Abstract

Volumetric Layer Segmentation Using a Generic Shape Constraint

with Applications to Cortical Shape Analysis

Xiaolan Zeng

Yale University

2000

A novel approach has been developed in this thesis for the problem of segmenting volumetric layers, a type of structure often encountered in medical image analysis. This approach is aimed towards the use of structural information to enhance the performance of the segmentation process. While some organs have more consistent global shape and can be characterized using a specific shape model, other anatomical structures possess much more complex shape with possibly high variability which needs a more generic shape constraint. The three-dimensional(3D) nature of anatomical structures necessitates the use of volumetric approaches that exploit complete spatial information and therefore are far superior to the non-optimal and often-biased 2D methods. Our method takes a volumetric approach, and incorporates a generic shape constraint – in particular, a thickness constraint. The resulting coupled surfaces algorithm with a level set implementation not only offers segmentation with the advantages of minimal user interaction, robustness to initialization and computational efficiency, but also facilitates the extraction and measurement of many geometric features of the volumetric layer.

The algorithm was applied to 3D Magnetic Resonance (MR) brain images for skull-stripping, cortex segmentation and various feature measurements including cortical surface shape and cortical thickness. Validation of the model was done through both

synthetic images with "ground truth" and a wide range of real MR images with expert tracing results. As a natural follow up to the segmentation work, a new approach was developed for the extraction of sulcal ribbon surfaces which are distinctive cortical landmarks of the brain. This effective and efficient 3D method of sulcal ribbon extraction has potential in a variety of applications such as the automatic parcellation of cortical regions and the problem of geometry-constrained brain atlas building. The tools of cortical and sulcal shape analysis developed in this work are of great importance to studies of neuroanatomy through medical imaging, and are bringing about new understanding of brain anatomy and function.

Volumetric Layer Segmentation Using a Generic Shape Constraint with Applications to Cortical Shape Analysis

A Dissertation

Presented to the Faculty of the Graduate School

of

Yale University

in Candidacy for the Degree of

Doctor of Philosophy

by

Xiaolan Zeng

Dissertation Director: James Scott Duncan

May 2000

Acknowledgments

I owe heartfelt thanks to my advisor James Duncan. This work would not have been possible without his continued guidance, encouragement and support. I am especially grateful to Larry Staib for his numerous offers of help and for always being ready to discuss any question I have. I thank Roman Kuc and Peter Belhumeur for serving on my committee. I thank Bob Schultz for our collaboration, and Hemant Tagare for many stimulating technical discussions and his generous advice. Thanks to Arye Nehorai for guidance during my first year at Yale and his continued encouragement. Thanks also to Eric Grimson for being my external reader.

Financial support for this work was provided, in part, by NIH grant NINDS R01 NS35 193-01 and NSF grant IRI-9530768. Thanks to the McConnell Brain Imaging Center at Montreal Neurological Institute and the Center for Morphometric Analysis at Massachusetts General Hospital for providing test images.

I thank all the past and present members of the Image Processing and Analysis Group for making my work here enjoyable: Ravi Bansal, Haili Chui, Don Dione, Francios Meyer, Xenios Papademetris, Suguna Pappu, Jim Rambo, Anand Rangarajan, Pengcheng Shi, Oskar Skrinjar, Rik Stokking, Colin Studholme, Yongmei Wang, Larry Win and George Zubal. Carolyn Meloling and Mary Lally provided administrative assistance and encouragement.

A special thanks goes to Xiangyi Huang, Jun Gu, and many others, for their friendship.

Finally, I thank Xiaoyi Wu for his love, support and unwavering confidence in me. I dedicate this thesis to my parents. Their love is my constant source of comfort and happiness.

List of Figures

List of Tables

Contents

Chapter 1

Introduction

1.1 Introduction

Segmentation and quantitative analysis of objects depicted in an image have tremendous applications in the field of computer vision. For example, modern medical imaging technology provides exceptional views of internal anatomy by peering noninvasively into the human body [92]. Medical imaging is playing an increasingly important role in diagnosis, as well as in the planning and treatment of disease. In order to fully realize the value of medical imaging in both clinical and research settings, information about anatomical structures must be extracted and quantified from medical images with accuracy, efficiency and reproducibility.

There are many sources of information to assist in the recovery of anatomical structures of interest. Medical imaging techniques depend on differentiating properties among tissue types, for example their density and water content. As a result, homogeneity of features in an image, such as similarity of gray level and texture pattern, is displayed by tissue belonging to the same anatomical region. A change in such characteristics is usually

evidence of the presence of a boundary between regions. Image gradient, a measure of the intensity difference between neighboring voxels, and image curvature, an indicator of corner information, are among the most commonly used features for estimating such changes from images. While such image-derived sources of information are intuitive and direct tools to use, they are often inconsistent due to the ambiguity in image content caused by a host of artifacts, and are incomplete for defining coherent structure due to the limitation that they provide only local information.

Structural information which reflects our knowledge of anatomy can be introduced to help solve these problems. This knowledge can be the perception of the particular shape of a certain structure, or the neighboring relationships between structures. This type of information can be used in various ways. One simple thing to do is to use it as a guideline for grouping or other post-processing after image classification. More sophisticated methods investigate its integration into the mechanism of shape recovery as a direct constraint. While some organs have more consistent global shape and can be characterized using a specific shape model, other anatomical structures possess much more complex shape with possibly high variability where a more generic shape constraint is helpful or necessary. The successful incorporation of structural information depends on the nature and purpose of a specific task. The challenge lies in having a model that not only describes the structure to certain precision, but also permits variations in the reflecting characteristics that are observed across different subjects [55].

The 3D analysis of anatomical structures that are volumetric in nature is made possible by modern imagery such as MR and Computed Tomography (CT). Processing of such volumetric images on a slice by slice basis is a labor intensive process, and requires further processing to stack and interpolate results into 3D. Moreover, such processing can

make incorrect assumptions by ignoring inter-slice correlation in a 3D image, and therefore is often non-optimal or even biased. To obtain an accurate and robust segmentation of 3D anatomical structures, volumetric approaches are of particular importance.

1.2 Overview

This thesis is aimed at using structural information to assist 3D segmentation and analysis of one type of structure that is encountered in medical image analysis – namely volumetric layers of finite thickness bounded by two surfaces. Key examples include the myocardium of the left ventricle (LV) and especially the cortical gray matter of the brain (see Figure 1.1). Others include vessel walls, the lining of the membrane inside the liver, and the sigmoid wall of the bowel lumen.

Chapter 2 reviews related image segmentation techniques, concerning the use of region-based methods, deformable model-based methods, the incorporation of shape priors and the integration of the above, with an emphasis on 3D approaches. Background knowledge of the level set method, which is the mechanism of surface propagation in our work, is also given in this chapter.

In Chapter 3, we discuss the nature of the problem of segmenting a volumetric layer, as well as the motivation behind using a coupled surfaces approach with a thickness constraint. We discuss a level set narrow band technique that was developed especially for the coupled surfaces propagation. With this narrow band implementation of our algorithm, the correspondence between the points on two bounding surfaces automatically falls out of the rebuilding of narrow bands at each iteration. This correspondence is essential for the coupling between the two surfaces through the thickness constraint.

We then apply the method to MR brain images to segment cortical gray matter.

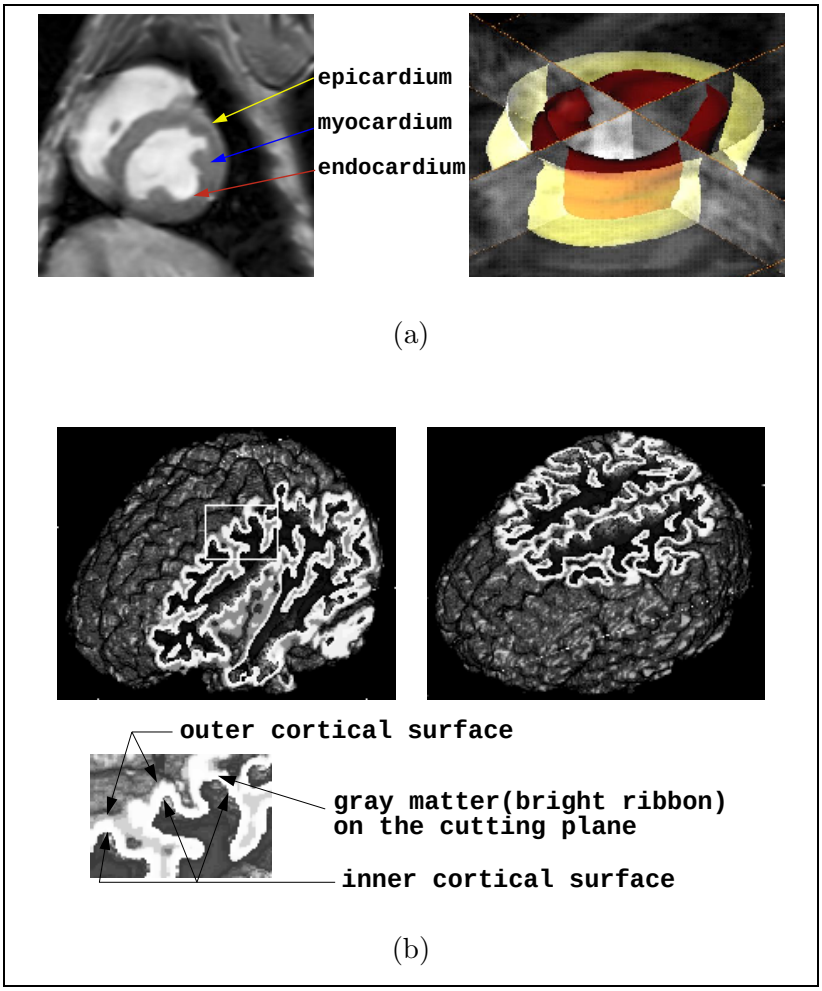

Figure 1.1: (a): Myocardium of the left ventricle bounded by the epicardium and the endocardium. (b): 3D volume rendering of the cortical gray matter from our 3D algorithm with oblique cutting planes. The convoluted thin bright ribbons are the cortical gray matter captured on the cutting plane, the darker parts are the rendering of the out-of-plane gray matter.

Chapter 4 describes the design of speed terms for the coupled surfaces propagation which is key to the segmentation algorithm. The speed terms for the coupled surfaces are influenced by two factors – an appropriate image gray level transition captured by a local operator specifically designed to exploit image homogeneity, and a thickness constraint emulating the nearly constant thickness of human cerebrum cortex. Validation of the segmentation algorithm is presented through experiments on both simulated images with "ground truth" information, and real images with expert tracing results. The results from our algorithm on real images with a range of quality are also compared to those from various classification-based methods, to demonstrate the advantage of the thickness shape constraint.

Chapter 5 showcases a number of cortical geometric measurements made possible by the coupled surfaces algorithm, such as the cortical shape index and cortical thickness. Several studies are shown to demonstrate the usage of such measurements in the study of neuroanatomy.

In Chapter 6, we present a 3D method of intra-sulcal ribbon extraction. This work is a natural follow-up to the coupled surfaces algorithm for cortex segmentation, however, it can be used with general brain segmentation methods. We also discuss how this sulcal ribbon extraction work can be applied in many ways to facilitate neuroanatomy studies.

1.3 Main Contributions

A novel coupled surfaces approach has been developed for segmenting and measuring volumetric layers bounded by two surfaces [100, 99, 101] (see Figures 4.4, 7.1). Surface propagation is driven by the search for boundaries of interest which are characterized by appropriate image gray level transitions. A local operator exploiting image homogeneity

is designed specifically to capture such transitions. A generic shape constraint, namely a thickness constraint, is also incorporated into the system to guide the recovery process. This is motivated by the coupled nature of the inner and outer bounding surfaces of such layers. Through this coupling, image information on one bounding surface is made available to the partner surface, thus achieving a more robust segmentation.

The algorithm is implemented using a level set method with a specifically-designed coupled narrow band technique. This is one of the first times shape constraints have been incorporated into level set methods. As a result of the level set implementation, our system not only performs automatic segmentation with the advantage of minimal user interaction, robustness to initialization and computational efficiency, but also offers a compact representation of the volumetric layer that easily facilitates a variety of geometric measurements.

The system is applied to the problems of cortical segmentation and shape analysis, which are significant to the study of brain morphology. Our segmentation algorithm is extensively validated through both simulated images with "ground truth" and a wide range of real MR images with expert tracing results. Several cortical geometric measurements, such as the cortical surface shape and cortical thickness, can be derived directly from the coupled surfaces algorithm. Sulcal ribbons, which are distinctive cortical landmarks, can be efficiently extracted using a 3D method derived from the coupled surfaces algorithm. The extraction and measurement of these cortical and sulcal geometric features made possible or made easy by our work are offering new understanding of neuroanatomy, and have the potential of facilitating new studies of brain structure and function.

Chapter 2

Related Work

In this chapter, we review image segmentation techniques that are relevant to our work, with an emphasis on 3D approaches. Despite their importance in medical image analysis, full volumetric approaches are still underdeveloped compared to those in 2D.

Detailed discussions of traditional low-level image processing techniques, such as thresholding, region growing, edge detection and mathematical morphology operations, are not provided here since they can be found in a number of textbooks [31, 70]. Overall, these low-level techniques have considerable limitations when applied to medical image segmentation, largely due to their under-constrained nature. Excellent reviews of recent development of image segmentation techniques can be found in [55, 62].

2.1 Region-Based Methods

Region-based methods exploit homogeneity in images. They primarily depend on the underlying consistency of any relevant feature in different regions, e.g. gray level values and texture properties etc, to assign individual voxels into classes.

2.1.1 MRF-based approach

Markov Random Fields (MRF) are a natural and powerful way of incorporating spatial correlations into a segmentation process. Following the work by Geman & Geman [25], MRF-based methods have been widely used for image restoration and segmentation. The underlying assumption is that for a range of degradation, the posterior distribution of an image is an MRF with a structure akin to the image model. The task becomes to compute the maximum *a posteriori* estimate of the original image given a degraded realization. Geman & Geman's work draws an analogy between images and statistical mechanics systems. The assignment of an energy function in the physical system determines its Gibbs distribution. Because of the Markov-Gibbs equivalence established by the Hammersley-Clifford Theorem, the assigned energy function determines the MRF image model. This energy function is a more convenient and natural mechanism for embodying image attributes than are the local characteristics of the MRF. Now by gradual temperature deduction in the physical system, low energy states, i.e. the most probable states under the Gibbs distribution, are recovered. The resulting method is a highly parallel "relaxation" algorithm for MAP estimation with the convergence to a global maximum.

When applied to image segmentation, MRF-based methods label each pixel or voxel with different classes [11, 12, 32, 41, 42, 51, 98]. There are two things that need to be estimated simultaneously for this type of application: the parameters of the underlying Gibbs Random Field, and the labels for each pixel or voxel. Lakshmanan and Derin [41] developed an adaptive algorithm which recursively yielded segmentations and parameter estimates that converged to the MAP segmentation, and the ML estimates of the Gibbsian model parameters.

Leahy *et al.* applied MRF's as a statistical image model to medical imaging. In

medical images, with the exception of the boundary between different anatomical regions, the image intensity tends to vary slowly. MRFs well capture this property with a suitable choice of the Gibbs energy function to model the high degree of local correlation between neighboring voxels in regions of a single tissue type.

MRF models have been extended into 3D for volumetric image segmentation [12, 32, 98]. The 3D MRF models better capture the spatial correlation, and are demonstrated to outperform the 2D models. When applied to brain data [32], 3D MRF-based segmentation algorithm was able to capture features that are of special importance for MR images, i.e., distributions of tissue intensities, neighborhood correlations, and signal inhomogeneities.

2.1.2 Mixture models and the EM algorithm

Finite mixture models have been very popular recently especially in MR brain image segmentation [34, 43, 93, 96], because they are capable of modeling the volume averaging effect in medical imaging, i.e., the voxel intensity reflects contributions from different types of tissues that make up the voxel. Instead of assigning each voxel into a single class, mixture models label each voxel with several classes probabilistically, thus achieving a "soft" segmentation. The Expectation-Maximization (EM) algorithm is widely used to approximate ML estimates for such incomplete data problems [56, 68]. Liang *et al.* [43] also employed the information criteria of Akaike (AIC) and Schwarz & Rissanen (MDL) to determine the number of classes in the image.

Overall, the advantage of region-based methods lies in their relatively insensitivity to noise [10]. However, they typically require further processing to group segmented regions into coherent structures. Moreover, quantitative measurement of geometric features of objects of interest does not follow immediately.

2.2 Deformable Model-based Methods

Deformable models have been a very popular technique in medical image analysis in recent years. Combining geometry, physics and approximation theory, they overcome the local limitation of traditional low-level image processing techniques, and offer a unique and powerful approach to segmentation problems [55]. Geometry serves to represent object shape, physics imposes constraints on how the shape may vary over space and time, and optimal approximation theory provides the underlying mechanism for fitting the model to image data. In addition, their compact representation of objects segmented greatly facilitates the quantification of measurements which are of importance in medical imaging applications.

2.2.1 Snakes, Balloons and T-snakes

One of the most commonly used deformable model-based methods is the snakes approach due to Kass *et al.* [36]. A snake is an energy-minimizing spline guided by internal forces governing the smoothness of the spline, and external forces that attract the spline onto nearby lines and edges. The minimization process can be achieved by using a variational method that involves solving the corresponding Euler equation. One concern regarding this method is that a close initialization has to be provided in order to obtain good final results. Otherwise, the spline has the tendency to shrink.

A balloon model with a pressure force outward was then introduced as a way to generalize and solve some of the problems encountered with the above snake method [15]. This internal "inflation" force expands a snakes model past spurious edges towards the real edges of a structure, making the snake less sensitive to initial conditions. Extensions to 3D were made, and finite-element methods were adopted to solve the model yielding

greater stability and fast convergence.

Others have also made efforts to offer solutions to the problems associated with the original snake methods. Radeva *et al.* [64] used both the magnitude and direction of image gradient in the attracting force. As a result, the snake has a selective behavior according to edge orientation.

However, despite these improvements, the snake method still has the problem of having too many parameters to optimize over. The internal energy constraints can limit its geometric flexibility and prevent it from representing shapes with significant protrusions. Furthermore, the topology of the object of interest must be known in advance since the above snake models are parametric and are incapable of topological transformation without additional machinery.

To improve over the above stated problems, McInerney and Terzopoulos developed topologically adaptable snakes (T-snakes) [53], and further extended it to 3D [54]. Embedded in an Affine Cell Decomposition (ACD) framework, their model maintains the traditional parametric physics-based formulation, and handles complex geometry and topologies. The computational load is very heavy for 3D T-snakes, however it can potentially be alleviated on parallel machines.

2.2.2 Fourier surface

Staib and Duncan [81] described a 3D surface model with a Fourier parameterization for deformable boundary finding. The model decomposes surfaces into a weighted sum of sinusoidal basis functions. Since the higher-index basis functions in the sum represent higher spatial variations, the series can be truncated to allow relatively smooth surfaces to be described with a small set of parameters. Surface finding can then be formulated as an optimization problem using gradient ascent to fit surfaces to places with high image

gradient. One concern about this model lies in its limitation in capturing convoluted surfaces.

2.2.3 Free-Form deformation

Free-Form Deformation (FFD) was a generalization by Sederberg and Parry [75], following the work of Barr [5] on deformation by creating operators for stretching, twisting, bending and tapering surfaces around a central axis. This method imposes an initial deformation lattice on a parallelepiped, and defines the deformable space as the trivariate splines volume defined by the lattice points. The parallelepiped form of the lattice allows points of an embedded object to be quickly parameterized in the space of the lattice, and as the lattice is deformed, the deformed points can be calculated by simple substitution into the defining equations of the trivariate volume. This method is widely used because of its power to create many types of deformations with little user-interaction. Recently, the FFD model is starting to show promise in applications to medical imaging. Bardinet *et al.* used FFDs to reconstruct the myocardium of the Left Ventricle [4].

2.2.4 Level Set Method

While the above listed deformable models have been widely used in various medical imaging applications, they have severe limitations: they are unable to handle complex geometry and changing topology without extra machinery and heavy computational load. To circumvent this, a new family of deformable models based on the level set techniques of Osher & Sethian [77] have recently been introduced [47, 48, 49, 60, 77, 78].

Level set methods are powerful numerical techniques concerning the motion of a closed interface $\gamma(t)$ propagating along its normal direction. This closed interface $\gamma(t) : [0, \infty) \rightarrow R^N$, can either be a curve in 2D space ($N = 2$) or a surface in 3D space ($N = 3$).

The essential idea is to represent the front $\gamma(t)$ as the zero level set of a high dimensional function Ψ. Let $\Psi(\mathbf{x}, t = 0)$, where $\mathbf{x} \in R^N$, be defined by

$$\Psi(\mathbf{x}, t = 0) = d, \tag{2.1}$$

where d is the signed distance from position x to $\gamma(0)$, and plus(minus) sign is chosen if the point x is outside(inside) the initial front $\gamma(0)$. Thus we have an initial function $\Psi(\mathbf{x}, t = 0) : R^N \to R$ with the property that

$$\gamma(t = 0) = (\mathbf{x} \mid \Psi(\mathbf{x}, t = 0) = 0) \tag{2.2}$$

The goal now is to produce an equation for the evolving function $\Psi(\mathbf{x}, t)$ so that Ψ always remains zero on the propagating interface. Let $\mathbf{x}(t), t \in [0, \infty)$ be the path of a point on the propagation front. That is, $\mathbf{x}(t = 0)$ is a point on the initial front $\gamma(t = 0)$, and $\mathbf{x}_t = F(\mathbf{x}(t))$ with the vector $\mathbf{x}_t$ normal to the front at $\mathbf{x}(t)$. Since the evolving function Ψ is always zero on the propagating front, we have

$$\Psi(\mathbf{x}(t), t) = 0. \tag{2.3}$$

By the chain rule, we have

$$\Psi_t + \sum_{i=1}^{N} \Psi_{x_i} x_{i_t} = 0 \tag{2.4}$$

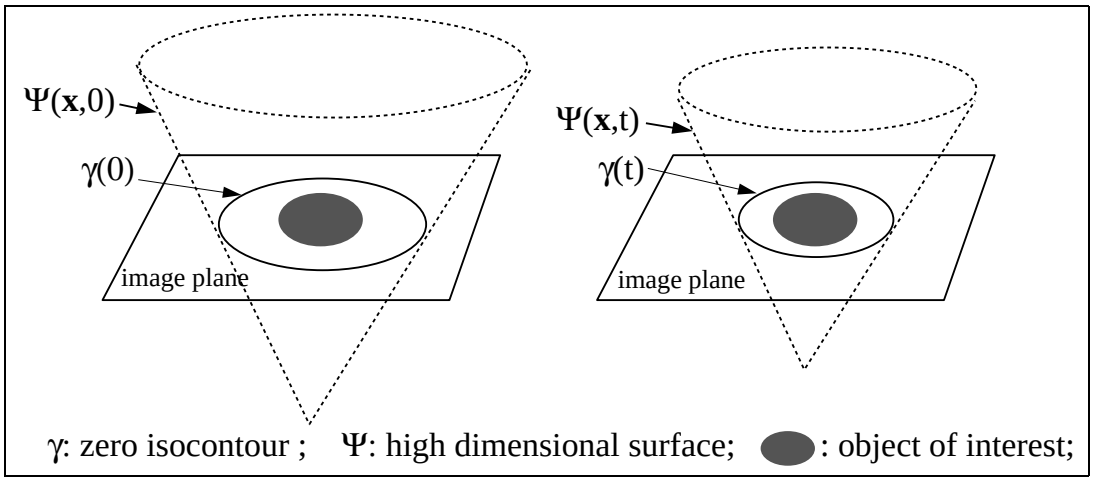

Figure 2.1: Example of 2D curve propagation with the level set method. In this case, a curve contracts to capture the oval object of interest on the image plane. Notice the evolution of the signed distance function Ψ, and the resulting contraction of its zero level set γ.

where x_i is the ith component of $\mathbf{x}$. Since

$$\sum_{i=1}^{N} \Psi_{x_i} x_{i_t} = (\Psi_{x_1}, \Psi_{x_2}, ..., \Psi_{x_N}) \cdot (x_{1_t}, x_{2_t}, ..., x_{N_t}) = F(\mathbf{x}(t)) \mid \nabla \psi \mid, \tag{2.5}$$

the evolution equation for ψ is then:

$$\Psi_t + F \mid \nabla \Psi \mid = 0 \tag{2.6}$$

Figure 2.1 shows a diagram of curve propagation with the level set method. F can be a function of the front characteristics (such as the curvature, normal direction

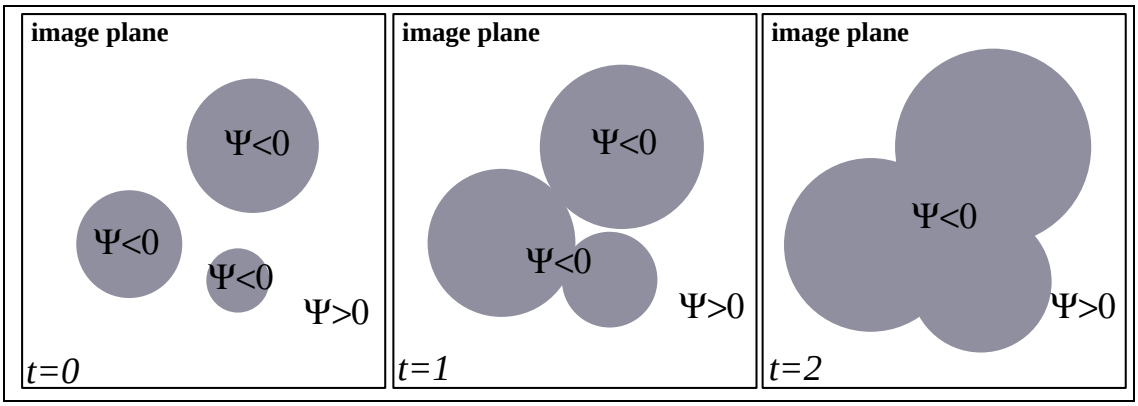

Figure 2.2: Example of three circles propagating with constant speed along normal direction. At $t = 0$, Ψ takes minus sign within the circles in the shaded area, and plus sign outside the circles. Thus the circles are the zero level set of $\Psi(t = 0)$. As time goes on, the shaded areas where Ψ takes negative values start to merge. As a result, the zero level set of Ψ, which corresponds to the current front, changes topology.

etc.) and the image characteristics (e.g. gray level and gradient etc.). When applied to segmentation problems, the speed term F of the propagation front is designed to be dependent on image information, and stop at desired places by having a zero value there. Since the level set method concerns the motion of closed fronts, when used to extract open curves or surfaces in an image, problems of extracting closed fronts need to constructed where the desired open curves and surfaces are part of the closed fronts. This can usually be done by adding artificial boundaries at the opening of the desired curves or surfaces.

The major advantages of using this method over other active contour strategies include the following. First, although the evolving level function $\Psi(\mathbf{x}, t)$ remains a function, the embedded propagating front $\gamma(t)$ may change topology, break, merge and form sharp corners as Ψ evolves. This ability of handling flexible topology change and complex

geometry is taken care of without additional machinery as in traditional parameterization models such as the splines. As a result, initializing curves or surfaces for segmentation in an image do not need to be placed close to the boundary of interest and in similar topological form, thus relieving user-interaction in the stage of initialization. Figure 2.2 illustrates the merging of three expanding circles with the level set method.

The second advantage of this Eulerian formulation concerns numerical approximation. Because $\Psi(\mathbf{x}, t)$ remains a function as it evolves, we may use a discrete grid in the domain of $\mathbf{x}$ and substitute finite difference approximations for the spatial and temporal derivatives [77]. For example, using the image grid with nodes ij and spacing h, and employing the standard notation that Ψ_{ij}^n is the approximation to the solution $\Psi(ih, jh, n\triangle t)$, where $\triangle t$ is the time step, we can write:

$$\frac{\Psi_{ij}^{n+1} - \Psi_{ij}^n}{\triangle t} + F(\nabla_{ij}\Psi_{ij}^n) = 0. \tag{2.7}$$

Third, the intrinsic geometric properties of the front may be easily determined from Ψ. For example, at any point of the front, the normal vector is given by $\vec{n} = \nabla\Psi$.

Finally, There are no significant differences in handling fronts in 3D space and fronts in 2D space. Surface propagation works in similar fashion as curve propagation, only differing in that Ψ changes to a function in 4D.

This original formulation, as proposed by Malladi *et al.* [77] is not derived from an energy minimization point of view. When applied to image segmentation, it requires careful design of the stopping criterion for the propagating fronts.

A modified approach, which considers the deformation of a level set as a gradient flow to a state of minimal energy [47, 78] performs better provided that objects can clearly

be identified from their boundaries [60]. This approach elegantly connects energy and curve evolution approaches of active contours. However, it has drawbacks in its sensitivity to initializations due to the possibility of being trapped at a local minimum of energy.

2.3 Incorporating a Shape Prior

In medical image analysis, the use of implicit or explicit anatomical knowledge to guide shape recovery is of great importance for automatic and robust image interpretation. There have been a number of efforts aimed at the integration of shape information into deformable models as shown below. Other efforts can be found in [16, 34].

2.3.1 Model-based Snake

The first approach uses local shape information integrated with the original snake method. To relieve the initialization problem, Gunn and Nixon [30] used a dual active contour model, with one contour expanding from inside the target feature and the other contracting from the outside, to improve parameterization. A local shape model with constraints on contour curvature was imposed to favor circular shape.

A similar effort was made by Radeva *et al.* [64] by using internal forces that included structural information about the expected snake shape. In particular, a term reflecting the difference between the derivatives of the current active contour and the model contour was added to the internal force. As a result, changes from the model shape were penalized, and the snake model had a tendency to keep the model shape and resist shrinkage.

Though an improvement over the original snake model to some extent, these methods are still limited to handle only simple shapes. In addition, extension to 3D is

not trivial.

2.3.2 Deformable shape templates

Deformable shape templates are another form of using *a priori* knowledge to constrain the shape variations in objects of interest.

Superquadrics are examples that are often used for the recovery of compact volumetric models [79, 87, 90]. Superquadrics allow a small number of intuitive global shape parameters to govern the gross shape of a target anatomic structure. Furthermore, the global parameters can often be coupled with local shape parameterizations such as the splines to reconstruct the fine details of complex shapes.

Tagare [84] proposed a formulation of a template matching algorithm using orthogonal curves. This method used a lower-dimensional search space than the conventional methods by precomputing extensions of the deformable template along the orthogonal curves. The reduction in search space allowed the use of dynamic programming to obtain globally optimal solutions and reduce the sensitivity of the algorithm to initial placement of the template. Although successfully applied to 2D case, this model has yet to be extended to solve 3D problems.

2.3.3 Using statistical shape priors

Other researchers use deformable models in a probabilistic framework to include statistical shape information, i.e. the prior probability distribution on shape variables.

Staib and Duncan [80] used elliptic Fourier descriptors as model parameters to represent open and closed boundaries. *A priori* shape information was included as a spatial probability expressed through the likelihood of each model parameter. Thus, the Fourier coefficients served to bias toward a range of shapes favored by the shape prior

distribution. A Bayesian approach [57] was used to solve the estimation.

Cootes *et al.* [18] described a method of building models of image objects from a training set of example images. Each model consisted of a flexible shape template describing how important points of the object could vary, and a statistical model of the expected gray levels in regions around each model point. The deformations were modeled using the eigenvectors of the variation from the mean shape, to allow deformation reflecting the variations obtained in the training set.

Szekely *et al.* [83] proposed a segmentation technique combining shape representation by Fourier parameterization and modeling of natural shape variability. Their flexible parametric shape models were represented by a parameter vector describing the mean contour and by a set of eigenmodes of the parameters characterizing the shape variation. Elastic fit of the mean model in the subspace of eigenmodes restricted possible deformations, and found an optimal match between the model surface and boundary candidates.

Wang and Staib [91] exploited the similar idea of investigating the statistical variation of object boundary points in a training set, but used it in a more sophisticated way. A Bayesian formulation, based on this prior knowledge and the edge information of the input image, was employed to find the object boundary. The Bayesian framework allowed the model to adjust the weighting between the statistical shape model and the image information based on the image quality and the reliability of the training set. The structure of interest was delineated and the spatial correspondence of these points to the model was established when the *a posteriori* probability was maximized using conjugate gradient optimization.

The deformable models listed above have already shown the power of statistical

shape priors in image segmentation. However, extensions to 3D for recovering complex shapes are yet to be fully realized, due to difficulties in 3D parameterization, volumetric model building and computational requirements.

2.4 Integrated Methods

Since different segmentation techniques have different advantages when applied to images, it is only natural to consider an integration of methods to achieve more robust and accurate results. Some effort has been made toward this goal in recent years. Many of these efforts are customized for a specific task either by using a pipeline of different segmentation techniques [35, 74], or by optimizing a single global objective function that combines different objectives [10]. Others aim at a more elegant solution.

One piece of notable work was done by Zhu and Yuille on region competition with a statistical and variational approach [104]. This algorithm was derived by minimizing a generalized Bayes/MDL criterion using the variational principle. The algorithm combined aspects of snakes/balloons and region growing, and was guaranteed to converge to a local minimum. Indeed the classic snakes/balloons and region growing algorithm could be directly derived from the region competition approach. Theoretical analysis of accuracy of boundary location, criteria for initial conditions, and the relationship to edge detection using filters were also provided.

Other effort includes that of Bozma and Duncan [9]. Image analysis systems aimed at robust segmentation often require integration of a variety of modules. Concerning the importance of adhering to the multiple coexisting nature of the objectives of different modules, they developed a segmentation system aimed at module integration using a game-theoretic approach. The existence of a Nash equilibrium which corresponds

to a most rational segmentation was investigated. Chakraborty [10] further developed the system to integrate region-based segmentation and boundary finding with Fourier parameterization.

The search for an optimal integration of information is far from complete. Each of the various region-based methods and deformable models listed above can be used as a candidate component in the integrated segmentation system. The selection of such components depends heavily on the nature of a particular segmentation problem. And the choice of integration technique can be greatly influenced by both the region-based and boundary-based components, and the nature of shape information of the particular anatomy of interest.

In solving challenging 3D problems in segmenting complex medical images, we expect that models which use region-based and boundary-based information, as well as incorporate anatomical knowledge, are required. As described in detail in the following chapters, the work in this thesis makes an initial effort in combining shape information and region-based information using a level set approach.

Chapter 3

Volumetric Layer Segmentation with Coupled Surfaces Propagation

3.1 Motivation and Model

In medical image analysis, we often encounter the problem of segmenting a volumetric layer of finite thickness. For example, in 3D cardiac image analysis, the myocardium of the left ventricle is a thick structure bounded by the endocardial and epicardial walls (see Figure 1.1(a)). In the analysis of neuroanatomical structures from 3D MR images, the cortex, which is the outermost layer of gray matter in the brain (see Figure 1.1(b)), is bounded by the outer cortical surface (Cerebral Spinal Fluid(CSF)/gray matter interface) and the inner cortical surface (gray/white matter interface). The two surfaces which bound the layer can be viewed as coupled in several ways, ranging from loose coupling (e.g. the endocardial and the epicardial boundaries) to tight coupling where the thickness

of the layer is nearly constant(e.g. the outer and the inner cortical surfaces).

This nature of the coupling between the inner and outer bounding surfaces is characteristic to the volumetric layer. If incorporated into the segmentation process, this coupling would greatly enhance accuracy and robustness. More specifically, because of the limitations of the imaging technique used and the volume averaging effect, it is often observed that in some regions, there is not enough information from the image data to clearly define either the outer or the inner bounding surface. When applying a single surface approach, we may very well end up with error in such a region. While through the coupling, information on one surface is available to the partner surface, thus improving boundary finding.

The volumetric layer is bounded by two surfaces. Across each surface, there is a local difference in the gray scale values, while in between the two surfaces there is a homogeneity of certain voxel statistics. The layer is defined completely by its bounding surfaces and the homogeneity in between. Based on this definition, we propose a new approach of coupled surfaces propagation, which takes into account the coupling information through the distance between the two surfaces. By evolving two embedded surfaces simultaneously, each driven by its own image-derived information while maintaining the coupling, we are able to achieve a more robust segmentation on this special set of volumes. The information used to drive the surface propagation should distinguish the inner and outer boundary features, and reward the homogeneity inside the layer.

Take for example the MR brain images. Due to volume averaging, in some regions the boundary between white matter and gray matter is not well shown, while the CSF appears clearly. The single surface approach may hence have the inner cortical surface collapse into CSF. However, with the coupled surfaces approach, we can maintain some

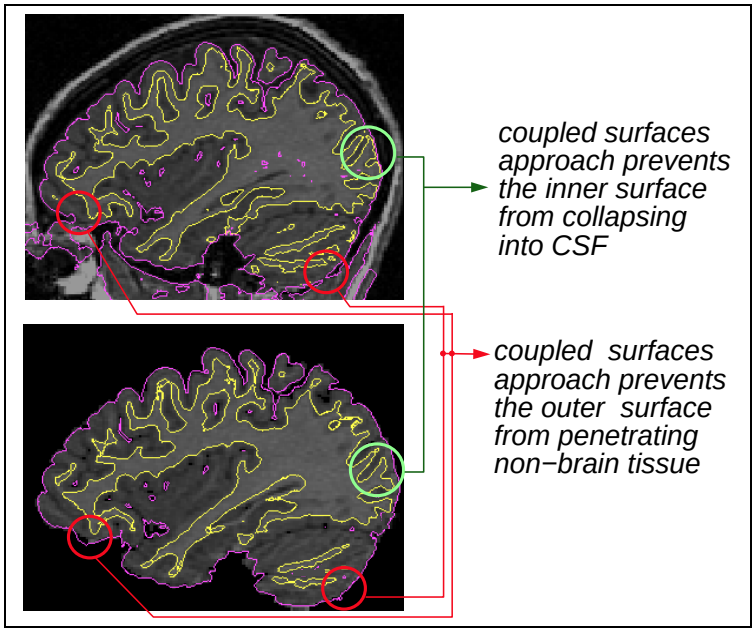

Figure 3.1: Single vs. coupled surfaces approach on cortex segmentation. Top: cut view of surfaces resulting from the single surface approach shown on a sagittal slice of the original image (finding the inner and outer cortical surfaces separately); bottom: cut view of surfaces resulting from the coupled surfaces approach shown on the same sagittal slice of the skull-stripped data (skull-stripping was done by our expert with manual tracing). Both the single surface approach and the coupled surfaces approach were run on the original 3D image. The result from the coupled surfaces approach is shown with the skull-stripped brain to demonstrate that the outer cortical surface resulting from the algorithm nicely fits the boundary from expert tracing.

minimal distance between the inner cortical surface and CSF, thus preventing the inner cortical surface from going into CSF. There are also places where structures such as the eye sockets obscure the CSF in the image. With the coupled surfaces approach, the white/gray matter boundary information is then used to stop the propagation of the outer cortical surface before it penetrates non-brain structures. Figure 3.1 shows examples of the above mentioned cases where the coupled surfaces approach outperforms the single surface approach, and illustrates the motivation of using the coupled surfaces approach.

3.2 Coupled Surfaces Propagation with Level Set Method

We chose the level set method as the surface propagation mechanism of our coupled surfaces approach for a number of reasons: its ability to handle complex geometry and topological changes, its computational efficiency, its natural correspondence between the bounding surfaces, and its easy measurements of volume and many geometric features of the layer.

In our coupled surfaces formulation, embedding each surface as the zero level set in its own level function, we have two equations:

$$\begin{aligned} \Psi_{in_t} + F_{in} \mid \nabla \Psi_{in} \mid &= 0 \\ \Psi_{out_t} + F_{out} \mid \nabla \Psi_{out} \mid &= 0 \end{aligned} \tag{3.1}$$

The coupling of the two surfaces is embedded in the design of F_{in} and F_{out}. In particular, the speed of the outer surface will depend on information of the current position of the inner surface, and vice versa. To ensure that the coupled surfaces stop at

desired places in images to perform satisfactory segmentation of volumetric layers, F_{in} and F_{out} also depend heavily on characteristic image features. In Chapter 4, we propose the use of a local operator to capture inner and outer boundary features, and reward the homogeneity inside the layer. More discussions on speed term design can also be found in Chapter 4.

Here, we first discuss the details of algorithm construction with the level set method for coupled surfaces propagation. The above equations are discretized on the image grid (see Equation 2.7), and solved iteratively.

3.2.1 Initialization

The ease of initialization has always been a desirable feature of any segmentation algorithm. Since the level set method handles the change of topology, our initializing surfaces do not need to be placed close to the boundaries of interest and in similar topological form. We thus use concentric spheres as initialization for the inner and outer surfaces. Such concentric spheres can be specified with simple mouse clicks on images. We choose to place the concentric spheres inside the inner bounding surface, and have them propagate outward to capture the volumetric layer. In this way, with the signed distance function Ψ_{in} and Ψ_{out}, we can easily track the inside and outside of the volumetric layer. Examples of concentric spheres as initialization are shown in Figure 4.4. Studies of the robustness to initialization will be discussed in Chapter 4 as well.

3.2.2 Marching Cubes algorithm for iso-surface construction

In order to construct the current position of the propagating surfaces, we use the Marching Cubes algorithm to get the zero level set of Ψ functions. The Marching Cubes algorithm was designed by Lorensen and Cline to extract surface information from a 3D field of

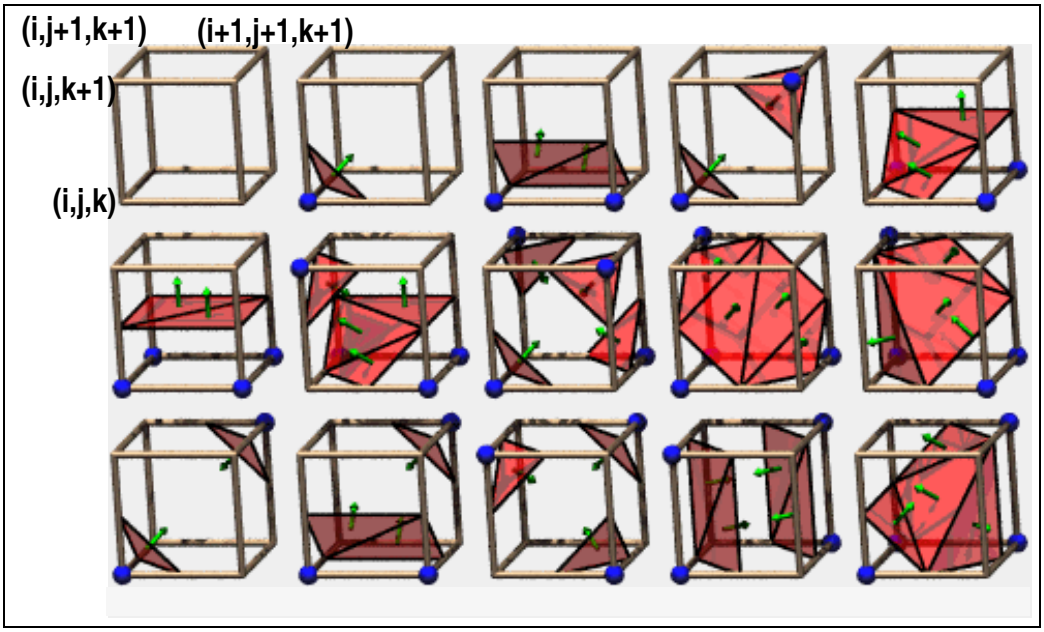

Figure 3.2: The 15 cube combinations used in the Marching Cubes algorithm (Courtesy of J. Sharman, http://exaflop.org/docs/marchcubes/). In our case of getting iso-zero level set of Ψ functions, each cube unit has 8 corners at positions $(i, j, k), (i + 1, j, k), (i, j + 1, k), (i + 1, j + 1, k), (i, j, k + 1), (i + 1, j, k + 1), (i, j + 1, k + 1)$, and $(i + 1, j + 1, k + 1)$ on the 3D image grid. The blue spheres denote corners that have tested as inside the zero level set, and the green arrows denote the surface normals of the relevant triangles of zero Ψ value.

values [14]. In our case, the 3D field of values are Ψ_{in} and Ψ_{out} for the inner and outer surfaces respectively.

The basic principle behind the Marching Cubes algorithm is to subdivide space into a series of small cubes. In our case of getting the iso-zero level set of function Ψ, each cube unit has 8 corners at positions $(i,j,k), (i+1,j,k), (i,j+1,k), (i+1,j+1,k), (i,j,k+1), (i+1,j,k+1), (i,j+1,k+1), (i+1,j+1,k+1)$ on the 3D image grid. The algorithm then marches through each of the cubes testing the Ψ value of corner points, and records an appropriate set of polygons at the interpolated iso-zero intersection.

With 8 corners for each cube unit, there are potentially 256 possible combinations of corner status of either being inside or outside the zero level set surface. This complexity can be reduced to 15 cases (see Figure 3.2) by taking into account the duplication caused by reflection and rotational symmetry.

As illustrated in Figure 3.2, the result of the current zero level set of function Ψ constructed using the Marching Cubes algorithm is thus a set of triangles in 3D space. In the following discussion, we refer to the vertices of such triangles as points on the constructed iso-zero surface.

3.2.3 Narrow band implementation

The updating of level set function Ψ on the entire image grid at each time step is computationally very expensive. Upon careful examination, it is observed that to move the front, it is only necessary to update Ψ at a small set of points in the neighborhood of the zero level set. Algorithms can thus be implemented using a narrow band method [77], which modifies the level set method so that it only explicitly updates the points close to the current propagating fronts. Our implementation of the narrow band method uses this idea, but is designed specifically for coupled level sets so that the distance between the

two embedded surfaces (necessary for being used as a thickness constraint) is available with no further computation after the narrow band rebuilding.

Based on the fact that any point b in the narrow band of the current surface should be within some neighborhood of a certain point a on the current surface, the narrow band is constructed dynamically in the 3D neighborhood of each point on the current surface by including points that lie within a certain distance range (*i.e.* bandwidth) away from that particular point. Also, since a point b in the narrow band can be within the neighborhood of several points $a_1, ..., a_l$ on the current surface, we update the value of the level function Ψ at b to be

$$sign(\Psi(b)) \cdot (Min_{i=1,...,l} dist(b, a_i))$$

where function $dist$ gives the positive Euclidean distance.

The steps for rebuilding the narrow band and reinitializing Ψ inside the band at each time step are as follows:

```
for every point a on the current front {
    for every point b in the neighborhood of a {
        if b is not already in the narrow band,
           then add b to the narrow band;
        if dist(b,a) is less than the absolute value of the current Ψ(b),
           then update Ψ(b) to be sign(Ψ(b)) · dist(b,a);
        }
}
```

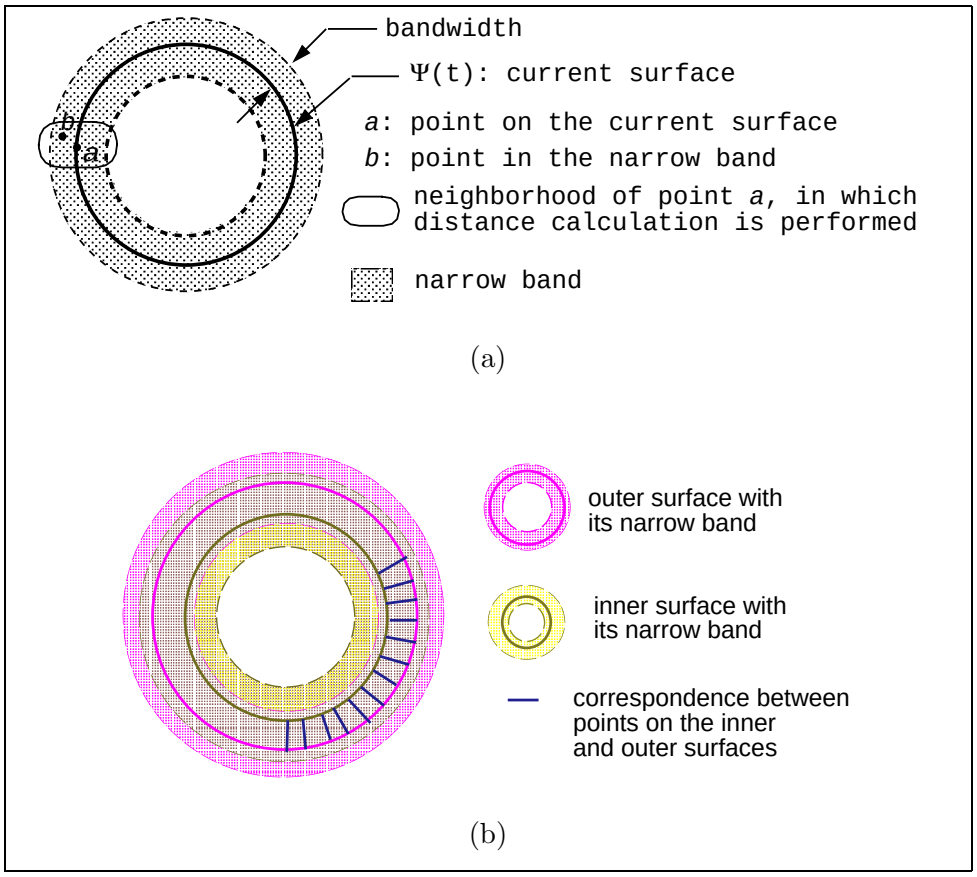

Figure 3.3: Schematic of the narrow band implementation for the case of 2D curves (same argument holds in 3D). (a): dynamic construction of the narrow band and the update of the level function Ψ within are performed in the neighborhood of the current surface. (b): inner and outer surfaces with their narrow bands. Notice the inner surface lies within the narrow band of the outer surface, and vice versa.

The size of the neighborhood depends on the allowed bandwidth, and therefore is fixed. Thus, for a surface with N points (as mentioned before, points here refer to the triangular vertices from the Marching Cubes algorithm), the construction of its narrow band and the update of Ψ in the narrow band is an $\mathcal{O}(\mathcal{N})$ calculation. We constrain the neighborhood to be a cylinder-shaped area centered at points on the propagation surface, with the axis of the cylinder pointing in the surface normal direction (for a schematic drawing, see Figure 3.3(a)). To further improve efficiency, we also pre-calculate and store the set of possible offsets to surface points in the cylinder-shaped neighborhood depending on its axial direction. Thus when rebuilding a narrow band, we loop over every point on the current surface, and use the pre-calculated offsets depending on the surface norm to visit through its neighborhood.

In our application, two different narrow bands are computed for the inner and outer interfaces, Ψ_{in} and Ψ_{out}, respectively. As shown in Figure 3.3(b), to ensure that the distance-based correspondence between the coupled surfaces falls out automatically, the two bandwidth ranges (for the inner and outer narrow bands separately) are chosen such that the inner surface lies within the narrow band of the outer surface and vice versa.

3.2.4 Extending speed function

Based on the fact that the speed terms are designed to force the propagating level set to stop at the desired boundary, the image dependent speed terms have meaning only on the front, *i.e.* the zero level set of Ψ. However the level set equation of motion is written for the function Ψ defined over the entire image grid. We thus extend the speed terms from the zero level set to the whole narrow band at each iteration as described by [77], *i.e.* point b takes on the speed of point a which is the closest point to b and lies on the

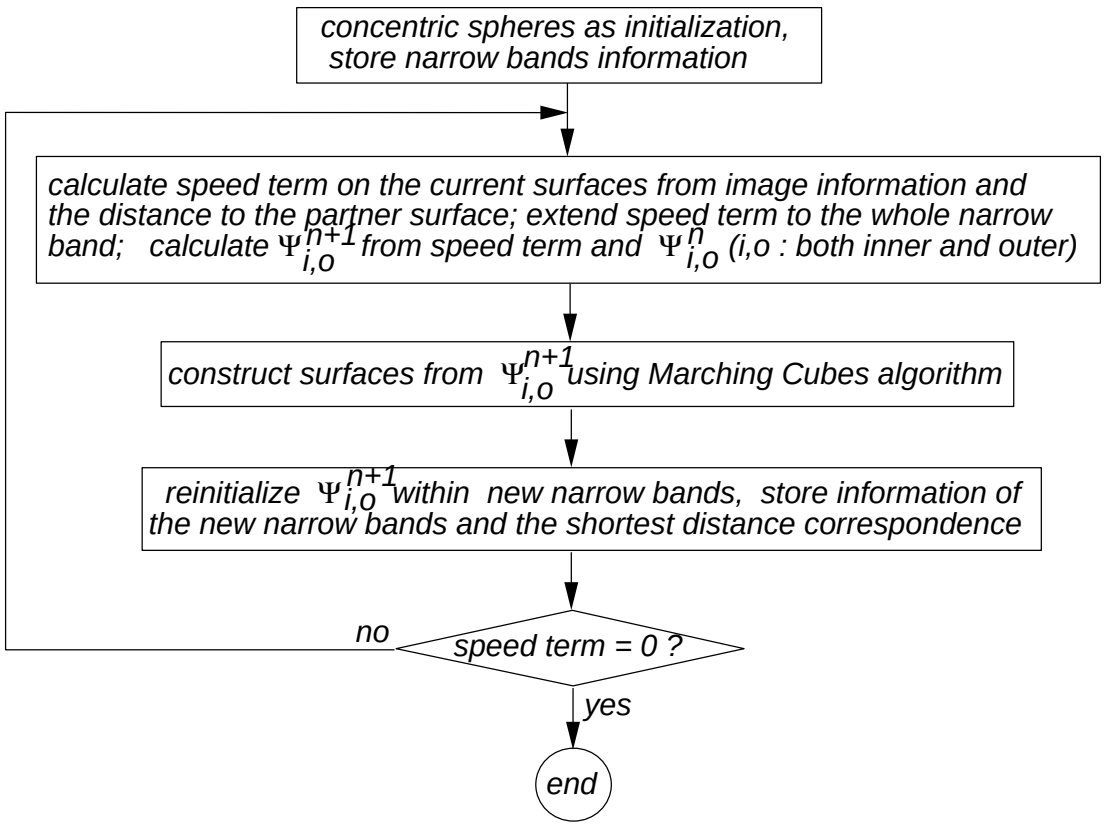

Figure 3.4: Diagram of the coupled surfaces algorithm

zero level set. The correspondence between any point b in the narrowband and its closest point a on the current surface is determined during the reconstruction of the narrowband described above.

3.2.5 Summary of the coupled surfaces algorithm

To summarize, the algorithm diagram is shown in Figure 3.4.

We begin with concentric spheres as initialization for the coupled surfaces. Ψ_{in}

and Ψ_{out} at $t = 0$ are calculated as the narrow bands around the inner and outer surfaces are constructed. In the process of building the narrow bands, the shortest distance correspondence between the points in the narrow bands and the points on the surfaces are determined.

We then calculate the speed term for all the points on the current inner and outer surfaces, and extend speeds to the whole narrow bands using the stored shortest distance correspondence. Ψ_{in} and Ψ_{out} are then updated in their own narrow bands using the discretized version of level set Equation 2.7.

The new position of inner and outer surfaces are then constructed from the updated Ψ_{in} and Ψ_{out} using the Marching Cubes algorithm. New narrow bands for the inner and outer surface are rebuilt respectively, and Ψ_{in} and Ψ_{out} are reinitialized simultaneously within the narrow bands.

We repeat the above process of calculating speed term, updating level functions within the narrow bands, constructing the new position of propagating surfaces, and rebuilding narrow bands. The whole process stops when the speed terms for both the inner and outer surfaces reach a zero value everywhere.

Chapter 4

Application to Cortex Segmentation

4.1 Introduction

In this chapter, we apply the coupled surfaces algorithm to the problem of segmenting cerebral cortex, a focus of a number of recent neuroanatomy studies. As the outer most aspect of the brain, the cerebral cortex is responsible for human cognitive function, and abnormalities of the cortex are believed to be associated with a number of neurological disorders, such as schizophrenia and Alzheimer's. Despite recent advances in neuroimaging technology that facilitate the visualization of detailed cortical features, accurate quantitative measurement has yet to be fully realized, largely due to the lack of a reliable and efficient means of quantifying information from brain images. Segmentation is the first step toward achieving such goals.

The cerebral cortex is characterized by its convoluted surface (see Figure 1.1(b)). The narrow groove between adjacent convolutions is called a sulcus, and the ridge between

two sulci is call a gyrus. Due to its convoluted nature, the segmentation of the cortex must be considered in 3D. For example, although the cerebral cortical layer is nearly $3mm$ thick [7] everywhere on the cortex, an oblique 2D slice that happens to be approximately parallel to a particular sulcus will give the appearance of a much thicker structure. Also, a slice can show several disconnected regions of cortical gray matter despite the fact that the cortical gray matter is all connected in 3D. Only by going through the neighboring slices can we get complete information to perform segmentation. Slice by slice manual tracing of the cortex is extremely tedious and labor intensive, hence automatic, reliable and relatively efficient segmentation which enables automated measurement is a highly desirable goal.

4.2 Related Work

There have been a number of cortical segmentation efforts. The first group to consider are region-based methods, which exploit homogeneity in images. Cline *et al.* used a multi-spectral voxel classification method [14] in conjunction with connectivity to segment the brain into different tissue types from 3D MR images. Liang *et al.* [43] applied a material mixture model. Kapur *et al.* [35] used an MRF model, together with prior encoding of the relative geometry of structures, and gave a probabilistic segmentation of different tissue types in MR brain images. As discussed in Chapter 2, region-based methods need additional machinery for region grouping. In addition, they do not directly provide a compact representation of segmented regions for boundary analysis and incorporating shape information.

Other techniques use deformable models. MacDonald *et al.* presented an iterative algorithm for simultaneous deformation of multiple surfaces with inter-surface proximity

constraints and self-intersection avoidance, where the deformation was formulated as a cost function minimization problem [45, 46]. This method was applied to 3D MR brain data to extract surface models for the skull and the cortical surfaces. This approach takes advantage of the information of the interrelation between the surfaces of interest. However, drawbacks lie in its extremely high computational expense, and the difficulty of tuning weighting factors in the cost function due to the complexity of the problem.

Teo *et al.* [86] used a system that exploited knowledge of cortical anatomy, in which white matter and CSF regions were first segmented. After the connectivity of the white matter was verified in regions of interest, a representation of the gray/CSF surface was created by a constrained growing-out from the gray/white matter boundary. The algorithm was quite complex with a set of rules for connectivity checking. The focus of this work was to create a representation of cortical gray matter for functional MRI visualization. It was argued that the goal of this work was to achieve similar qualitative results to manual segmentation.

Davatzikos *et al.* introduced the concept of a ribbon for modeling the outer cortex in cross-sectional brain images [21] and then extended the model into 3D [20]. A deformable surface algorithm was constructed to find the central layer of the cortex. Based on this parameterization, the cortical structure was characterized through its depth map and curvature map. This method explicitly used the structural information of the cortex. However, close initialization and significant human interaction were needed to force the ribbon into sulcal folds. To compensate for this, Xu *et al.* further extended the method by using a new external force model called gradient vector flow for surface deformation [97], and solved the problem to some extent. However, the method was still subject to the problem of surface self-intersection.

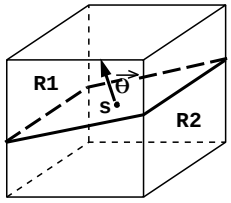

Figure 4.1: A local operator to derive image information.

4.3 Image Information Derivation

To apply the coupled surfaces algorithm, we first need to derive image features that could select the boundary points on both the inner and outer cortical surfaces, and keep them well distinguished.

Medical images consist of a number of different anatomical regions. The homogeneity of each region can usually be characterized by various voxel statistics inside. Thus, by using gradient features (information of gray level difference between neighboring voxels) alone, we actually lose important pieces of information. Here in our approach, instead of using gradient features, we design a local operator which makes use of the gray level information, and gives a measure of the likelihood of a voxel lying on the boundary between tissue A and tissue B. This model can also be extended to make use of a vector of registered parametric images (such as T1, T2 and PD MR images) or images from different modalities.

At each voxel site s, a small neighborhood around s is drawn (see Figure 4.1). Now given a possible boundary with normal direction $\vec{\theta}$, dividing the neighborhood into parts $R1$ and $R2$, the probability that s lies on the boundary between tissue A and tissue

B is:

$$p_{AB}(\vec{\theta}) = p(R1 \in TissueA) \cdot p(R2 \in TissueB) \tag{4.1}$$

Given an estimation $\vec{\theta}^*$ of $\vec{\theta}$, we can use $p(\vec{\theta}^*)$ as a measure of the likelihood that s lies on the boundary between tissue A and tissue B.

One way of estimating $\vec{\theta}^*$ is to first generate the vector $P = [p(\vec{\theta}_1), p(\vec{\theta}_2), ..., p(\vec{\theta}_k)]^T$ where k is the number of possible directions corresponding to the 26 first order neighbors. Then, $\vec{\theta}^*$ is the direction which corresponds to the element in vector P that has the largest magnitude. Here we make the assumption of one single parametric image X, in which voxels belonging to tissue A are independently drawn from a Gaussian distribution $G(\mu_A, \sigma_A)$, and voxels belonging to tissue B are independently drawn from $G(\mu_B, \sigma_B)$. Thus, we have

$$p_{AB}(\vec{\theta}) = \prod_{r \in R1} \frac{1}{\sqrt{2\pi}\sigma_A} e^{-\frac{(X_r - \mu_A)^2}{\sigma_A^2}} \cdot \prod_{t \in R2} \frac{1}{\sqrt{2\pi}\sigma_B} e^{-\frac{(X_t - \mu_B)^2}{\sigma_B^2}} \tag{4.2}$$

In our laboratory, brain images were first processed with MR inhomogeneity correction followed by a histogram normalization procedure [73]. The MR inhomogeneity effect was reduced by correcting using a simple fixed map, which was determined manually by sampling tissue types throughout the field to decide the average inhomogeneity. More sophisticated approach of MR inhomogeneity correction that combines ideas of multi-modality registration criteria, non-rigid registration and models of geometric and intensity distortion in MR image formation is currently being developed [82]. The histogram normalization procedure set the modal white matter and gray matter peaks to

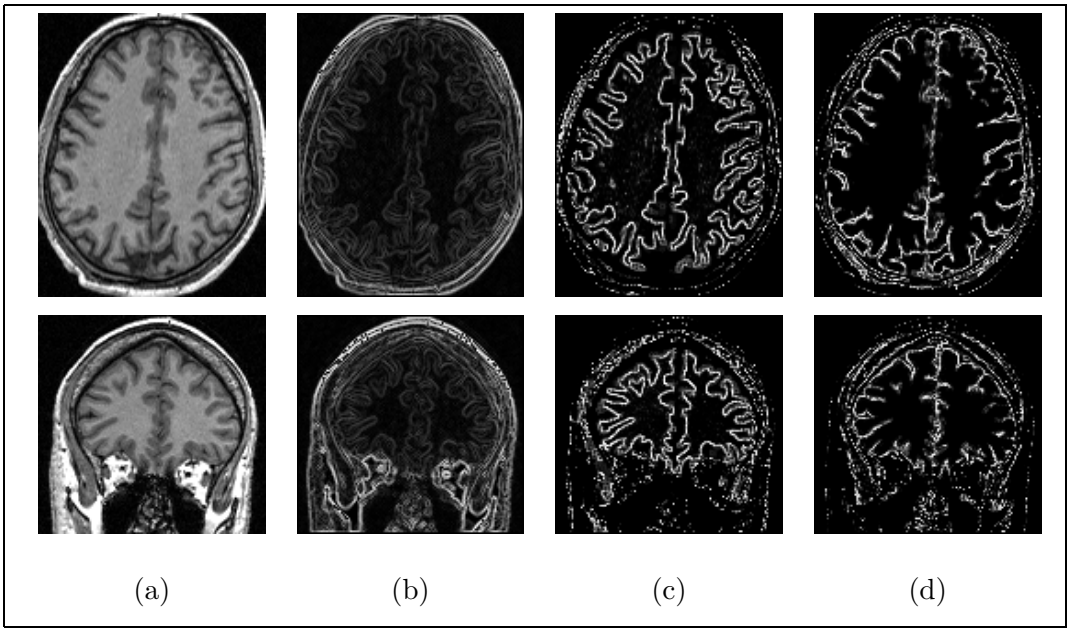

Figure 4.2: Results from our local operator compared to the magnitude of image gradient. (a): axial(top) and coronal(bottom) slices from original 3D brain images; (b): the magnitude of image gradient; (c): $p_{BC}(\vec{\theta}^*)$ from our local operator, B= gray matter, C= white matter; (d)$p_{AB}(\vec{\theta}^*)$, A= CSF, B= gray matter.

predefined gray level values (200 and 300, respectively). This has the benefit of normalizing the relative brightness of the images across subjects, so that difference in pixel intensity cannot contribute to error in manual tracing. Our algorithm starts with such processed data, and takes the fixed distribution parameters as input to the local operator. This gives us the convenience of processing large numbers of images without tuning parameters, however, more elaborate adaptive estimation of tissue distribution parameters would further enhance the algorithm performance.

In our implementation, $R1$ and $R2$ are now set to include one voxel each. A limited expansion to several voxels could potentially further enhance the capability of capturing

homogeneity. In Figure 4.2, we show examples of results from our local operator showing how well it selects the appropriate gray level transition, which is crucial for subsequent processing. Note that more complicated MR image models [25, 41, 42] can be used to calculate $p(\vec{\theta})$.

4.4 Speed Term Design

To apply the coupled surface algorithm to the cortical segmentation problem, we consider two moving interfaces describing the inner and outer cortical bounding surfaces respectively. Starting from inside the inner cortical surface (*i.e.* inside the white matter), with an offset in between (see Figure 4.4), the interfaces propagate along the outward normal direction stopping at the desired place, while maintaining a certain distance between them.

Embedding each surface as the zero level set in its own level function, we have two equations:

$$\Psi_{in_t} + F_{in} \mid \nabla \Psi_{in} \mid \;=\; 0 \tag{4.3}$$

$$\Psi_{out_t} + F_{out} \mid \nabla \Psi_{out} \mid \;=\; 0 \tag{4.4}$$

where F_{in} and F_{out} are functions of the surface normal direction, image-derived information and distance between the two surfaces. The coupling is embedded in the design of F_{in} and F_{out}. At places where the distance between the two surfaces is within a normal range, the two surfaces propagate according to the image-based information. Where the distance between the two surfaces is out of the normal range, the distance imposes a constraint on the propagation of the surfaces.

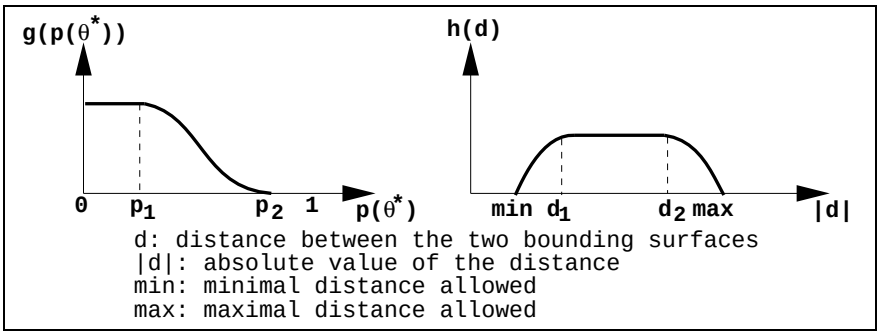

Figure 4.3: Functions g and h used in speed term design.

With the level set implementation, we have a natural way to establish a correspondence between the points on the two evolving surfaces through distance, which is evaluated with little extra computational expense. Recall that the value of the level function of a front at any point is simply the distance from this point to the current front, which is calculated as the shortest distance from this point to all the points on the front [77]. In our case of two moving surfaces, for any point on the inner moving surface, the distance to the outer moving surface is the value Ψ_{out} at this point, and vice versa for the point on the outer moving surface. Hence, we write

$$F_{in} = g(p_{BC}(\vec{\theta}^*))h(\Psi_{out}) \tag{4.5}$$

$$F_{out} = g(p_{AB}(\vec{\theta}^*))h(\Psi_{in}) \tag{4.6}$$

where g and h are the functions whose shapes are as shown in Figure 4.3, and A, B, C denote CSF, gray matter and white matter respectively.

Function g maps larger likelihood to slower speed, *i.e.*, as the likelihood gets

larger, g tends to zero, while as the likelihood gets to near zero, g tends to a constant. Function h penalizes the distance out of the normal range. As the distance goes out of normal range, h goes to zero. Thus, each surface moves with constant speed along its normal direction, and slows down when either the image-based information becomes strong or the distance to the other surface moves away from the normal range. Each surface finally stops when the image-derived information is strong enough or the distance to the other surface is out of the normal range. For smooth speeds, g and h can take form of the following functions with their slopes being Hermite polynomials to have shapes shown in Figure 4.3:

$$g(p) = \begin{cases} 1 & 0 \leq p \leq p_1; \\ 1 - 3(\frac{p-p_1}{p_2-p_1})^2 + 2(\frac{p-p_1}{p_2-p_1})^3 & p_1 < p < p_2; \\ 0 & p_2 \leq p \leq 1 \end{cases}$$

$$h(d) = \begin{cases} (\frac{|d|-d_{min}}{d_1-d_{min}})^2(3 - 2\frac{|d|-d_{min}}{d_1-d_{min}}) & d_{min} \leq \mid d \mid \leq d_1 \\ 1 & d_1 < \mid d \mid < d_2 \\ 1 - 3(\frac{|d|-d_2}{d_{max}-d_2})^2 + 2(\frac{|d|-d_2}{d_{max}-d_2})^3 & d_2 \leq \mid d \mid \leq d_{max} \end{cases}$$

In our experiments, we have found that the slopes in the functions g and h have little effect on the algorithm results, since the surface propagation only stops when the speed value reaches zero. For computational efficiency, we keep p_1 close to p_2 to ensure relatively fast propagation speed when surfaces are away from desired boundary. The time step is set to a small value so that the propagating surfaces move less than half a voxel size at each iteration, thus avoiding stepping over desired boundaries. The allowed distance between the inner and outer surfaces is set to a range from $d_{min} = 1.5mm$

to $d_{max} = 5.5mm$, based on knowledge from reported *post mortem* studies [7] to cover typical variations of cortical thickness. Parameter p_2 is set to allow the surface to stop at a strong intensity change. Weak intensity changes would cause propagating surfaces to leak out, however the leakage can partially be controlled by the thickness constraint imposed.

Due to the level set formulation, we have a notion of the inside and outside of the current moving front, which is embedded in the outward normal direction $\vec{n}$. This information can be used to reduce the feasible space of possible $\vec{\theta}$s, or $\vec{n}$ can be used directly as an estimate of $\vec{\theta}^*$, thus obtaining a better result. In our implementation, we directly use $\vec{n}$ as an estimate of $\vec{\theta}^*$, which contributes to computational efficiency as well.

4.5 Verification of Cortical Volume Connectivity

With the signed distance function Ψ, the level set formulation keeps track of the inside and outside of the current moving front. Once the evolution of the coupled surfaces is completed, the cortical gray matter voxels are those that lie inside the outer (gray/CSF) cortical surface while outside the inner (gray/matter) cortical surface. In the same fashion, non-brain tissue voxels will be the ones that are outside the outer cortical surface, and voxels of white matter will lie inside the inner cortical surface in addition to sub-cortical gray matter and ventricles.

We then perform a connectivity check on the estimated gray matter volume based on the fact that cortical gray matter is a single sheet encasing white matter. Starting from a candidate gray matter voxel, we find all the gray matter voxels that are connected to the candidate using a second order neighborhood system. We keep the largest connected group as cortical gray matter, and remove isolated gray matter voxels which usually

correspond to sub-cortical structures in deep brain or are caused by erroneous edges. Even after this removal process, there are chances that our gray matter volume contains sub-cortical structures that are actually connected to cortical gray matter, such as the hippocampus. The delineation of such structures can be done relatively easily either in pre-processing or post-processing stage with manual or other automatic methods.

Because the signed distance-based measures have sub-voxel accuracy, we can obtain a sub-voxel segmentation instead of a binary segmentation on the data set. In other words, if the distance from a voxel to the zero level set surface is less than the voxel size in width, the voxel is considered to contain multiple tissue types.

4.6 Segmentation Results

In this section, we show validations of our approach using various simulated and real MR data, as well as applications to specific cortical studies. We use only T1-weighted images because they provide the best gray/white contrast [72] and are therefore commonly used for neuroanatomical analysis. Figure 4.4 shows the propagation of the coupled inner and outer cortical surfaces.

4.6.1 Validation using simulated MR data with ground truth

We first present our segmentation results using simulated MR brain images provided by the McConnell Brain Imaging Center at the Montreal Neurological Institute [52]. The images are generated using an MRI simulator [40] which allows users to independently control various acquisition parameters to obtain realistic MR images. The ground truth of the phantom is provided in the form of membership functions of each voxel belonging to different tissue types, such as the skull, CSF, gray matter and white matter.

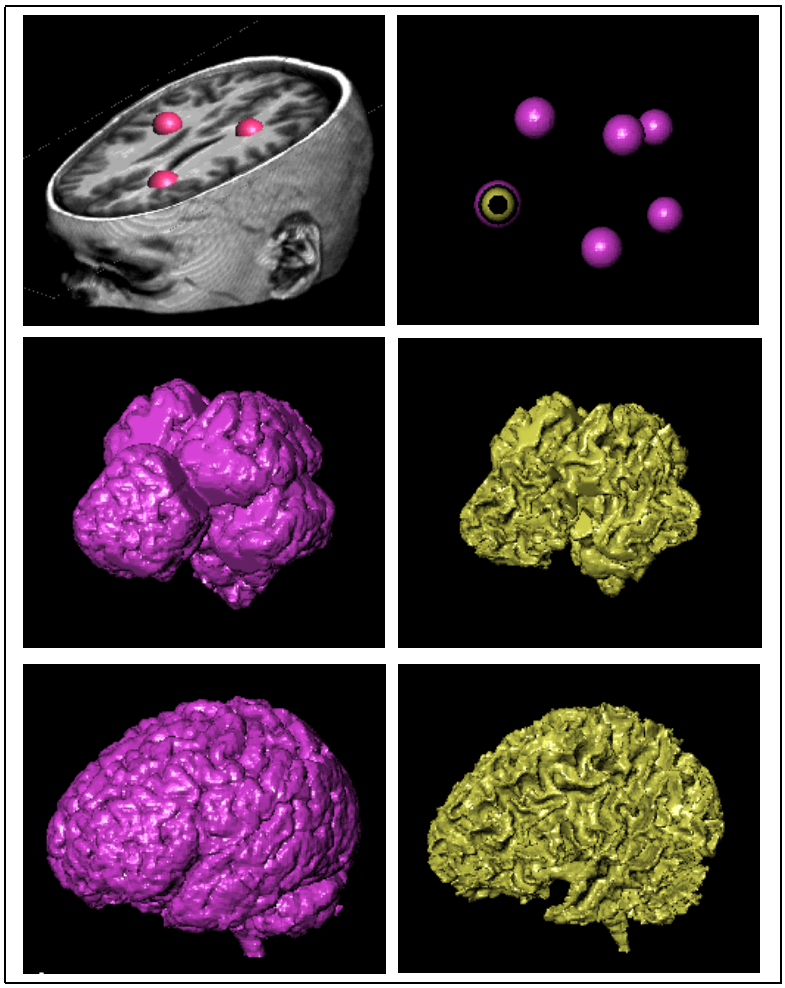

Figure 4.4: Propagation of the outer (pink) and inner (yellow) bounding surfaces. Top: pairs of concentric spheres (only the outer ones are shown on the left, both are shown in one pair with a cutting plane on the right) as initialization in unedited 3D MR brain images; middle: intermediate step in surface propagation; bottom: final result of the outer and inner cortical surfaces.

The simulated data we tested our algorithm on were T1 images of a normal brain, with the following parameter settings: voxel size= $1mm^3$, noise= 3%, intensity non-uniformity= 0%. Starting from the unedited images, no further user interaction is needed after initialization using several pairs of concentric spheres. The spheres grow out and automatically lock onto the inner and outer cortical surfaces. As long as the spheres are placed inside the white matter, the algorithm is robust to starting position (for more quantitative analysis, see section 4.6). Measurement of the volume is then performed as described above; we use a binary segmentation in this experiment. In our implementation of cortex segmentation, the allowed distance between the inner and outer surfaces is set to range from $1.5mm$ to $5.5mm$ based on knowledge from reported *post mortem* studies [7]. Therefore, to ensure the proper overlapping of the inner and outer narrow bands, the bandwidth ranges for the inner and outer interfaces are chosen to be $(-3mm, 6mm)$ and $(-6mm, 3mm)$ respectively.

To evaluate the segmentation result, we apply several measures defined as follows. For any tissue type T in the region of interest, we denote the voxels of tissue type T recovered from our 3D algorithm as V_a and the voxels that are mostly of tissue type T according to the phantom (*i.e.* the value of tissue T membership function is greater than 0.5) as V_e. We denote the overlap of V_a and V_e as V_{ae}, and the part that is in V_a but not in V_e as $V_{ae'}$. A *true positive(TP) rate* is then defined to be the size of V_{ae} relative to the size of V_e, while the *false positive(FP) rate* is defined to be the ratio of the size of $V_{ae'}$ to the size of V_e. We also define the *volume ratio* to be the volume of all the voxels segmented as of tissue type T by our algorithm to the total sub-voxel volume of tissue type T) specified by the phantom (sub-voxels contribute in only part of the voxel volume).

%	whole brain	cortical gray matter *	white matter
TP rate	92.3	92.8	92.4
FP rate	2.0	6.0	3.3
volume ratio	96.3	103.2	98.1

Table 4.1: Comparison of our volume measurements with the phantom ground truth. whole brain: total brain tissue (white+gray matter) on all slices; cortical gray matter *: cortical gray matter on the frontal 49 coronal slices and the top 56 axial slices;

Table 4.1 shows our measurement results over 3 types: total brain tissue (including white matter and gray matter), cortical gray matter in selected slices, and white matter. Since the algorithm is designed specifically for the nearly constant thickness of the cerebral cortex, it recovers only part of the gray matter in the brain stem and the cerebellum where the constant thickness constraint is not well satisfied. These regions account for most of the errors in the *TP* rate and volume ratio for whole brain tissue. For the same reason that the algorithm is specifically tailored for the cerebral cortex, we would compare the cortical gray matter volume only in the cerebrum. Since the phantom data does not provide the information related to partitioning cerebrum, the cerebellum and the brain stem, we only compare the cortical gray matter volume on selected slices where cerebellum and brain stem are not present: frontal 49 coronal slices and top 56 axial slices. The resulting average error of the *TP* and *FP* rate is around 6% to 7% , and the volume ratio error is within 4%. For the white matter, the errors for the *TP*, *FP* rate and volume ratio are also low. These results show that our algorithm performs well in isolating the brain from non-brain tissues and in segmenting the cortex.

4.6.2 Validation on 20 normal brains

To further evaluate our segmentation approach under a wide range of imaging conditions, we tested the algorithm on real MR data and compared the results obtained with gray segmentation by manual experts. Since for 3D data it is a very labor intensive job to segment gray and white matter, we utilized the data provided by the Internet Brain Segmentation Repository (IBSR) of the Center for Morphometric Analysis (CMA) at Massachusetts General Hospital [33].

The purpose of IBSR is to encourage the development and evaluation of segmentation methods by providing test image data, human expert segmentation results, and methods for comparing segmentation results. It is one of the first efforts to offer solutions to the problem of validating and comparing new algorithms in this rapidly growing medical image analysis field. The test image data sets provided in this repository permit a standardized mechanism for evaluation of the sensitivity of a given analysis method to signal to noise ratio, contrast to noise ratio, shape complexity, degree of partial volume effect, etc.

We obtained 20 normal MR brain data sets and their manual segmentations from IBSR. These 20 coronal 3D T1-weighted spoiled gradient echo MRI scans were performed on two different imaging systems. Ten FLASH scans on four males and six females were performed on a 1.5 tesla Siemens Magnetom MR System (Iselin, NJ) with the following parameters: TR = 40 msec, TE = 8 msec, flip angle = 50 degrees, field of view = 30 cm, slice thickness = contiguous 3.1 mm, matrix = 256x256, and averages = 1. Ten 3D-CAPRY scans on six males and four females were performed on a 1.5 tesla General Electric Signa MR System (Milwaukee, WI), with the following parameters: TR = 50 msec, TE = 9 msec, flip angle = 50 degrees, field of view = 24 cm, slice thickness =

contiguous 3.0mm, matrix = 256x256, and averages = 1.

All data sets were positionally normalized at CMA by imposing a standard 3D brain coordinate system on each 3D MR scan using the midpoints of the decussations of the anterior and posterior commissures and the mid-sagittal plane at the level of the posterior commissure as points of reference for rotation and (non-deformation) transformation [85, 23]. The repositioned scans were then resliced into normalized $3.0mm$ coronal, $1.0mm$ axial, and $1.0mm$ sagittal scans which were used for subsequent analysis.

Manual segmentation was performed on the normalized scans by trained investigators at CMA using a semi-automated intensity contour mapping algorithm [37, 33]. Once the external border was determined by intensity contour mapping, grey-white matter borders were demarcated using signal intensity histograms. Using this technique, borders were defined as the midpoint between the peaks of the bimodal histogram for a given structure and its adjacent tissue. Other neuroanatomical structures were segmented similarly [24].

An overlap metric is used by IBSR to compare results from automatic segmentation and manual segmentation. While manual segmentations are not "ground truth", they provide a reasonable way to compare automated segmentation methods. The overlap metric is defined for a given voxel class assignment as the number of voxels that have the class assignment in both segmentations divided by the number of voxels where either segmentation has the class assignment, which is equivalent to $TP/(1+FP)$. This metric ranges from 1.0 , for perfect agreement, to 0.0, for no agreement of classified voxels.

We interpolated and resampled the image data into $1mm$ thick coronal slices, and then ran our coupled surfaces algorithm. Figure 4.5 shows the overlap metric for gray matter segmentation on 20 normal brains from the manual method, various automatic

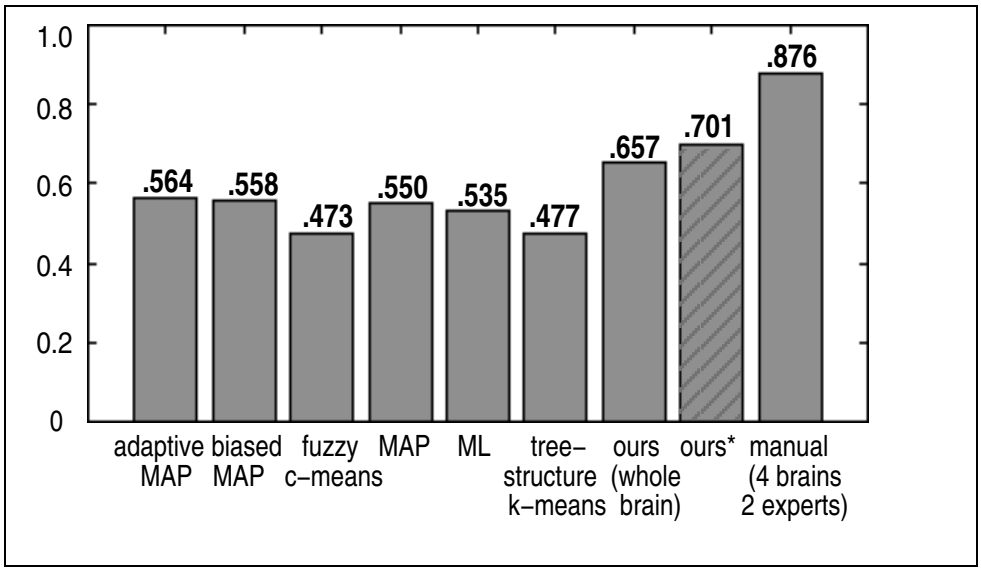

Figure 4.5: Average overlap metric (with expert tracing results) for gray matter segmentation on 20 normal brains from various segmentation methods. The results of automatic segmentation methods provided by IBSR were from work done by Rajapakse [65]. MAP: Maximum A Posteriori Probability. ML: Maximum-Likelihood. *: using frontal 13 coronal slices and upper 50 axial slices of each brain to exclude brain stem and cerebellum.

segmentation methods and our coupled surfaces algorithm. The results from the automatic segmentation provided by IBSR were from work done by Rajapakse, and partially based on the methods described in Rajapakse *et al.* [65]. The gray matter overlap metric for our algorithm on the whole brain is 0.657, which is well above those from the other 6 listed automatic methods ranging from 0.473 to 0.564 (shown in columns 1-6 in Figure 4.5). Since our algorithm is designed specifically for the cerebral cortex, we compute an improved overlap metric on the upper and frontal part of the brain (to exclude brain stem and cerebellum) of 0.701. Moreover, considering that the other 6 listed automatic

methods started with brain-only data sets, while the coupled surfaces algorithm started with the original un-stripped brain images, the advantage of our method with geometric structural constraints is clear.

These 20 brain scans were chosen by IBSR because they have been used in published studies [65], and cover a range of image quality [33] with the worst ones having low contrast and relatively large intensity inhomogeneities. The overlap scores shown in Figure 4.5 from the automatic classification methods may appear low, however they need to be taken into the context of a wide range of image quality, and should not be compared with numbers from different studies. More recently acquired (i.e. better quality) data should result in far better results from the automatic classification methods, which holds for our coupled surfaces algorithm as well. In fact, as shown in the section above, the overlap metric for our phantom cortical segmentation is $0.928/(1 + 0.060) = 0.875$, which compares well with the manual overlap metric of 0.876 showing inter-operator reproducibility from tests on 4 brains averaged over 2 experts (see Figure 4.5).

4.6.3 Results on real MR data for frontal lobe study

We also tested our algorithm on the frontal lobes of 7 high resolution MRI data sets (SPGR, 2NEX, TR=24 msec, TE=5 msec , flip angle= 45 degrees, field of view = 30 cm, $1.2 \times 1.2 \times 1.2mm^3$ voxels) from a randomly chosen subset of young adult autistic and control subjects from ongoing studies to measure frontal lobe volume. After preprocessing to reduce the effects of MR bias field inhomogeneity using a simple standard nonlinear map (this is also a step before expert manual tracing), we ran the coupled surfaces algorithm to isolate the brain tissue and segment the cortex (see Figure 4.4). Figure 4.6 shows 2D orthogonal slices of original image data, results from expert tracing and results from our algorithm. The frontal lobe was then manually defined independently in the left

and right hemispheres as all tissue anterior to the central sulcus, excluding sub-cortical nuclei [67]. We then create a mask of the frontal lobe, and use it to exclude the posterior part of the volume.

As shown in Table 4.2, over the 7 frontal lobes, the *TP* and *FP* rate (compared to manual tracing by our neuroanatomy specialist) of the whole frontal lobe averaged 94.1% and 2.1% respectively, which demonstrated that our algorithm nicely isolated the brain tissue from the non-brain tissue. The average *TP* and *FP* rate for the cortical gray matter (measured on 2 orthogonal slices, one coronal and one axial, to cover the entire range of the frontal lobe) in the frontal lobe were 86.7% and 20.8%. Note that the *FP* rate for gray/white segmentation is a very sensitive measure, especially considering the fact that manually drawing a boundary between gray and white matter to some extent depends on subjective individual judgment. However, in quantifying the difference between populations, despite the FP rates, the volume measurements would still yield useful information as long as they are consistent.

The volume of the constituent parts of the brain is often the measurement of interest for comparison among different subjects in studies of neuroanatomy. Thus, as a second way to analyze the utility of our algorithm, we compute reliability statistics on the volume measurements using the methods described in Schultz and Chakraborty [72] (see also [100]). There was strong agreement between the algorithm and the expert on the volume of the frontal lobe (Pearson r = .991; intraclass correlation coefficient [ICC] = .901). The algorithm systematically estimated the frontal lobe volume to be less than the expert tracer (mean difference = 4%), and this accounts for the lower ICC than the Pearson coefficient. Similarly, for gray matter volume of the frontal lobe there was also good agreement (Pearson r = .96). Thus, for both whole frontal lobe volume and frontal

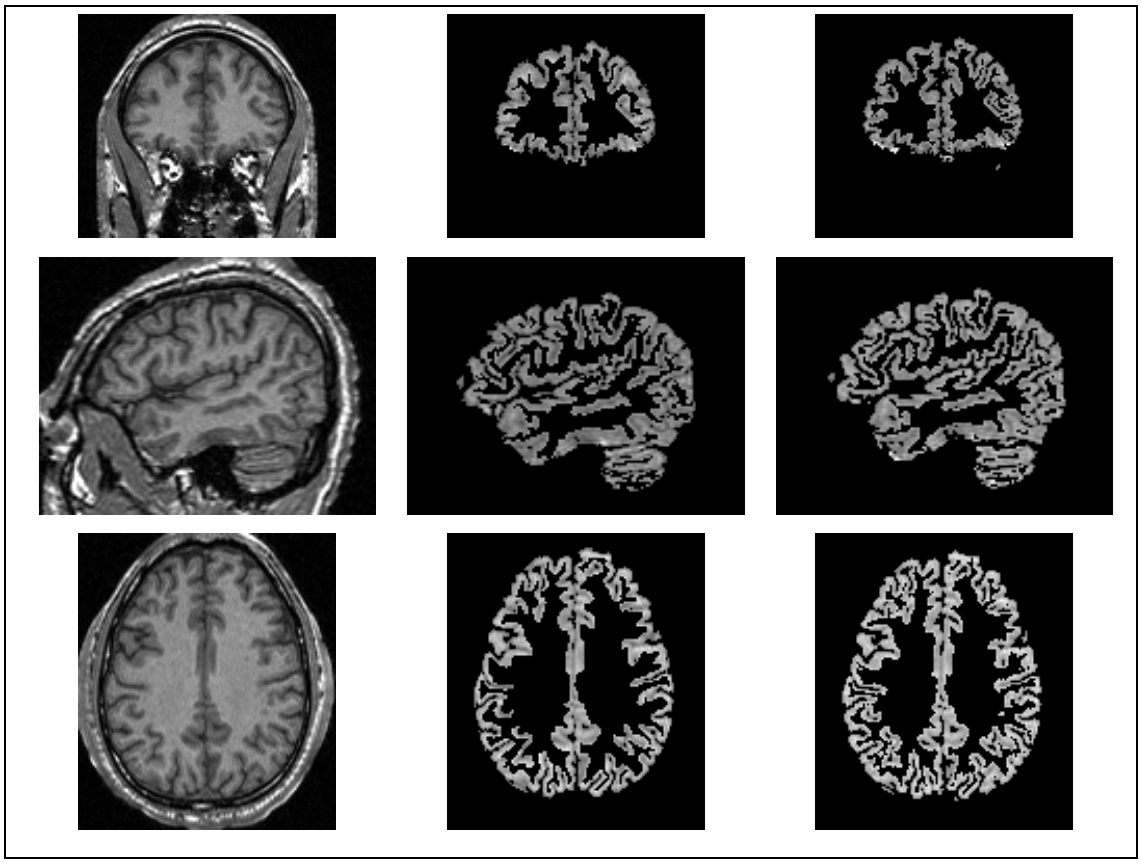

Figure 4.6: Results of cortical gray matter segmentation shown on orthogonal slices from 3D images. Left: original image slices; middle: gray matter from manual tracing of those slices; right: gray matter from our 3D algorithm shown on the same slices.

frontal lobe		frontal lobe cortex	
TP(%)	FP(%)	TP(%)	FP(%)
93.8	3.4	83.6	25.5
93.9	1.9	86.2	20.1
95.2	2.9	86.5	24.4
93.7	1.7	86.7	24.5
94.5	1.5	88.9	21.2
94.1	1.7	87.0	20.5
94.1	1.4	89.0	19.5

Table 4.2: Our measurements on 7 frontal lobes compared with expert tracing results

gray matter volume, the coupled surfaces algorithm produced measurements that were very similar to expert tracings.

4.7 User Interaction and Speed Issues

In addition to robustness and accuracy, minimum user interaction and computational efficiency have always been two important issues in the problem of segmenting and measuring the cortex. For an expert to manually isolate non-brain tissue (using a combination of image thresholding, region growing, and fine editing with manual tracing slice by slice to carefully remove any non-brain voxels such as the CSF within sulci and the dura) alone can take about 2 hours. (Structures such as the dura and the CSF in sulci can only be removed by careful slice-by-slice inspection. Therefore, considering the thoroughness and obsessiveness of the fine editing, we believe 2 hours is a fair estimate of the processing

time.) The manual tracing of cortical gray matter is even more time consuming. MacDonald *et al.* deformed two ellipsoids with inter-surface constraints to approximate the inner and the outer cortical surfaces. Their processing time for such segmentation on each subject was reported to be 100 hours on a SGI Origin 200 R10000 processor running at 180 MHz [46]. Davatzikos and Bryan [20] reported that the "ribbon" algorithm was a fairly computationally demanding iterative procedure; while manual placement of the initial cortical surface and a multi-scale formulation could decrease the computational load. The processing time per subject for Xu's method was reported to vary between 4.5 to 6.5 hours on a SGI O2 workstation with a 174MHz R10000 processor [97].

The initialization for our algorithm only requires the user to specify several pairs of concentric spheres in the unedited images, which can be done with several mouse clicks within seconds. It should be emphasized that neither the number nor the placement of the spheres (within a broad range of acceptable values) affects the accuracy or the reproducibility of the final result. To illustrate this, Figure 4.7 shows the coupled surfaces propagation on the same brain as in Figure 4.4 but from a different set (and different number) of initializing spheres. The final results of the surfaces show little visual difference. Quantitatively the TP rate of the gray matter volume from one with respect to that from the other is over 99.5%, and FP rate is less than 0.5%.

For a 3D image ($1.2 \times 1.2 \times 1.2mm^3$ in voxel size) of the whole brain, our algorithm runs in about 1 hour on a SGI Indigo2 machine with a 195MHz R10000 processor. **Skull-stripping, segmentation and measurement of the cortex are done simultaneously.** (Geometric measurements are discussed in detail in the next Chapter.) Comparatively, to our knowledge, our algorithm outperforms other related techniques with respect to user interaction and computational efficiency.

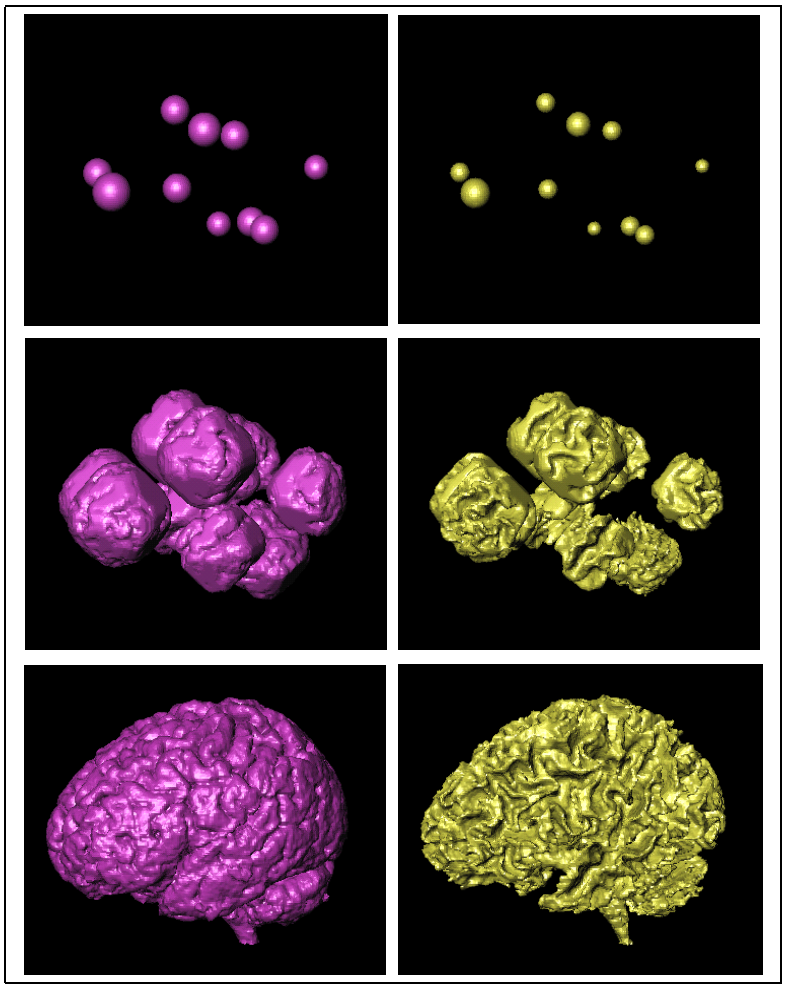

Figure 4.7: Coupled surfaces propagation on the same brain image as in Figure 4.4 but with a different set of initializing spheres. For the two final results from different initializations, the *TP* rate of the gray matter segmentation from one initialization with respect to that from the other is over 99.5%, and *FP* rate is less than 0.5%, which demonstrates our algorithm's robustness to initialization.

Chapter 5

Cortical Geometric Measurement and Analysis

5.1 Introduction

In this chapter, we discuss the derivation of several geometric measurements facilitated by our coupled surfaces algorithm, and present some results on cortical feature analysis. This type of analysis is important in the study of neuroanatomy and related disorders.

For example, cortical convolution is believed to be associated with the development of functionally distinct regions. In particular, during the development of the brain in embryo, connections between specific cortical regions and connections between cortical regions and sub-cortical structures, in conjunction with an overall growth process constrained by the skull, are believed to induce the inward folding of the cortex, resulting in the formation of sulci [50]. The curvature and shape index measurements proposed below offer convenient tools for the automatic identification and analysis of cortical folds.

In addition, the cerebral cortex is composed of columns of neurons, aligned per-

pendicularly to the cortical surface, that serve as basic units of information processing. Cortical surface area is likely to be proportional to column number and therefore surface area could be related to functional capacities. Moreover, regional cortical thickness may relate to functional capacity, and alteration in each of these features has been suspected in specific neuropsychiatric disorders [76].

5.2 Curvature and Shape Index Measurements

One advantage of the level set implementation is that geometric properties of the propagation front are easily calculated [77]. In our case of surfaces propagating in 3D space, there are many choices of surface curvatures of the front (for formal definitions of the curvatures, refer to [22]), including mean curvature, κ_M, and Gaussian curvature, κ_G. Both may be conveniently expressed [77] in terms of the level set function Ψ:

$$\kappa_M = \frac{\sum_{(i,j,k)\in C}((\Psi_{ii}+\Psi_{jj})\Psi_k^2 - 2\Psi_i\Psi_j\Psi_{ij})}{2(\Psi_x^2+\Psi_y^2+\Psi_z^2)^{3/2}} \tag{5.1}$$

$$\kappa_G = \frac{\sum_{(i,j,k)\in C}(\Psi_i^2(\Psi_{jj}\Psi_{kk}-\Psi_{jk}^2) + 2\Psi_i\Psi_j(\Psi_{ik}\Psi_{jk}-\Psi_{ij}\Psi_{kk}))}{(\Psi_x^2+\Psi_y^2+\Psi_z^2)^2} \tag{5.2}$$

where $C = \{(x,y,z),(y,z,x),(z,x,y)\}$ is the set of circular shifts of (x,y,z).

The maximum principle curvature, κ_1, and the minimum principle curvature, κ_2, are related to Gaussian and mean curvatures through the following formulas:

$$\kappa_1 = \kappa_M + \sqrt{\kappa_M^2 - \kappa_G}; \qquad \kappa_2 = \kappa_M - \sqrt{\kappa_M^2 - \kappa_G}; \tag{5.3}$$

We also adopt the classification of surfaces by Koenderink [39] using the numerical relationship between the two principal curvatures. A shape index function is defined as:

$$si = \frac{2}{\pi} arctan((\kappa_1 + \kappa_2)/(\kappa_1 - \kappa_2)) \tag{5.4}$$

which classifies surfaces into nine types as shown in Figure 5.1. With the shape index, gyri (mostly ridges) and sulci (mostly ruts) are automatically identified. Further potential use of the shape index includes the definition of an atrophy index (sulci widen with age). This information can also be used to study the percentage of the cortical surfaces that are buried in sulcal valleys.

Figure 5.1 shows the outer and inner cortical surfaces of a frontal lobe colored with their shape indices. As we see, most parts of the gyri are automatically identified as ridges while most parts of the sulci are identified as ruts, which coincides with our knowledge of the cortical structure.

5.3 Surface Area

The marching cubes algorithm [44] is performed on the signed distance functions, Ψ_{in} and Ψ_{out}, to extract the embedded zero level sets. The resulting surfaces are realized using a triangular representation. Surface area is then calculated as the sum of the areas of the composing triangles.

5.4 Thickness Measurement

As discussed earlier, the value of the level function of a front at any point is the distance from this point to the current front. Also recall that the inner and outer surfaces are the

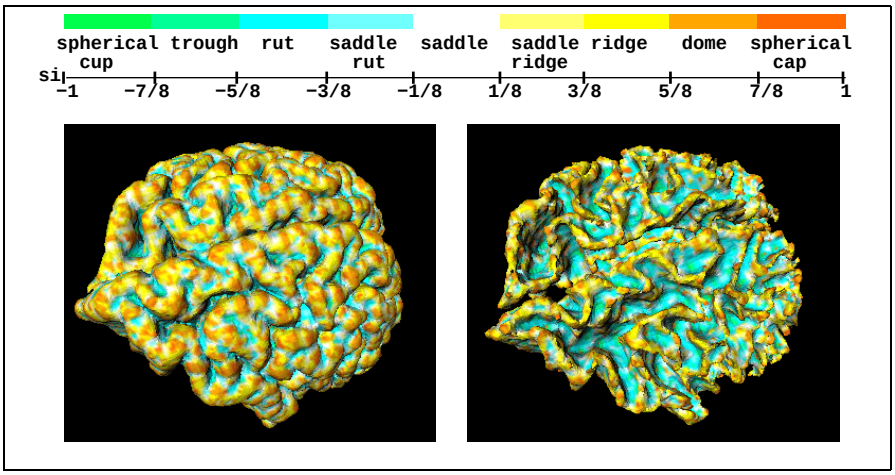

Figure 5.1: The outer and inner cortical surfaces of a frontal lobe colored with the specified spectrum representing shape index si.

zero level sets of Ψ_{in} and Ψ_{out}. Thus, for any point on the outer surface, the absolute value of Ψ_{in} at the point is simply the distance from the point to the inner surface. Using this measure, we obtain a thickness map between the inner and outer cortical surfaces, which can be used to study the normal thickness variations across different regions of the brain, and also abnormalities in brain structures.

We applied our algorithm to more than 50 high resolution MRI data sets (SPGR, 2NEX, $1.2 \times 1.2 \times 1.2 mm^3$ voxels) to obtain cortical thickness maps. Figure 5.2 showcases a sample of such maps with red color indicating cortical thickening and green indicating cortical thinning. Though the cortical thickness and pattern vary from subject to subject, we have consistently observed the pattern of thinning behind the central sulcus in the gyral area, as well as the thinning in the occipital lobe.

We further computed regional cortical thickness by lobe and in 2 specific regions (see Figure 5.3(a)) of 10 normal males (average IQ = 109), and compared findings with

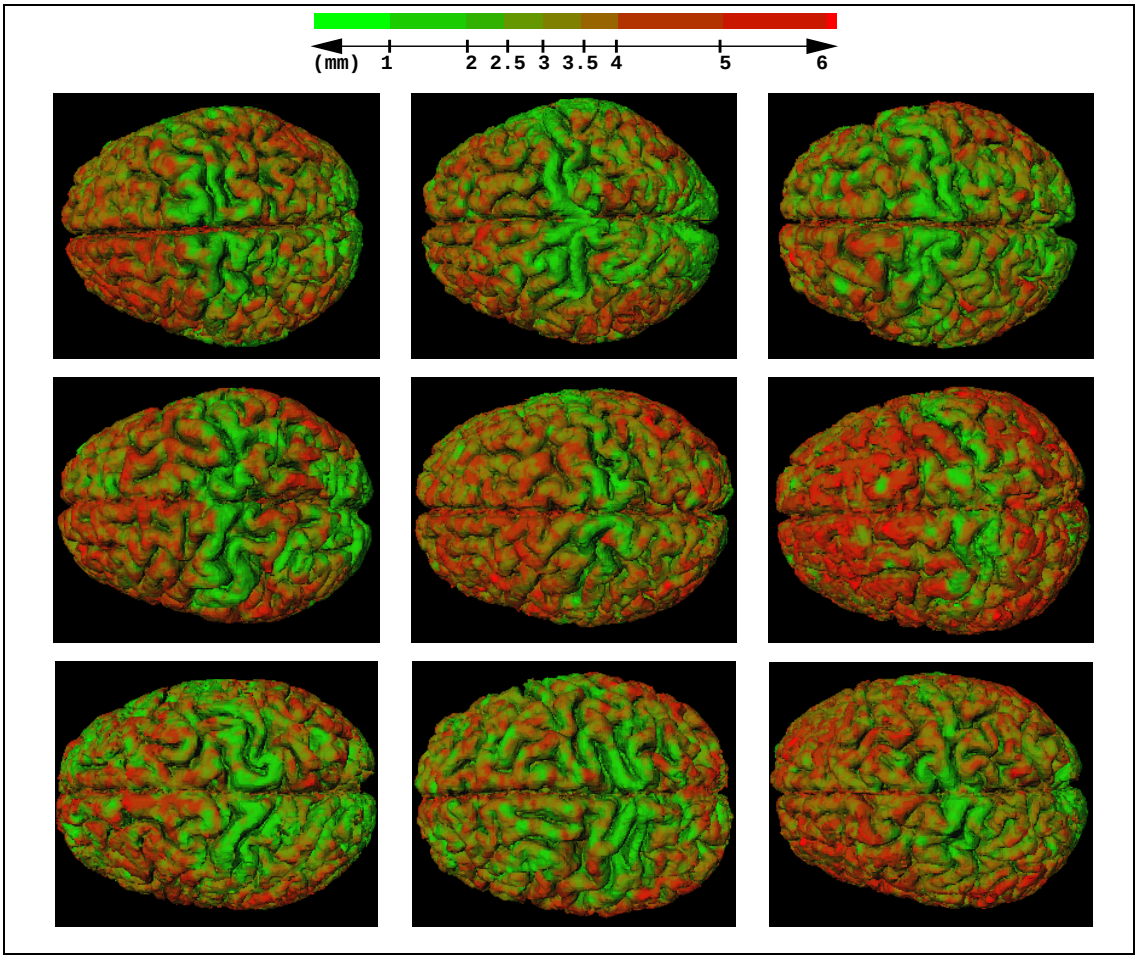

Figure 5.2: Cortical thickness map of 9 normal subjects obtained through our 3D coupled surfaces algorithm. Notice the variation in cortical thickness and its pattern across different subjects, as well as the consistent thinning behind the central sulcus in the gyral area.

that from *post mortem* studies. The lobes of the brain were labeled using locally developed software [66] in conjunction with the ANALYZE software package [69]. The frontal lobe was segmented by tracing the central sulci directly on 3D renderings of the brain, and then in successive 2D slices extending the traces to the depth of the sulci and through the white matter to the mid-line at an angle perpendicular to the inter-hemispheric fissure. Next, the temporal lobes were segmented by tracing the sylvian fissure on 3D renderings until the point where the fissure arched upward into the parietal lobe. At that point of inflection, a plane parallel to the AC-PC was used to segment the temporal and parietal lobes. The occipital-parietal boundary was set at mid-line by placing a oblique plane through the pariet-occipital sulcus, and a coronal plane at the intersection of the pariet-occipital sulcus and the calcarine fissure. The pre-central gyral area was then defined as the portion in the frontal lobe that was between the pre-central sulcus and the central sulcus. In a similar fashion, the post-central area was defined to be the part of the parietal lobe that was between the central sulcus and the post-central sulcus. Figure 5.3(a) shows the parcellation of the 4 lobes and the 2 specific areas of a brain, as described above.

Shown in Figure 5.3(b) are the side and back views of an outer cortical surface colored with cortical thickness. Figure 5.4 shows the regional cortical thickness measurements in 4 lobes and 2 specific areas over the 10 subjects. We compared the mean thickness of each lobe to the data on 63 males from the *post mortem* study of Pakkenberg and Gundersen [61], and found the same rank ordering of thickness; the frontal cortex was the thickest and the occipital cortex the thinnest (see Figure 5.4). The *post mortem* data measurements were 5 to 14% thinner by lobe than our *in vivo* data. This might be due to both the older age of the subjects, tissue shrinkage in the *post mortem* study, and volume averaging with our MRI data. However, it is important to note that the variability of

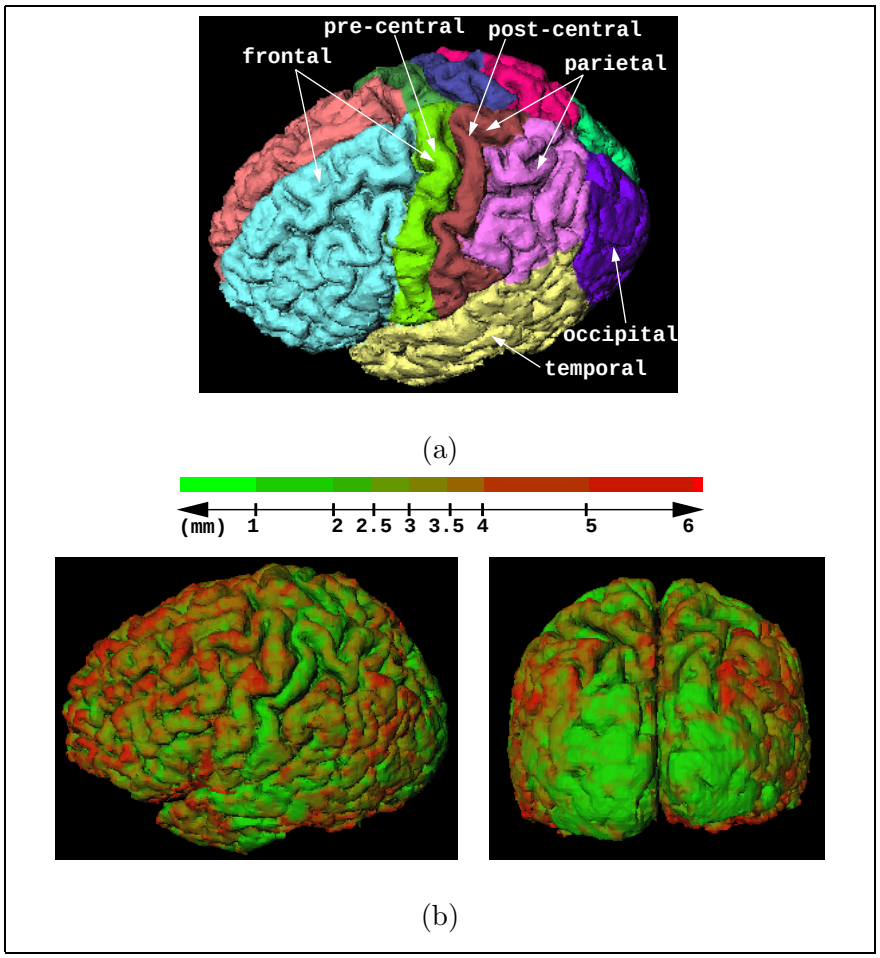

Figure 5.3: Brain partitioning and regional cortical thickness. (a): Parcellation of lobes and specific areas of interest where regional cortical thickness was measured. (b): Side and back views of an outer cortical surface colored with cortical thickness.

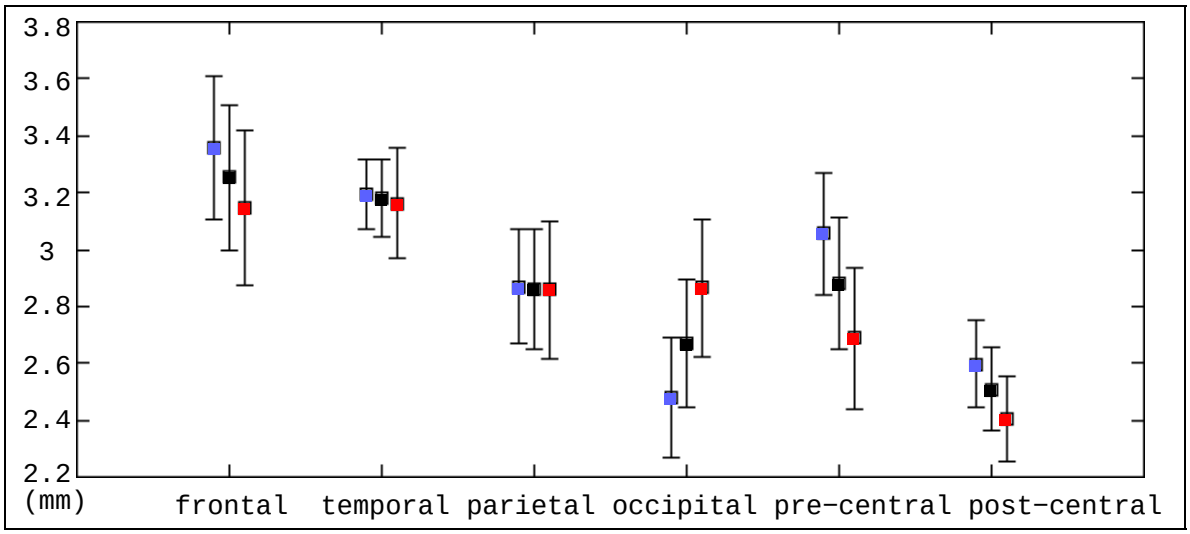

Figure 5.4: Plot of the mean and standard deviation of regional cortical thickness in 4 lobes (frontal, temporal, parietal and occipital) and 2 specific areas (pre-central and post-central gyral area) averaged over 10 normal male subjects. For each lobe or area, the left blue column and the right red column show the result on the left and right hemispheres respectively, while the middle black column shows the result over the two hemispheres.

thickness was the same for both samples (about 0.15 mm). This gradient of thickness from front to back in the brain is well known and due to the greater number of large pyramidal neurons in the anterior as compared to the posterior cortices. In addition, the post-central gyral area was the thinnest among different regions, as we observed over a larger sample size consistently. This too was in good agreement with known cortical thickness variations [7].

A repeated measures analysis of variance (ANOVA) tested whether cortical thickness differed by lobe, and found significant differences between the 4 lobes ($F[3, 27] = 56.3$, $p < .0001$). Post hoc paired t-tests showed that the frontal and temporal lobes

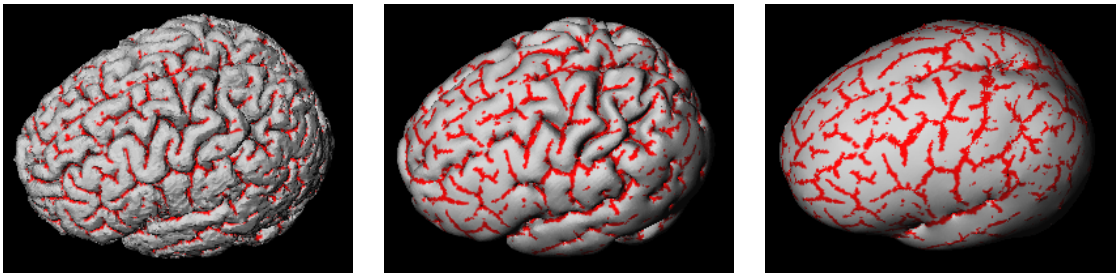

Figure 5.5: Flattening of an outer cortical surface using surface relaxation that simulates a mean curvature flow (iteration=0,50,250). Flattening is useful for the visualization of cortical features, especially within deep sulci. Here, high maximum principal curvature points are shown in red.

were each significantly thicker than either the parietal or occipital lobe (p's $< .001$), but they did not differ in thickness from one another. Likewise, the thickness in parietal and occipital lobes was not significantly different. Differences between the left and right hemispheres were also present in the thickness measurements, which could suggest interesting findings related to anatomical and functional asymmetry of the brain.

5.5 Cortical Surface Flattening

A flattened outer cortical surface map provides us with a better visualization of cortical features buried in sulcal ruts. Figure 5.5 shows an example of the visualization of maximum principal curvature on a cortical map. As the surface flattens, more and more details are shown from any single viewing angle.

We have implemented a straightforward method that simulates a mean curvature flow to realize the flattening. Our cortical surface is represented as a set of triangles

resulting from the Marching Cubes algorithm. For each triangular vertex v on the cortical surface, its position is updated with an average of the positions of its first order neighbors on the triangular net. Thus, the concave and convex parts of the surface flatten out. The flattening was carried out iteratively over all surface vertices, and stopped till the maximum move of all surface vertices was within a preset threshold.

Note that in this straightforward flattening method, the preservation of surface features, such as surface area and curve angles, is not our goal since we do not intend to carry out quantitative analysis of the flattened surface map. However, the connectivity of the triangles on the surface net does not change in our operation, so our method serves a second purpose of checking the topology of the original unflattened surface from our segmentation algorithm. Work on cortical surface flattening with emphasis on surface feature preservation can be found in [2, 19].

Chapter 6

Sulcal Ribbon Extraction

In this chapter, we present a new approach to 3D automatic intra-sulcal ribbon finding using dynamic programming and surface deformation [102], following the work on cortical segmentation with coupled surfaces propagation .

6.1 Introduction

Sulcal medial surfaces are 3D thin convoluted ribbons embedded in cortical sulci (see Figure 6.1 for a schematic drawing). The deepest part of a sulcus is called the fundus, which often demarcates the boundary between cortical regions with observable differences in their cytoarchitecture (the packing density and laminar distribution of different neuron types) and function [63]. Sulci thus provide important gross anatomical features for distinguishing different functional regions of the cortex, and thereby can serve as geometric guidelines in brain volume warping and cortical atlas building. In functional brain mapping, numerous efforts are in progress to propose cortical functional maps, as well as to correlate the gyral/sulcal anatomy of one individual and his/her neuro-functional

recordings. Relevant applications include the study of cerebral metabolism, hemodynamics or neurochemical and neurophysiological function [28]. All of this work depends on the development of efficient, fast and objective methods of extracting and representing sulcal patterns from brain images.

6.2 Related Work

Because of the importance of sulcal boundaries in brain structural and functional analysis, a number of recent efforts have begun to deal with the automatic extraction and representation of sulci [88, 28, 103], the probabilistic study of sulcal geometry and configuration [29], and automatic sulcal labeling [50].

Our work shares the same aim of previous work in [88, 28, 103], which is to automatically extract a sulcal ribbon surface and provide a parametric representation, thereby further facilitating quantitative shape analysis and cortical-constrained brain matching and warping.

Zhou *et al.* [103] described sulci as a collection of 2D contours generated from successive image slices, and proposed an algorithm that included two phases: contour generation and sulcal surface generation. The 2D sulcal contours were generated by using a voxel-coding method to extract 2D skeletons from segmented regions. In phase two, the integration between slices for the surface parametric representation resulted from testing adjacency relationships of the contour end points within adjacent slices. This method does not exploit the 3D nature of such sulcal ribbons, and the adjacency testing is rather *ad hoc*. As a result, the ambiguity in contour connections prevents the method from handling complicated cases that are common to sulcal patterns.

Vaillant and Davatzikos [88] used a physical model to retrieve sulcal ribbons.

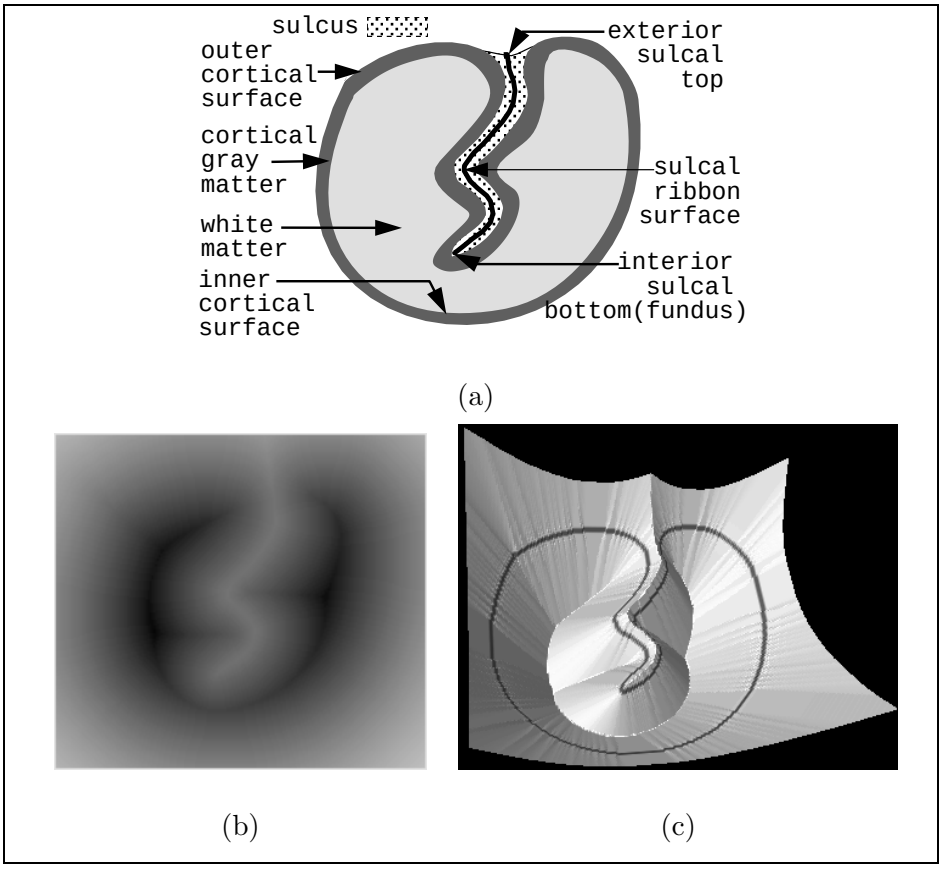

Figure 6.1: (a): 2D schematic representation of a sulcus; Gray scale view (b) and surface view (c) of the signed distance function corresponding to the outer cortical boundary shown in (a).

They started by initializing an active contour at the exterior part of a sulcus (see Figure 6.1). A parametric representation of the sulcal medial surface was then obtained as the active contour slides down toward the deep sulcal bottom under the influence of weighted external forces, including a center-of-mass force and an inward force. The latter force is based on a combination of surface normals, curve sliding speed and sliding acceleration. This deformable model uses characteristics of the cortical shape, and has been successfully applied to MR brain images. However, the manual placement of the initializing curve is a limitation, and it is likely that the tuning of weights on the external forces is as well.

Goualher's algorithm [28] used a very similar underlying mechanism of physical modeling of multiple forces to drive a curve transversing through the entire sulcal area, and differed only in its particular choice of the forces and its B-Spline curve parameterization.

6.3 Sulcal Ribbon Extraction Using Dynamic Programming and Deformable Surface Models

Following our earlier work on cortex segmentation with coupled surfaces using a level set implementation (see earlier chapters), we propose a new approach to automatic sulcal ribbon finding. Through the utilization of the signed distance function in which the outer cortical surface is embedded as its zero level set, we are able to formulate the sulcal ribbon finding problem as one of surface deformation, avoiding possible control problems of tuning weights on external forces in the sliding contour methods. This formulation provides our algorithm with the capability of modeling the complexity of sulcal patterns. Our sulcal ribbon finding algorithm starts from the outer cortical surface S_{out} and its associated level function Ψ_{out}, and requires several steps as shown in Figure 6.2.

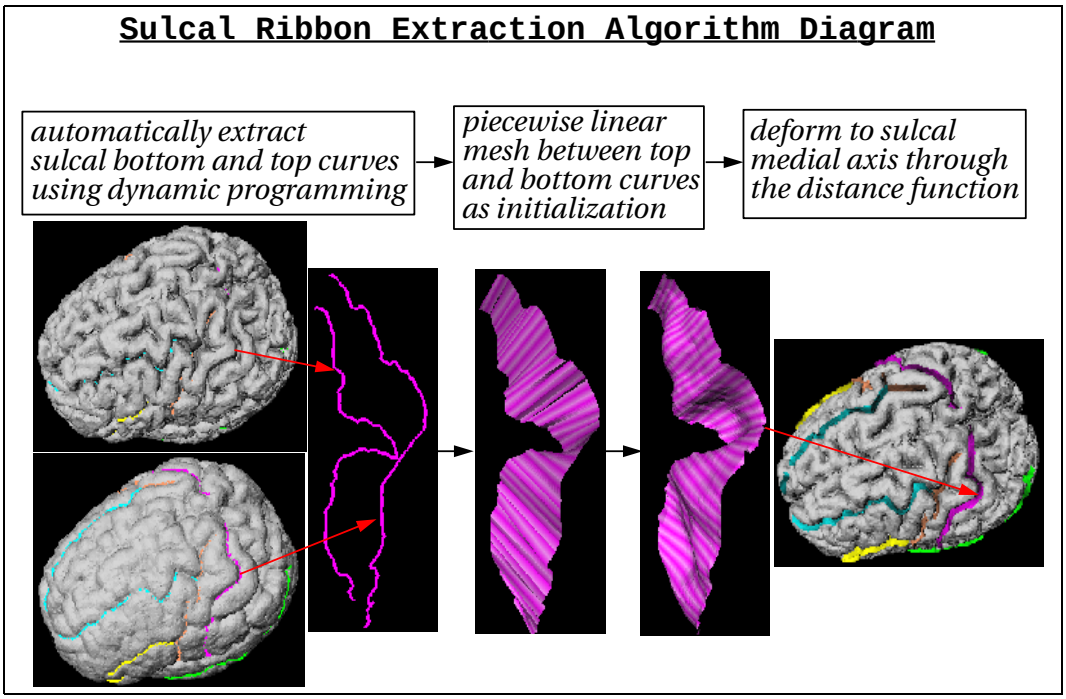

Figure 6.2: Diagram of our intra-sulcal ribbon extraction algorithm.

6.3.1 Calculation of a signed distance function

After running our coupled surfaces algorithm for cortical segmentation, we have two level functions Ψ_{out} and Ψ_{in} in which the outer and inner cortical surfaces S_{out} and S_{in} are embedded as zero level sets respectively.

Because of the level set implementation, our segmentation algorithm has the advantage of handling highly convoluted structures. As a result, S_{out} captures the deep sulcal folds and gets down to the interior sulcal bottom (see Figure 6.1 and 6.4) rather than staying at the exterior sulcal top. Therefore, we can use this surface to extract a sulcal bottom curve at the fundus, which greatly facilitates sulcal ribbon finding. This is an important difference between our sulcal ribbon finding and the approach due to

Vallaint and Davatzikos [88].

Though the sulcal ribbon finding algorithm proposed here is a natural follow up to our earlier segmentation work, it easily applies in general settings. S_{out} can be extracted from any segmented brain volume using the Marching Cubes algorithm [44]. To obtain a smoother S_{out} from the brain volume, other iso-surface based methods with specific constraints to preserve high curvature areas, such as that proposed by Gibson [27], can be used. Ψ_{out} can then be calculated as the signed distance function:

$$\Psi_{out}(x) = \begin{array}{ll} -dist(x, S_{out}) & if \quad x \in brainvolume, \\ dist(x, S_{out}) & else. \end{array}$$

where $dist$ is the Euclidean distance from position x on 3D image grid to S_{out}. Narrow band techniques (see Section 3.3) can be used here as well to limit the calculation of Ψ_{out} to positions close to S_{out} for computational efficiency.

6.3.2 Automatic extraction of sulcal curves

Our first step in sulcal ribbon finding is to define the interior sulcal bottom curve at the fundus which resembles crest lines corresponding to points where maximum principal curvature holds a local maximum [38]. The complex anatomy of different sulcal regions often exhibits small bumps and dents on the cortical surface within sulcal valleys. To overcome distractions from these, we pose the extraction of sulcal bottom curves given end points as a control problem where $\int (Curv_{max} - Curv(x))^2 dx$ needs to be optimized over all possible curves on cortical surface. $Curv(x)$ here denotes the largest maximum principal

curvature on the surface. On our discretized outer cortical surface net, the maximum principal curvature on each vertex $Curv(v_i)$ on S_{out} is calculated directly from Ψ_{out} as in Equation 5.3. Figure 5.5 shows an example of high maximum principal curvature points on an outer cortical surface.

We introduce the following notation for the description of our automatic curve extraction method.

- $V = \{v_i \mid 0 \leq i < M\}$: the set of all vertices on the surface net, where M is the number of vertices;
- $T = \{t_j\}$: the set of all triangles on the surface net;
- $E = \{e_{i,j}\}$: the set of all edges on the surface net, where $e_{i,j}$ is an edge if for some triangle $t \in T$, v_i, v_j are vertices of t. $e_{i,j}$ is a degenerate edge if $i = j$.
- $N(v_i) = \{v_j \mid e_{i,j} \in E\}$: the set of neighbors of vertex v_i;
- $Cost(e_{i,j}) = cost(Curv(v_i), Curv(v_j)) \cdot dist(v_i, v_j)$: the cost of stepping through edge $e_{i,j}$, where function $cost(,)$ penalizes small maximum principal curvatures, and $dist(v_i, v_j)$ gives the Euclidean distance between vertices v_i and v_j. Function $cost(Curv(v_i), Curv(v_j))$ takes on the form of $\frac{(Curv_{max}-Curv(v_i))^2+(Curv_{max}-Curv(v_i))^2}{2}$, where $Curv_{max}$ is the largest maximum principal curvature of all surface vertices.
- $P_{i_0,i_1,...,i_K}$: a path from v_{i_0} to v_{i_K}, consisting of a sequence of edges e_{i_0,i_1}, e_{i_1,i_2}, ..., e_{i_{K-1},i_K} (degenerate edges are allowed), where K is the number of steps.

The problem of finding a sulcal curve given the starting point v_{start} and the ending

point v_{end} becomes finding the optimal path on our discrete surface net:

$$P^* = \arg \min_{P_{i_0,i_2,\ldots,i_K}} \Sigma_{k=0}^{K-1} Cost(e_{i_k,i_{k+1}}) \tag{6.1}$$

where $v_{i_0} = v_{start}, v_{i_K} = v_{end}$, and K is the number of steps taken in the optimal path.

Dynamic programming [3, 6] is a technique well suited for such an optimization problem. A similar application can be found in work by Khaneja *et al.* [38]. The basic idea is as follows. Suppose there are M vertices on the surface net, then the optimal path P^* takes no more than M steps, i.e. $K \leq M$. If we can find an optimal path of step $K-1$ to the neighbors of v_{end}, then P^* is just the optimal path to one particular neighbor of v_{end}, plus the edge from that particular neighbor to v_{end}.

The algorithm works in the following fashion. First, we initialize a *pathvalue* for v_{end} to be 0, and $+\infty$ for the rest of the vertices. Let PV_t denote the set of possible vertices at the end of step $t (0 \leq t \leq K-1)$. Since v_{end} is the end of step K, it is obvious that $PV_{K-1} = N(v_{end})$. So for each $v_i \in PV_{K-1}$, we store *pathvalue* to be $Cost(e_{end,i})$, and assign v_{end} to be the *successor* of such a v_i. By similar reasoning, we have $PV_t = \{N(v_i) \mid v_i \in PV_{t+1}\}$. Now for $t = K-1$ down to 0, for each $v_i \in PV_t$, we compare $Pathvalue(v_j) + Cost(e_{ji})$ where $v_j \in N(v_i)$, to find the optimal *pathvalue* for v_i and assign its corresponding *successor*. When the operation is done for $t = 0$, starting from v_{start}, we trace back from the *successor* all the way to v_{end}, which gives us the optimal path.

Note that such a discrete optimal path is only an approximation of the weighted geodesic curve on the continuous surface. However, since our triangulation of the surface is done at the level of voxel size (the surface net of each whole brain has about 500,000

triangles), we have found the discrete path to be a fine enough approximation, as verified by expert inspection. In our implementation, the starting and ending points of a particular sulcus are specified by the user to allow flexibility, which only takes two mouse clicks on a surface rendering. The maximal step number M is set to be 300 (large enough for the possible steps for each sulcus) to provide nearly real-time operation – the extraction of each sulcal curve takes 2-3 seconds on a SGI R10000 processor running at 250 MHz. In the case of sulcal bifurcation, segments of a sulcus can be traced out separately and assigned the same sulcal label. Example automatic traces are shown in Figure 6.3.

6.3.3 Compute a brain wrapper

After extracting the interior sulcal bottom curve, we then define the exterior sulcal top. One simple way is to use the signed distance level function Ψ_{out}. While the outer cortical surface S_{out} is the zero level set of Ψ_{out}, a positive value ϵ can be chosen so that the ϵ level set of Ψ_{out} provides a brain wrapper – a surface that wraps around the brain volume, while following indentations at exterior sulcal tops. In our implementation, ϵ is chosen to be $3mm$, which results in a consistent brain wrapper suitable for subsequent sulcal ribbon extraction. After the brain wrapper surface is extracted, sulcal top curves are automatically traced out on the brain wrapper surface in the same fashion that sulcal bottom curves are on the cortical surface (see Figure 6.4). End points of the sulcal top curves are specified by simple mouse clicks as well.

6.3.4 Sulcal surface extraction and representation

The sulcal ribbon surface corresponds to the medial axis of a particular sulcus. There are different ways of extracting the medial axis of 3D structures, such as those using Voronoi diagrams [8, 59, 58]. The drawbacks of 3D Voronoi methods lie in their algorithmic

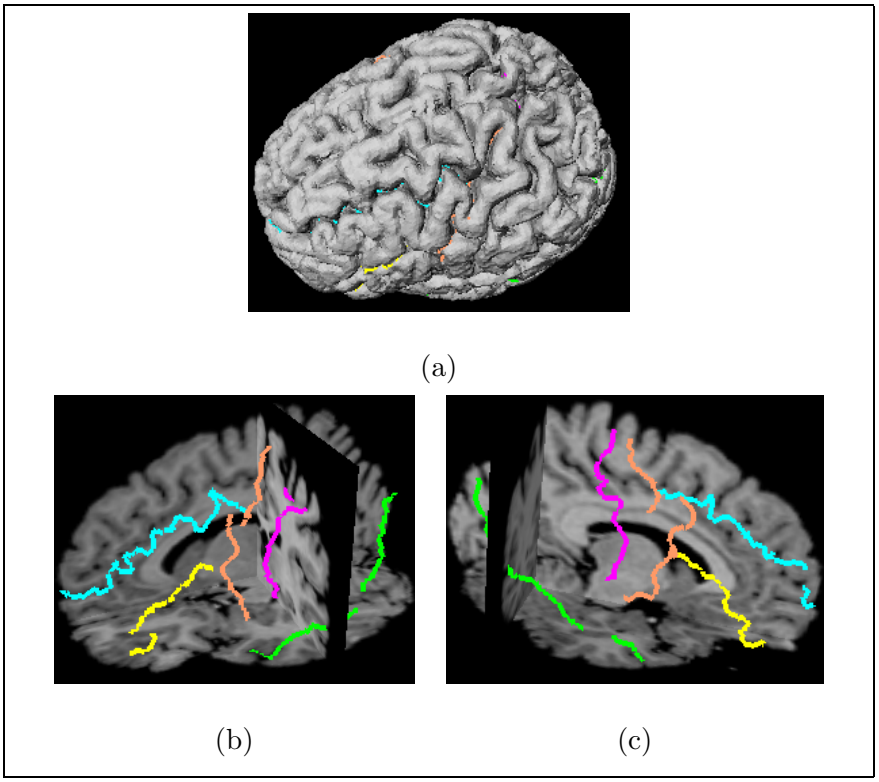

Figure 6.3: Examples of nearly real-time automatic sulcal tracing on an outer cortical surface. Volume rendering of the sulci shown with the cortical surface (a) and orthogonal image cards (b,c showing the left and right sides of the brain) are central sulci (magenta), superior frontal sulci (blue), inferior frontal sulci (yellow), superior temporal sulci (green) and pre-central sulci (tan). The traces are not clearly visible in (a) due to the convolution of the cortical surface. Also see Figure 6.5.

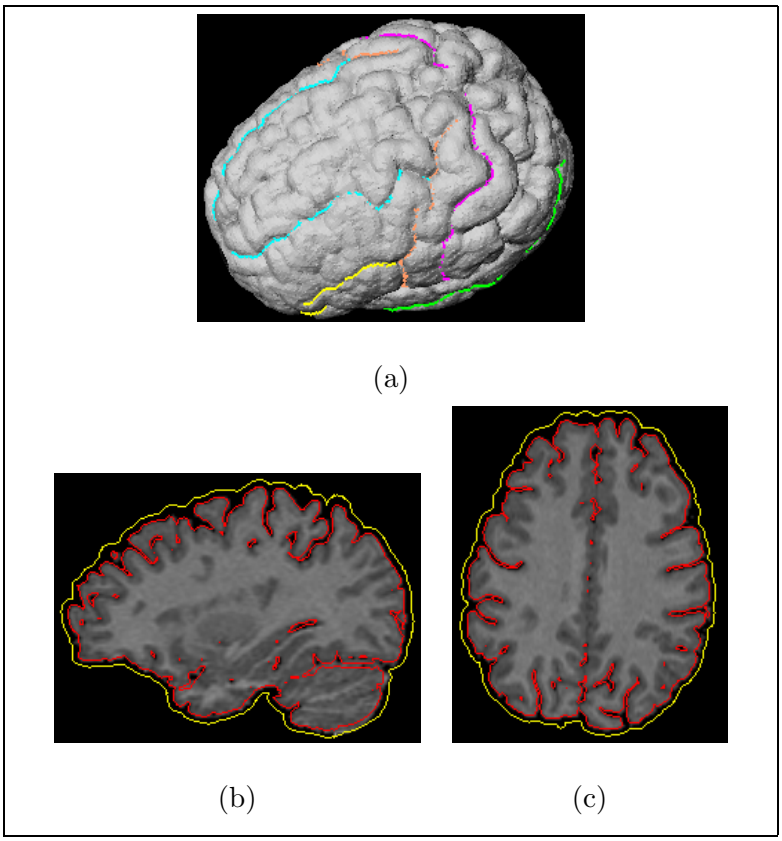

Figure 6.4: (a): Traces of sulcal top curves on the brain wrapper corresponding to the cortical surface in Figure 6.3. Cuts through the brain wrapper surface (yellow) and outer cortical surface (red) on (b):sagittal and (c):axial brain image slices.

difficulties caused by sampling problems and pruning procedures. Moreover, a parametric representation of the sulcal ribbon does not follow immediately from these methods.

From the level function Ψ_{out} and the sulcal bottom and top curves traced out automatically, we have a simple and natural way of defining the entire sulcal ribbon. Our method is based on the fact that the medial axis of a sulcus corresponds to directional local maximum of the signed distance function Ψ_{out}. This idea is illustrated by the 2D schematic drawing in Figure 6.1. Figure 6.1(b,c) shows the image and surface view of the signed distance function of the outer cortical boundary in Figure 6.1(a). The ridge curve in the sulcal region shown in Figure 6.1(c) is the medial axis of the sulcus, and has the property of being located at a local maximum of the distance function along its normal direction. The 3D case is similar, only differing in that the medial axial ridge curve becomes a surface that has a local maximum of Ψ_{out} along the surface normal direction. Our goal now is to capture such a surface through Ψ_{out}.

We will define a sulcal surface mesh $\vec{R}(u, v)$ on the domain $[0, 1] \times [0, 1]$, so that the parameter u runs in the direction parallel to the sulcal top and bottom curves, while parameter v runs across the sulcal depth. Figure 6.6 helps in illustrating this process. To start, we map $\vec{R}(u, 0)$ to be the interior sulcal bottom curve, and $\vec{R}(u, 1)$ to be the exterior sulcal top curve. We parameterize both $\vec{R}(u, 1)$ and $\vec{R}(u, 0)$ in u from $[0, 1]$. This parameterization is chosen for the convenience for measurement purposes. In this way, a correspondence between the points on the sulcal top and bottom curves is also set up, which helps to offer an intuitive and reasonable concept of sulcal depth discussed later in Section 6.4.1.

We then realize a piece-wise linear triangulation between the sulcal bottom and

top curves to generate the entire mesh as an initialization of the sulcal ribbon surface:

$$\vec{R}(u,v) = (1-v)\vec{R}(u,0) + v\vec{R}(u,1) \tag{6.2}$$

The number of v iso-parametric curves is chosen to be 20 in our implementation, so that the triangulation of sulcal ribbon surfaces are fine enough to be on the order of a voxel or less. Figure 6.6(a,b) shows such an initialization of a central sulcal surface and a superior frontal sulcal surface.

We then deform the surface according to the following equation while fixing the sulcal bottom curve $\vec{R}(u,0)$:

$$\begin{aligned} \frac{\partial \vec{R}(u,v)}{\partial t} &= F_{smooth} + F_{image} \qquad (6.3) \\ &= a(\vec{R}_{uu} + \vec{R}_{vv}) + (\nabla \Psi_{out} \cdot N_R(u,v))N_R(u,v), \quad (u,v) \in [0,1] \times (0,1] \end{aligned}$$

where $N_R(u,v)$ denotes the unit normal of the sulcal ribbon surface. The first term on the right hand side of the equation guarantees the smoothness of the sulcal ribbon surface, while the second force drives the sulcal surface along its normal direction towards the local maximum of Ψ_{out} which corresponds to the sulcal medial axis. Thus, the surface deforms to the sulcal medial axis while maintaining a certain smoothness decided by parameter a. This equation is discretized on the surface mesh and solved iteratively. The iteration stops when the increase in the value of $\int\int_{\vec{R}(u,v)} \Psi_{out} du dv$ falls below a certain threshold, which is set to be 0.05.

Figure 6.6(3,4) shows the captured central and superior frontal sulcal surfaces

with their iso-parametric u and v curves superimposed. The iso-u parametric curves show the correspondence between the points on sulcal top and bottom curves which is used in the sulcal depth measurement, and the iso-v parametric curves help with the visualization of the convolution across the sulcal depth.

The advantage of our approach is that the information defining the sulcal medial axis is implicit in the signed distance function Ψ_{out}, and by using Ψ_{out}, the problem of finding the sulcal ribbon becomes explicitly one of surface deformation. This formulation avoids the difficulty in tuning the weights of multiple inward force components in the sliding contour model.

6.4 Applications

In this section, we present some results of our sulcal ribbon finding algorithm, and discuss how they can be used for structural and functional analysis of sulci. The MR images we tested our algorithm on were acquired on a 1.5 T GE scanner, using a 3D SPGR volume acquisition. The image resolution was $1.2 \times 1.2 \times 1.2mm^3$ in voxel size.

We first ran our coupled surfaces algorithm to segment cortical gray matter from white matter and non-brain tissues, which resulted in the outer cortical surface S_{out} and its level function Ψ_{out}. The segmented cortical gray matter volume was then inspected by an expert, and corrections were made. Accordingly, S_{out} was modified and Ψ_{out} recomputed locally at the places of correction. This completed the pre-processing step of the sulcal ribbon finding algorithm.

As described in section 6.3.3, a brain wrapper surface S_{wrap} was extracted based on Ψ_{out}, and the maximum principal curvature was calculated on both S_{out} and S_{wrap}. Our expert then dropped starting and ending points of sulcal top and bottom curves on

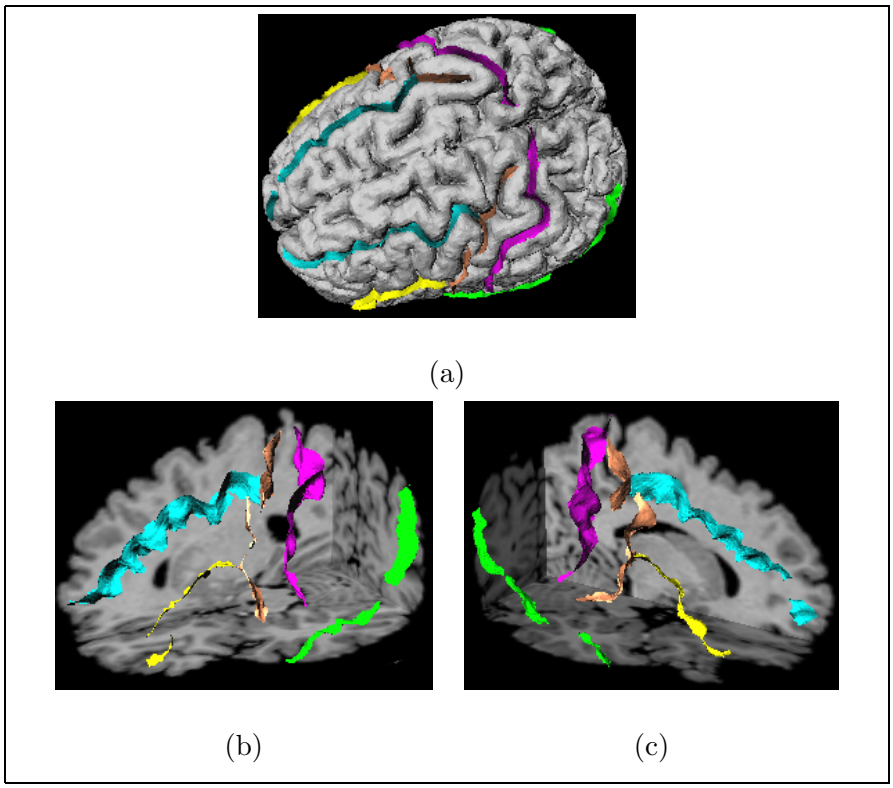

Figure 6.5: Sulcal ribbon surfaces corresponding to sulcal curves in Figure 6.3. shown with (a):cortical surface and with (b,c):3D orthogonal image cards. Notice the convolution captured by our ribbon extraction algorithm.

S_{wrap} and S_{out} respectively. After the sulcal curves were extracted automatically using dynamic programming, a piecewise linear triangulation between them was realized as an initialization of the sulcal ribbon. Finally, the sulcal ribbon was deformed to the sulcal medial axis through Ψ_{out} using Equation 6.3.

Software written in C++/Open Inventor was used for these steps on a SGI Octane machine with a 250MHz R10000 processor. The automatic tracing of sulcal curves can be done in nearly real-time (about 2-3 seconds), and the deformation of each ribbon surface

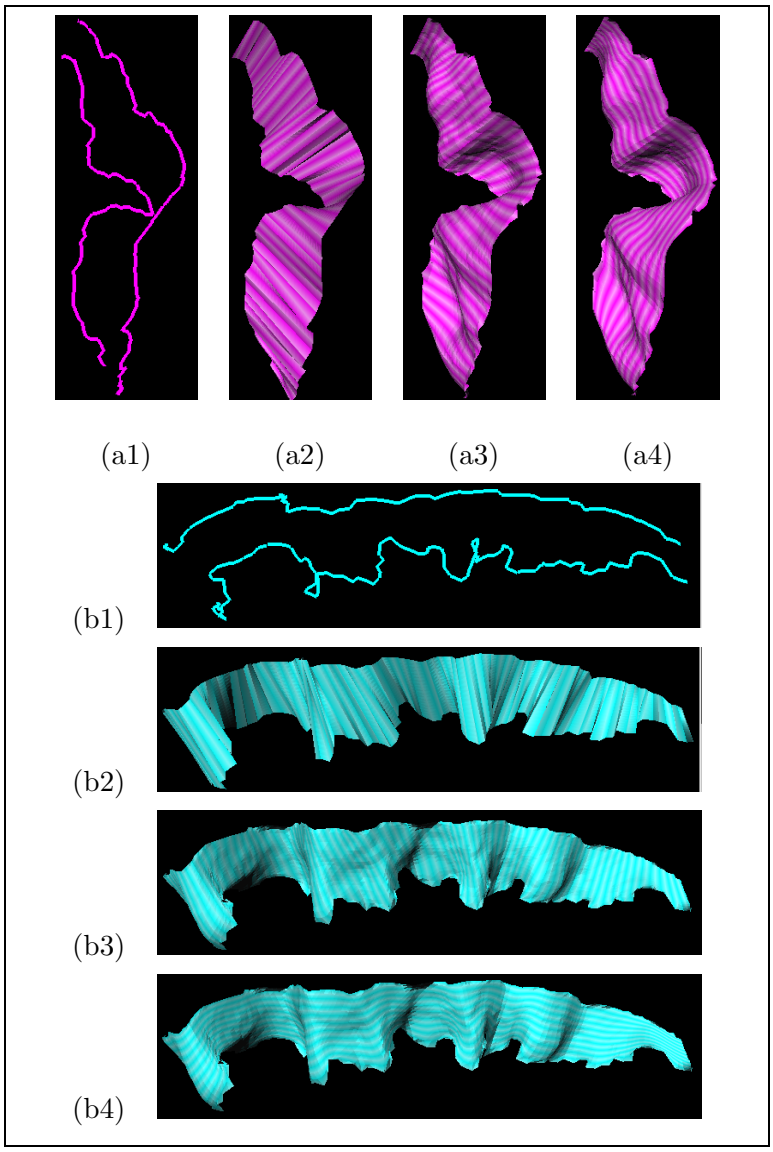

Figure 6.6: Deformation of a (a):central sulcal ribbon and a (b):superior frontal sulcal ribbon. (1): Sulcal top and bottom curves traced out automatically; (2): A piece-wise linear mesh between the two curves serves as an initialization of the sulcal ribbon. The final sulcal ribbon surface with iso-parametric (3):u and (4):v curves superimposed.

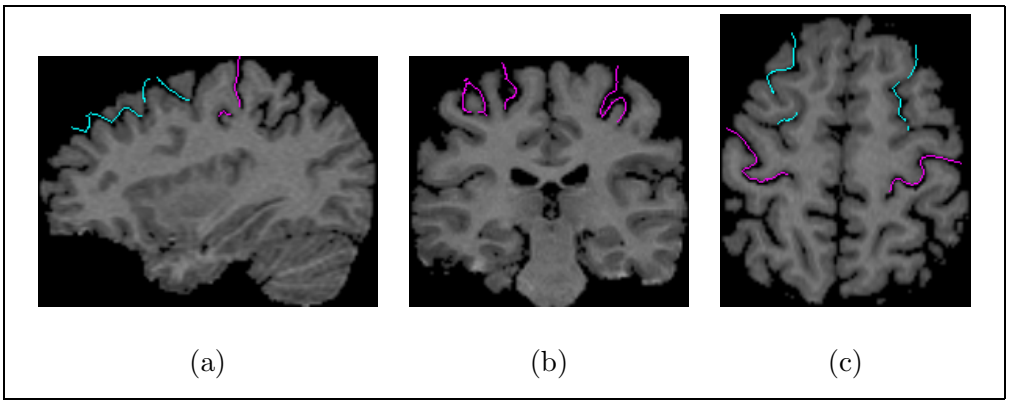

Figure 6.7: Left and right central sulcal ribbons (magenta) and superior frontal sulcal ribbons (blue) shown on (a):sagittal, (b):coronal and (c):axial images slices. Notice that our sulcal ribbons capture the in-plane and out-of-plane 3D convolutions.

takes about $3-5$ seconds.

Shown in Figure 6.5(a) are the ribbon surfaces of central sulci (magenta), superior frontal sulci (blue), inferior frontal sulci (yellow), superior temporal sulci (green) and pre-central sulci (tan) on the outer cortical surface. Figure 6.5(b,c) shows sulcal ribbons of the left and right hemispheres with orthogonal image cards. The cut views of central sulci and superior frontal sulci on orthogonal image slices shown in Figure 6.7 demonstrate that the sulcal ribbons from our algorithm capture the complexity of 3D sulcal convolution.

6.4.1 Quantitative measurement of central sulcal ribbons

Once the sulcal ribbons are captured in their parametric form, we can make quantitative measurements such as surface area, sulcal depth and sulcal curvature etc. Sulcal ribbon surface area is calculated as the sum of the area of all the triangles used to compose the surface. A reasonable and consistent way to measure sulcal depth is to measure the

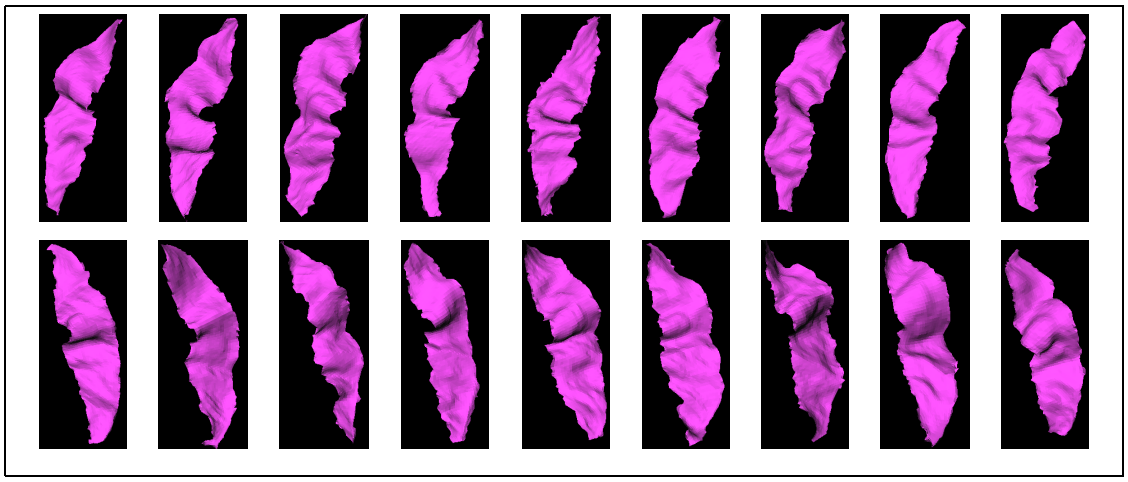

Figure 6.8: Left(top) and right(bottom) central sulcal ribbons of 9 normal subjects extracted from MR images using our algorithm.

geodesics on the sulcal surface to the sulcal fundus from each point on the top curve, in other words, the geodesics between $R(u,1)$ to $R(u,0)$ for all $u \in [0,1]$. The geodesics are computed using the dynamic programming technique described in Section 6.3.2, with the cost function set to be simply the length of a step, i.e. $dist(v_{i_k}, v_{i_{k+1}})$. Our experiments suggest that most of time the geodesic from $R(u_0,1)$ to $R(u_0,0)$ for a particular u_0 coincides with the iso-u curve $R(u_0)$.

We measured the depth and intra-sulcal ribbon surface area of the complete course of the central sulcus in both hemispheres across 15 right handed subjects to demonstrate our methods in an area of interest to neuroscientists. The fundus of the central sulcus is the dividing point for the primary motor region (Brodmann area 4) on the anterior bank and the primary somatosensory strip (Brodmann area 3b) on the posterior bank of the sulcus. In addition to serving as boundaries, the depth of sulci and the total intra-sulcal surface area may bear some relationship to the functional capacity of that region. There is

a somatotopic mapping of the body on the primary motor and somatosensory regions, such that different regions of these cortical strips process information from different regions of the body. There are larger cortical representations for body parts with more complex duties (e.g., hands vs. back, [26]) and greater signal resolution in areas containing more tissue [71]. Recent work has suggested that anatomic asymmetry in the depth and surface area of the region of the central sulcus which maps to the hand and arm may be associated with asymmetry in motor function; right handers have a deeper central sulcus in this region in the left hemisphere, while the opposite is true in left handers [95, 1, 94].

Figure 6.8 shows the central sulcal ribbons of 9 of the right-handed subjects to demonstrate the sulcal variability captured by our algorithm. Although not presented here, ongoing work in our lab is comparing regional measurements corresponding to the representation of the hand in this group and a matched comparison group of left handers, in order to test for structure function relationships. Results for the total central sulcus in this group of 15 neurologically healthy young adults of normal general intelligence (IQ mean $\pm SD = 108 \pm 15$) including 7 males and 8 females between the ages 9 and 41 years (23.8 ± 8.0 years) are as follows. Average sulcal depth of the complete central sulcal ribbon was $18.12(\pm 1.66)mm$ on the left and $18.08(\pm 1.67)mm$ on the right, with no significant right-left difference. Total surface area of the cortical ribbon (corresponding to the surface area of one bank, not both sulcal banks) was $1724(\pm 202)mm^2$ on the left and $1764(\pm 205)mm^2$ on the right. These measurements are roughly commensurate with *post mortem* measurements [94] and prior *in vivo* morphometry [1], especially considering the differences in measurement procedure (the cited methods were based on interpolation between traces on 2D slices).

6.4.2 Brain matching with cortical constraints

Another potential use of the extracted sulcal surfaces lies in non-rigid brain warping and cortical atlas building. As distinctive features of the brain, sulcal surfaces can be used as geometric guidelines in shape transformation methodologies. There has already been work toward this direction [89, 17, 13], and our method offers an alternative for obtaining a starting point.

6.5 Discussion

In this chapter we have presented a new approach to automatic 3D sulcal ribbon finding. Dynamic programming is used to automatically extract interior sulcal bottom curves on the outer cortical surface, and exterior sulcal top curves on a brain wrapper computed from the distance function Ψ_{out} associated with the outer cortical surface. A sulcal ribbon surface is then initialized through a piecewise linear triangulation between the sulcal top and bottom curves, and deformed to the sulcal medial axis through the distance function Ψ_{out}. The use of Ψ_{out} makes the information defining the sulcal medial axis implicit, and the resulting surface deformation formulation is simpler without multiple forces to tune. Though a natural follow up to our segmentation method, our sulcal ribbon finding algorithm can be adapted to follow other segmentation procedures. By allowing the user to define a sulcal ribbon with a few mouse clicks, our method offers automation, flexibility and nearly real-time operation.

All the sulcal ribbon surfaces captured by our algorithm are evaluated slice by slice on axial image slices by an expert. The positions of the ribbon surfaces are always within one voxel's distance from the sulcal medial axis by visual inspection. However, since any type of expert tracing of the sulcal ribbon suffers from its own limitations in

capturing the 3D nature of sulci, we feel the best way to do full quantitative analysis of our algorithm is to create a phantom with known convolutions, and experiment on its images.

Our method for extracting the brain wrapper has potential for further improvement. There are no well-defined methods to locate the exterior top edge of a sulcus. Although our way of extracting a fixed ϵ level set surface is consistent in its own right, ideally different values of ϵ need to be chosen for different sulci in order to obtain a more geometry-specific definition. Other directions of future research include localized sulcal measurement and shape analysis, and the study of structure and function relationships in sulcal regions by connecting our sulcal representation with functional data.

Chapter 7

Summary

This thesis presents a 3D approach for the problem of segmenting and measuring volumetric layers - a type of structure often encountered in medical image analysis. Key examples include the cortical gray matter of the brain and the left ventricular myocardium of the heart. We adopt a definition of the volumetric layers that includes the inner and outer bounding surfaces, as well as the homogeneity in between. Motivated by the nature of the coupling between the bounding surfaces, we have developed a coupled surfaces approach where a generic shape constraint – more specifically, a thickness constraint, is incorporated. By evolving two embedded surfaces simultaneously, each driven by its own image-derived information while maintaining the coupling, a final representation of the bounding surfaces and the automatic segmentation of the layer in between are achieved. The image-derived information is obtained by using a local operator based on gray-level information rather than image gradient alone, which gives our algorithm the ability of capturing the homogeneity of the tissue inside the volumetric layer.

The algorithm was implemented using a level set method, a powerful numeric technique for computing interface motion. A coupled narrow band algorithm was cus-

tomized for the coupled surfaces propagation. The correspondence between points on the two bounding surfaces falls out automatically during the narrow band rebuilding which is required for surface propagation at each iteration. This shortest distance-based correspondence is essential in imposing the coupling between the two bounding surfaces through the thickness constraint. The level set implementation of the coupled surfaces propagation offers the advantages of easy initialization, computational efficiency and the ability to handle complex geometry and topology change, as well as the ready evaluation of several characteristics of the layer, such as surface curvature and a thickness map.

Our approach was applied successfully to the problem of cortical analysis. We presented extensive validation and experiments on both simulated MR brain images, as well as real data with a range of image qualities. Comparisons were also given between our algorithm and several classification-based techniques to show the power of the shape constraint incorporated. With our method, skull-stripping, cortical gray matter segmentation and measurements are done simultaneously. The geometric measurements made easy by our algorithm, such as the cortical surface shape index and the cortical thickness, are offering new insights into the understanding of brain structure and function through neuroimaging. The 3D intra-sulcal ribbon algorithm derived from the coupled surfaces work provides a new mechanism for the efficient analysis of sulcal patterns which are also of great interest in the study of neuroanatomy and related disorders.

As a starting point, the coupled surfaces algorithm has shown promise in Left Ventricular (LV) myocardium segmentation (see Figure 7.1 for an example). The advantage of its easy initialization, computational efficiency, etc., translates into cardiac image analysis as well.

There are, of course, areas of potential improvement that this work could benefit

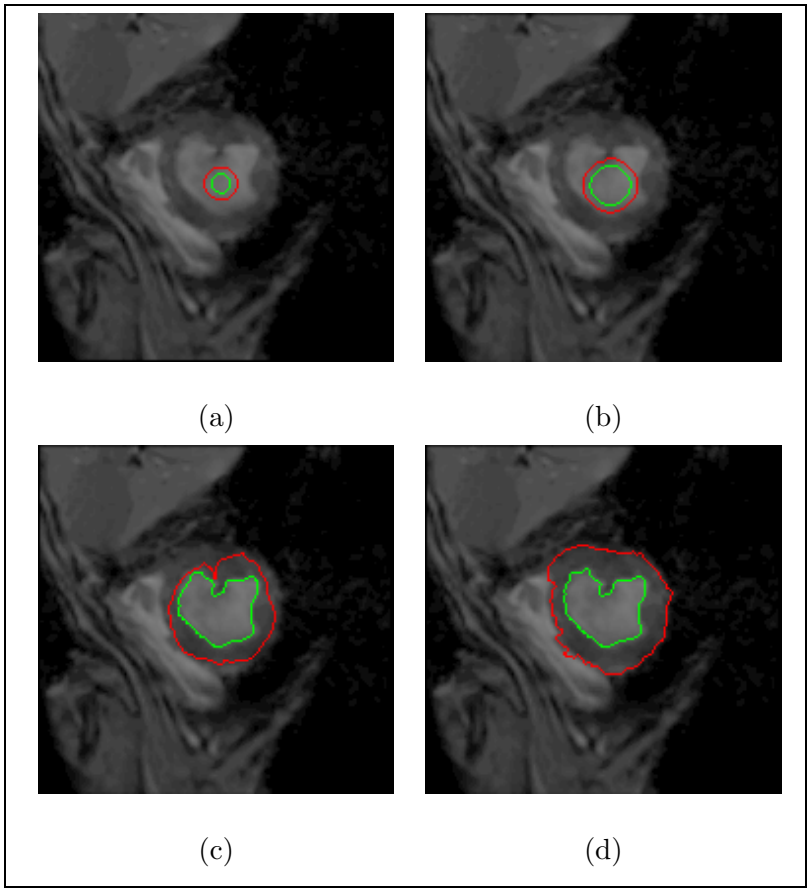

Figure 7.1: Example of LV Myocardium Segmentation. (a): initialization of curves representing epicardium (red) and endocardium (green) boundary; (b,c): intermediate steps in propagation; (c): final results of the captured epicardium and endocardium boundary.

from. Computationally, even though the speed of our algorithm already compares favorably with other methods available in the neuroscience community, further speed up could help to realize its full analytic power in neuroimaging studies. This could be done with more efficient use of data structures, and with the possibility of using multi-processors. In the current implementation, the narrow band rebuilding at each iteration takes up most of the processing time. Since updating the distance within the neighborhood of one point on surface does not affect the rest places, this rebuilding process is perfectly suitable for parallel programming.

The second possible improvement concerns better capturing the intensity homogeneity inside the volumetric layer. Currently, we simply check the immediately-surrounding voxels in the regions around the propagating front to calculate the image feature that partially determines the speed of the front. An investigation of voxels statistics within a strip of potential cortical gray along the current propagating surface and in between the coupled surfaces would help to make the segmentation process more robust to image noise and poor contrast.

The third potential improvement is to incorporate additional information that will help in trying to further capture of the full depth of each and every sulcus. Currently, there is some residual inaccuracy in imperfect sulcal delineation, especially in the case of tight sulci of young subjects. We may take advantage of the automatically identified high maximum principal curvature points that are indicators of sulci, and "sink" them locally into the brain volume with the help of information on the underlying cortical gray matter and inner cortical surface as well as its associated distance function.

Other directions of future research include localized sulcal measurement and shape analysis, and the study of structure and function relationships in sulcal regions.

Finally we discuss the generalization of this work. The core idea of this thesis work is using generic constraints to couple neighboring structures (in our case, the two bounding surfaces that are part of the definition of the volumetric layer) for the purpose of more robust and accurate segmentation. This idea can be further developed and generalized for applications to other interesting problems in medical image analysis. For example, the LV segmentation can potentially be helped by the information of the coupling between its epicardium and endocardium, as well as the coupling between LV and the right ventricle. Also, the segmentation of deep brain structures, such as the caudate, thalamus and lateral ventricle etc., could benefit from their simultaneous recovery with constraints emulating their neighboring relationships. This type of coupling offers an alternative to traditional template-based approaches by using more generic constraints, and therefore has the advantage of flexibility.

Bibliography

[1] K. Amunts, G. Schlaug, A. Schleicher, H. Steinmetz, A. Dabringhaus, P. E. Roland, and K. Zilles. Asymmetry in the human motor cortex and handedness. *Neuroimage*, 4(3 Pt 1):216–22, Dec 1996.

[2] S. Angenent, S. Haker, A. Tannenbaum, and R. Kikinis. Conformal geometry and brain flattening. *Proceedings of Medical Image Computing and Computer-Assisted Intervention*, pages 271–278, 1999.

[3] D. H. Ballard and C. M. Brown. *Computer Vision*, chapter 4. Prentice-Hall Inc., New Jersey, 1982.

[4] E. Bardinet, L. Cohen, and N. Ayache. Superquadrics and free-form deformation: a global model to fit and track 3D medical data. In *Proceedings of CVRMed*, pages 319–326, 1995.

[5] A. H. Barr. Global and local deformation of solid primitives. In *Computer Graphics (Proceedings of SIGGRAPH)*, volume 18, pages 21–30, 1984.

[6] R. Bellman and R. Kalaba. *Dynamic Programming and Modern Control Theory*. Academic Press Inc, 1965.

[7] S. M. Blinkov and I. I. Glezer. *The Human Brain In Figures and Tables. A Quantitative Handbook.*, chapter X. Basic Books, Inc. Publishers, Plenum Press, 1968.

[8] H. Blum. *Models for the Perception of Speech and Visual Form.* MIT Press, 1967.

[9] H. I. Bozma and J. S. Duncan. A game theoretic approach to integration of modules. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16:1074–1086, 1994.

[10] A. Chakraborty. *Feature and Module Integration for Image Segmentation.* PhD thesis, Yale University, 1996.

[11] H. Choi, D. Haynor, and Y. Kim. Multivariate tissue classification of MRI images for 3D volume reconstruction - a statistical approach. In *SPIE Medical Imaging XVII*, 1989.

[12] S. M. Choi, J. E. Lee, J. Kim, and M. H. Kim. Volumetric object reconstruction using the 3D-MRF model-based segmentation. *IEEE Transactions on Medical Imaging*, 16(6), 1997.

[13] H. Chui, J. Rambo, R. Schultz, J. Duncan, and A. Rangarajan. Registration of cortical anatomical structures via robust 3D point matching. In *Proceedings of Information Processing in Medical Imaging*, 1999.

[14] H. E. Cline, W. E. Lorensen, R. Kikinis, and F. Jolesz. Three-Dimensional segmentation of MR images of the head using probability and connectivity. *Journal of Computer Assisted Tomography*, 14(6):1037–1045, 1990.

[15] L. D. Cohen and I. Cohen. Finite-Element methods for active contour models and balloons for 2D and 3D images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1131–1147, 1993.

[16] D. Collins, T. Peters, W. Dai, and A. Evans. Model-based segmentation of individual brain structures from MRI data. In *Proceedings of the First Conference on Visualization in Biomedical Computing*, pages 10–23, 1992.

[17] D. L. Collins, G. L. Goualher, and A. C. Evans. Non-linear cerebral registration with sulcal constraints. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention*, pages 974–984, MIT, 1998.

[18] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models - their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995.

[19] A. M. Dale, B. Fischl, and M. I. Sereno. Cortical surface-based analysis. *NeuroImage*, pages 179–194, 1999.

[20] C. A. Davatzikos and R. N. Bryan. Using a deformable surface model to obtain a shape representation of cortex. *IEEE Transactions on Medical Imaging*, 15:785–795, 1996.

[21] C. A. Davatzikos and J. Prince. An active contour model for mapping the cortex. *IEEE Transactions on Medical Imaging*, 14:65–80, 1995.

[22] M. P. DoCarmo. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, New Jersey, 1976.

[23] P. A. Filipek, D. N. Kennedy, and V. S. Caviness. Volumetric analysis of central nervous system neoplasms based on MRI. *Pediatric Neurology*, pages 347–351, 1991.

[24] P. A. Filipek, C. Richelme, D. N. Kennedy, and V. S. Caviness. The young adult human brain: an MRI-based morphometric analysis. *Cerebral Cortex*, 4(4):344–360, 1994.

[25] D. Geman and S. Geman. Stochastic relaxation, Gibbs distribution and Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:721–741, 1984.

[26] N. Geschwind. Specializations of the human brain. *Scientific American*, pages 180–199, 1979.

[27] S. Gibson. Constrained elastic surface nets: Generating smooth surfaces from binary segmented data. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention*, pages 888–898, MIT, 1998.

[28] G. L. Goualher, C. Barillot, and Y. Bizais. Modeling cortical sulci using active ribbons. *International Journal of Pattern Recognition and Artificial Intelligence*, 11(8):1295–1315, 1997.

[29] G. L. Goualher, D. L. Collins, C. Barillot, and A. C. Evans. Automatic identification of cortical sulci using a 3D probabilistic atlas. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention*, pages 509–518, 1998.

[30] S. R. Gunn and M. S. Nixon. A robust snake implementation: a dual active contour. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(1):63–68, 1997.

[31] R. M. Haralick and L. G. Shapiro. *Computer and Robot Vision*. Addison-Wesley Publishing Company, Inc, 1993.

[32] K. Held, E. R. Kops, V. J. Krause, W. M. Wells III, R. Kikinis, and H. Muller-Gartner. Markov random field segmentation of brain MR images. *IEEE Transactions on Medical Imaging*, 16(6):878–886, 1997.

[33] The Internet Brain Segmentation Repository, Center for Morphometric Analysis, Massachusetts General Hospital. *http://neuro-www.mgh.harvard.edu/cma/ibsr*.

[34] T. Kapur, W. E. Grimson, R. Kikinis, and W. M. Wells. Enhanced spatial prior for segmentation of MRI. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention*, pages 457–468, MIT, 1998.

[35] T. Kapur, W. L. Grimson, W. M. Wells III, and R. Kikinis. Segmentation of brain tissue from magnetic resonance images. *Medical Image Analysis*, 1(2):109–127, 1996.

[36] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: active contour models. *International Journal of Computer Vision*, 1:312–331, 1988.

[37] D. N. Kennedy, P. A. Filipek, and V. S. Caviness. Anatomic segmentation and volumetric calculations in nuclear Magnetic Resonance imaging. *IEEE Transactions on Medical Imaging*, 8:1–7, 1989.

[38] N. Khaneja, M. I. Miller, and U. Grenander. Dynamic programming generation of curves on brain surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(11):1260–1265, 1998.

[39] J. J. Koenderink and A. J. Van Doorn. Surface shape and curvature scale. *Image and Vision Computing*, 10:557–565, 1992.

[40] R. K.-S. Kwan, A. C. Evans, and G. B. Pike. An extensible MRI simulator for post-processing evaluation. In *Visualization in Biomedical Computing*, volume 1131, pages 135–140. Springer-Verlag, 1996.

[41] S. Lakshmanan and H. Derin. Simultaneous parameter estimation and segmentation of Gibbs random fields using simulated annealing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11:799–810, 1989.

[42] R. Leahy, T. Hebert, and R. Lee. Applications of Markov Random Fields in medical imaging. In *Information Processing in Medical Imaging*, pages 1–14. Wiley-Liss Inc, 1991.

[43] Z. Liang, R. F. Jaszczak, and R. E. Coleman. Parameter estimation of finite mixtures using the EM algorithm and information criteria with application to medical image processing. *IEEE Transactions on Nuclear Science*, 39:1126–1133, 1992.

[44] W. Lorenson and H. Cline. Marching Cubes: a high resolution 3D surface construction algorithm. In *Proceedings of SIGGRAPH*, volume 21, pages 163–169, July 1987.

[45] D. MacDonald, D. Avis, and A. E. Evans. Multiple surface identification and matching in Magnetic Resonance images. In *SPIE*, volume 2359, pages 160–169, 1994.

[46] D. MacDonald, D. Avis, and A. E. Evans. Proximity constraints in deformable models for cortical surface identification. In *Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 650–659, 1998.

[47] R. Malladi, R. Kimmel, D. Adalsteinsson, G. Sapiro, V. Caselles, and J. A. Sethian. A geometric approach to segmentation and analysis of 3D medical images. In *MMBIA*, 1996.

[48] R. Malladi and J. A. Sethian. A real-time algorithm for medical shape recovery. In *International Conference on Computer Vision*, pages 304–310, 1998.

[49] R. Malladi, J. A. Sethian, and B. C. Vemuri. Shape modeling with front propagation: a level set approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17:158–174, 1995.

[50] J. Mangin, J. Regis, I. Bloch, V. Frouin, Y. Samson, and J. LopezKrahe. A MRF based random graph modeling the human cortical topography. In *Proceedings of CVRMed*, pages 177–183, Nice, 1995.

[51] B. S. Manjunath and R. Chellappa. Unsupervised texture segmentation using markov random field models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13:478–482, 1991.

[52] McConnell Brain Imaging Center, Montreal Neurological Institute, McGill University. *http://www.bic.mni.mcgill.ca/*.

[53] T. McInerney and D. Terzopoulos. Topologically adaptable snakes. In *Proceedings of the Fifth International Conference on Computer Vision*, pages 840–845, Cambridge, MA, 1995.

[54] T. McInerney and D. Terzopoulos. Medical image segmentation using topologically adaptable surfaces. In *Proceedings of CVRMed*, pages 23–32, 1997.

[55] Tim McInerney and Demetri Terzopoulos. Deformable models in medical image analysis. In *Proceedings of MMBIA*, pages 171–180, 1996.

[56] X. Meng and D. Dyk. The EM algorithm - an old folk song sung to a fast new tune. *J. R. Statis. Soc. B.*, 59(3):511–567, 1997.

[57] D. Mumford. Bayesian rational for the variational formulation. In B. H. Romeny, editor, *Geometry Driven Diffusion in Computer Vision*. Kluwer, 1994.

[58] L. R. Nackman and S. M. Pizer. Three-Dimensional shape description using the symmetric axis transformation I: Theory. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 7(2):187–202, 1985.

[59] M. Naf, O. Kubler, R. Kikinis, M. E. Shenton, and G. Szekely. Characterization and recognition of 3D organ shape in medical image analysis using skeletonization. In *Proceedings of MMBIA*, pages 139–150, 1996.

[60] W. Niessen, B. Romeny, and M. A. Viergever. Geodesic deformable models for medical image analysis. *IEEE Transactions on Medical Imaging*, 17:634–641, 1998.

[61] B. Pakkenberg and H. J. Gundersen. Neocortical neuron number in humans: effect of sex and age. *Journal of Comparative Neurology*, 384:312–320, 1997.

[62] N. R. Pal and S. K. Pal. A review of image segmentation techniques. *Pattern Recognition*, 26(9):1277–1294, 1993.

[63] J. Rademacher, V. S. Caviness Jr, H. Steinmetz, and A. M. Galaburda. Topographical variation of the human primary cortices: implications for neuroimaging, brain mapping, and neurobiology.. *Cerebral Cortex*, 3(4):313–329, 1993.

[64] P. Radeva, J. Serrat, and E. Marti. A snake for model-based segmentation. In *Proceedings of the Fifth International Conference on Computer Vision*, pages 816–821, 1995.

[65] J. C. Rajapakse and F. Fruggel. Segmentation of MR images with intensity inhomogeneities. *Image and Vision Computing*, 16(3):165–180, 1998.

[66] J. V. Rambo, R. T. Schultz, L. H. Staib, L. Win, X. Zeng, and J. S. Duncan. Fast interactive volumetric segmentation of MR brain images. *Submitted to NeuroImage*, 1999.

[67] J. V. Rambo, X. Zeng, R. T. Schultz, L. Win, L. H. Staib, and J.S.Duncan. Platform for visualization and measurement of gray matter volume and surface area within discrete cortical regions from MR images. In *Proceedings of the 4th International Conference on Functional Mapping of the Human Brain*, page 795, 1998.

[68] R. A. Redner and H. F. Walker. Mixture densities, Maximum Likelihood and the EM algorithm. *SIAM Review*, 26(2):195–239, 1984.

[69] R. A. Robb. *Analyze software, Biomedical Imaging Resource.* Mayo Foundation/Clinic, Rochester, MN, 1990.

[70] A. Rosenfeld and A. C. Kak. *Digital Picture Processing.* Academic Press, Inc, 1982.

[71] G. Schlaug, L. Jancke, Y. Huang, and H. Steinmetz. In vivo evidence of structural brain asymmetry in musicians. *Science*, pages 699–701, 1995.

[72] R. T. Schultz and A. Chakraborty. Magnetic resonance image analysis. In E. Bigler, editor, *Handbook of Human Brain Function: Neuroimaging*, pages 9–51. Plenum Press, 1996.

[73] R. T. Schultz, L. Win, L. Staib, X. Zeng, J. Rambo, R. Bronen, A. Rangarajan, and J. Duncan. *MRI Brain Morphometry Laboratory Manual*. Yale University, 1999.

[74] T. B. Sebastian, H. Tek, J. J. Crisco, S. W. Wolfe, and B. B. Kimia. Segmentation of carpal bones from 3D ct images using skeletally coupled deformable models. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention*, pages 1184–1193, 1998.

[75] T. W. Sederberg and S. R. Parry. Free-form deformation of solid geometric models. In *Computer Graphics (Proceedings of SIGGRAPH)*, volume 20, pages 151–160, 1986.

[76] L. D. Selemon, G. Rajkowska, and P. S. Goldman-Rakic. Abnormally high neuronal density in the schizophremic cortex, a morphometric analysis of Prefrontal Area 9 and Occipital Area 17. *Arch Gen Psychiatry*, 52:805–828, 1995.

[77] J. A. Sethian. *Level set methods: evolving interfaces in geometry, fluid mechanics, computer vision and materials science*. Cambridge University Press, 1996.

[78] K. Siddiqi, Y. Berube, A. Tannenbaum, and S. W. Zucker. Area and length minimizing flows for shape segmentation. *IEEE Transactions on Image Processing*, 7(3):433–443, 1998.

[79] F. Solina and R. Bajcsy. Recovery of parametric models from range images: the case for superquadrics with global deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(2):131–147, 1990.

[80] L. H. Staib and J. S. Duncan. Boundary finding with parametrically deformable models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(11):1061–1075, 1992.

[81] L. H. Staib and J. S. Duncan. Model-based deformable surface finding for medical images. *IEEE Transactions on Medical Imaging*, 16:720–731, 1996.

[82] C. Studholme, R. T. Constable, and J. S. Duncan. Incorporating an image distortion model in non-rigid alignment of EPI with conventional MRI. *Proceedings of Information Processing in Medical Imaging*, pages 454–459, 1999.

[83] G. Szekely, A. Lelemen, C. Brechbuhler, and G. Gerig. Segmentation of 2D and 3D objects from mri volume data using constrained elastic deformation of flexible Fourier surface models. *Medical Image Analysis*, 1(1), 1996.

[84] H. D. Tagare. Deformable 2-d template matching using orthogonal curves. *IEEE Transactions on Medical Imaging*, 16(1):108–117, 1997.

[85] J. Talairach and P. Tournoux. *Co-planar Stereotaxic Atlas of the Human Brain.* Thieme, New York, 1988.

[86] P. C. Teo, G. Sapiro, and B. A. Wandell. Creating connected representations of cortical gray matter for functional MRI visualization. *IEEE Transactions on Medical Imaging*, 16:852–863, 1997.

[87] D. Terzopoulos and D. Metaxas. Dynamic 3D models with local and global deformation: deformable superquadrics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7):703–714, 1991.

[88] M. Vaillant and C. Davatzikos. Finding parametric representations of the cortical sulci using an active contour model. *Medical Image Analysis*, pages 295–315, Sep 1997.

[89] M. Vaillant and C. Davatzikos. Mapping the cerebral sulci: application to morphological analysis of cortex and to non-rigid registration. In *Proceedings of Information Processing in Medical Imaging*, pages 141–154, 1997.

[90] B. Vemuri and A. Radisavljevic. Multiresolution stochastic hybrid shape models with fractal priors. *ACM Transactions on Graphics*, 13(2), 1994.

[91] Y. Wang and L. H. Staib. Boundary finding with correspondence using statistical shape models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 338–345, 1998.

[92] S. Webb, editor. *The Physics of Medical Imaging.* Institute of Physics Publishing, 1995.

[93] W. M. Wells, W. Grimson, R. Kikinis, and F. A. Jolesz. Adaptive segmentation of MRI data. *IEEE Transactions on Medical Imaging*, 15:429–442, 1996.

[94] L. E. White, T. J. Andrews, C. Hulette, A. Richards, M. Groelle, J. Paydarfar, and D. Purves. Structure of the human sensorimotor system. II: lateral symmetry.. *Cerebral Cortex*, 7(1):31–47, 1997.

[95] L. E. White, G. Lucas, A. Richards, and D. Purves. Cerebral asymmetry and handedness. *Nature*, 368(6468):197–8, 1994.

[96] D. Wilson and J. A. Noble. Segmentation of cerebral vessels and aneurysms from MR angiography data. In *Proceedings of Information Processing in Medical Imaging*, pages 423–428, 1997.

[97] C. Xu, D. L. Pham, J. L. Prince, M. E. Etemad, and D. N. Yu. Reconstruction of the central layer of the human cerebral cortex from MR images. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention*, pages 481–488, MIT, 1998.

[98] M. X.H. Yan and J. S. Karp. An adaptive bayesian approach to 3D mr brain segmentation. In *Proceedings of Information Processing in Medical Imaging*, pages 201–214, 1995.

[99] X. Zeng, L. H. Staib, R. T. Schultz, and J. S. Duncan. Segmentation and measurement of the cortex from 3D MR images. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention*, pages 519–530, MIT, 1998.

[100] X. Zeng, L. H. Staib, R. T. Schultz, and J. S. Duncan. Volumetric layer segmentation using coupled surfaces propagation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 708–715, Santa Barbara, CA, 1998.

[101] X. Zeng, L. H. Staib, R. T. Schultz, and J. S. Duncan. Segmentation and measurement of the cortex from 3D MR images using coupled surfaces propagation. *IEEE Transactions on Medical Imaging*, 1999.

[102] X. Zeng, L. H. Staib, R. T. Schultz, H. Tagare, L. Win, and J. S. Duncan. A new approach to 3D sulcal ribbon finding from MR images. In *Proceedings of Medical*

Image Computing and Computer-Assisted Intervention, pages 148–157, Cambridge, UK, 1999.

[103] Y. Zhou, P. M. Thompson, and A. W. Toga. Extracting and representing the cortical sulci. *IEEE Computer Graphics and Applications*, pages 49–54, 1999.

[104] S. C. Zhu and A. Yuille. Region competition: unifying snakes, region growing, and Bayes/MDL for multiband image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(9):884–899, 1996.