

Cortical Surface-Based Analysis

I. Segmentation and Surface Reconstruction

Anders M. Dale,^{*,1} Bruce Fischl,^{*} and Martin I. Sereno[†]

^{*}Nuclear Magnetic Resonance Center, Massachusetts General Hosp/Harvard Medical School, Building 149, 13th Street, Charlestown, Massachusetts 02129; and [†]Department of Cognitive Science, University of California at San Diego, Mail Code 0515, 9500 Gilman Drive, La Jolla, California 92093-0515

Received May 22, 1998

Several properties of the cerebral cortex, including its columnar and laminar organization, as well as the topographic organization of cortical areas, can only be properly understood in the context of the intrinsic two-dimensional structure of the cortical surface. In order to study such cortical properties in humans, it is necessary to obtain an accurate and explicit representation of the cortical surface in individual subjects. Here we describe a set of automated procedures for obtaining accurate reconstructions of the cortical surface, which have been applied to data from more than 100 subjects, requiring little or no manual intervention. Automated routines for unfolding and flattening the cortical surface are described in a companion paper. These procedures allow for the routine use of cortical surface-based analysis and visualization methods in functional brain imaging. © 1999 Academic Press

Key Words: segmentation; cortical surface reconstruction.

1. INTRODUCTION

The cerebral cortex, which makes up the largest part of the human brain, has the topology of a 2-D sheet and a highly folded geometry. The human cortex, like that of other primates, is divided into a large number of different areas (Zeki and Shipp, 1988; Felleman and Van Essen, 1991; Kaas and Krubitzer, 1991; Sereno and Allman, 1991). Most of the features that distinguish these cortical areas can only be measured relative to the local orientation of the cortical surface. A partial list of these includes laminar features (e.g., cortical thickness), as well as retinotopic, tonotopic, and somatotopic maps. There is no principled way to measure these quantities in an unlabeled 3-D data set. For this reason, very few studies of the cortex in nonhuman

primates have relied exclusively on a 3-D (stereotactic) system; instead, surface-based methods have been routinely applied. Nevertheless, most approaches to analyzing and displaying human brain imaging data have relied exclusively upon 3-D approaches. For instance, functional activations are commonly displayed using slices of a 3-D volume or volume renderings, and the locations of these activations are typically reported in Talairach coordinates.

A major impediment to the routine use of cortical surface-based approaches in human neuroimaging has been the difficulty of automating the surface reconstruction, alignment, painting, and display processes. The traditional method for reconstructing complex, folded, 2-D surfaces from thin sections was to trace contours from each section onto thin wax sheets, and then align and stack the sheets for viewing. Computer-based versions of such techniques have been used to analyze and visualize primate electrophysiological data (Schwartz, 1990) and human neuroimaging data (DeYoe *et al.*, 1996; Van Essen and Drury, 1997). However, since such manual tracing methods are quite time-consuming, very few studies have made extensive use of surface-based display of functional activation data.

Here we describe a largely automated set of computational tools that we have developed in order to facilitate the routine use of cortical surface-based analysis and visualization methods in human neuroimaging studies. These tools, which are extensions of the methods described by Dale and Sereno (1993), have been routinely used for mapping the detailed topographic organization of human retinotopic visual areas (Sereno *et al.*, 1995; Tootell *et al.*, 1996a,b; Tootell *et al.*, 1997, 1998a,b; Hadjikhani *et al.*, 1998) primary motor, somatosensory, and auditory areas (Talavage *et al.*, 1996; Moore *et al.*, 1997; Talavage *et al.*, 1997a,b; Moore *et al.*, 1998; Talavage, 1998a,b), as well as areas involved in the processing of visual motion (Tootell *et al.*, 1995a,b; Reppas *et al.*, 1997; Culham *et al.*, 1998), color (Hadjikhani *et al.*, 1998), perception of faces and objects

¹ To whom correspondence and reprint requests should be addressed. Fax: (617) 726-7422. E-mail: dale@nmr.mgh.harvard.edu.

(Halgren *et al.*, 1998), and processing the meaning of words (Halgren *et al.*, 1998). Part of the complexity of the tools has arisen due to modifications made in the course of many surface reconstructions. The resulting procedure is quite robust and has been used to automatically reconstruct well over 100 brains without requiring any parameter modifications.

Although methods for accomplishing several of the subtasks involved in the surface reconstruction and unfolding process have been described by others, such as bias-field correction (Wells *et al.*, 1994; Meyer *et al.*, 1995; Brechbühler *et al.*, 1996), removal of skull and extrameningeal tissues (Atkins and Mackiewicz, 1996), nonlinear smoothing of MRI data to preserve edges (Whitaker and Gerig, 1994), gray/white matter segmentation (Gerig and Kikinis, 1990; Gerig *et al.*, 1991; Bullmore *et al.*, 1995; Székely *et al.*, 1995; Atkins and Mackiewicz, 1996; Wells *et al.*, 1996; Teo *et al.*, 1997; Wagner *et al.*, 1998), cortical surface reconstruction (Davatzikos and Bryan, 1996; MacDonald, 1998), surface unfolding/flattening (Schwartz and Merker, 1986; Schwartz *et al.*, 1989; Carman *et al.*, 1995; Drury *et al.*, 1996a), and morphing (Drury *et al.*, 1996b; Thompson *et al.*, 1996; Thompson and Toga, 1996; Davatzikos, 1997; Drury *et al.*, 1997; Van Essen and Drury, 1997), this represents the first complete, automated procedure capable of reliably performing all these functions based on standard structural MRI data.

The automaticity is evidenced by the fact that these techniques have been successfully applied to well over 100 datasets taken with a variety of scanners (e.g., GE, Siemens, Picker) and pulse sequences (e.g., 3D SPGR, MP-RAGE²), without modification of any of the parameters. The entire procedure, from intensity normalization, through segmentation, tessellation, and inflation, runs completely automatically in about 1.5 h on current PC architectures, and, for the vast majority of subjects, requires no manual intervention beyond specifying input file names and formats. The inflation process typically reveals a small number of topological defects, mostly located in subcortical regions, which are not noticeable in the standard views of the folded surface. These defects are removed by manual editing of the labeled volume, typically requiring 30 min of manual interaction per subject.

In the following, we describe the various steps involved in the surface reconstruction and deformation process. The entire procedure is designed to be modu-

lar, with each component being relatively insensitive to variations in the rest of the process. This modularity should facilitate the modification or replacement of the individual components with improved algorithms.

2. CORTICAL SURFACE RECONSTRUCTION

The reconstruction of the cortical surface is a complex procedure which is broken into a number of subtasks. First, intensity variations due to magnetic field inhomogeneities are corrected and a normalized intensity image is created from a high resolution, T1-weighted, anatomical 3-D MRI dataset. Next, extracerebral voxels are removed, using a “skull-stripping” procedure. The intensity normalized, skull-stripped image is then operated on by a segmentation procedure based on the geometric structure of the gray–white interface. Cutting planes are then computed which separate the cerebral hemispheres and disconnect subcortical structures from the cortical component. This generates a preliminary segmentation which is partitioned using a connected components algorithm. Any interior holes in the components representing white matter are filled, resulting in a single filled volume for each cortical hemisphere. Finally, the resulting volume is covered with a triangular tessellation and deformed to produce an accurate and smooth representation of the gray/white interface as well as the pial surface. This surface departs from a simple spherical topology due to subcortical gray matter as well as various midbrain structures. These topological “defects” are removed through a manual editing procedure that results in a surface that is both geometrically accurate and topologically spherical. Each of these steps is explained in detail below.

2.1 Talairach Registration

We use the automated Talairach registration procedure developed and distributed by the Montreal Neurological Institute (Talairach and Tournoux, 1988; Collins *et al.*, 1994) to compute the transformation matrix from a high resolution T1-weighted scan taken using a short 3-D acquisition sequence (e.g., GE SPGR, Siemens MP-RAGE). In brief, the procedure computes the transformation parameters by using gradient descent at multiple scales to maximize the correlation between the individual volume and an average volume composed of a large number of previously aligned brains. We store this transformation matrix which takes image coordinates into Talairach coordinates for use in subsequent processing stages.

2.2 Intensity Normalization

The high resolution T1-weighted images generated by an MR scanner are typically corrupted by magnetic

² Typical MRI pulse sequence parameters used are: GE 3D SPGR (Fast, IR-Prepped, Flip = 25°, TE = 2.1 ms, TI = 300 ms, Plane = Sagittal, FOV = 25 cm, Thk = 1.2 mm, Matrix = 192 × 256, NEX = 1), Siemens 3D MP-RAGE (Flip = 10°, TR = 9.7 ms, TE = 4 ms, TI = 20 ms, TD = 0 ms, Plane = Sagittal, FOV = 25 cm, Thk = 1.0 mm, Matrix = 256 × 256, NEX = 1). Two to four repetitions of these sequences are usually acquired sequentially in a given session and averaged off-line after motion correction.

susceptibility artifacts and RF-field inhomogeneities, resulting in variations in both intensity and contrast across the image. That is, identical tissue types (e.g., cortical gray and white matter) will give rise to varying intensities as a function of their spatial location. This is obviously undesirable for any segmentation procedure which utilizes intensity information in order to classify voxel data into different tissue types. Here we describe a fast and relatively simple procedure for correcting such image intensity variations (or “bias fields”), which we have found results in accurate surface reconstructions when applied to data from a variety of scanners and MRI protocols.

The procedure relies on the assumption that in any given slice parallel to the x - y plane in the magnet coordinate system (perpendicular to the long axis of the bore) the highest intensity tissue type of any significance will be white matter. This assumption is used to center the mean white matter intensity in every x - y slice at a desired value. The remaining bias field is then corrected by automatically detecting a set of “control points” which are determined to be in white matter. The intensity value of these points represents a sampling of the remaining bias field. A Voronoi partitioning algorithm is then used to compute the bias field at non-control points, that is, non-control points are assigned the value of the closest control point. More specifically, the algorithm proceeds as follows:

1. Construct a set of histograms from overlapping slices parallel to the x - y plane in the magnet coordinate system. We use 24 10-mm-thick slices, with 50% overlap between adjacent slices. The stack of slices are centered at Talairach coordinates ($x = 0$, $y = -10$, $z = 10$), and oriented perpendicular to the long axis of the magnet bore (z axis in the magnet coordinate system).

2. Smooth the resulting histograms using a fairly broad ($\sigma = 2$) Gaussian to eliminate noise-induced peaks (note that the intensity values of the entire volume have originally been scaled to fit within the range of [0,255]).

3. Use a peak-finding algorithm to determine the mean white matter intensity, specified as the rightmost (i.e., highest intensity) peak in the histogram containing more than 15% of the voxels within an intensity range of [30,225]. This serves to exclude background voxels, as well as those containing CSF or fat. The peak itself is constrained to be the largest histogram value in a seven bin neighborhood. If no peak meets these criteria, the slice is not used in subsequent spline interpolation calculations.

4. Discard outliers from the array of detected mean white matter intensities. This step makes use of the fact that the variation in intensity due to magnetic field inhomogeneities is smooth across space, and that there-

fore the variation in detected mean white matter intensity should be small in any two adjacent slices. Specifically, we use the median detected white matter value in a set of slices around the Talairach origin as a starting point, then add slices for which the gradient of the intensity change over space with respect to the nearest valid slice is below a prespecified threshold (0.4/mm). This serves to exclude slices containing few white matter voxels (e.g., towards the very top of the head, or below the temporal lobes) from the bias-field calculations.

5. Fit a set of cubic splines to the resulting coefficients of the valid slices.

6. Use the splines to interpolate the coefficients for each point along the z axis.

7. Adjust each intensity value by the coefficient at its z coordinate, so that the white matter peak at that coordinate is normalized to a value of 110.

8. Find all points in the volume that are at the center of a $5 \times 5 \times 5$ neighborhood of intensity values that all lie within 10% of the white matter peak. This assures that the detected points (“control-points”) are not partial-volumed, and therefore accurately represent white matter intensity at those points. Set the bias field correction at these values to the ratio of the desired white matter intensity to the local average.

9. Build a Voronoi diagram and set all voxels unassigned in step 8 to the correction value of the nearest control point.

10. Perform a few iterations of “soap-bubble” smoothing. That is, replace each non-control point by the average of it and its 26 neighbors. This serves to smooth the boundaries between Voronoi regions.

11. Scale the intensity at each voxel in the volume by the computed correction field.

Steps 8–11 are typically iterated 5–10 times, with the number of control points at each step, and hence the spatial extent of the correction field, increasing at each iteration. We have found this procedure to be quite robust due to the minimal assumptions on which it relies (i.e., no parametric form) as well as the built-in cross-validation in step (4). Any minor intensity variations not removed by this procedure, have minimal effect on the rest of the procedure, due to the adaptive nature of the segmentation process described in section 2.4.

2.3 Skull-Stripping

The next step in the reconstruction process is the automated stripping of the skull from the intensity-normalized image. This procedure involves deforming a tessellated ellipsoidal template into the shape of the inner surface of the skull. The deformation process is driven by two kinds of “forces”: (1) an MRI-based force, designed to drive the template outward from the brain,

and (2) a curvature reducing force, enforcing a smoothness constraint on the deformed template. This latter force can be seen as encoding *a priori* knowledge about the smoothness of the inner surface of the skull. The current approach is similar to the “shrink-wrapping” method for refining the cortical surface estimates, described by Dale and Sereno (1993), which is closely related to the active contour methods described (Kaas *et al.*, 1988; Terzopoulos and Fleischer, 1988), and subsequently used in medical imaging applications (MacDonald *et al.*, 1994; Collins *et al.*, 1996; Davatzikos and Bryan, 1996; Davatzikos, 1997; MacDonald, 1998)—see also Sections 2.4 and 2.7. However, in order to improve the robustness of the algorithm, the MRI-based force used here is calculated based on nonlocal information, obtained by sampling the MRI data along the surface normal from each vertex of the template tessellation. The force acting on each vertex is then designed to drive the surface to regions consistent with CSF (i.e., with low MRI values), repelling the surface outwards from contiguous regions consistent with brain (i.e., with higher MRI intensities). More precisely, the procedure proceeds as follows:

We first center the template, based on a supertessellated unit icosahedron with 2562 vertices and 5120 triangles, at Talairach coordinates ($x = 0$, $y = -10$, $z = 10$), and scale its x , y , and z dimensions by $0.35 \cdot (\max_x - \min_x)$, $0.45 \cdot (\max_x - \min_x)$, $0.35 \cdot (\max_x - \min_x)$, respectively, where $\min_x$ and $\max_x$ are the minimum and maximum x coordinates of MRI voxels exceeding a threshold of 40. We base our initial estimates of the size of the brain on the x range as it is the least variable of the three cardinal axes, while the y and z extent of the MRI data has considerable variability due to the neck and the nose.

Once the initial template has been positioned, we gradually deform it through a series of iterative steps. On each iteration t , we update the coordinate $\mathbf{x}_k$ of each vertex according to the “forces” mentioned above, i.e.,

$$\mathbf{x}_k(t+1) = \mathbf{x}_k(t) + \mathbf{F}_S(t) + \mathbf{F}_M(t), \quad (1)$$

where the smoothness force $\mathbf{F}_S$ is given by

$$\mathbf{F}_S = \lambda_T \sum_{j \in N_k} (\mathbf{I} - \mathbf{n}_k \mathbf{n}'_k) \cdot (\mathbf{x}_j - \mathbf{x}_k) + \lambda_N \left(\sum_{j \in N_k} (\mathbf{n}_k \mathbf{n}'_k) \cdot (\mathbf{x}_j - \mathbf{x}_k) - \frac{1}{V} \sum_{i \in N_i} \sum_{j \in N_i} (\mathbf{n}_i \mathbf{n}'_i) \cdot (\mathbf{x}_j - \mathbf{x}_i) \right), \quad (2)$$

and the MRI-based force $\mathbf{F}_M$ is given by

$$\mathbf{F}_M = \lambda_M \mathbf{n}_k \prod_{d=1}^{30} \max(0, \tanh(I(\mathbf{x}_k - d\mathbf{n}_k) - I_{\text{thresh}})), \quad (3)$$

where λ_T and λ_N specify the strengths of the tangential and normal components of the smoothness force, λ_M specifies the strength of the MRI-based force, $\mathbf{I}$ is the 3-by-3 identity matrix, N_k denotes the set of vertices neighboring the k^{th} vertex, V is the number of vertices in the template tessellation, $I(\mathbf{x})$ is the MRI value at location $\mathbf{x}$, while $\mathbf{n}_k$ and $\mathbf{n}'_k$ denote the surface normal at location k and its transpose, respectively. Note that we have omitted the functional dependence on the iteration number t in Eqs. (2) and (3), to avoid notational clutter. For MRI intensity values normalized as discussed in Section 2.2, we use the following parameter values: $\lambda_T = 0.5$, $\lambda_N = 0.1$, $\lambda_M = 1.0$, and $I_{\text{thresh}} = 40$.

The evolution of the template is shown in Fig. 1, with the leftmost image depicting the initial position of the ellipsoidal template and the rightmost image showing the surface after 50 iterations. The MRI-based force was turned off (by setting $\lambda_M = 0$) during the last 10 iterations to smooth the final surface.

The deformed template is then used to strip the skull from the 3D MRI volume, by removing all voxels outside the tessellated surface (Fig. 2). The entire procedure requires about 30 s of computation time on a standard PC architecture.

2.4 White Matter Labeling

A common approach to MRI segmentation is the use of global gray-scale-based thresholds or classes to label different tissue types (Teo *et al.*, 1997; Wagner *et al.*, 1998). Such methods are intrinsically sensitive to variations in image intensity or contrast due to partial voluming, magnetic susceptibility artifacts, or RF-field inhomogeneities. Some of these difficulties can be overcome by modeling the inhomogeneities in the spatial distribution of tissue class intensities as a slowly varying bias field and solving for it concurrently with the tissue classification (Wells *et al.*, 1994). This helps minimize the effect of RF-field inhomogeneities, but does not address the problems of partial voluming and magnetic susceptibility artifacts which cannot be modeled as low frequency distortions. In addition, these procedures do not use any prior information regarding the proper local geometric structure of the cortex and may therefore generate surfaces with many topological defects that can prevent successful inflation and flattening of the cortical surface. The current approach overcomes many of these problems by exploiting the laminar structure of the cortical gray matter. Specifically, we know that the cortical surface is smooth, with finite curvature everywhere, resulting in a locally planar structure where cortical gray matter borders other tissue types such as white matter or CSF. The segmentation procedure employs this information by detecting the plane-of-least-variance and using intensity information in this plane as a basis for classifying regions in






FIG. 1. Snapshots of template deformation process.

FIG. 4. Cutting planes overlaid on sections through white matter-labeled volume.

FIG. 7. Intersection of the tessellated white matter surface with the skull-stripped MRI volume.

FIG. 8. Intersection of the tessellated pial surface with the skull-stripped MRI volume.




FIG. 2. Intensity normalized (left) and skull-stripped images (right).

which intensity information alone is insufficient for accurate tissue classification.

Another approach to cortical surface reconstruction is to deform a template using elastic transformations based on MRI intensity information to force it to conform to the shape of the cortex (Davatzikos and Bryan, 1996; MacDonald, 1998). One advantage of this approach, as noted in Dale and Sereno (1993), is that the topology of the resulting surface is known to be identical to that of the template, i.e., spherical. This facilitates unfolding and flattening of the resulting surface and can reduce or eliminate the need for manual intervention. However, the topological correctness of the surface is no guarantee of geometric accuracy. In fact, such methods often fail to accurately represent deep sulci, as noted by Manceaux-Demiau *et al.* (1998), requiring additional measures to force the surface into the depths of cerebral sulci (Davatzikos and Bryan, 1996). This problem is overcome to some extent by using the white matter, rather than the gray matter, as the target of the deformation process (MacDonald, 1998)—see also Dale and Sereno (1993). It should also be noted that segmentation errors, such as a single-voxel “bridge” across a sulcus, which would result in *topological* errors in the tissue classification process, may instead result in large *geometrical* errors in the surface generated by a deformation procedure which enforces a given topology. In the surface reconstruction process described here, such classification errors result in topological defects in the reconstructed surface, which can be readily detected by surface inflation and removed through manual intervention.

Recently, many techniques have been proposed in the computer vision literature to perform oriented filtering

of noisy data as a preprocessing step to facilitate segmentation (Perona and Malik, 1987; Nördstrom, 1990; Perona and Malik, 1990; Saint-Marc *et al.*, 1991; Alvarez *et al.*, 1992; Catté *et al.*, 1992; Nitzberg and Shiotani, 1992; Whitaker and Gerig, 1994; Fischl and Schwartz, 1997; Weickert, 1997; Fischl and Schwartz, 1998). Many of these are based on the work of Perona and Malik (1987, 1990), who used a nonlinear anisotropic version of the diffusion equation to smooth images within regions, but not between regions. This type of approach can also be formulated as a nonlinear filtering operation, in which the filter shape is gradually constructed, through the numerical integration of the nonlinear diffusion equation, to conform to local image structure (Fischl and Schwartz, 1997; Fischl and Schwartz, 1998). In uniform regions, away from boundaries, the diffusion equation becomes approximately linear, and the filter kernels become Gaussian. Near edges, the kernels take on the orientation and shape of the border in an effort to avoid averaging data from two different objects or classes. This approach has been made explicit by Nitzberg and Shiotani (1992), who construct a nonlinear filter with a shape, orientation, and offset directly based on local image derivatives. Our approach has a similar spirit. However, in the case of gray/white segmentation we know the shape of the desired filter a priori. Within a given tissue class, the shape is unimportant as all directions are functionally equivalent. However, on the borders between classes, the shape should always be planar due to the laminar structure of the cortex. Thus, we directly enforce this shape and formulate an adaptive procedure to compute the correct orientation. This approach has some notable advantages within this problem domain: (1) it pre-

serves single voxel thick structures, which prevents thin white matter strands from being blurred with surrounding gray matter or CSF; (2) it avoids the time-consuming iterative process inherent in most implementations of diffusion filtering; and (3) it integrates filtering and model-based knowledge as part of the labeling process, which results in a robust segmentation procedure.

The segmentation process is a two-step procedure. First, a preliminary classification is performed based solely on intensity information. Next, this volume is examined and regions which contain more than one tissue type are marked for further processing. The planar orientation which minimizes within-plane intensity variance is then computed for each of these “border” voxels. Finally, if the in-plane intensities indicate a mislabeling and the intensity of the voxel itself is borderline, the segmentation decision is reversed, and the label of the voxel is changed to reflect its assignment to a new tissue class. With this overview in mind, we now proceed to a more detailed accounting of the full segmentation procedure.

First, three intensity-based constants WM_LOW (90), WM_HIGH (140), and GRAY_HIGH (100) are defined. These specify the lowest possible intensity for white matter, the highest possible intensity for white matter, and the highest possible intensity for gray matter, respectively. Overlap in the intensity distribution of white and gray matter is accounted for by allowing GRAY_HIGH to be greater than WM_LOW. This defines a range of ambiguous intensities which are in and of themselves insufficient for classification of white matter versus non-white matter.

Next, a preliminary segmentation is performed based solely on the WM_LOW and WM_HIGH constants. In this step every voxel in the range [WM_LOW, WM_HIGH] is tentatively labeled as white matter, while all other voxels are labeled as non-white. Denoting the output of the skull-stripping as $I_0(\mathbf{x})$, we have

$$I_1(\mathbf{x}) = \begin{cases} 1, & I_0(\mathbf{x}) \in [\text{WM_LOW}, \text{WM_HIGH}] \\ 0, & \text{otherwise} \end{cases}. \quad (4)$$

This volume is then analyzed to locate voxels whose labels are inconsistent with that of their neighborhood. Specifically, voxels that lie in a $3 \times 3 \times 3$ neighborhood which contains a significant number of voxels with labels that differ from that of the central voxel are marked for further processing. The threshold for further examination is a fairly liberal one, usually in the range of 0.2. That is, if more than 20% of a voxel's nearest neighbors have a different label than the voxel in question, it is marked as ambiguous. Note that the bulk of these voxels lie on the border between two regions. Formally, the set of ambiguous voxels A is

given by

$$A = \left\{ \mathbf{x} \mid \frac{\#\{\mathbf{x}' \in N(\mathbf{x}) \mid (I_1(\mathbf{x}') \neq I_1(\mathbf{x}))\}}{\#N(\mathbf{x})} > 0.2 \right\}, \quad (5)$$

where $N(\mathbf{x})$ specifies the (26) nearest neighbors of the voxel at location $\mathbf{x}$. This preliminary sorting into ambiguous and unambiguous voxels is done to reduce the need for the more computationally intensive geometry-based calculations described below.

At this point in the segmentation we use local geometric information to guide the remaining classification. Specifically, we identify the maximally planar orientation of local intensity values and use the values and preliminary labels in the plane of this orientation containing the central voxel to guide the classification procedure. If the local intensities have no obvious planar structure, then the location is typically not near the gray–white border, and all orientations are functionally equivalent for the purposes of segmentation.

For each location in A , we determine the plane-of-least-variance by calculating the intensity variance within slabs of voxels with different orientations, distributed uniformly over the unit hemisphere. More precisely, for each orientation, $\mathbf{n}$, specified by a vertex coordinate of a supertessellated icosahedron, the set of voxels $P_{\mathbf{n}}(\mathbf{x})$ within the slab is given by

$$P_{\mathbf{n}}(\mathbf{x}) = \{\mathbf{x}' \in N_2(\mathbf{x}) \mid \text{abs}((\mathbf{x}' - \mathbf{x}) \cdot \mathbf{n}) \leq p/2\}, \quad (6)$$

where p denotes the thickness of the slab (1 voxel), and $N_2(\mathbf{x})$ denotes the set of voxels within a $5 \times 5 \times 5$ neighborhood. The in-plane mean ($\mu_{\mathbf{n}}$) and variance ($\sigma_{\mathbf{n}}^2$) within the slab orthogonal to the candidate direction are given by

$$\mu_{\mathbf{n}} = \frac{1}{\#P_{\mathbf{n}}(\mathbf{x})} \sum_{\mathbf{x}' \in P_{\mathbf{n}}(\mathbf{x})} I(\mathbf{x}'), \quad (7)$$

$$\sigma_{\mathbf{n}}^2 = \frac{1}{\#P_{\mathbf{n}}(\mathbf{x})} \sum_{\mathbf{x}' \in P_{\mathbf{n}}(\mathbf{x})} (I(\mathbf{x}') - \mu_{\mathbf{n}})^2.$$

We can then determine the orientation which minimizes the within-slab variance, i.e.,

$$\mathbf{n} = \min_{\mathbf{n}'} \sigma_{\mathbf{n}'}^2. \quad (8)$$

Next, we perform an order statistic filtering (similar to a median filter, but with hysteresis) within this “slab-of-least-variance.” In this process the preliminary classification of ambiguous voxels is reversed if more than 60% of the within-slab voxels are labeled as a




FIG. 3. Skull-stripped (left) and white matter-labeled images (right).

different class than the voxel in question. That is, we form two more sets of ambiguous voxels, B and C given by

$$B = \{\mathbf{x} | I_0(\mathbf{x}) \in [\text{WM_LOW}, \text{GRAY_HIGH}]\},$$

$$C = \left\{ \mathbf{x} \left| \frac{\#\{\mathbf{x}' \in P_n(\mathbf{x}) | I_1(\mathbf{x}') \neq I_1(\mathbf{x})\}}{\#P_n(\mathbf{x})} > 0.6 \right. \right\}, \quad (9)$$

where the threshold of 0.6 provides a hysteresis effect—i.e., the classification of the central voxel is changed only if indicated by at least 60% of the voxels within the slab-of-least-variance.

Finally, we form the completed labeling by reversing the preliminary classification decision for voxels which are members of all three of the ambiguous sets:

$$I_{\text{final}}(\mathbf{x}) = \begin{cases} \bar{I}_1(\mathbf{x}), & \mathbf{x} \in A \cap B \cap C \\ I_1(\mathbf{x}), & \text{otherwise,} \end{cases} \quad (10)$$

where $\bar{I}_1(\mathbf{x})$ indicates logical negation of the binary classification $I_1(\mathbf{x})$ of a given voxel $\mathbf{x}$.

Thus, the final labeling is identical to the preliminary intensity-based one except in cases where a voxel and its neighborhood meet three criteria: (1) the voxel's intensity is ambiguous, (2) some of its nearest neighbors have labels which differ from the voxel in question (i.e., their intensity indicates a different tissue class), and (3) a significant number of voxels in the plane of least intensity variance also have differing labels. The result of this white matter labeling procedure is illustrated in Fig. 3.

2.5 Cutting Planes

Once the MR image has been segmented, we wish to generate representations of each cortical hemisphere.

For this purpose, we automatically establish two cutting planes³ which prevent connectivity across them. The first is a sagittal cut along the corpus callosum which serves to separate the two hemispheres, while the second is horizontal through the pons, removing subcortical structures, allowing us to generate two topologically closed surfaces. Both cutting planes are established using an initial seed point specified by the Talairach coordinate of the relevant structures. Next, the in-plane connected components are computed (horizontal for the pons, and sagittal for the corpus callosum) in a region around each seed point, and the slice which contains the minimal cross-sectional area is identified as the location for the desired cutting plane. Finally, the bounding box of the connected component in this slice is used to specify the coordinates of the entire cutting plane, with the sagittal cutting plane extended to the bottom of the slice. The locations of the resulting cutting planes are shown in Fig. 4.

2.6 Connected Components

After the segmentation is complete, we wish to generate a single connected mass representing the white matter structure of each hemisphere. This is accomplished using a connected components procedure. Briefly, starting with seed points in the white matter of each hemisphere, we partition the segmented data into six-connected components, discarding all components with no connectivity to the seed points. Next, we partition the complement of the set into connected components, representing all regions classified as non-white. Isolated components of this partitioning represent regions interior to the white matter components

³“Cutting plane” is actually a misnomer, as the regions which prevent connectivity are rectangular solids.

are then filled. After each phase of the connected components procedure, an iterative $3 \times 3 \times 3$ order statistic filter is applied, filling (or unfilling) a voxel if at least 66% of its neighboring voxels are filled (or unfilled). This is done to eliminate invaginations or protuberances with a 1-D structure. The procedure results in two solid masses of connected voxels, each representing a single cortical hemisphere (Fig. 5).

2.7 Surface Tessellation, Refinement, and Deformation

Once the two white matter volumes have been generated, a surface tessellation is constructed for each hemisphere. This is done by using two triangles to represent each (square) face separating voxels classified as a white matter of a given hemisphere from differently classified voxels. Note that since the resulting tessellations faithfully represent the boundary of the mass of connected voxels for each hemisphere, they must also have the same topology as the surface of the corresponding volume.

However, since the voxel faces are necessarily orthogonal to one of the cardinal axes, the resulting tessellation tend to be jagged at the single-voxel scale. To alleviate this effect, we smooth the initial tessellation using a deformable surface algorithm guided by local MRI intensity values (Dale and Sereno, 1993), resulting in two smoothly tessellated cortical hemispheres. Since the connectivity is explicitly maintained, and self-intersection is prevented (see below), the topology of the surface cannot change during this deformation process.

Deformable templates and related algorithms have been widely used in the computer vision and medical imaging community (Terzopoulos and Fleischer, 1988; Sethian, 1990; Yuille, 1991; Christensen *et al.*, 1996; Davatzikos and Bryan, 1996; McInerney and Terzopoulos, 1996; Sethian, 1996; Davatzikos, 1997; MacDonald, 1998). Many of the difficulties in employing these procedures arise due to factors such as lack of smooth-

ness in the target object, variable or unknown topology, and sensitivity to appropriate initialization. Fortunately, many of these factors are not issues in our problem domain. The topology of the structures we intend to reconstruct is fixed and simple. Furthermore, the initial tessellation of the gray/white surface is quite close to the desired configuration, and thus local minima are not an issue. Finally, the structures we reconstruct are all smooth on the scale of the tessellation, which obviates the need to incorporate techniques for dealing with “shocks” or nondifferentiable locations in the surface.

The surface repositioning we employ is similar in spirit to a number of current techniques for reconstructing a cortical surface by “shrink-wrapping” a predefined shape driven by the intensity values of a T1-weighted MRI volume (Davatzikos and Bryan, 1996; MacDonald, 1998). Specifically, the surface repositioning is accomplished by constructing an energy functional based on the tessellated white-matter surface, the minimization of which results in a final surface that is both smoother than the initial tessellation and more accurate. The first two terms in the energy functional give the surface a spring-like property that acts to smooth the surface and regularize the tessellation. The spring-property is decomposed into terms normal and tangential to the surface, allowing the smoothing and regularization properties to be weighted unequally. These terms are given by:

$$J_n = \frac{1}{2V} \left(\sum_{i=1}^V \sum_{j \in N_1(i)} (\mathbf{n}(i) \cdot (\mathbf{x}_i - \mathbf{x}_j))^2 \right), \quad (5)$$

$$J_t = \frac{1}{2V} \left(\sum_{i=1}^V \sum_{j \in N_1(i)} (\mathbf{e}_0(i) \cdot (\mathbf{x}_i - \mathbf{x}_j))^2 + (\mathbf{e}_1(i) \cdot (\mathbf{x}_i - \mathbf{x}_j))^2 \right), \quad (6)$$

where $N_1(i)$ denotes the set of nearest neighbors of the i^{th}




FIG. 5. Result of calculating connected components (right) for white matter-labeled volume (left). Connected right and left hemisphere voxels are labeled dark and light, respectively.

vertex, V is the total number of vertices in the tessellation, $\mathbf{n}(i)$ is the unit normal vector to the surface at the i^{th} vertex, $\{\mathbf{e}_0(i), \mathbf{e}_1(i)\}$ is an orthonormal basis for the tangent plane at the i^{th} vertex, and $\mathbf{x}_k$ refers to the (x, y, z) position of the k^{th} vertex in the tessellation. The J_t term acts to redistribute vertices to regions where they are needed (i.e., where there is local expansion). This serves to encourage a uniform spacing of vertices without requiring a tessellation with a prohibitively large number of elements. The J_n term imposes a smoothness constraint on the surface, as each node is penalized for being too far in the direction normal to the surface from its neighbors. This has an effect similar to curvature terms frequently employed in deformable surfaces, in that it prevents the surface from becoming too jagged. Note that if these two terms are given equal weighting, they become identical to the “spring” term used in the surface inflation detailed in (Dale and Sereno, 1993; Fischl *et al.*, 1998).

The third term in the energy functional is an intensity-based one. Denoting the volume intensity at position $\mathbf{x}_i$ by $I(\mathbf{x}_i)$, the intensity (or data term) of the energy functional is given by:

$$J_I = \frac{1}{2V} \sum_{i=1}^V (T(i) - I(\mathbf{x}_i))^2, \quad (7)$$

where $T(i)$ is the target intensity value for the i^{th} vertex in the surface tessellation. The value of this term is computed differently in order to construct the gray/white interface as opposed to the pial surface. For finding the gray/white boundary, we use the previously labeled volume to estimate the mean white matter value of border voxels in a 5-mm neighborhood of each vertex, which then serves as the target intensity value. For the pial surface deformation, we have no local information regarding the intensity value of the gray matter/CSF border, as neither voxels containing gray matter nor CSF are explicitly labeled. However, the contrast between CSF and gray matter is sufficiently large that using a predefined global target is a simple and viable alternative to a local estimation of the target values. For both procedures, the value of $I(\mathbf{x}_i)$ is computed on a subvoxel basis using trilinear interpolation. The surface deformation is then accomplished by minimizing an energy functional which is a weighted sum of these three terms:

$$J = J_t + \lambda_n J_n + \lambda_I J_I, \quad (8)$$

where the coefficients λ_n and λ_I specify the relative strength of the smoothness and regularization constraints vis-à-vis the intensity term. The gradient of this functional specifies the direction of movement for the tessellation. That is, the direction of movement of

the k^{th} vertex is given by the negative of the directional derivative with respect to $\mathbf{x}_k$ and is given by:

$$-\frac{\partial J}{\partial \mathbf{x}_k} = \lambda_I (T(k) - I(\mathbf{x}_k)) \nabla I(\mathbf{x}_k) + \sum_{j \in N_1(k)} (\lambda_n (\mathbf{n}(k) \cdot \mathbf{x}_j) + \mathbf{e}_0(k) \cdot \mathbf{x}_j + \mathbf{e}_1(k) \cdot \mathbf{x}_j), \quad (9)$$

where the volume gradient $\nabla I(\mathbf{x}_k)$ is computed using a Gaussian blurred ($\sigma = 1$) version of the MRI volume. For MRI intensity values normalized as discussed in Section 2.2, we weight the components of the gradient $\lambda_n = 0.25$ and $\lambda_I = 0.075$.

The use of the deformable surface procedure necessitates the implementation of an algorithm to prevent self-intersections from occurring in the surface. A straightforward implementation of such an algorithm would have computational complexity F^2 , where F is the number of faces (i.e., triangles) in the tessellation. A single time-step using such an exhaustive approach would require hours on a reasonable sized tessellation, rendering the procedure enormously computationally intensive. Instead, we dramatically reduce the number of triangles to be tested by maintaining a coarsely discretized (4-mm voxels) spatial lookup table (LUT). Each point in this volume contains a list of faces which intersect it. As the surface is deformed, the gradient of the energy functional is computed as detailed above, and each vertex is tentatively moved a short distance along the negative of the gradient direction. Next, the faces attached to the vertex are examined for self-intersection. This is accomplished by constructing a list of discretized voxels that each face intersects, then limiting the triangle intersection computation to those triangles which also intersect that discretized voxel, using the highly optimized triangle-triangle intersection algorithm described by Möller (1997). If self-intersection is detected, the movement delta is cropped to a point where the self-intersection no longer takes place. A surface rendering of the result of these procedures is given in Fig. 6.

Figure 7 shows coronal, sagittal, and horizontal slices (from left to right) with the smoothed reconstructed white matter surface overlaid in as a lighter yellow line. Note the accuracy with which the estimated surface follows all the folds in the cortical surface. Figure 8 contains an analogous set of images displaying the intersection of the pial surface with the same slices. Some care must be taken in evaluating surfaces based on their projection onto a two-dimensional slice, as regions in which the surface appears to vanish may simply be offset from the surface by a millimeter or less in the through-plane direction.



FIG. 6. Original (left), gray/white boundary (middle), and pial surface (right) reconstructions of a left hemisphere.

3. VALIDATION OF SURFACE RECONSTRUCTION ACCURACY

In order to be useful in a variety of contexts (e.g., for computation, visualization, etc.), cortical surface reconstructions must faithfully represent both the topology and the geometry of the true cortical surface. Enforcing the correct topology without respect for the cortical geometry may result in a surface that does not follow that of the actual cortex. Conversely, a surface that represents many of the cortical folds but does not preserve the spherical topology⁴ of the cortex cannot be accurately inflated nor unfolded.

3.1 Validation of Surface Topology

The reconstructed cortical surface is internally represented as a polygonal mesh. The topology of such a polygonal tessellation can be easily determined by computing the Euler number of the surface. The number of holes, g , in the manifold is related to the Euler number by Euler–Lhuillier’s formula: $v - e + f = 2 - 2g$, where v , e , and f are the number of vertices, edges, and faces, respectively (Pont, 1974; Preparata and Shamos, 1985). A topologically perfect reconstruction should result in a surface with no holes ($g = 0$), and thus with Euler number 2. However, the Euler number itself provides no information as to the scale of the defect and is thus of somewhat limited use. For our purposes, small-scale topological defects are not a concern, as they have no effect on the usability of the surface reconstruction. However, larger-scale defects, which either connect or detach vertices across substantial distances, prevent the surface from being accurately flattened or inflated. These types of defects occur occasionally in the surface reconstruction and are removed through a manual editing process that makes use of the surface inflation to highlight the presence and location of the defects. This is graphically illustrated in Fig. 9, which displays lateral and medial views of an inflated surface before (left) and after

(right) editing. This particular surface has no cortical defects prior to editing, although some defects are apparent in the surface due to subcortical gray matter, such as the basal ganglia (1), or due to subcortical white matter tracts (2). The basic topological correctness of the final surfaces is attested to by the fact that we are able to successfully flatten and inflate them, as detailed in the companion to this paper (Fischl *et al.*, 1998).

3.2 Validation of Surface Geometry

Validating the geometry of the surface reconstruction is more difficult than validating the topology, since the geometry is variable across individuals, and is not known *a priori*. Qualitative validation can be accomplished in a number of ways. Visual validation can be achieved by inspecting the intersection of the reconstructed surface with the underlying MRI intensity data, as depicted in Figs. 7 and 8. Note the accuracy with which the surface follows the gray/white interface into the deep sulci. A further indication of the robustness and accuracy of the surface reconstruction process is illustrated in Fig. 10, showing coronal views of the intersection of the pial surface of the left hemispheres of 24 subjects, using data from different scanners and protocols. Each of the brains was reconstructed using the same parameter settings.

A more quantitative means for validating the accuracy of the surface reconstruction is by examining the test–retest reliability of the procedure. That is, will different scans of the same subject result in similar reconstructions? Figure 11 shows estimates of the gray/white matter surface reconstructed from different scans during the same session (center vs right) and from different sessions (left vs center/right). The correlation between the segmented white matter volumes is greater than 96%.

Another means of qualitatively validating the geometry of the surface reconstruction is through the use of functional data. A great deal is known about the location and response-properties of a variety of modality-specific cortical sensory areas in humans. These areas occur in stereotyped locations relative to major sulci and gyri, although some intersubject variability re-

⁴ The actual topology of the cortex is that of a disk. However, we “cap” the midbrain in order to generate a closed surface with spherical topology.





FIG. 9. Lateral (top) and medial (bottom) views of an inflated surface before (left) and after (right) manual editing.

mains. Robust functional activation in these geometrically defined regions on the surface reconstruction therefore provides a qualitative validation of the accuracy of the surface geometry. If there were large geometric inaccuracies in the surface reconstructions, activation would either be assigned to an incorrect location or no statistically significant activation would be detected. The surface reconstruction procedure we have detailed in this paper has been used to analyze data from many cortical areas, attesting to the geometric accuracy of the surfaces obtained from data taken from many subjects and a variety of scanners. These areas include primary visual cortex (Serenó *et al.*, 1995; Tootell *et al.*, 1998a), which localizes to the calcarine sulcus, primary motor, and somatosensory cortices on the anterior and posterior banks of the central sulcus, respectively (Moore *et al.*, 1997; Moore *et al.*, 1998), and primary auditory cortex in Heschl's gyrus, a secondary fold within the sylvian fissure (Talavage *et al.*, 1996, 1997a,b; Talavage 1998a,b). The retinotopic mapping is demonstrated in Fig. 12, which displays the visual field-sign maps (Serenó *et al.*, 1995) of 4 subjects. These maps indicate the borders of the early human retinotopic areas which are known to be localized around the

calcarine sulcus. The robustness and reliability of these maps across subjects, despite large differences in individual folding patterns, provides evidence of the accuracy with which the reconstructed surface follows the true cortical surface.

4. FUTURE WORK


There are a number of unresolved issues that we plan to address in the surface reconstruction procedure. The first modification we are experimenting with is employing curvature information along the normal direction $\mathbf{n}$ as defined by Eq. (8). This allows us to modify the intensity threshold used in the labeling procedure, based on the prior knowledge that cortical white matter does not border any higher intensity tissue classes in T1-weighted MR images. In this way, voxels which are higher intensity than voxels in the directions parallel and antiparallel to $\mathbf{n}$ will be more likely to be classified as white matter, and those with lower intensity are more likely to be classified as nonwhite. This technique will prove particularly useful for white matter voxels contained in thin temporal lobe strands that border the tip of the lateral ventricle. Due to the low intensity



10



11



12

FIG. 10. Intersection of pial surface with the skull-stripped MRI volume in 24 subjects.

FIG. 11. Test-retest reliability of the surface reconstruction from MRI scans taken on different runs in the same session (center vs right) and different sessions on different days (left vs center/right).

FIG. 12. Visual field-sign maps of 4 subjects painted onto flattened representations of the occipital pole, which have been cut down the fundus of the calcarine sulcus at the right of each map.

values of CSF in T1-weighted MR images, partial volume effects frequently reduce the intensity of voxels in these regions, which mainly contain white-matter, to a low enough value that they are misclassified. The use of normal curvature information has succeeded in classifying these voxels correctly, although further work is required before the technique can be put into widespread use.

In addition, we intend to make use of the surface deformation procedure outlined in section 2.7 in order to “shrink-wrap” a shape of known topology onto the segmented filled volume, allowing us to enforce the proper topology on the surface. This type of approach has been employed with some success elsewhere (Davatzikos and Bryan, 1996; MacDonald, 1998), using the MRI intensity volume itself to guide the surface deformation. In contrast, we intend to use the labeled volume, which has been demonstrated to be geometrically correct, as the target of the deformation, thus simplifying the process considerably. This should allow usable surfaces to be generated with no manual intervention, although such intervention will still be useful in order to correct minor geometric inaccuracies in the surface.

Finally, it has not escaped our attention that having explicit representations of both the gray/white and the pial surfaces allows us to compute a variety of morphometric properties of the cortex, such as gray matter volume, thickness, degree of folding, etc. When such measures are combined with automated procedures for morphing individual subjects into a standardized space, or surface-based coordinate system (Drury *et al.*, 1996; Sereno *et al.*, 1996; Thompson *et al.*, 1996; Thompson and Toga, 1996; Davatzikos, 1997; Fischl *et al.*, 1998), it becomes possible to create statistical maps of morphometric differences between individuals and/or groups. The largely automated nature of the methods described here facilitates the routine use of such analysis techniques in a wide range of applications, including the study of psychiatric and neurodegenerative diseases.

5. CONCLUSION

In this paper, we have presented a set of procedures which automatically reconstruct the gray/white and pial surface of the cerebral cortex from a high resolution T1-weighted MRI dataset. Such surface reconstructions are essential for the analysis of features of the cortical surface, including structural properties such as cortical thickness, as well as functional characteristics such as topographic structure (e.g., tonotopy, retinotopy, somatotopy). The reconstruction is a complex task that requires the solution of a number of subtasks such as intensity normalization, skull-stripping, filtering, segmentation, and surface deformation. The tools described here constitute a largely automated process,

which has been used to reconstruct more than 100 brains. Together with the surface transformation methods described in the companion paper (Fischl *et al.*, 1998), these tools provide a means for routinely carrying out cortical surface-based analysis of structural and functional data.

ACKNOWLEDGMENTS

Thanks to Christophe Destrieux for helping with the validation of the surface reconstruction, to Kevin Hall for generating most of the surfaces shown in the document (and many more!), and to Arthur Liu for many useful suggestions. In addition, the authors extend their thanks to an anonymous reviewer who provided many helpful suggestions.

REFERENCES

- Alvarez, L., Lions, P.-L., *et al.* 1992. Image selective smoothing and edge detection by nonlinear diffusion. II. *SIAM J. Numer. Anal.* **29**(3):845–866.
- Atkins, M. S., and Mackiewicz, B. T. 1996. *Automated Segmentation of the Brain in MRI*. The 4th International Conference on Visualization in Biomedical Computing, Hamburg, Germany.
- Brechtbühler, C., Gerig, G., *et al.* 1996. *Compensation of Spatial Inhomogeneity in MRI Based on a parametric Bias Estimate*. The 4th International Conference on Visualization in Biomedical Computing, Hamburg, Germany.
- Carman, G. J., Drury, H. A., *et al.* 1995. Computational Methods for reconstructing and unfolding the cerebral cortex. *Cerebral Cortex*. **5**(6):506–517.
- Catté, F., Lions, P.-L., *et al.* 1992. Image selective smoothing and edge detection by nonlinear diffusion. *SIAM J. Numer. Anal.* **29**(1):182–193.
- Christensen, G. E., Rabbitt, R. D., *et al.* 1996. Deformable Templates Using Large Deformation Kinematics. *IEEE Trans. Image Proc.* **5**(10):1435.
- Collins, D. L., Le Boualher, G., *et al.* 1996. Cortical constraints for non-linear cortical registration. *Visualization in Biomedical Computing* (K. H. Hohne and R. Kikinis, Eds.), pp. 307–316. Springer, Berlin.
- Collins, D. L., Neelin, P., *et al.* 1994. Data in Standardized Talairach Space. *J. Comput. Assist. Tomogr.* **18**(2):292–205.
- Culham, J. C., Brandt, S. A., *et al.* 1998. Cortical fMRI activation produced by attentive tracking of moving targets. *J. Neurophysiol.*, in press.
- 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**(2):162–176.
- Davatzikos, C. 1997. Spatial transformation and registration of brain images using elastically deformable models. *Comput. Vision Image Understand.* **66**(2):207–222.
- Davatzikos, C., and Bryan, R. N. 1996. Using a Deformable Surface Model to Obtain a Shape Representation of the Cortex. *IEEE Trans. Med. Imaging* **15**:785–795.
- DeYoe, E. A., Carman, G. J., *et al.* 1996. Mapping striate and extrastriate visual areas in human cerebral cortex. *Proc. Natl. Acad. Sci. USA* **93**(6):2382–2386.
- Drury, H. A., Van Essen, D. C., *et al.* 1996a. Computerized Mappings of the Cerebral Cortex: A Multiresolution Flattening Method and a Surface-Based Coordinate System. *J. Cogn. Neurosci.* **8**(1):1–28.
- Drury, H. A., Van Essen, D. C., *et al.* 1996b. Analysis and comparison

- of areal partitioning schemes using two-dimensional fluid deformations. *NeuroImage* **3** (S130).
- Drury, H. A., Van Essen, D. C., *et al.* 1997. Warping fMRI activation patterns onto the visible man atlas using fluid deformations of cortical flat maps. *NeuroImage* **5** (S421).
- Felleman, D., and Van Essen, D. C. 1991. Distributed hierarchical processing in primate cerebral cortex. *Cerebral Cortex* **1**:1–47.
- Fischl, B., Dale, A. M., *et al.* 1998. A coordinate system for the cortical surface. *NeuroImage* **7**(4):S740.
- Fischl, B., and Schwartz, E. 1997. Learning an integral equation approximation to nonlinear anisotropic diffusion in image processing. *IEEE Trans. Pattern Anal. Machine Intel.* **19**(4):342–352.
- Fischl, B., and Schwartz, E. 1998. Adaptive nonlocal filtering—a fast alternative to anisotropic diffusion for image enhancement. *IEEE Trans. Pattern Anal. Machine Intel.*, in press.
- Fischl, B., Sereno, M. I., *et al.* 1998. Cortical Surface-Based Analysis II: Inflation, Flattening, a Surface-Based Coordinate System. *NeuroImage*, in press.
- Hadjikhani, N. K., Liu, A. K., *et al.* 1998. Retinotopy and color sensitivity in human visual cortical area V8. *Nature Neurosci.* **1**:235–241.
- Halgren, E., Dale, A. M., *et al.* 1998a. Cortical location of implicit repetition effects in a size-judgment task to visually-presented words. *NeuroImage*, submitted.
- Halgren, E., Dale, A. M., *et al.* 1998b. Location of human face-selective cortex with respect to retinotopic areas. *Human Brain Mapping*, submitted.
- Kaas, J. H., and Krubitzer, L. A. 1991. The organization of extrastriate visual cortex. *Neuroanatomy of Visual Pathways and their Retinotopic Organization* (B. Dreher and S. R. Robinson, Eds.), pp. 302–359. MacMillan, London.
- Kaas, M., Witkin, A., *et al.* 1988. Snakes: Active contour models. *Int. J. Comput. Vision* **1**:321–331.
- MacDonald, D. 1998. *A Method for Identifying Geometrically Simple Surfaces from Three Dimensional Images*. Montreal Neurological Institute, McGill University, Montreal.
- MacDonald, D., Avis, D., *et al.* 1994. Multiple surface identification and matching in magnetic resonance imaging. *Proc. Vis. Biomed. Comp.* **2359**:160–169.
- Manceaux-Demiau, A., Bryan, R. N., *et al.* 1998. A probabilistic ribbon model for shape analysis of the cerebral sulci: Application to the central sulcus. *J. Comput. Assist. Tomogr.*, in press.
- McInerney, T., and Terzopoulos, D. 1996. Deformable Models in Medical Image Analysis: A Survey. *Med. Image Anal.* **1**(2).
- Meyer, C. R., Peyton, H. B., *et al.* 1995. Retrospective correction of Intensity Inhomogeneities in MRI. *IEEE Trans. Med. Imaging* **14**:36–41.
- Möller, T. 1997. A Fast Triangle-Triangle Intersection Test. *J. Graphics Tools* **2**(2):25–30.
- Moore, C. I., Gehi, A., *et al.* 1997. *Basic and Fine Somatotopy in Human SI*. Cogn. Neurosci. Soc., San Francisco, CA.
- Moore, C. I., Stern, C. E., *et al.* 1998. Segregation of Multiple Somatosensory Maps within the Human Postcentral Gyrus Using fMRI. *NeuroImage* **7**(4).
- Nitzberg, M., and Shiota, T. 1992. Nonlinear image filtering with edge and corner enhancement. *IEEE Trans. Pattern Anal. Machine Intel.* **16**(8):826–833.
- Nördstrom, N. K. 1990. Biased anisotropic diffusion: a unified regularization and diffusion approach to edge detection. *Image Vision Comput.* **8**(4):318–327.
- Perona, P., and Malik, J. 1987. *Scale-Space and Edge Detection Using Anisotropic Diffusion*. Proc. IEEE Comput. Soc. Workshop on Computer Vision, Miami, FL.
- Perona, P., and Malik, J. 1990. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Machine Intel.* **12**(7):629–639.
- Pont, J.-C. 1974. *La Topologie Algebrique: Des Origines a Poincare*. Presses Universitaires de France, Paris.
- Preparata, F. P., and Shamos, M. I. 1985. *Computational Geometry: An Introduction*. Springer-Verlag, Berlin.
- Reppas, J. B., Niyogi, S., *et al.* 1997. Representation of motion boundaries in retinotopic human visual cortical areas. *Nature* **388**:175–179.
- Saint-Marc, P., Chen, J. S., *et al.* 1991. Adaptive Smoothing: A General Tool for Early Vision. *IEEE Trans. Pattern Anal. Machine Vision* **13**(6):514.
- Schwartz, E. L. 1990. Computer-aided neuroanatomy of macaque visual cortex. *Comput. Neuroanat.* (E. L. Schwartz, Ed.), pp. 295–315. MIT Press, Cambridge.
- Schwartz, E. L., and Merker, B. 1986. Computer-Aided Neuroanatomy: Differential Geometry of Cortical Surfaces and an Optimal Flattening Algorithm. *IEEE Comp. Graph. Appl.* **6**:36–44.
- Schwartz, E. L., Shaw, A., *et al.* 1989. A numerical solution to the generalized mapmaker's problem: Flattening nonconvex polyhedral surfaces. *IEEE Trans. Pattern Anal. Machine Intel.* **11**:1005–1008.
- Sereno, M. I., and Allman, J. M. 1991. Cortical visual areas in mammals. In *The Neural Basis of Visual Function* (A. G. Leventhal, Ed.), pp. 160–172. MacMillan, London.
- Sereno, M. I., Dale, A. M., *et al.* 1996. A Surface-based Coordinate System for a Canonical Cortex. *NeuroImage* **3**:S252.
- Sereno, M. I., Dale, A. M., *et al.* 1995. Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* **268**:889–893.
- Sethian, J. A. 1990. Numerical Algorithms for propagating interfaces: Hamilton-Jacobi equations and conservation laws. *J. Differ. Geom.* **31**:131–161.
- Sethian, J. A. 1996. *Level Set Methods. Evolving Interfaces in Geometry, Fluid Mechanics, Computer Vision, and Materials Science*. Cambridge University Press, Cambridge.
- Talairach, J., and Tournoux, P. 1988. *Co-Planar Stereotaxic Atlas of the Human Brain*. Thieme Medical Publishers, New York.
- Talavage 1998a. *Primary Auditory Cortex: HST*. MIT, Boston.
- Talavage, T., Ledden, P., *et al.* 1997a. Multiple phase-encoded tonotopic maps in human auditory cortex. *NeuroImage* **5**(S8).
- Talavage, T., Ledden, P., *et al.* 1997b. Phase-encoded tonotopic maps in human auditory cortex. *ISMRM* **5**:6.
- Talavage, T., Ledden, P., *et al.* 1996. Preliminary fMRI evidence for tonotopicity in human auditory cortex. *NeuroImage* **3**(S355).
- Talavage, T. M. 1998b. Functional magnetic resonance imaging of the frequency organization of human auditory cortex. In *Health Science and Technology*. MIT, Boston.
- Teo, P. C., Sapiro, G., *et al.* 1997. Creating Connected Representations of Cortical Gray Matter for Functional MRI Visualization. *IEEE Trans. Med. Imag.* **16**(6):852–863.
- Terzopoulos, D., and Fleischer, K. 1988. Deformable models. *Visual Comp.* **4**(6):306–331.
- Thompson, P., Schwartz, C., *et al.* 1996. Three-dimensional statistical analysis of sulcal variability in the human brain. *J. Neurosci.* **16**(13):4261–4274.
- Thompson, P. M., and Toga, A. W. 1996. A surface-based technique for warping 3-dimensional images of the brain. *IEEE Trans. Med. Imag.* **15**:1–16.

- Tootell, R. B. H., Dale, A. M., *et al.* 1996a. fMRI analysis of human visual cortical area V3A. *NeuroImage* **3**:S358.
- Tootell, R. B. H., Hadjikhani, N. K., *et al.* 1998a. Functional analysis of primary visual cortex (V1) in humans. *Proc. Natl. Acad. Sci. USA* **95**:811–817.
- Tootell, R. B. H., Mendola, J. D., *et al.* 1997. Functional analysis of V3A and related areas in human visual cortex. *J. Neurosci.* **17**:7060–7078.
- Tootell, R. B. H., Mendola, J. D., *et al.* 1998b. The representation of the ipsilateral visual field in human cerebral cortex. *Proc. Natl. Acad. Sci. USA* **95**:818–824.
- Tootell, R. B. H., Reppas, J. B., *et al.* 1995a. Visual motion aftereffect in human cortical area MT revealed by functional magnetic resonance imaging. *Nature* **375**(May 11):139–141.
- Tootell, R. B. H., Reppas, J. B., *et al.* 1995b. Functional analysis of human MT and related visual cortical areas using magnetic resonance imaging. *J. Neurosci.* **15**(4):3215–3230.
- Tootell, R. B. H., Dale, A. M., *et al.* 1996b. New images from human visual cortex. *Trends Neurosci.* **19**:481–489.
- 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.
- Wagner, M., Fuchs, M., *et al.* 1998. Cortex segmentation from 3D MR images for MEG reconstructions. In *Biomagnetism: Fundamental Research and Clinical Applications*, pp. 433–438. IOS Press, Amsterdam.
- Weickert, J. 1997. *Recursive Separable Schemes for Nonlinear Diffusion Filters*. First International Conference on Scale-Space Theory in Computer Vision, Springer, Utrecht, The Netherlands.
- Wells, W., Kikinis, R., *et al.* 1994. *Statistical Intensity Correction and Segmentation of Magnetic Resonance Image Data*. Proceedings of the Third Conference on Visualization in Biomedical Computing VBC'94.
- Whitaker, R., and Gerig, G. 1994. Vector-valued Diffusion. In *Geometry-Driven Diffusion in Computer Vision* (B. M. t. H. Romeny, Ed.), pp. 93–134. Kluwer Academic, Dordrecht.
- Yuille, A. L. 1991. Deformable Templates for Face Recognition. *J. Cogn. Neurosci.* **3**(1):59–70.
- Zeki, S., and Shipp, S. 1988. The functional logic of cortical connections. *Nature* **335**(6188):311–317.