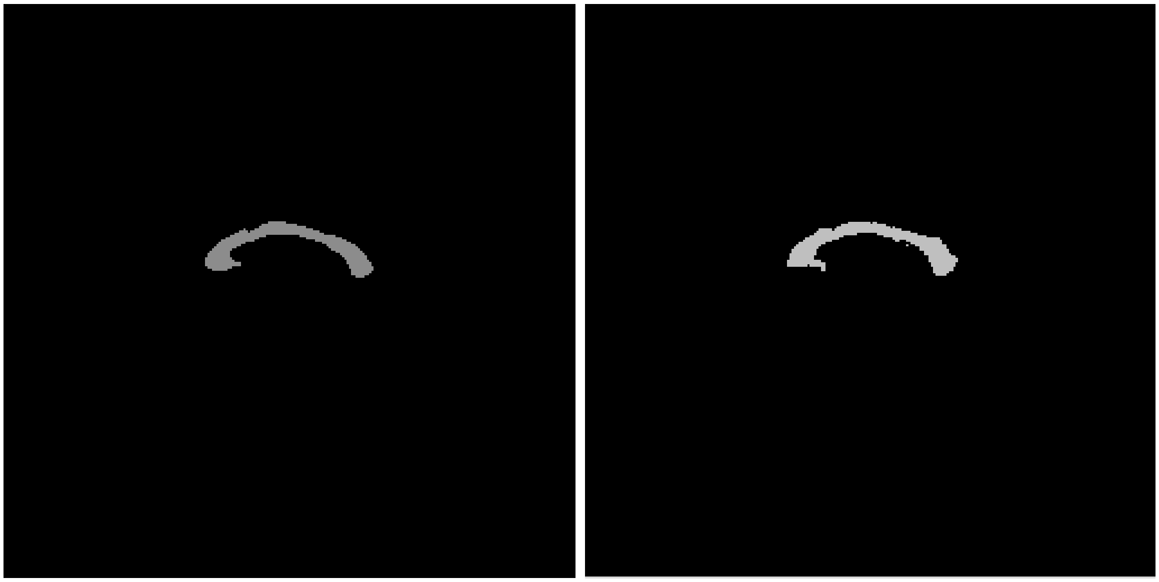

Error Ratio = 2.35%

Expert Segmentation

Automatic Segmentation

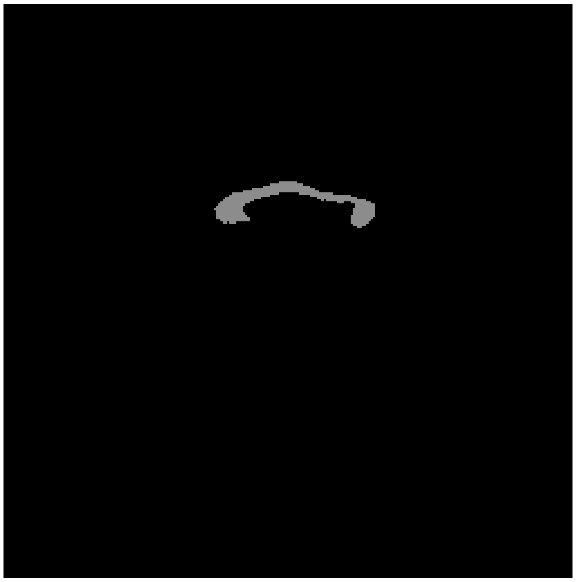

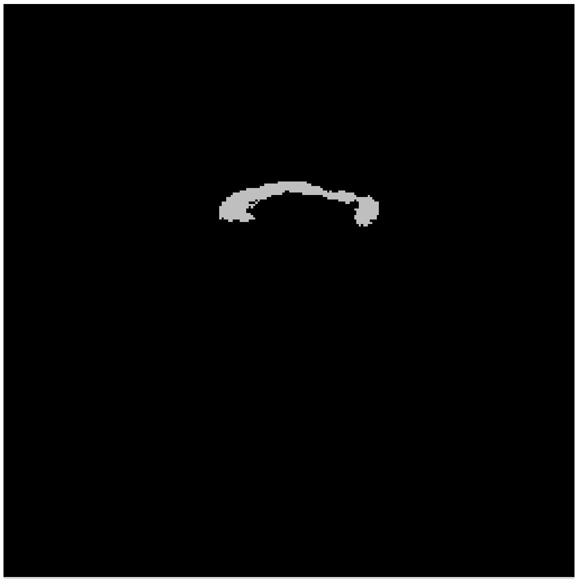

Error Ratio = 0.85%

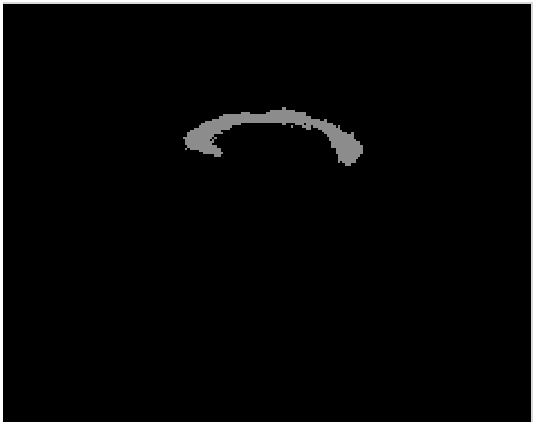

Error Ratio = 2.51%

Expert Segmentation Automatic Segmentation

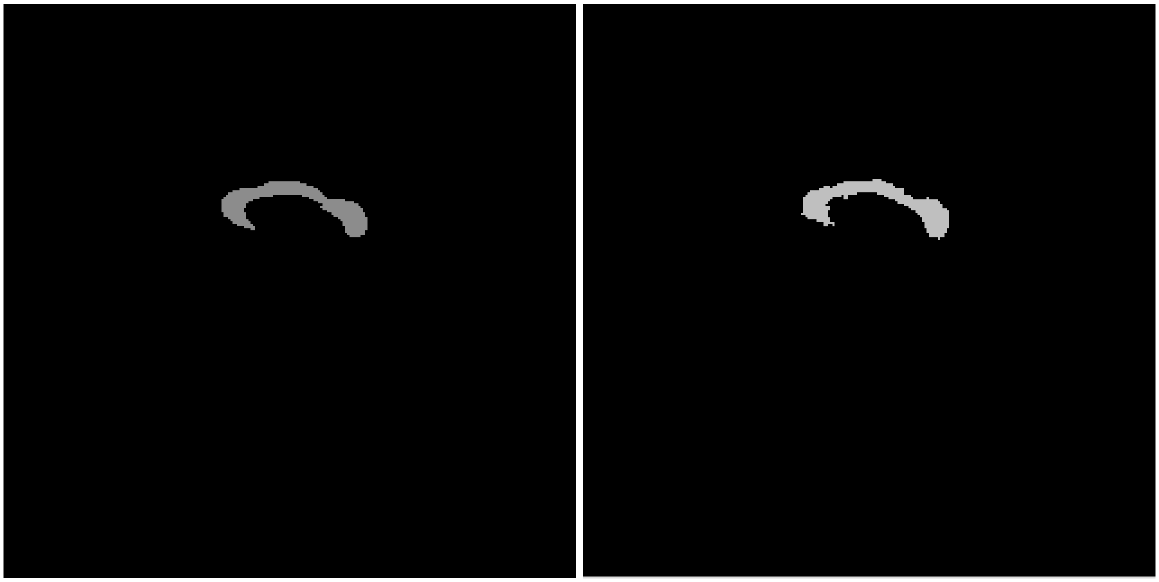

Error Ratio = 2.45%

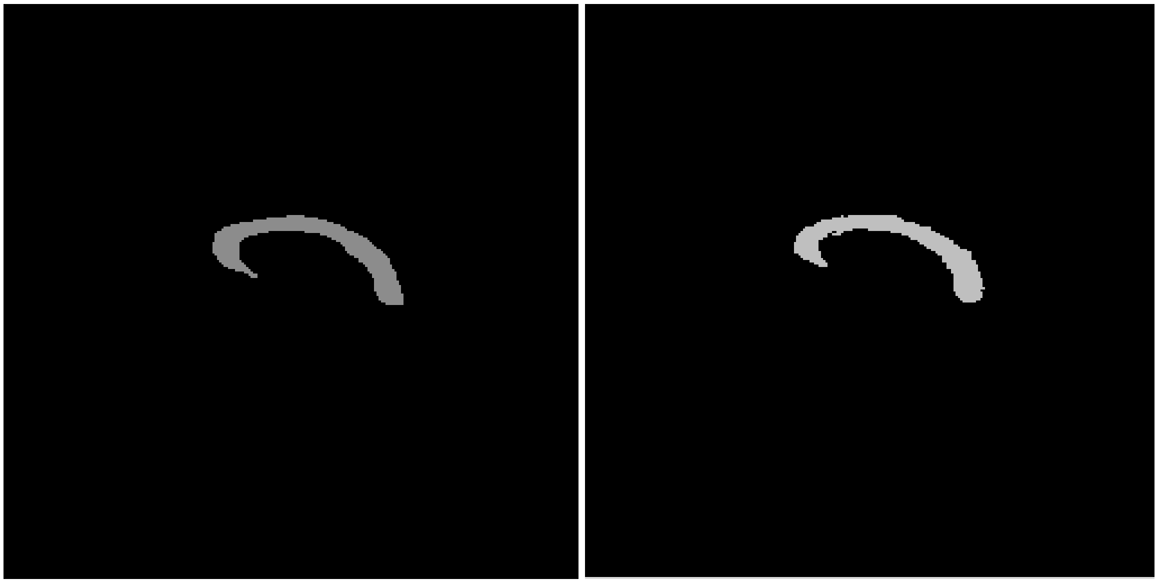

Error Ratio = 2.05%

Expert Segmentation

Automatic Segmentation

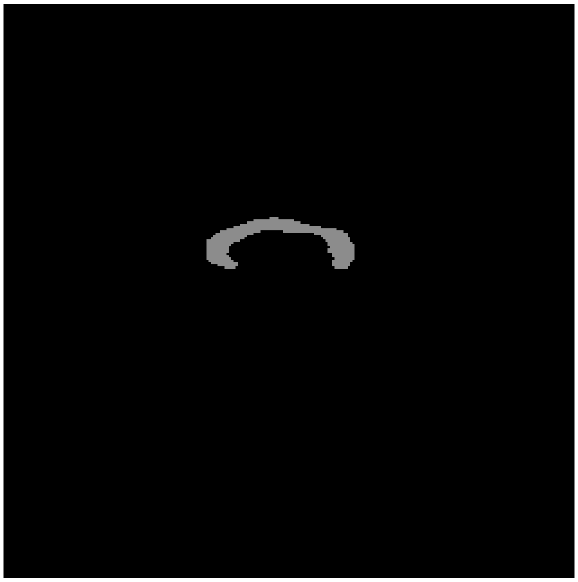

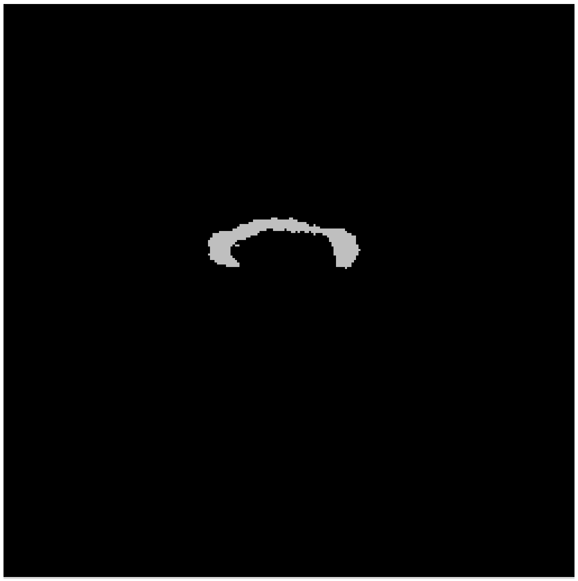

Error Ratio = 1.25%

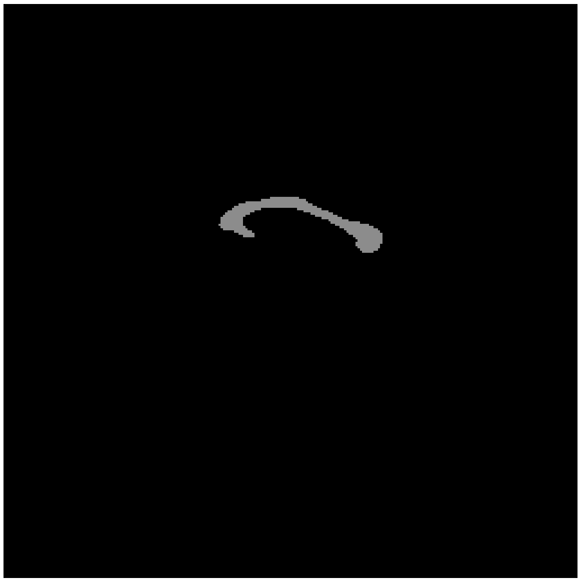

Error Ratio = 7.96%

Expert Segmentation

Automatic Segmentation

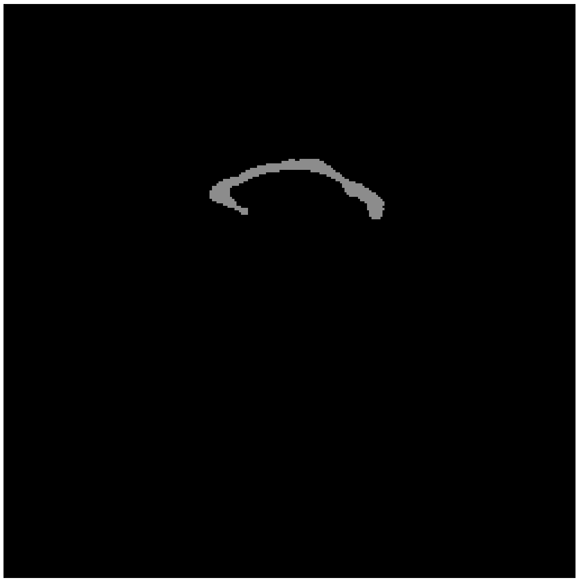

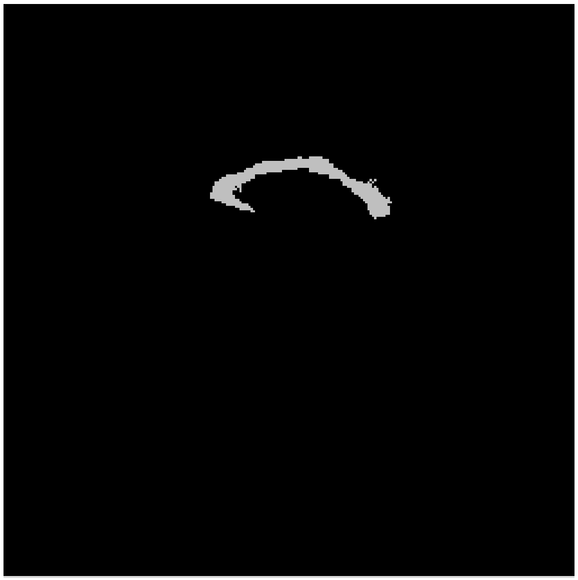

Error Ratio = 5.09%

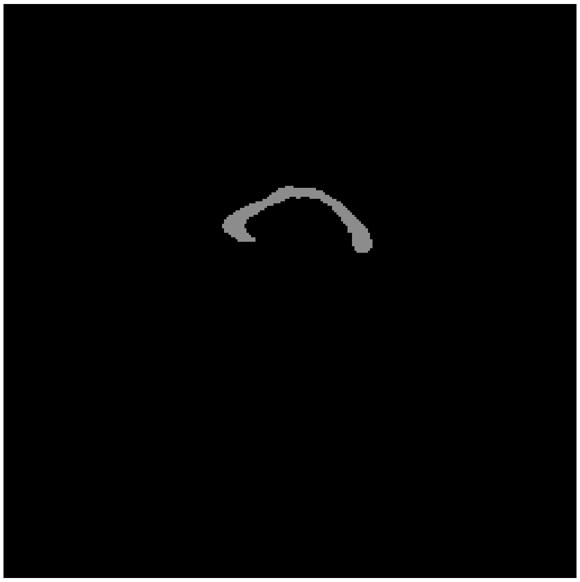

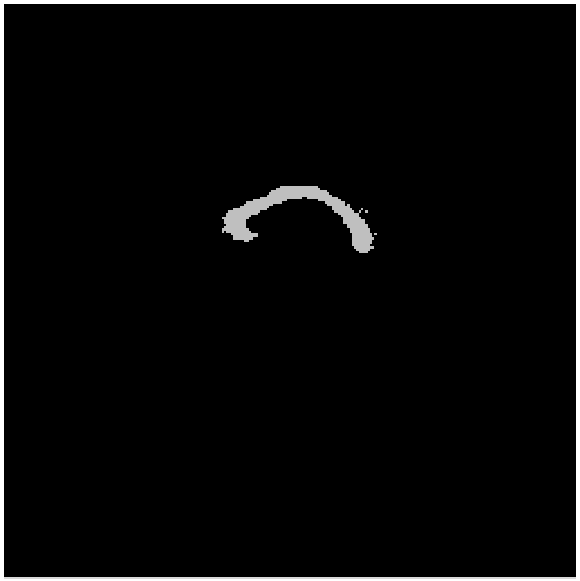

Error Ratio = 1.02%

Expert Segmentation

Automatic Segmentation

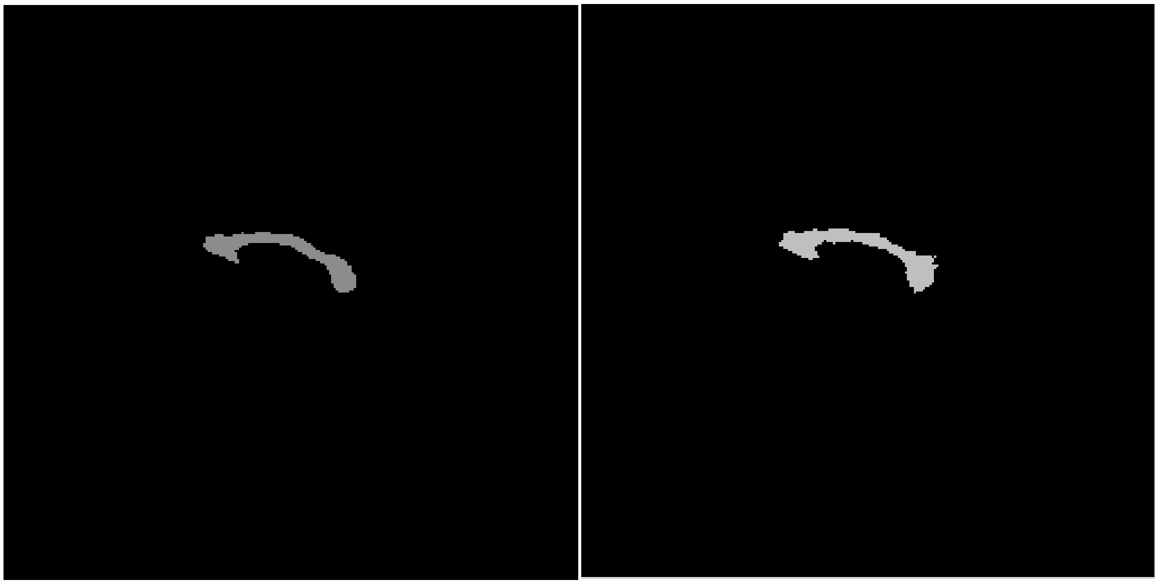

Error Ratio = 4.4%

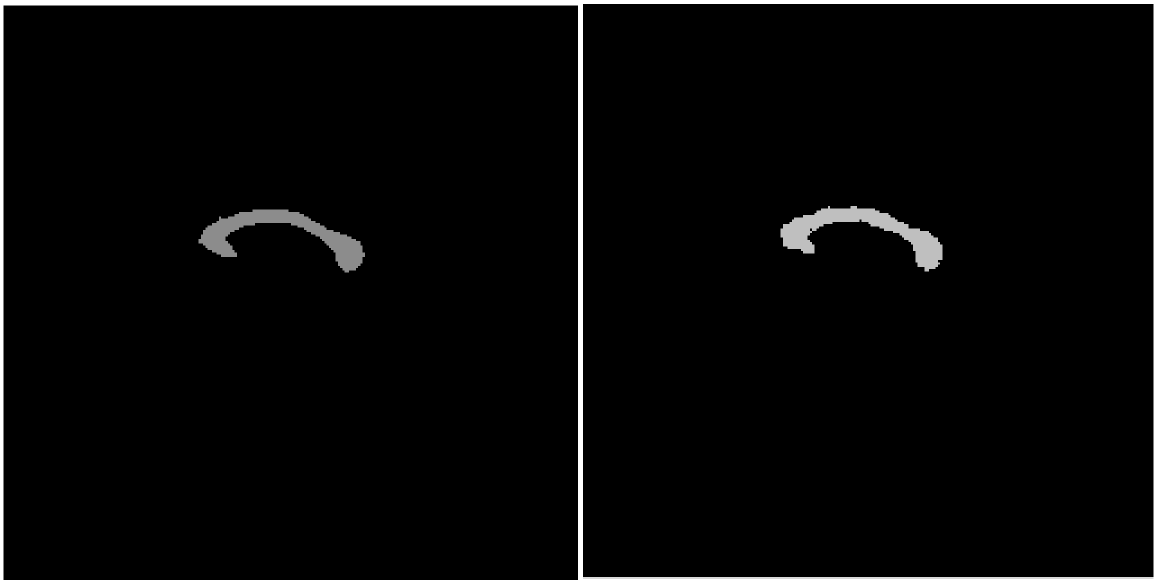

Error Ratio = 0.62%

Expert Segmentation Automatic Segmentation

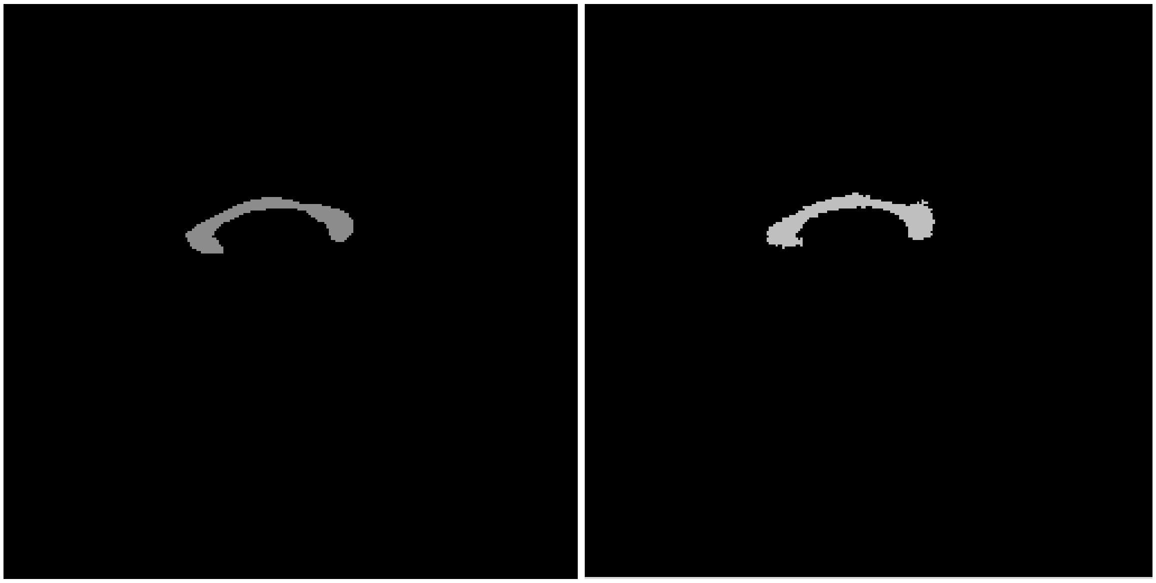

Error Ratio = 3.87%

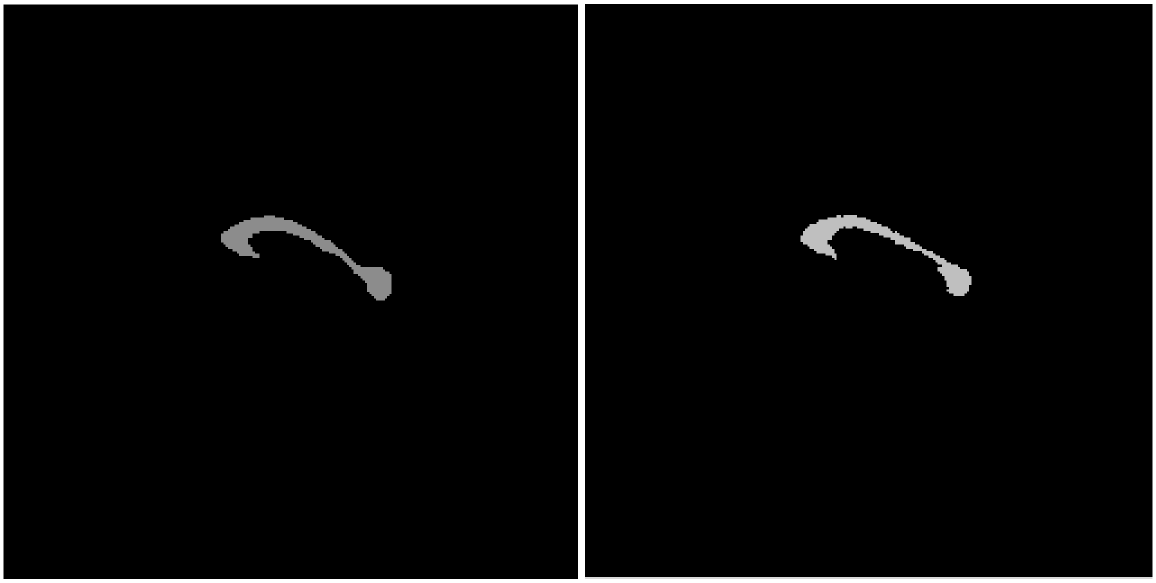

Error Ratio = 3.22%

Expert Segmentation

Automatic Segmentation

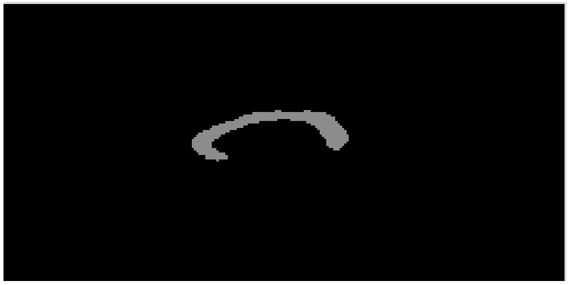

Error Ratio = 4.13%

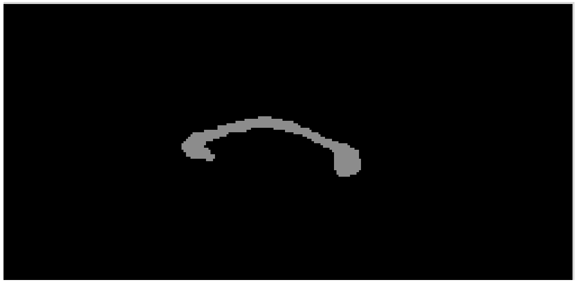

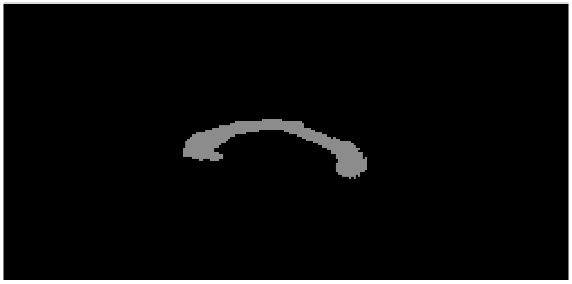

Error Ratio = 3.51%

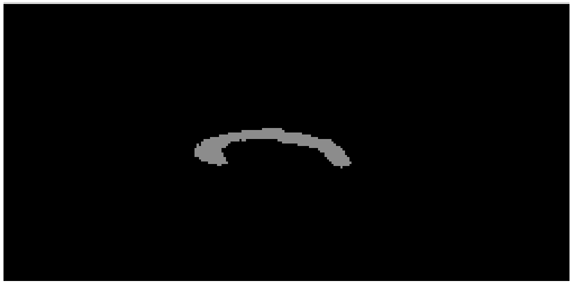

Error Ratio = 0.99%

Expert Segmentation Automatic Segmentation

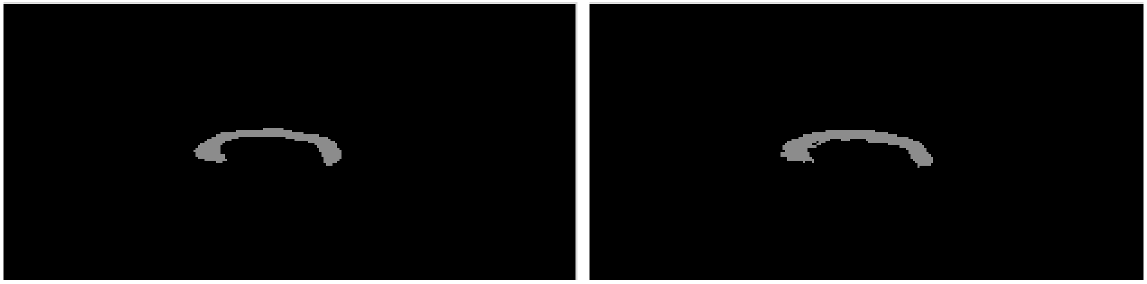

Error Ratio = 9.78%

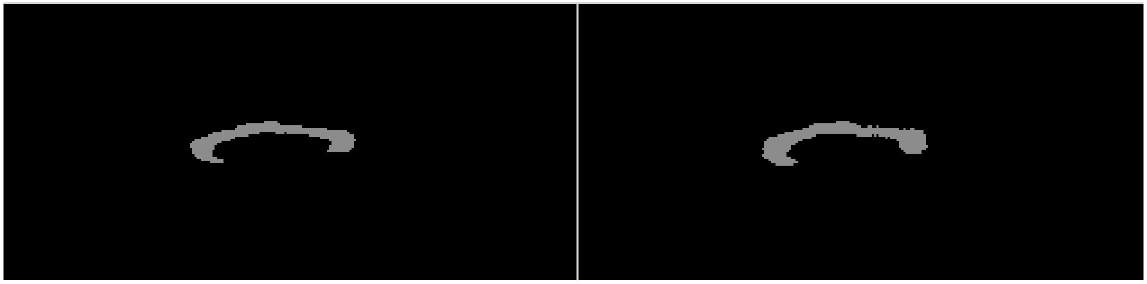

Error Ratio = 4.05%

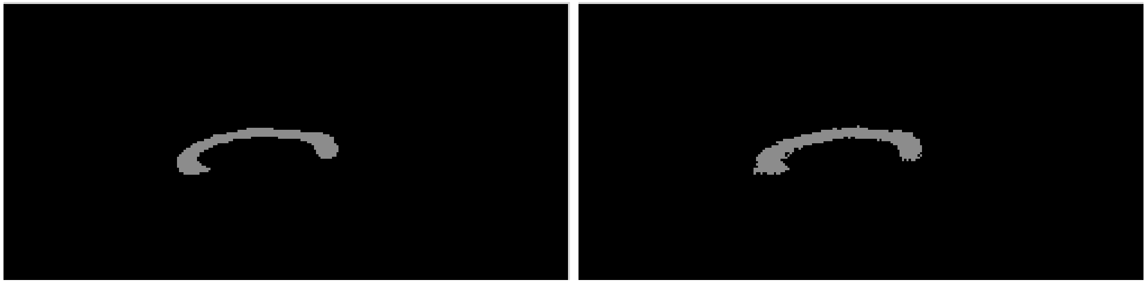

Error Ratio = 3.86%

Expert Segmentation | Automatic Segmentation

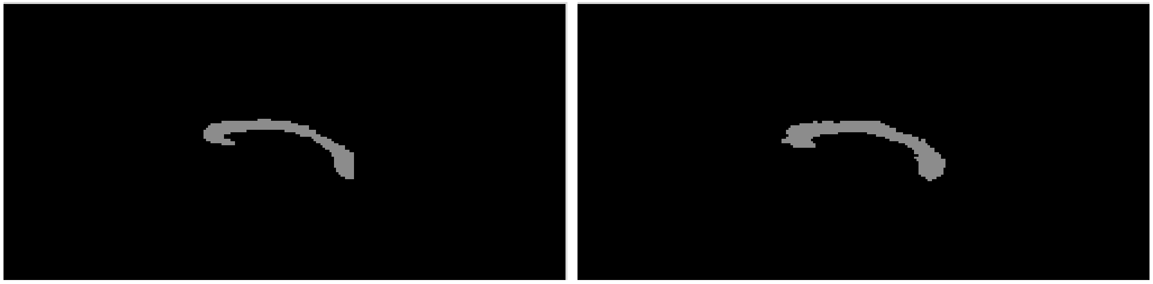

Error Ratio = 8.08%

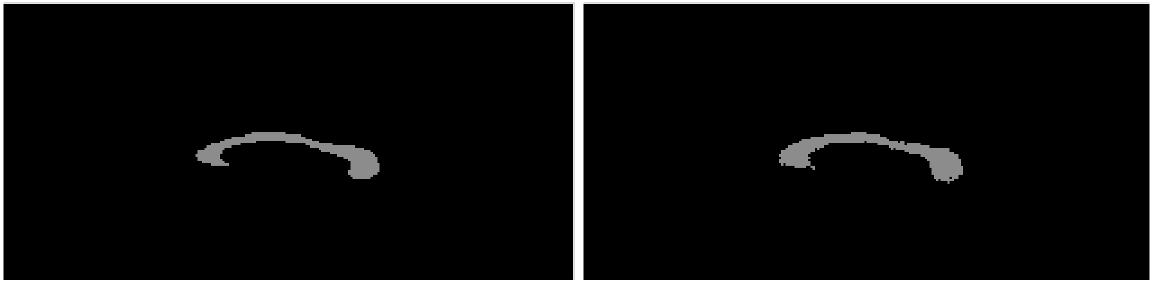

Error Ratio = 4.15%

Error Ratio = 9.55%

Expert Segmentation Automatic Segmentation

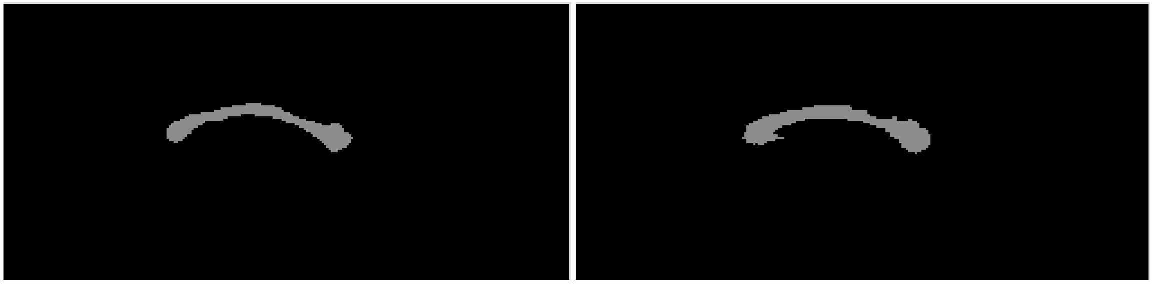

Error Ratio = 5.84%

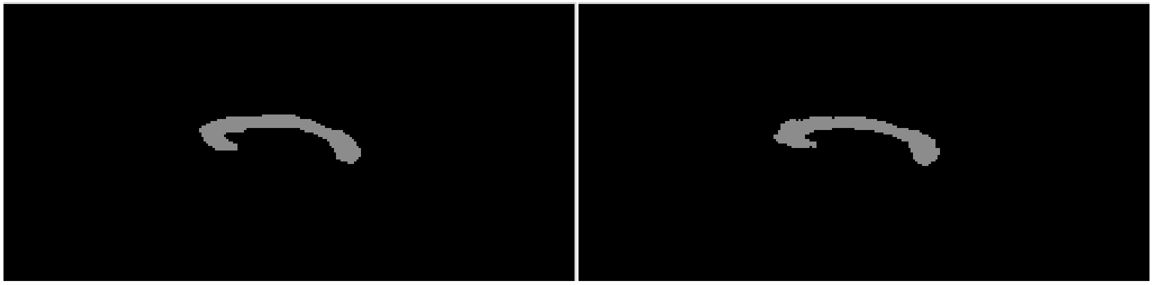

Error Ratio = 1.08%

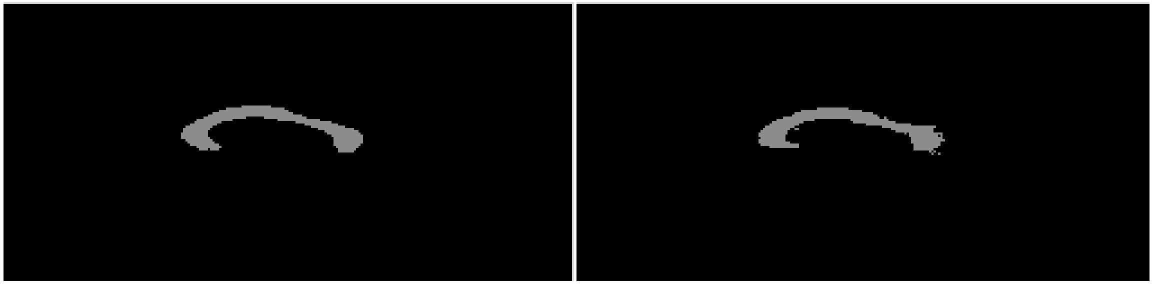

Error Ratio = 3.78%

Expert Segmentation

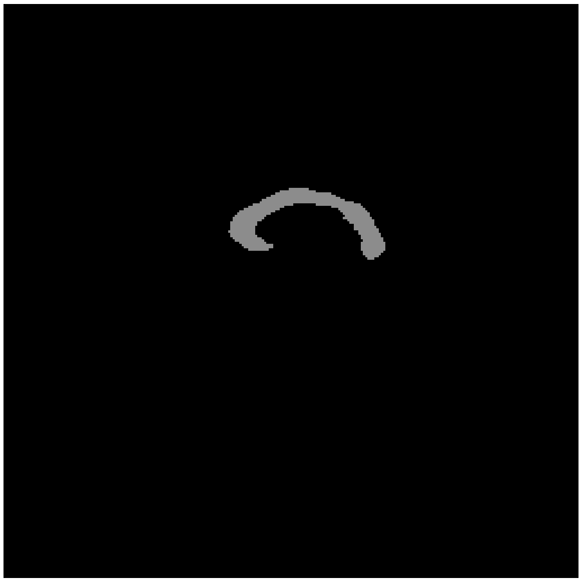

Automatic Segmentation

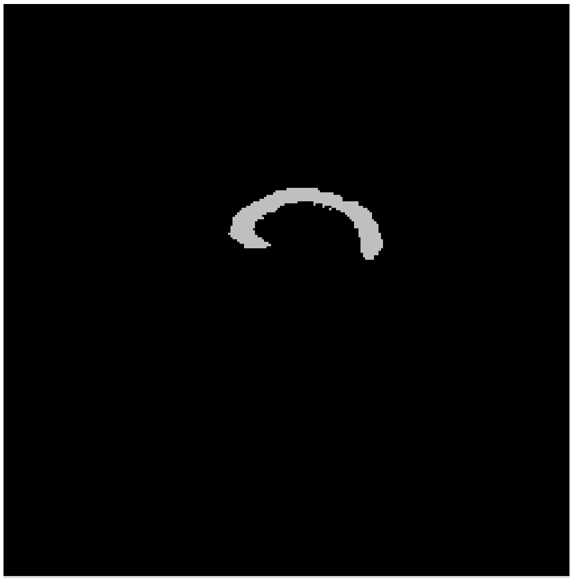

Error Ratio = 0.68%

Bibliography

[1] Anandan, “A computational framework and an algorithm for the measurement of visual motion”, International Journal of Computer Vision, vol. 2, No. 3, pp 283-310, January, 1989.

[2] Bajcsy and Kovacic, “Multiresolution Elastic Matching”, Computer Vision, Graphics, and Image Processing, vol. 46, pp 1-21, 1989.

[3] Benes et al., “Normal ventricles in young schizophrenics”, British Journal of Psychiatry, vol. 141, pp. 90-93, 1982.

[4] Berrut et al., “Modeling and indexing medical images: the RIME approach”, Proceedings of Hypertext-Information Retrieval-Multimedia, pp. 105-115, April, 1995.

[5] Besag, “On the statistical analysis of dirty pictures”, RoyalStatistics, B-48(3), pp. 259-302, 1986.

[6] Black and Anandan, “Robust dynamic motion estimation over time”, Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, pp. 296-302, June, 1991.

[7] Bookstein, “Shape and the information in Medical Images: A decade of the Morphometric Synthesis”, Computer Vision And Image Understanding, vol. 66, No. 2, pp. 97-118, 1997.

[8] Bookstein, “Landmark methods for forms without landmarks: morphometrics of group differences in outline shape”, Medical Image Analysis, vol. 1, No. 3, pp. 225-243.

[9] Chen et al., “Anomaly detection through registration”, Pattern Recognition, vol. 32, pp 113-128, 1999.

[10] Chen et al., “3-D deformable registration of medical images using a statistical atlas”, Proceedings of the second international conference on Medical Image Computing and Computer-Assisted Intervention, pp. 621-630, September, 1999.

[11] Christensen et al., “Volumetric transformation of brain anatomy”, IEEE Transactions on Medical Imaging, vol. 16, No. 6, December, 1997.

[12] Christensen et al., “Deformable templates using large deformation kinematics”, IEEE Transactions on Image Processing, September, 1996.

[13] Christensen et al., “Individualizing Neuroanatomical Atlases Using A Massively Parallel Computer”, IEEE Computer, pp. 32-38, January 1996.

[14] Collins et al., “Automatic 3-D model-based neuroanatomical segmentation”, Human Brain Mapping, vol. 3, pp 190-208, 1995.

[15] Collins et al., “Automatic 3D intersubject registration of MR volumetric data in standardized Talairach space”, Journal of Computer Assisted Tomography, vol. 18, No. 2, pp 192-205, 1994.

[16] Cook et al., “Analysis of cortical patterns”, Magnetic Resonance Scanning and Epilepsy, pp. 263-274, Plenum, New York, 1994.

[17] Cootes et al., “Trainable method of parametric shape description”, Image and Vision Computing, vol. 10, No. 5, pp. 289-294, June, 1992.

[18] Cootes, “Combining point distribution models with shape models based on finite element analysis”, Proceedings of British Machine Vision Conference, 1994.

[19] Dann et al., “Evaluation of elastic matching systems for anatomic (CT, MR) and functional (PET) cerebral images”, Journal of Computer Assisted Tomography, vol. 13, No. 4, pp 603-611, 1989.

[20] Davatzikos, “Spatial transformation and registration of brain images using elastically deformable models”, Computer Vision and Image Processing, Special Issue on Medical Imaging, May, 1997.

[21] Davatzikos et al., “A computerized approach for morphological analysis of the corpus callosum”, Journal of Computer Assisted Tomography, vol. 20, pp. 88-97, 1996.

[22] Davidson and Hugdahl, “Brain Symmetry”, A Bradford Book, 1994.

[23] Declerk, et. al., “Automatic retrieval of anatomical structures in 3-D medical images”, Technical Report 2485, INRIA, Sophia-Antipolis, France, 1995.

[24] DeLisi et al., “A prospective follow-up study of brain morphology and cognition in first-episode schizophrenic patients: Preliminary findings”, Biological Psychiatry, vol. 38, pp. 349-360, 1995.

[25] Evans et al., “Computational approaches to quantitative human neuroanatomical variability”, Brain Mapping: the methods, New York, NY, Academic Press, 1996.

[26] Evans et al., “3D statistical neuroanatomical models from 305 MRI volumes”, Proceedings of IEEE Nuclear Science Symposium and Medical Imaging Conference, pp 1813-1817, 1993.

[27] Evans et al., “An MRI-based stereotactic atlas from 250 young normal subjects”, Society of Neuroscience Abstracts, vol. 18, pp 408, 1992.

[28] Evans et al., “Warping of Computerized 3D Atlas to Match Brain Image Volumes for Quantitative Neuroanatomical and Functional Analysis. Proceedings of SPIE Medical Imaging, vol. 1445, pp. 236-246, 1991.

[29] Evans et al., “MRI-PET correlation in three dimensions using a volume-of-interest (VOI) atlas”, Journal of Cerebral Blood Flow and Metabolism, vol. 11, A69-A78, 1991.

[30] Feldmar and Ayache, “Locally affine registration of free-form surfaces”, Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, pp 496-501, 1994.

[31] Gee and Le Briquer, “An empirical model of brain shape”, Maximum Entropy and Bayesian Methods, August 4-8, 1997.

[32] Gee et al., “Bayesian Approach to the Brain Image Matching Problem”, SPIE Medical Imaging, 1995.

[33] Gee et al., “Elastically Deforming 3D Atlas to Match Anatomical Brain Images”, Journal of Computer Assisted Tomography, vol. 17, No. 2, pp 225-236, 1993.

[34] Geiger and Girosi, “Parallel and deterministic algorithms from MRFs: surface reconstruction”, IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 13, No. 5, pp. 401-412, May, 1991.

[35] Geman and Geman, “Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images”, IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 6, No. 6, pp. 721-741, November, 1994.

[36] Gleicher and Witkin, “Through-the-lens camera control”, Computer Graphics, vol. 26, No. 2, pp 331-340, July, 1992.

[37] Guimond and Subsol, “Automatic MRI database exploration and application”, Pattern Recognition and Artificial Intelligence, vol. 11, No. 8, December, 1997.

[38] Guimond et. al., “Automatic Computation of Average Brain Models”, Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention, pp. 631-640, 1998.

[39] Haller et al., “Three-dimensional hippocampal MR morphometry with high-dimensional transformation of a neuroanatomic atlas”, Radiology, No. 202, pp. 504-510, 1997.

[40] Hill et al., “Model-based interpretation of 3D medical images”, Proceedings of the 4th British Machine Vision Conference, pp. 339-348, 1993.

[41] Hillman et al., “Measurement of brain compartment volumes in MR using composition calculations”, Journal of Computer Assisted Tomography, vol. 15, pp. 640-646, 1991.

[42] Hommer et al., “Decreased corpus callosum size among alcoholic women”, Archives of Neurology, vol. 53, pp. 359-363, April, 1996.

[43] Horn, “Closed-form solution of absolute orientation using unit quaternions”, Journal of Optical Society of America, vol. 4, No. 4, April, 1987.

[44] Horn and Schunck, “Determining optical flow”, Artificial Intelligence, vol. 17, pp 185-203, 1981.

[45] Jacobson, “The contributions of sex and drinking history to the CT brain scan changes in alcoholics”, Psychological Medicine, vol. 16, pp. 547-559, 1986.

[46] Joshi et al., “On the geometry and shape of brain sub-manifolds”, International Journal of Pattern Recognition and Artificial Intelligence, vol. 11, No. 8, pp. 1317-1343, 1997.

[47] Joshi et al., “Hierarchical brain mapping via a generalized Dirichlet solution for mapping brain manifolds”,

[48] Kamber et al., “Model-based 3-D segmentation of multiple sclerosos lesions in magnetic resonance brain images”, IEEE Transactions on Medical Imaging, vol. 14, No. 3, pp. 442-453, September, 1995.

[49] Kapur et al., “Enhanced spatial priors for segmentation of magnetic resonance imagery”, Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention, pp. 457-468, 1998.

[50] Kelly and Cannon, “CANDID: Comparison algorithm for navigating digital image databases”, Proceedings of 7th International Working Conference on Scientific and Statistical Database Management, pp. 252-258, September, 1994.

[51] Keshavan et al., “Brain structural alterations in early Schizophrenia: do they progress?”, Journal of Psychiatric Research, 1996.

[52] Kikinis et al., “Routine quantitative analysis of brain and cerebrospinal fluid spaces with MR imaging”, Journal of Magnetic Resonance Imaging, vol. 2, pp. 619-629, 1992.

[53] Kikinis et al., “A digital brain atlas for surgical planning, model driven segmentation and teaching”, IEEE Transactions on Visualization and Computer Graphics, vol. 2, No. 3, 1996.

[54] Kischell et al., “Classification of brain compartments and head injury lesions by neural networks applied to MRI”, Diagnostic Neuroradiology, vol. 37, pp. 535-541, 1995.

[55] Korn et al., “Fast nearest neighbor search in medical image databases”, Proceedings of Visual Libraries and DataBases, 1996.

[56] Kroft et al., “Brain ventricular size in female alcoholism: an MRI study”, Alcoholism, vol. 8, pp. 31-34, 1991.

[57] Kyriacou and Davatzikos, “A biomedical model of soft tissue deformation with applications to non-rigid registration of brain images with tumor pathology”, Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention, pp. 531-538, 1998

[58] Lirio et al., “Localization error in biomedical imaging”, Computers in Biology and Medicine, vol. 22, No. 4, pp 277-286, 1992.

[59] Liu et al., “Automatic bilateral symmetry (mid-sagittal) plane extraction from pathological 3D neuroradiological images”, Proceedings of SPIE International Symposium on Medical Imaging, February, 1998.

[60] Liu and Dellaert, “Classification-driven medical image retrieval”, Image Understanding Workshop, November, 1998.

[61] Loncaric et al., “3-D image analysis of intra-cerebral brain hemorrhage from digitized CT films”, Computer Methods and Programs in Biomedicine, vol. 46, pp. 207-216, 1995.

[62] Maes et al., “Multimodality Image Registration by Maximization of Mutual Information”, IEEE Transactions on Medical Imaging, vol. 16, No. 2, April, 1997.

[63] Maintz et al., “Evaluation of ridge seeking operators for multimodality medical image matching”, IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 18, No. 4, April, 1996.

[64] Mann et al., “Do women develop alcoholic brain damage more readily than men?”, Alcohol Clin Exp Res., vol. 16 pp. 1052-1058, 1992.

[65] Marias et al., “Visualizing cerebral asymmetry”, Visualization in Biomedical Computing, pp. 411-416, Springer, Hamburg, September, 1996.

[66] Martin et. al., “Characterization of Neuropathological Shape Deformations”, IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 20, No. 2, 1998.

[67] McInerney and Terzopoulos, “Deformable models in medical image analysis: a survey”, Medical Image Analysis, vol. 1, No. 2, pp 91-108, 1996.

[68] Metropolis et al., “Equation of state calculations by fast computing machines”, Journal of Chemistry and Physics, vol. 21, No. 6, p. 1087-1092, 1953.

[69] Meyer et al., “Demonstrating the accuracy and clinical versatility of mutual information for automatic multimodality image fusion using affine and thin-plate spline warped geometric deformation”, Medical Image Analysis, vol. 1, No. 3, pp 195-206, 1997.

[70] Miller et al., “Magnetic resonance imaging in monitoring the treatment of multiple sclerosis: concerted action guideline”, Journal of Neurology, Neurosurgery, and Psychiatry, vol. 54, pp. 683-688, 1991.

[71] Moghaddam et al., “A Bayesian similarity measure for direct image matching”, M.I.T. Media Laboratory Perceptual Computing Section Technical Report No. 393.

[72] Pennec and Thirion, “Validation of 3-D registration methods based on points and frames”, INRIA, Technical Report, No. 2470, January, 1995.

[73] Pentland and Sclaroff, “Closed-form solutions for physically-based shape modeling and recognition”, IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 13, No. 7, July, 1991.

[74] Pokrandt, “Fast Non-Supervised Matching: A Probabilistic Approach”, internal technical report, Institute for Real-Time Computer Systems and Robotics, University of Karlsruhe, 1996.

[75] Pokrandt, “Medical image registration: fast feature space evaluation with Gaussian entropy function”, internal technical report, Institute for Real-Time Computer Systems and Robotics, University of Karlsruhe, 1996.

[76] Press et. al., “Numerical Recipes in C”, Cambridge University Press, Cambridge, England, second edition, 1992.

[77] Quam, “Hierarchical warp stereo”, Image Understanding Workshop, pp 149-155, December, 1984.

[78] Rademacher et al., “Topographical variation of the human primary cortices: implications for neuroimaging, brain mapping and neurobiology”, Cerebral Cortex, vol. 3, pp. 313-329, 1993.

[79] Rangarajan and Duncan, “Matching point features using mutual information”, Internal Technical Report, Department of Diagnostic Radiology and Electrical Engineering, Yale University, 1998.

[80] Sandor and Leahy, “Surface-based labelling of cortical anatomy using a deformable atlas”, IEEE Transactions on Medical Imaging, vol. 16, No. 1, February, 1997.

[81] Schulz et al., “Ventricular enlargement in teenage patients with schizophrenia spectrum disorder”, American Journal of Psychiatry, vol. 140, No. 12, pp. 1592-1595, December, 1983.

[82] Sclaroff and Pentland, “On modal modeling for medical images: underconstrained shape description and data compression”, Proceedings of IEEE Workshop on Biomedical Image Analysis, pp. 70-79, Seattle, 1994.

[83] Shenton et al., “Harvard brain atlas: a teaching and visualization tool”, Proceedings of Biomedical Visualization, pp. 10-17, 1995.

[84] Simon et al., “Techniques for fast and accurate intrasurgical registration”, Journal of Image Guided Surgery, vol. 1, 17-29, 1995.

[85] Sponheim et al., “Stability of ventricular size after the onset of psychosis in Schizophrenia”, Psychiatry Research: Neuroimaging, vol. 40, pp. 21-29, 1991.

[86] Stensaas et al., “The topography and variability of the primary visual cortex in man”, Journal of Neurosurgery, vol. 40, pp. 747-755, 1974.

[87] Szekely et. al., “Segmentation of 2-D and 3-D objects from MRI volume data using constrained elastic deformations of flexible Fourier contour and surface models”, Medical Image Analysis, vol. 1, No. 1, pp. 19-34.

[88] Szeliski, “Bayesian modeling of uncertainty in low-level vision”, Technical Report, Computer Science Department, Carnegie Mellon University, CMU-CS-88-169, August, 1988.

[89] Szeliski, “Matching 3-D Anatomical Surfaces with Non-Rigid Deformations using Octree-Splines”, International Journal of Computer Vision, vol. 18, No. 2, pp 171-186, 1996.

[90] Szeliski and Coughlan, “Spline-Based Image Registration”, Digital Equipment Corporation, Cambridge Research Lab, Technical Report 94/1.

[91] Thevenaz et al., “A pyramid approach to subpixel registration based on intensity”, IEEE Transactions on Image Processing, vol. 7, No. 1, January, 1998.

[92] Thirion, “Fast non-rigid matching of 3D medical images”, INRIA, Technical Report No. 2547, May, 1995.

[93] Thirion et al., “Statistical analysis of normal and abnormal dissymmetry in volumetric medical images”, Technical Report No. 3178, INRIA, June, 1997.

[94] Thompson and Toga, “Detection, visualization and animation of abnormal anatomic structure with a deformable probabilistic atlas based on random vector field transformations”, Medical Image Analysis, vol. 1, No. 4, pp 271-294.

[95] Thompson et al., “High-resolution random mesh algorithms for creating a probabilistic 3-D surface atlas of the human brain”, NeuroImage, vol. 3, No. 1, pp 9-34, February, 1996.

[96] Thompson and Toga, “A surface-based technique for warping three-dimensional images of the brain”, IEEE Transactions on Medical Imaging, vol. 15, No. 4, pp 402-417, August, 1996.

[97] Thompson et al., “Mapping structural alterations of the corpus callosum during brain development and degeneration”, Proceedings of the NATO ASI on the Corpus Callosum, Kluwer Academic Press.

[98] Thompson et al., “Tracking tumor growth rates in patients with malignant gliomas”, Annual Meeting of the Society for Neuro-Oncology, Scottsdale, Arizona, November 18-21, 1999.

[99] Thompson et al., “Three-dimensional statistical analysis of sulcal variability in the human brain”, Journal of Neuroscience, vol. 16, No. 13, pp. 4261-4274, July, 1996.

[100] Ullman,

[101] van den Elsen, Pol, and Viergever, “Medical Image Matching--A Review with Classification”, IEEE Engineering in Medicine and Biology, vol. 12, No. 1, pp 26-39, 1993.

[102] Vemuri et al., “An Efficient Motion Estimator with Applications to Medical Image Registration”, Medical Image Analysis.

[103] Viola and Wells III, “Alignment by Maximization of Mutual Information”, Proceedings of the fifth International Conference on Computer Vision, pp 16-23, 1995.

[104] Wang and Staib, “Elastic Model Based Non-rigid Registration Incorporating Statistical Shape Information”, Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention, pp. 1162-1173, 1998.

[105] Warfield et al., “Automatic identification of grey matter structures from MRI to improve the segmentation of white matter lesions’, Journal of Image Guided Surgery, vol. 1, No. 6, pp. 326-338, 1995.

[106] Wells et al., “Adaptive segmentation of MRI data”, IEEE Transactions on Medical Imaging, 1996.

[107] West et al., “Comparison and evaluation of retrospective intermodality brain image registration techniques”, Journal of Computer Assisted Tomography, vol. 21, pp 554-566, 1997.

[108] Wheeler and Ikeuchi, “Iterative estimation of rotation and translation using the quaternions”, Technical Report, Computer Science Department, Carnegie Mellon University, CMU-CS-95-215, December, 1995.

[109] Witkin et. al., “Signal matching through scale space”, International Journal of Computer Vision, vol. 1, pp 133-144, 1987.

[110] Woods et al., “MRI-PET registration with automated algorithm”, Journal of Computer Assisted Tomography, vol. 17, No. 4, pp 536-546, 1993.

[111] Woods et al., “Rapid automated algorithm for aligning and reslicing PET images”, Journal of Computer Assisted Tomography, vol. 16, No. 4, pp 620-633, 1992.

[112] Zhu and Yuille, “Forms: a flexible object recognition and modeling system”, Harvard Robotics Laboratory, No. 94-1, 1993.