

# Automated 3-D Extraction of Inner and Outer Surfaces of Cerebral Cortex from MRI

David MacDonald,\* Noor Kabani,\* David Avis,† and Alan C. Evans\*

\*McConnell Brain Imaging Centre, Webster 2B, Montréal Neurological Institute, 3801 University, Montréal, Québec, Canada H3A 2B4; and †School of Computer Science, McGill University, Montréal, Québec, Canada

Received June 22, 1999

**Automatic computer processing of large multidimensional images such as those produced by magnetic resonance imaging (MRI) is greatly aided by deformable models, which are used to extract, identify, and quantify specific neuroanatomic structures. A general method of deforming polyhedra is presented here, with two novel features. First, explicit prevention of self-intersecting surface geometries is provided, unlike conventional deformable models, which use regularization constraints to discourage but not necessarily prevent such behavior. Second, deformation of multiple surfaces with intersurface proximity constraints allows each surface to help guide other surfaces into place using model-based constraints such as expected thickness of an anatomic surface. These two features are used advantageously to identify automatically the total surface of the outer and inner boundaries of cerebral cortical gray matter from normal human MR images, accurately locating the depths of the sulci, even where noise and partial volume artifacts in the image obscure the visibility of sulci. The extracted surfaces are enforced to be simple two-dimensional manifolds (having the topology of a sphere), even though the data may have topological holes. This automatic 3-D cortex segmentation technique has been applied to 150 normal subjects, simultaneously extracting both the gray/white and gray/cerebrospinal fluid interface from each individual. The collection of surfaces has been used to create a spatial map of the mean and standard deviation for the location and the thickness of cortical gray matter. Three alternative criteria for defining cortical thickness at each cortical location were developed and compared. These results are shown to corroborate published postmortem and *in vivo* measurements of cortical thickness.** © 2000

Academic Press

## 1. INTRODUCTION

In recent years, there has been a rapid proliferation of algorithms for segmentation of MRI data into identifiable anatomical objects that can be visualized, ma-

nipulated, and measured. For brain imaging, segmentation has found widespread application in basic neuroscience, neurological research, clinical neurology, and neurosurgery planning. However, the inherent limitations of image under-sampling, lack of contrast, intensity bias, and noise have made the segmentation of fine anatomic details a difficult task. This problem is nowhere more evident than in the highly convoluted human neocortex. Until recently, the accurate extraction of the cortical surface has been an exceedingly difficult task. Partial volume effects blur the distinction between closely adjacent surfaces in deep sulci, leading to a well-known segmentation error in which the deeper reaches of sulci are not penetrated by the putative surface model. As a result, any attempt to extract morphometric information (e.g., cortical area, percentage of buried cortex, or cortical thickness) is meaningless. Similarly, any attempt to relate structural anatomy and functional data from positron emission tomography (PET) or functional magnetic resonance imaging (fMRI) activation studies with cortical morphology in 3-D or 2-D (as flattened surfaces) is a hazardous undertaking. We present here a method for overcoming these problems, using a multiple-surface deformation algorithm, called ASP, which stands for anatomic segmentation using proximities. Topology constraints ensure that the cortical mantle, defined by its exterior and interior boundaries, is properly identified in areas susceptible to partial volume errors. This results in a more accurate sulcal penetration and a measure for gray matter thickness at every point over the whole cortical mantle. The method for extracting a 3-D geometric model of human cortical surface is presented, as well as results from validation experiments and techniques for combining information from large population datasets. Finally, measures of cortical gray matter thickness are presented and compared to related imaging and postmortem studies.

The underlying method of cortex identification presented here is a deformable surface technique. Deformable models provide solutions for registration, segmentation, and matching tasks in computational

neuroanatomy, by combining the bottom-up approach of edge detection with the top-down approach of model-based constraints. Purely bottom-up approaches such as the "Marching Cubes" method (Lorenson and Cline, 1987) provide high resolution representations of data but suffer from their inability to consider information that is not present in the data. One consequence is that bottom-up methods have limited means to compensate for noise. Another deficiency is that the topology of the reconstructed surface cannot be controlled, even though the correct topology is often known. Furthermore, it is difficult to incorporate domain-specific constraints on the reconstruction process, information that could help resolve ambiguous regions. The active contour method of Kass *et al.* (1988), commonly referred to as "SNAKES," attempts to address these issues and has been the foundation upon which many deformation methods have been based. Essentially, a 2-D or 3-D spline function is assigned an energy function that consists of a constraint on stretching and bending based on first and second derivatives and an image term that decreases in energy as the spline moves closer to image boundaries defined by a change in intensity. Numerical integration techniques deform the spline from a starting position to a minimum energy configuration, which represents a compromise between the shape constraints of the model and the edge features of the image. Shape constraints have been used to reduce sensitivity to noise, impose an expected class of shapes on data, and to help formulate mappings among models and individual datasets. As well, other complex constraints based on *a priori* models are readily incorporated into the deformation process. For these reasons, deformable models have been chosen as the framework upon which to build the novel segmentation algorithm described in this paper.

Many adaptations of deformable models in medical imaging have since been presented (Collins, 1992; Davatzikos and Bryan, 1995; Sandor and Leahy, 1997; Cohen and Cohen, 1993; Staib and Duncan, 1992; Xu *et al.*, 1998; Zeng *et al.*, 1998) and a survey of these methods is presented in (McInerney and Terzopoulos, 1996). Some of these methods specifically address the problem of locating the cortical gray matter of human brains in three-dimensional images. This is a particularly difficult task due to the high degree of geometric complexity of the cortex, combined with the noise, under-sampling, inhomogeneities, and other sampling artifacts, which are often present in the images. Several approaches have used information about related anatomical structures, notably white matter, to improve the segmentation of the cortical gray surface. The solution of Dale and Sereno (1993) performs morphological operations to improve the topology of a binary white matter volume and then expands a deformable surface from the white matter boundary toward the outer gray matter boundary. The use of the white mat-

ter information is an important advance, but the surfaces produced have two drawbacks. First, the algorithm as described does not prevent self-intersecting topologies, which occur as opposite banks of a sulcus are expanded toward each other, and, second, it is difficult to guarantee that the resulting surface has the topology of a sphere by using morphological operators.

Davatzikos and Bryan (1995) present an active contour method that models the cortical gray matter as a finite thickness sheet. The sheet has a two-dimensional parameterization, making it suitable for subsequent morphometric analysis. Rather than parameterizing the complete folded cortical surface, the method parameterizes the outer boundary of the cortex, only entering the upper portion of each sulcus. Location of the depths of the sulcus is performed as a second step, where a curve is initialized at the top of each sulcus and pushed down into the depths by a method similar to the Snakes algorithm. While this method is novel in its attempts to locate the depths of the sulci, the total surface of the cortex is not contained in a single model. Coordinating the two models used, the outer surface model and the set of deep curves, into a single model that faithfully represents the true folded cortex is a nontrivial task. The curves represent the center of the sulcus, providing no information about the width of the sulcus, which would be necessary to reconstruct the folded surface.

A more recent method by Zeng *et al.* (1998) uses a level set method to represent the deformable surface as the zero set of an implicit function. Two concentric surfaces are deformed with intersurface distance constraints to impose a predicted range of thickness on the gray matter, resulting in a more robust identification of the cortical gray surface. Again, the limitation of this algorithm is that the resulting surface may not be a two-dimensional manifold. The correspondence between points on the gray and white matter surfaces is computed as the nearest point, which is one of the three measures of cortical thickness that will be evaluated here. The deficiencies of this measure of thickness relative to the other two will be shown.

In order to address the deficiencies in the above segmentation methods, we present a novel deformable surface method (ASP—*anatomic segmentation using proximities*), which creates simple (non-self-intersecting) cortical gray matter surfaces with topologies isomorphic to the sphere (two-dimensional manifolds). Most deformable surfaces typically use stretching and bending constraints for a regularization effect, which penalizes but does not fully prevent nonsimple (self-intersecting) surface geometries. The deformable method presented here *explicitly* prevents self-intersecting configurations. This allows the deforming surface to be less stiff and thus better able to capture fine detail, while guaranteeing that the topology of the surface will not be changed by arbitrary self-intersections.

This will be shown in a simple validation experiment where conventional methods produce a nonsimple surface, whereas the ASP algorithm does not. The idea of using other anatomical structures to guide the identification of the outer gray matter surface is extended here to model two concentric surfaces with explicit proximity links between a dense field of corresponding points on both surfaces. A large number of sample points is used on each surface to produce a faithful approximation to the complex in-folded neuroanatomical surfaces.

The following section describes the deformable method as a minimization problem, consisting of an objective function and a procedure for numerical minimization. Some simple validation experiments are presented and then the application of ASP to identifying cortical surfaces from normal human MR images is outlined. The results of applying this to 150 subjects are presented in a variety of forms, including a mean surface, a mean thickness map, and a mean curvature map. Discussion of measures of cortical thickness obtained with ASP and the relevance of these results to conventional neuroanatomic knowledge of the cortex is presented.

## 2. METHOD

The essence of the ASP algorithm is the formulation of an objective function that when minimized deforms a set of surfaces to fit some models and data. The domain of the function is the set of vertex coordinates describing one or more polyhedra to be deformed, and the range is a scalar value representing a goodness of fit of the polyhedra to the target data.

### 2.1. Objective Function

The objective function to be minimized,  $O(S)$ , may be defined generally as a weighted sum of  $N_t$  terms, each of which may be thought of as a data or model term, depending on whether it constrains the *deforming* polyhedra to match image data or some model-based *a priori* information:

$$O(S) = \sum_{k=1}^{N_t} T_k(S),$$

where each term  $T_k(S)$  measures some aspect of  $S$ , a set of  $N_s$  deforming polyhedral surfaces,

$$S = \{S_i; S_i \text{ is a polyhedral surface, } 1 \leq i \leq N_s\}.$$

Each term,  $T_k(S)$ , is formulated as:

$$T_k(S) = W(D_k(S)),$$


where  $D_k(S)$  is a signed scalar measure of deviation from some ideal, and  $W(x)$  is a general weighting function. Usually this is just a squaring function,  $W(x) = ax^2$ , where  $a$  is a constant indicating the weight of this term relative to the other terms. However, by making  $a$  a function of  $x$ , one can introduce constraints into the system by use of functions that increase rapidly when  $x$  passes a certain threshold value. For the sake of simplicity, the objective terms are presented here as equations that are in the form of a simple squaring of a deviation from ideal, with the weight  $a$  equal to 1. Where appropriate it will be indicated how an appropriate  $a(x)$  weighting function can be used to enforce hard constraints in the system.

**2.1.1. Definitions.** We have previously defined the set of deforming surface,  $S$ . For the model terms, we also define a set of companion surfaces,  $\hat{S}$ , which is a set of  $N_s$  model polyhedral surfaces where each  $\hat{S}_i$  has the same topology as  $S_i$ , i.e., the number of nodes and indices of connected nodes are equivalent. These are used to define shape-based ideals. We now present some definitions that are used in the objective functions that follow:  $\bar{x}_v = (x_v, y_v, z_v)$ , 3-D position of vertex  $v$  in a deforming polyhedral mesh,  $S_i$ ;  $\hat{x}_v = (\hat{x}_v, \hat{y}_v, \hat{z}_v)$ , 3-D position of vertex  $v$  in a static model polyhedral mesh,  $\hat{S}_i$ ;  $n_v$ , the number of vertices in a polyhedral mesh;  $n_e$ , the number of edges in a polyhedral mesh;  $n_p$ , the number of polygons in a polyhedral mesh;  $m_v$ , the number of neighbors of vertex  $v$ ;  $d(\bar{x}, \bar{y})$ , the 3-D Euclidean distance between two points;  $n_{v,j}$ , the  $j$ th neighbor of vertex  $v$ ;  $\bar{N}_v$ , the surface normal at vertex  $v$ , defined as the unit normal to the polygon consisting of the counterclockwise ordered neighbors of the vertex,  $v$ .

**2.1.2. Image term.** Many conventional methods use an image term based on the image value at a point on the deforming surface, implicitly using the local gradient of the image to “push” the surface toward the correct edge. However, if the deforming surface is far from an edge, or the gradient direction is misleading, there may be difficulties locating the “correct” edge. For this reason, the image term presented here is based on the distance from a vertex on the deforming surface to the nearest image boundary in the direction of the local surface normal (see Fig. 1) and is expressed as

$$T_{\text{boundary-dist}} = \sum_{v=1}^{n_v} d_B(\bar{x}_v, \bar{N}_v, t)^2,$$

where  $d_B(\bar{x}_v, \bar{N}_v, t)$  is the distance to the nearest image contour of the threshold,  $t$ , from the vertex,  $v$ , along the line defined by the surface normal,  $\bar{N}_v$  (see Fig. 2). The explicit search in both directions along the surface normal increases the power of locating image bound-



**FIG. 1.** Searching along the surface normal provides improved edge finding as compared to searches based on local gradient.

aries that are relatively far from the current surface position, because it is less sensitive to local variations in gradient. However, the search is truncated after a certain distance in order to prevent very distant boundaries from dominating the objective function. This term may be modified to search for other image features, such as the maximum gradient magnitude or zero crossings of the Laplacian. In addition, the image term may be oversampled between vertices to match the sampling of the data.

**2.1.3. Stretch term.** The stretch term increases as lengths between vertices are stretched or compressed relative to a user-defined model surface representing the ideal lengths,

$$T_{\text{stretch}} = \sum_{v=1}^{n_v} \sum_{j=1}^{m_v} \left( \frac{d(\bar{x}_v, \bar{x}_{n_{v,j}}) - L_{v,j}}{L_{v,j}} \right)^2,$$

where  $L_{v,j}$ , the ideal length of an edge, is defined as the corresponding length in the model polyhedron:

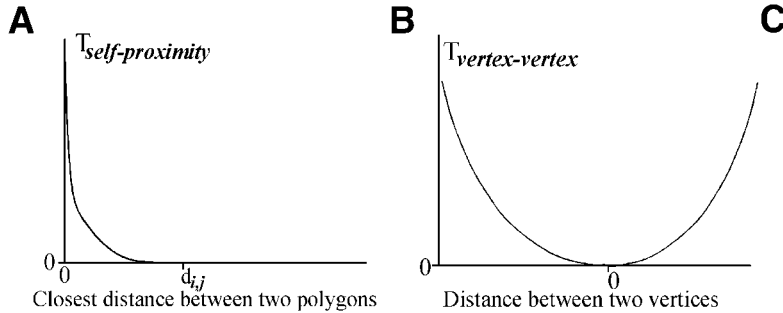

$$L_{v,j} = d(\hat{x}_v, \hat{x}_{n_{v,j}}).$$

The intended effect of this term is to make distances between corresponding pairs of vertices on the model and deformed surface roughly equivalent and is analogous to the term involving the magnitude of the first derivative of the spline in the original Snakes formulation (Kass *et al.*, 1988). One can consider this term a regularization force, maintaining equivalence of corresponding edge lengths between the model and the deforming surface.

**2.1.4. Bending term.** The bending term provides a measure of deviation from a model shape based on an estimate of local curvature and is analogous to the second derivative term in the Snakes formulation (Kass *et al.*, 1988),

$$T_{\text{bend}} = \sum_{e=1}^{n_e} (a(S, e) - a(\hat{S}, e))^2,$$

where  $a(S, e)$  is the signed angle between the two polygons adjacent to the edge,  $e$ . This term can be considered a regularization force like the stretching term, but is also intended to be used for shape-based matching and segmentation. By keeping the deforming surface similar to a preferred model shape, points on the de-



**FIG. 2.** Shapes of terms: (A)  $T_{\text{boundary-dist}}$ , (B)  $T_{\text{stretch}}$ , (C)  $T_{\text{bend}}$ , (D)  $T_{\text{self-proximity}}$ , (E)  $T_{\text{vertex-vertex}}$ .

forming surface will be encouraged to move toward places in the image that have the same shape as the model, thus automatically improving the match between the model and the data.

*2.1.5. Self-proximity term and intersurface proximity term.* The previous three terms are found in some form in most conventional deformable models. Here we introduce the novel self-proximity term, which measures the proximity of pairs of nonadjacent polygons in a surface,

$$T_{\text{self-proximity}} = \sum_{i=1}^{n_p-1} \sum_{j=i+1}^{n_p} \begin{cases} (d_{\min}(P_i, P_j) - d_{i,j})^2, & \text{if } d_{\min}(P_i, P_j) < d_{i,j} \\ 0, & \text{otherwise,} \end{cases}$$

where  $d_{\min}(P_i, P_j)$  is the smallest Euclidean distance between the  $i$ th polygon,  $P_i$ , and the  $j$ th polygon,  $P_j$ , and  $d_{i,j}$  is a distance threshold. Pairs of adjacent polygons are not included in the above equation, as the distance  $d_{\min}(P_i, P_j)$  is a constant zero value for any deformation of the polyhedra. The self-proximity term is used to explicitly prevent nonsimple topologies by assigning a prohibitively high cost to self-intersecting topologies. As shown in Fig. 2d, the weighting function shown above must be modified to increase asymptotically toward infinity as the inter-polygon distance approaches zero, which will be described later under 2.3. The intersurface proximity term,  $T_{\text{surface-surface}}$ , is formulated in a similar fashion, and is used to prevent two surfaces from coming within a certain distance of each other.

*2.1.6. Vertex-vertex proximity constraints.* A related proximity term is used to designate that corresponding points on two surfaces prefer to be a certain distance apart. Assuming two polyhedra have the same topology, the term constraining the distance between two corresponding vertices is:

$$T_{\text{vertex-vertex}} = (d(\bar{x}_v, \bar{x}_w) - d_B)^2,$$


where  $d_B$  is the preferred distance between vertex  $v$  on one surface and vertex  $w$  on a second surface and can be a constant for all pairs of vertices, as shown here, or can vary across the surface. This term keeps specific points of two surfaces a preferred distance apart, but does not explicitly prevent intersurface intersection, which is achieved by the intersurface proximity term defined previously. Typically this term can be over-sampled uniformly over the polyhedra to match the sampling of the data.

## 2.2. Minimization of Objective Function

Having formulated an objective function for fitting surfaces to data, we now have to minimize it. For all but the most trivial examples, the size of the domain of

parameter space that must be searched prohibits any form of exhaustive search. For example, conventional polygonal representations of human cortical surface have ranged from a few hundred to half a million vertices, which is intractable for a grid search even with only two samples per dimension. The conventional approach of directed searching from an initial guess is employed here. From an initial configuration of polyhedra, deformation proceeds by minimization of the objective function using a conjugate gradient approach (Press *et al.*, 1988). This method involves iteratively computing line minimizations along directions computed from successive derivatives. In order to increase the chances of finding the global minimum, a multiscale approach is employed. Deformation begins with a low-resolution initial guess for each of the polyhedral surfaces being deformed, which may be a hand-crafted model or statistically generated approximation to the surfaces being identified. Low-resolution surfaces are deformed to fit the image data, resampled to contain more triangles, and the process is repeated several times. At the final resolution, triangles with sides of about 1 mm are sufficient to capture the surfaces in the 1-mm isotropic voxel MR data being segmented. This makes the algorithm less sensitive to noise at the earlier levels, where the low resolution nature of the surfaces does not admit noisy representations. Conventional multiscale algorithms also apply a decreasing blur to the image data during the iterative fitting, using the original image only at the final scale. However, we have chosen to use oversampling of the surface objective function to achieve the same noise reducing effect without blurring the data. One advantage is that the apparent topology of the data does not change during the fitting, a situation that can occur during blurring and cause the deforming model to have to readjust itself. This is particularly important when segmenting the cortex because large sulcal areas that could be represented early on in multiscale space might not be reached until later because blurring closes the thin opening providing access to the region. By using super-sampling of the objective function instead of data blurring, the surfaces represent the best fit to the original data at each scale space iteration and avoid having to take many steps later when it is more costly. Thus, the multiscale approach can also increase the speed of the algorithm significantly by taking large steps when the surfaces are at lower resolutions and are therefore less expensive to deform.

The multiscale deformation is implemented as a sequence of deformations where the weights and other parameters may have different values for each step. For instance, the stretching constraints are relaxed somewhat as the surface gets closer to a solution in order to better interpolate the data. There remains some investigation into the relative benefits of changing some of the other parameters as the surfaces converge.




**FIG. 3.** (A) Three-dimensional view of fit without self-intersection constraints. (B) Cross-section of image and surface from (A). (C) Three-dimensional view of fit with self-intersection constraints. (D) Cross-section of image and surface from (C).

2.3. Choice of Weights


One of the drawbacks of using a deformable surface method is the requirement for user-defined weights to control the tradeoffs among the various terms. In addition, it is desirable to choose weights that reduce curvature in the objective function, to allow speedy convergence to the solution, and reduce the number local minima. Therefore the choice of weighting functions is a critical issue. For the self-intersection and intersurface proximity constraints, the weighting functions are constructed so that the user can specify distances instead of weights. For instance, to prevent self-intersection, a small threshold distance is specified. If two triangles are greater than this distance

apart, the weight is zero, ignoring the pair. Below this distance, the weight increases exponentially from a small number ( $10^{-10}$ ) to a large number ( $10^{10}$ ) as the distance approaches zero, in an attempt to keep the objective function as smooth as possible. For example the self-proximity term of section 2.1.5 would now be written as:

$$T_{\text{self-proximity}} = \sum_{i=1}^{n_p-1} \sum_{j=i+1}^{n_p} \begin{cases} e^{\log 10^{-10} + (10^{10} - 10^{-10}) \frac{d_{i,j} - d_{\min}(P_i, P_j)}{d_{i,j}}} x^2, & \text{if } d_{\min}(P_i, P_j) < d_{i,j} \\ 0, & \text{otherwise.} \end{cases}$$



**FIG. 4.** (A) Cross-section of image representing a sulcus in which the distinction between opposing banks of the sulcus has been obscured by partial volume. (B) Cross-section of apparent gray/CSF boundary extracted by conventional single surface method. (C) Cross-section of apparent gray/white boundary. (D) Cross-section of dual-model deformation. (E) Dual-model gray/CSF surface. (F) Dual-model gray/white surface.



**FIG. 5.** Response of ASP to gray thickness outside normal range. Here the phantom's cortical thickness varied from 1 to 15 mm, i.e., far beyond normal values. The algorithm constrained the thickness to between 2 and 6 mm.

### 3. VALIDATION

The following sections present several experiments which examine various aspects of the ASP algorithm independently.

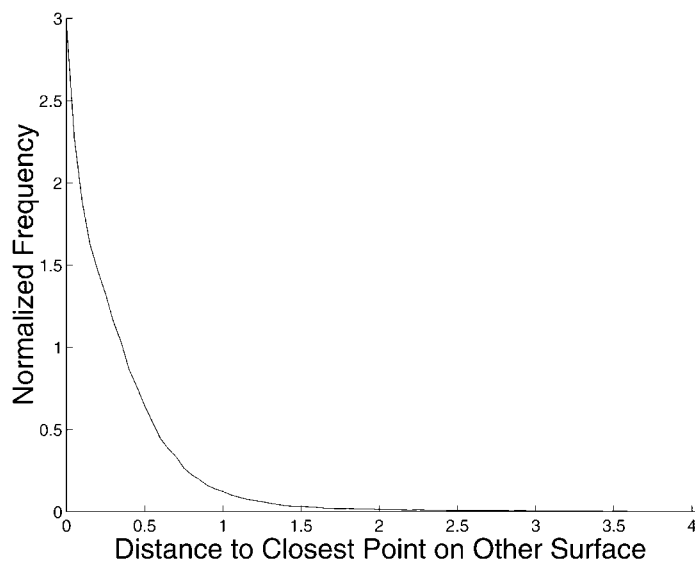
#### 3.1. Representing Surfaces with Enforced Topology

In order to verify that the surface deformation method correctly imposes a simple spherical topology on a data set with topological holes, a simple test was performed. A torus image was produced, and a sphere fit to this image, with and without self-intersection constraints. In both


cases an approximation to the torus was produced that spanned the hole of the torus (Fig. 3). Without self-intersection constraints, the resulting surface started to wrap itself around the torus a second time (Fig. 3B), stopped only by limits on stretching. With self-intersection constraints, a simple surface resulted (Fig. 3D). The use of explicit self-intersection avoidance is necessary when attempting to extract simple surfaces from images which may have topological holes, as is the case with the cortical mantle.

#### 3.2. Solving Partial Volume Effects with a Double Surface Model

Figure 4A illustrates a cross-section through a three-dimensional simulated sulcus phantom. Conventional single surface deformable methods find the sulcus in the white matter (Fig. 4C), but fail to find a sulcus in the gray matter (Fig. 4B) due to partial volume effects in the image. The ASP method successfully locates a reasonable approximation to the gray matter sulcus (shown in three dimensions in Fig. 4E), using a double surface formulation. The gray/CSF (Fig. 4E) and the gray/white (Fig. 4F) surfaces are simultaneously deformed to fit the image under an intersurface distance of  $4 \pm 2$  mm. The gray/CSF surface follows the gray/white surface deep into the sulci, and it is important to note that without self-proximity constraints, the opposite banks of the sulcus would cross over each other as they are moved away from the white matter. In this example the model cortex was of uniform thickness. Figure 5 shows how ASP applies the thickness constraints on a pathological phantom dataset, where the cortical thickness ranges from below to far above the a



**FIG. 6.** Histogram of distance between true surface and approximating surface, normalized so that area under curve equals unity. The histogram covers a grid of points on both surfaces and shows the minimum distance of each point to the closest point on the opposite surface. The actual histogram x axis continues to a maximum of 11.5 mm, but only the part that is sufficiently greater than zero is shown.



**FIG. 7.** (A) Deformation of sphere to fit white matter image model. The mean error between the surface and the “ground truth” was less than half a voxel. (B) sagittal section, (C) coronal section, and (D) transverse section of deformed surface (red) superimposed on a voxel map of white matter (black).


*priori* thickness. There are areas where the thickness constraints have prevented the surface from matching the edges in the data because doing so would violate the model. Expanding the range of allowable thicknesses would provide better faithfulness to the data, but reduce the ability of the algorithm to pull the outer surface into deep, narrow sulci, as in Fig. 4.

### 3.3. Ability to Approximate Complicated Structures

Since the ultimate goal is the application of ASP to MR images of human brains, an experiment was performed to determine how accurately the method can approximate a set of image boundaries with the complexity of the human brain. A dataset where the answer is known is required for this experiment. For the purpose of testing the ability of the ASP algorithm to deform surfaces into the highly convoluted shape of the human cortex, it is more important the data set have a known answer than that it be perfectly anatomically correct. Therefore, a reasonable data set to define as the “ground truth” for this experiment is the white


matter model previously devised for an MR simulator (Kwan *et al.*, 1996). The simulator models the properties of the various tissue types and creates a simulated MR image with user-specified amounts of noise, RF inhomogeneities, and other imaging artifacts, based on some models of human brain tissue. In this tissue model dataset, every voxel has been automatically labeled by tissue type (gray, white, CSF), using a neural network-based classification algorithm (Zijdenbos *et al.*, 1993) applied to an actual high resolution MR image. There is no explicit information that any given voxel is part of a particular anatomical object such as cortex. Given this volume of tissue types, we can define the “ground truth” or “true” surface as the set of all points on the boundary between a white pixel and a nonwhite pixel. A surface was extracted automatically by fitting a sphere to the voxelated data, using ASP.

The resulting surface was compared to the true surface by finding the mean of the distances between a grid of points on one surface and the nearest points on the other surface, and vice versa, interchanging the roles of the two



**FIG. 8.** Gray-CSF boundary automatically extracted with the dual-surface model.






**FIG. 9.** Gray-CSF boundary automatically extracted with the dual-surface model.

surfaces. The combined mean difference of this bidirectional distance measure was 0.34 mm (about a third of a voxel) with a standard deviation of 0.46. Figure 6 shows a histogram of this error, indicating that most points on either surface are within 1 mm of the other surface, with the maximum error being 11.5 mm. The larger errors occur predominantly on the under-surface of the cortex, where the deformation algorithm finds it difficult to stretch the surface sufficiently to represent this region. As well as being a sufficiently close approximation, the deformed surface has the advantage of an enforced simple topology and continuity constraints. The deformed surface and cross-sections of the surface and image model are shown in Fig. 7. The largest areas of discrepancy are in the inferior aspect, where the surface has reached up

into the ventricles. However, subsequent analysis is generally not concerned with these areas and this region can be masked off and ignored. Alternatively, the shape parameters of the model may be tuned in this area to force the deforming surface to simply make a smooth cut across the brain stem. This additional constraint was not imposed in these experiments in order to evaluate the results of the basic algorithm without extra model constraints.

#### 4. APPLICATION TO CORTICAL SURFACE SEGMENTATION

ASP was applied to a large number of normal human MR images to automatically identify the total surface



**FIG. 10.** Gray-white boundary automatically extracted with the dual-surface model.




FIG. 11. Gray-white boundary automatically extracted with dual-surface model.

of the cerebral cortical gray matter. Subjects were selected from a data base of 150 young normal subjects that were acquired as part of the ICBM project (Mazziotta *et al.*, 1995). These data were scanned on a Philips Gyroscan ACS 1.5 Tesla super-conducting magnet system at the Montreal Neurological Institute using a T1-weighted 3-D spoiled gradient-echo acquisition with sagittal volume excitation (TR = 18, TE = 10, flip angle = 30°, 1 mm isotropic voxels, 140–180 sagittal slices). T2-weighted and proton density-weighted images were acquired at a resolution of 1 mm in plane and 2 mm out of plane. The images were corrected for RF inhomogeneity artifacts (Sled *et al.*, 1997), linearly transformed into a stereotaxic coordinate system (Collins *et al.*, 1994), and classified into gray matter, white matter, and CSF (Zijdenbos *et al.*, 1993). It is important to note that these steps provide an advantageous frame of reference for subsequent analysis by reducing intensity biases and gross spatial

differences in all images. The identification of the cortical surface was accomplished in two steps. The first step uses high stretching and bending weights to rigidly deform a coarse cortex mask to fit the classified volume, thereby removing noncerebral white matter from the volume. Then the dual-surface deformation described previously was performed on the masked volume, in a multiscale fashion. The initial surfaces each contain 320 triangles, and after several iterations of deformation and subsampling the surfaces, the resulting surfaces contain 81,920 triangles each. At this surface sampling, the average intervertex distance was about 1 mm, equivalent to the sampling of the image. The segmentation for a single subject took about 30 h of time on a Silicon Graphics Origin 200 R10000 processor running at 180 megahertz. Much of this expense is attributable to the computation of the self-intersection constraints, but the reward is the guarantee of a simple surface. Figures 8 and 10 show the top, bottom,





FIG. 12. Loose consistency of vertex placement is illustrated by labeling the 1668th vertex on four different cortical surfaces.




**FIG. 13.** Six views of the average of 150 normal cortical surfaces. Note the prominence of major gyral and sulcal features such as the central sulcus, the superior temporal sulcus, and the superior frontal gyrus. (A) top, (B) left, (C) back, (D) bottom, (E) right, (F) front.

and right views of the outer and inner surfaces of a cerebral cortex, respectively, and Figs. 9 and 11 show cross-sections of the surface superimposed on the classified volume.

#### 4.1. Averaging and Mapping to Two Dimensions

The surfaces deformed have a topology based on a triangulation of a sphere, and there exists a one-to-one mapping between points on any two deformed surfaces

or between a deformed surface and a sphere. The mapping is simply that the  $i$ th vertex on one surface corresponds to the  $i$ th vertex on another surface, be it a cortex or the triangulated sphere, since all cortical surfaces were created from the same initial ellipsoid (a scaled sphere). It is important to note that the existence of such a mapping does not necessarily imply



**FIG. 14.** Curvature of cortex (blue = sulci, red = gyri) from Fig. 8a, mapped onto (A) the average surface and (B) a sphere.

**TABLE 1**

Studies of Anterior Versus Posterior Thickness of Central Sulcus Banks

Author	Year	Method	Anterior border (mm)	Posterior border (mm)	Significance
Brodmann	1908	Autopsy	3.94	1.86	
Von Economo	1925	Autopsy	3.75	1.82	
Glezer	1959	Autopsy	3.01	1.59	
Meyer	1996	MRI	2.70	1.76	
ASP- $T_{link}$	1999	MRI	3.22	2.73	$P < 0.0013$
ASP- $T_{nearest}$	1999	MRI	2.03	1.80	$P < 0.0099$
ASP- $T_{normal}$	1999	MRI	2.75	2.34	$P < 0.0084$

*Note.* For the present study, the significances of the differences are shown.




FIG. 15. The three different measures of thickness.

that there is any anatomical basis to the correspondence. In a worst case scenario, a point in the frontal lobe on one surface may correspond to a point on the occipital lobe of another surface. Having acknowledged that the deformation algorithm does not explicitly enforce homology among corresponding vertices of different brains, it is interesting to see if there is any consistency at all in the placement of vertices. Figure 12 shows the same vertex labeled on four different surfaces. While it is evident that these four points do not correspond exactly to each other, their consistent location near the superior aspect of the left precentral gyrus indicates somewhat of a trend. Another way to examine this is to perform a vertex by vertex average of the 150 deformed surfaces identified by ASP to create a mean surface; i.e., the four labeled points in Fig. 12 and their 146 counterparts would be averaged to give one vertex on the average surface. If there is no homology between points, one would expect a featureless cloud of points for the mean surface. However, the actual result, as shown in Fig. 13, is surprisingly full of anatomic detail, with the central sulcus and superior temporal sulcus clearly visible. It appears that performing the deformation in the standardized stereotaxic coordinate system mentioned earlier (Collins *et al.*, 1994), combined with stretch constraints that limit the movement of vertices, effectively enforces a relatively consistent placement of points on the cortical surface. By no means could one use this to locate an exact point on an individual cortex, for instance, on the fundus of a particular sulcus. However, this loose homology may be useful when dealing with averages of a large number of data sets or as the starting point for a post-processing step that attempts to find a more accurate homology among multiple brains. Further investigation is required to evaluate the errors in this homology based on some technique that explicitly maps corre-

sponding points; at a minimum, by using human observers to define homology of a subset of points across the 150 brains.

Figure 14 shows the average curvature of the 150 surfaces mapped onto both the average surface and a unit sphere, where blue areas correspond to sulci, and red areas to gyri. Due to the discrete nature of the polyhedral surfaces, curvature is approximated by the distance from a vertex to the centroid of its neighbors, relative to the diameter of the set of neighbours. Aside from the visual simplification aspect, morphometric analysis may be more easily performed in the two-dimensional parameter space of the sphere or on the average surface. Depending on the specific analysis requirements, it may be necessary to perform a further step of warping within the two-dimensional space, for instance, in order to make the mapping preserve distances, angles, or areas. Although approximating both hemispheres with a single closed surface is an oversimplification, the regions where this topology is inappropriate may be masked off. Alternatively, a surface model that matches the anatomical topology more closely may be chosen for the deformation process, but the self-intersection constraints remain important for preserving the desired topology. In the case of the cortical gray mantle, a more correct topology might include a pair of bilateral holes in the surface associated with the ventricles.

4.2. Cortical Thickness Measurements

Cerebral cortex is not of uniform thickness and the change in thickness is a reflection of change in cytoar-

TABLE 2

Mean and Standard Deviations of Three Thickness Measures over 150 Normal Cortical Surfaces

$T_{link}$ Mean (SD)	$T_{near}$ Mean (SD)	$T_{normal}$ Mean (SD)
4.20 (0.52)	2.60 (0.50)	4.43 (0.88)





FIG. 16. The trend for postcentral regions (top of image) to be thicker than inferior temporal regions (bottom) is clearly visible on this coronal slice, which cuts through these two regions at approximately a right angle.




**FIG. 17.** Top, left, and bottom views of normalized mean cortical thickness, measured three ways: (A)  $T_{\text{link}}$ , (B)  $T_{\text{nearest}}$ , (C)  $T_{\text{normal}}$ . The overall trend ranges from thin cortex in the postcentral regions to thicker in the inferior temporal regions.

chitecture. Cortical thickness in human isocortex has been reported to be 1.5 to 5.0 mm (Zilles, 1990), this thickness being lowest at the sulcal fundus and highest on the gyral ridge. The absolute values of cortical thickness may be influenced by the shrinkage of the tissue during fixation and embedding procedures, and it has been found that the shrinkage is greater in immature

brains that have a higher water content than mature ones (Blinkov and Glezer, 1968; Kretschmann *et al.*, 1979).

The most prominent evidence of change in cortical thickness comes from the studies done on the central sulcus. Several researchers have found that there is a difference in the cortical thickness in the two banks of



**FIG. 18.** Top, right, and bottom views of areas of statistical difference ( $P < 0.05$ ) between the three measures of thickness (red shows greater than, green shows less than): (A) between  $T_{\text{link}}$  and  $T_{\text{nearest}}$ , (B) between  $T_{\text{link}}$  and  $T_{\text{normal}}$ , and (C) between  $T_{\text{normal}}$  and  $T_{\text{nearest}}$ .

the central sulcus. This is attributed to the difference in cytoarchitecture of the two banks. The anterior bank is agranular cortex and forms the primary motor area 4 (Brodmann's division). This contains giant pyramidal cells of Betz in cortical layer V and relatively few granule cells. In contrast, the posterior bank forms the primary sensory cortex (areas 3, 1, and 2), which contains densely packed granule cells and is much thinner.

In 1996, Meyer *et al.* reported the mean ratio of cortical thickness in the MR images of pre- and post-central gyri in the two hemispheres of 10 brains. Using cortical thickness measurements across the central sulcus, this group provided a method for locating the primary motor and sensory cortices that correspond to the cytoarchitecture in these regions. Table 1 compares cortical thickness information from four previous stud-

ies with three measures based on the ASP method, which will be presented shortly.

Since the cytoarchitecture is not visible on the images acquired by MR scanning, the development of accurate methods for the estimation of cortical thickness is important for understanding the ontogeny of various cortical regions and for providing an indication of the hidden cytoarchitectural boundaries.

The cortical surface models described here lend themselves readily to measurements of cortical thickness. However, a definition of the thickness measure must be chosen. We present three such definitions and show the effectiveness of each for the task of comparing and relating thicknesses in various regions and across brains. Given a surface representing the gray/CSF boundary and another representing the gray/white boundary, the following three measures of thickness are defined:

$T_{\text{link}}$ : The distance from a vertex on the outer surface to the corresponding vertex on the inner surface, as defined by the linkages used in the two-surface deformation of the ASP algorithm.

$T_{\text{near}}$ : The distance from a vertex on the outer surface to the nearest point on the inner surface.

$T_{\text{normal}}$ : The distance from a vertex on the outer surface to the nearest point on the inner surface in the direction of the surface normal. The normal at a vertex is an average of the normals of the adjacent faces, smoothed.


A schematic of these three measures is presented in Fig. 15. For each measure of thickness, the mean and standard deviation over the entire surfaces of the 150 cortices was generated and shown in Table 2. In absolute terms, the three measures differ significantly but they nonetheless exhibit similar patterns of variability. To focus on this latter aspect, all three measures were normalized to a value that indicates number of standard deviations from the mean of the measure, i.e., the  $z$  score, and the resulting surfaces are shown in Fig. 17, color coded on a scale of  $-4$  to  $4$ . The three measures all indicate similar trends in thickness across the cortex, with the thinnest region being the post-central regions and the thickest being the inferior temporal regions, a trend that can be seen on most individual MR images (Fig. 16). Figure 18 shows the areas of the cortex where there is significant ( $P < 0.05$ ) difference among the three measures of thickness. If the measures provided the same information, there would be no areas of difference between any pair. However, the  $T_{\text{near}}$  measure is considerably dissimilar to the other two measures ( $T_{\text{link}}$  and  $T_{\text{normal}}$ ), which seem to be much more closely related to each other.

It is interesting to note that the  $T_{\text{normal}}$  has almost double the standard deviation of the other two. This can be attributed to the fact that a slight deviation in normal direction due to noise or imprecise surface fitting can cause the distance to the other surface to

change significantly. The mean value of  $T_{\text{near}}$  was much smaller than the other two. This can be attributed to the fact that with deep narrow sulci which have some degree of asymmetry, the nearest distance from the outer surface can miss the fundus of the sulcus, as suggested by Fig. 15. In contrast to these two measures, the  $T_{\text{link}}$  measure attempts to use the correspondence between points on the surface, which provides a measure of thickness that is less sensitive to fluctuations in surface normal and areas of high curvature. However, this measure is therefore dependent on the validity of the correspondence created by the deformation algorithm. For instance, if the points on one surface slide with respect to the other surface, such as seen on the right portion of Fig. 15, the thickness could be overestimated. Although the inner surface's tendency to pull the outer surface toward the fundus of each sulci helps provide the correct correspondence at each of the fundus points, it remains to investigate methods to improve the correspondence over the entire cortical surface.

The cortical surface thickness maps thus produced can be used readily to address issues of variability across the cortex. For example, the distribution of central sulcus thickness can be examined by comparing a point on the anterior bank with one on the posterior bank (Fig. 19). For all three methods of measuring cortical thickness, the anterior bank is significantly thicker than the posterior bank ( $P < 0.01$ ), which matches the previously published postmortem and imaging results (Table 1).

A more sophisticated analysis of this central sulcus wall thickness involves the use of a probabilistic atlas of neuroanatomic structure (Evans *et al.*, 1996; Mazziotta *et al.*, 1995). This atlas provides the probability of each of 90 structures at each voxel within a normalized stereotaxic space (Evans *et al.*, 1992; Collins *et al.*, 1994). Using the probabilistic precentral and postcen-



**FIG. 19.** Two sample points shown as red spheres, one each on the posterior bank and the anterior bank of the central sulcus.

**TABLE 3**

Thickness of Probabilistic Anterior and Posterior Bank of Central Sulcus

Structure	$T_{\text{link}}$ mm (SD)	$T_{\text{nearest}}$ mm (SD)	$T_{\text{normal}}$ mm (SD)
Anterior bank ( $N = 106998$ )	3.8 (1.3)	2.5 (1.0)	3.7 (3.2)
Posterior bank ( $N = 75684$ )	3.5 (1.3)	2.3 (0.96)	3.4 (3.2)

tral gyrus, the statistical test of two points is replaced with a similar test of about 180,000 points on these two regions. The result is that for all three thickness measures, the precentral cortex is thicker than postcentral cortex ( $P < 0.000001$ ). The mean and standard deviations of the thicknesses of these two regions are presented in Table 3. It is interesting to note that the absolute difference in thickness is much less than with the previous two-point test. This may indicate that while the mean thicknesses of the two entire gyri differ slightly but significantly, there may be smaller subregions of more substantial differences. The image pre-processing mentioned earlier removes some of the sources of bias in these thickness measurements through linear registration, RF inhomogeneity correction, intensity normalization, and tissue classification, but it needs to be examined what potential errors and bias in the cortical thickness measurements remain.

## 5. SUMMARY AND DISCUSSION

We have described ASP, a novel method of multiple surface deformation for cortical segmentation with the following features:

- a boundary search along the local surface normal is used to increase the range of attraction of edges;
- the use of proximity constraints with appropriate weights excludes self-intersecting surface configurations;
- some arbitrary weights are replaced by more intuitive geometric constraints;
- multiple surfaces, models, and datasets may be combined into a single objective function;
- automatic identification of the total cerebral cortical surface from MR images is achieved in a robust way with respect to partial volume effects;
- a preliminary map of cortical gray matter thickness has been produced and related to previous studies;
- and a higher resolution average brain surface has been created using the deeper sulcal penetration of ASP compared to earlier versions of this algorithm (MacDonald *et al.*, 1994).

The ability to automatically generate surface representations from images provides opportunities for so-

phisticated analysis of large populations of neuroanatomical data. Average surface models that incorporate improved descriptions of mean position, shape, and thickness have been generated and the results have been consistent with previous studies involving both postmortem and imaging data. Future directions include improving the mapping of cortical surface between subjects, as well as sulcal recognition and application to automated identification of cerebellum, brain stem, and other neuroanatomic structures.

A question arising from the use of digital models is to what extent does the imposed *a priori* model (i.e., thickness constraints) bias the resulting thickness maps. In order to minimize this bias, the current estimate of local cortical thickness should be used to define a second-generation model for subsequent surface deformation processes. Such refinement of the cortical thickness map may be repeated to investigate whether the process converges. More investigation is required to determine to what extent the various parameters of the ASP algorithm affect the perceived thickness of the cortex.

The advantage of the ASP algorithm in segmenting the human cortex arises from its attempt to use simple geometric constraints based on neuroanatomical knowledge. The resulting models provide valuable data in a form suitable for a variety of morphological experiments. While an important step, there remains significant opportunity to refine and extend these ideas to provide more detailed and accurate tools for understanding neuroanatomical relationships. For instance, modeling the forces that cause the smooth brain of the embryo to develop into a deeply folded adult brain could improve the thickness computations with a more accurate correspondence between surfaces. The most important feature of a technique such as ASP is its ability to refine and extend its underlying model as the questions asked by neuroanatomists become more sophisticated and more precise.

## ACKNOWLEDGMENTS

Special thanks to Dr. Louis Collins for insight and critique on the development of this paper. The authors express appreciation for funding from the Natural Sciences and Engineering Research Council of Canada, the McDonnell-Pew Program in Cognitive Neurosciences, the Medical Research Council of Canada, the Human Brain Map Project, NIMH, and NIDA.

## REFERENCES

- Bajcsy, R., and Kovačič. 1989. Multiresolution elastic matching. *Comput. Vis. Graph. Image Process.* **46**:1–21.
- Blinkov, S. M., and Glezer, I. I. 1968. *The Human Brain in Figures and Tables*. Plenum Press, New York.
- Brodman, K. 1908. Uber rindenmessugen. *Zentralbl Nervenheilkunde Psychiatrie*.



- Cohen, L. D., and Cohen, I. 1993. Finite-element methods for active contour models and balloons for 2-d and 3-d images. *IEEE Trans. Patt. Anal. Mach. Intell.* **15**(11):1131–1147.
- Collins, D. L., Neelin, P., Peters, T. M., and Evans, A. C. 1994. Automatic 3-d intersubject registration of mr volumetric data in standardized talairach space. *J. Comput. Assist. Tomogr.* **18**(2): 192–205.
- Collins, D. L., Peters, T. M., Dai, W., and Evans, A. C. 1992. Model-based segmentation of individual brain structures from mri data. In *Proceedings Visualization in Biomedical Computing 1992*, pp. 10–23.
- Dale, A. M., and Sereno, M. I. 1993. Improved localization of cortical activity by combining eeg and meg with mri cortical surface reconstruction: A linear approach. *J. Cogn. Neurosci.* **5**:162–176.
- Davatzikos, C., and Bryan, R. N. 1995. Using a deformable surface model to obtain a shape representation of the cortex. In *Proceedings of the IEEE International Conference on Computer Vision 1995*, pp. 212–217.
- Evans, A. C., Collins, D. L., and Holmes, C. J. 1996. Toward a probabilistic atlas of human neuroanatomy. In *Brain Mapping: The Methods* (J. C. Mazziotta and A. W. Toga, Eds.), pp. 343–361. Academic Press.
- Economo, C. V., and Koskinas, G. N. 1925. *Ctoarchitektonik der Hirnrinde des erwachsenen Menschen*. Springer-Verlag, Berlin.
- Evans, A. C., Marrett, S., Neelin, P., Gum, T., Dai, W., Milot, S., Meyer, E., and Bub, D. 1992. Anatomical mapping of functional activation in stereotactic coordinate space. *Neuroimage* **1**:43–53.
- Glissen, E., and Zilles, K. 1996. The calcarine sulcus as an estimate of the total volume of human striate cortex: A morphometric study of reliability and intersubject variability. *J. Hirnforschung* **37**(1): 57–66.
- Kwan, R. K.-S., Evans, A. C., and Pike, G. B. 1996. An extensible MRI simulator for post-processing evaluation. In *Proceedings of the International Conference on Visualization in Biomedical Computing*.
- Kretschmann, H. J., Schleicher, A., Wingert, F., Zilles, K., and Loblich, H. J. 1979. Human brain growth in the 19th and 20th century. *J. Neurol. Sci.* **40**(2–3):169–88.
- Kass, M., Witkin, A., and Terzopoulos, D. 1988. Snakes: Active contour models. *Int. J. Comput. Vis.* 321–331.
- Lorensen, W. E., and Cline, H. E. 1987. Marching cubes: A high resolution 3-d surface construction algorithm. *Comput. Graph. Quart. Rep. SIGGRAPH-ACM* **21**(4):163–169.
- MacDonald, J. D. 1998. *A Method for Identifying Geometrically Simple Surfaces from Three-Dimensional Images*. PhD thesis, School of Computer Science, McGill University, Montreal, Canada.
- MacDonald, J. D., Avis, D., and Evans, A. C. 1994. Multiple surface identification and matching in magnetic resonance images. In *Proceedings Visualization in Biomedical Computing 1994*, pp. 160–169.
- MacDonald, J. D., Avis, D., and Evans, A. C. 1998. Proximity constraints in deformable models for cortical surface identification. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention—MICCAI'98*.
- Meyer, J. R., Roychowdhury, S., Russell, E. J., Callahan, C., Gitelman, D., and Mesulam, M. M. 1996. Location of the central sulcus via cortical thickness of the precentral and postcentral gyri on mr. *Am. J. Neuroradiol.* **17**(9):1699–706.
- McInerney, T., and Terzopoulos, D. 1996. Deformable models in medical image analysis: A survey. *Med. Image Anal.* **1**(2).
- Mazziotta, J. C., Toga, A. W., Evans, A. C., Fox, P., and Lancaster, J. 1995. A probabilistic atlas of the human brain: Theory and rationale for its development. The international consortium for brain mapping. *NeuroImage* **2**(2):89–101.
- Nastar, C., and Ayache, N. 1993. Non-rigid motion analysis in medical images: A physically based approach. In *Proceedings of Information Processing in Medical Imaging*, pp. 17–32. Flagstaff, AZ.
- Press, W. H., Flannery, B. P., Teukolsky, S. A., and Vetterling, W. T. 1988. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge Univ. Press, Cambridge, England.
- Staib, L. H., and Duncan, J. S. 1992. Deformable fourier models for surface finding in 3d images. In *Proceedings Visualization in Biomedical Computing 1992*, pp. 90–104.
- Sandor, S., and Leahy, R. 1997. Surface-based labeling of cortical anatomy using a deformable atlas. *IEEE Trans. Med. Imag.* **16**(1): 41–54.
- Sled, J. G., Zijdenbos, A. P., and Evans, A. C. 1997. A comparison of retrospective intensity non-uniformity correction methods for MRI. In *Information Processing in Medical Imaging*, in press.
- Van Essen, D. C., and Drury, H. A. 1997. Structural and functional analyses of human cerebral cortex using a surface-based atlas. *J. Neurosci.* **17**(18):7079–7102.
- Xu, C., Pham, D. L., Prince, J. L., Etemad, M. E., and Yu, D. N. 1998. Reconstruction of the central layer of the human cerebral cortex from mr images. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention—MICCAI'98*, pp. 481–488.
- Zilles, K., Armstrong, E., Schleicher, A., and Kretschmann, H. J. 1988. The human pattern of gyrification in the cerebral cortex. *Anat. Embryol.* **179**(2):173–179.
- Zijdenbos, A. P., Dawant, B. M., and Margolin, R. A. 1993. Measurement reliability and reproducibility in manual and semi-automatic mri segmentation. *Proc. Annu. Conf. Eng. Med. Biol.* **15**:162–163.
- Zilles, K. 1990. *Cortex in the Human Nervous System*. Academic Press, San Diego, CA.
- Zeng, X., Staib, L. H., Schultz, R. T., and Duncan, J. S. 1998. Segmentation and measurement of the cortex from 3-d mr images. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention—MICCAI'98*, pp. 519–530.