Anatomically Consistent Segmentation of the Human Cortex for Functional MRI Visualization

Patrick C. Teo
Dept. of Computer Science
Stanford University
Stanford, CA 94305
teo@white.stanford.edu

Guillermo Sapiro
Imaging Technology Dept.
Hewlett-Packard Labs
Palo Alto, CA 94304
guille@hplgui.hpl.hp.com

Brian A. Wandell
Dept. of Psychology
Stanford University
Stanford, CA 94305
brian@white.stanford.edu

Abstract

We describe a system that is being used to segment gray matter from volumetric representations of the human cortex obtained using magnetic resonance imaging. The segmentation algorithm identifies gray matter voxels and computes their connectivity. The method differs from existing schemes in that it exploits knowledge of the anatomy of human cortex and produces anatomically consistent segmentations. The method is based on a novel and computationally efficient technique of incorporating structural constraints into the segmentation algorithm. Because the gray matter segmentation is anatomically consistent, it can be used together with functional magnetic resonance imaging measurements to visualize the spatial pattern of cortical activity within the gray matter.

Keywords: Segmentation, human cortex, structural MRI, functional MRI, visualization.

1 Introduction

Magnetic resonance scanners can be used to measure various aspects of a source material. In one important application magnetic resonance imaging (MRI) is used as a non-invasive method of visualizing biological structures (sMRI). Recently, functional magnetic resonance imaging (fMRI) has provided a method of visualizing a correlate of neural activity in the brain.[1] The ability to measure cortical activity non-invasively is an important breakthrough, providing us with a new opportunity to study the activity of single human brains at relatively high spatial resolution.

Among the various parts of the brain, the cortex is the most prominent, and one of the most intensely studied. The cortex is divided into two hemispheres connected by a massive set of nerve fibers. The cortex is composed largely of two types of tissue: gray matter and white matter. Gray matter forms the outer layer of the cortex, encasing the inner white matter almost completely. In humans, gray matter is, on average, 3 mm thick. Gray matter tissue contains a high density of neurons (approximately $10^5/\text{mm}^3$ [31]) which form the computational units of the cortex. White matter is made up of nerve fibers that connect different parts of the cortex, as well as the cortex with other parts of the brain. Functional magnetic resonance imaging measures the neural activity in the gray matter.

Despite its complex outward appearance, the structure of each hemisphere is quite straightforward and consistent across human brains. The cortical gray matter surface is the same as that of a crumpled sheet; i.e., it does not have any holes or self-intersections. To visualize the cortex as a three-dimensional structure from sMRI data, most algorithms require some prior segmentation. This segmentation may be simply discriminating between white matter, gray matter, and cerebral spinal fluid (CSF, the fluid that fills the cranial cavity), or the segmentation may involve a more elaborate labeling of different anatomical structures. Based on this segmentation, the brain is rendered as a three dimensional surface and the observer sees mainly those portions outside the folds.

Because it is important to see neural activity buried deep within the three-dimensional folds of the brain, visualizing fMRI data requires novel visualization techniques. An increasingly popular way of visualizing such mappings is to superimpose fMRI measurements on flattened representations of the cortical surface [8, 9, 10, 32]. One method of obtaining a

[1] This correlate is the relative amount of oxygen in the surrounding blood-flow. Because (a) the relative amounts of oxygen around active areas of the cortex are different from those around inactive areas of the cortex, and (b) the paramagnetic properties of oxygenated and deoxygenated blood differ, MRI can be used as an indirect measure of neural activity.

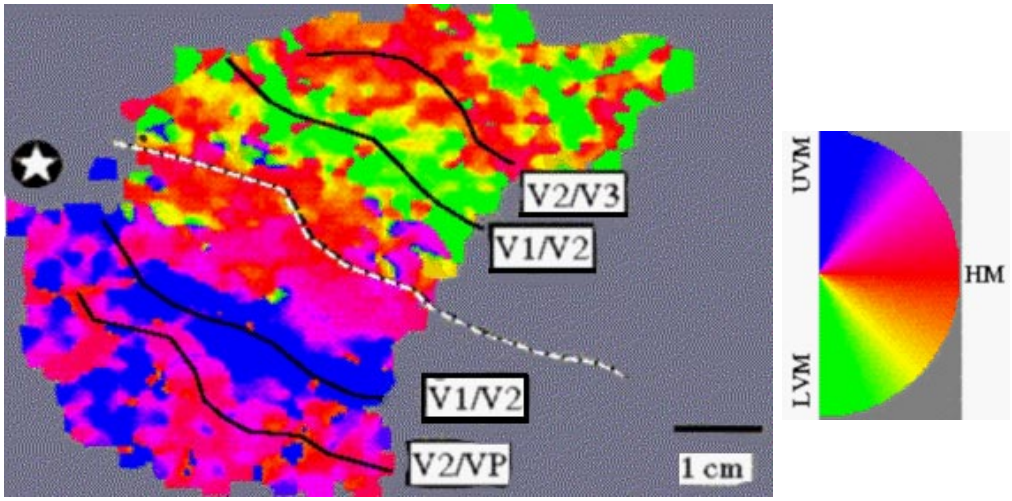

Figure 1: Functional MRI measurements in response to a visual stimulus overlaid on a flattened representation of the occipital lobe (source: Engel et al., 1997). The data represent the retinotopic, with respect to polar angle, of several visual areas in the occipital lobe. The color codes the stimulus angle at the moment each location within cortex was active. The angles are indicated by the key on the upper right: UVM is upper vertical meridian, HM is horizontal meridian, and LVM is lower vertical meridian. The left hemisphere represents only the right visual field. Reversals in the change of the polar angle representation identify the boundaries between several visual areas (V1,V2,V3,VP).

flattened representation is to compute the best planar representation of a region of gray matter such that distances on the plane are similar to the corresponding actual (geodesic) distances on the gray matter surface.

Figure 1 is an example of how fMRI measurements can be represented on a flattened region of the occipital lobe. Data from monkey and human studies show that neurons within area V1 are *retinotopically organized*: Neurons that are responsive to nearby regions of the visual field are located close to one another within the gray matter layer. Because of the retinotopic organization of visual areas, it is possible to create simple visual stimuli that generate continuous traveling waves of neural activity in visual cortex [31]. The figure illustrates how the spatial structure of these traveling waves, represented on the flattened cortical surface, can be used to determine the locations and boundaries of different areas of specialized processing.

In order to create flat maps, it is important not only to be able to identify regions in the MR data that correspond to gray matter but also to determine their connectivity. The most time-consuming aspect of creating flattened representations of visual cortex is to identify the layer of gray matter from the structural MRI of the brain. This task is often done manually, though as we discuss later, several algorithms have been proposed.

In this paper, we describe a semi-automatic system that is being used to segment gray

matter voxels in structural MR data. In addition to identifying gray matter voxels, the segmentation also computes the connectivity of gray matter voxels within the gray matter layer. Furthermore, the segmentation ensures that the topology of the connected gray matter voxels is anatomically correct; i.e., it cannot have holes or self-intersections. This is of particular importance if the gray matter segmentation will be used, together with fMR measurements, to visualize the spatial pattern of cortical activity within the gray matter layer. The presence of holes or self-intersections, caused by incorrect connections or erroneous segmentations, could have disastrous effects on these spatial patterns.

Gray matter segmentation is difficult for a variety of reasons. First, voxel intensities of gray matter tissue are spread over a large range. Second, given the current spatial resolution of MRI, regions of gray matter voxels can be as narrow as one or two voxels. This is thin compared to regions of white matter voxels which make up the bulk of the MR volume, so that a large percentage of gray matter voxels suffer from severe *partial volume* effects.[2] Partial volume reduces the effectiveness of intensity-based gray matter segmentation algorithms. Third, it is impossible to determine the connectivity of the gray matter voxels by only examining the segmented gray matter voxels. For example, the gray matter voxels on opposite sides of a sulcus (an infolding in the gray matter) could potentially be very close or even adjacent to one another on the sampling grid. Yet, the gray matter voxels on opposite sides of the sulcus should not be labeled as connected to each other. Because connectivity cannot be discerned from the intensity-based segmentation alone, ensuring that the topology of the segmented gray matter is anatomically correct is also impossible.

Various gray-matter segmentation techniques have been proposed and many of them use generic image segmentation techniques that do not fully take advantage of knowledge of the anatomy of the cortex [33, 29, 34, 30, 6, 3]. In our application, however, the gray matter segmentation is ultimately used to reveal spatial patterns of cortical activity within the gray matter layer so that the gray matter segmentation and its connectivity information needs to be anatomically consistent. There are several techniques that do benefit from some amount of anatomical knowledge, but the manner in which such knowledge is employed tends to be local and statistical [16, 28]. A notable exception is the method proposed by Joliot and Mazoyer [17] which will be discussed later in the paper.

Unlike generic image segmentation techniques, existing methods using deformable membranes (so called snakes or balloons) often do produce anatomically consistent segmenta-

[2] Partial volume effects occur when a voxel contains more than one tissue type. For example, the intensity of a voxel straddling the gray/white matter boundary or gray matter/CSF boundary would have a mean intensity value different from a voxel containing gray matter exclusively.

tions [5, 8, 18, 21, 25, 27]. If these membranes are initialized to be topologically equivalent to a sheet, they are automatically consistent with the topology of the gray matter layer. Unfortunately, the deformation of these membranes are quite prone to being captured by local minima. These minima occur when the membrane is required to deform into sulci that are deep but have narrow entrances. The region of human cortex near the occipital lobe, for example, is highly convoluted and has many such sulci. As a result, good initialization of such techniques is often required. On the other hand, these methods have the advantage that they incorporate smoothness as part of their segmentation criterion and are capable of producing sub-pixel classification (of the boundary between white and gray matter, for example). The segmentation method proposed in this paper could, quite straightforwardly, be used to initialize some of these algorithms.

2 Method

Our method consists of four stages. First, the white matter and CSF regions in the MR volume are segmented. Second, the desired cortical white matter component is selected by the user. Third, the consistency of the white matter topology with the actual anatomy is verified. Finally, gray matter is segmented by growing out from the white matter boundary while connectivity of the segmented gray matter is computed simultaneously.

2.1 Segmentation of White Matter and CSF

In the first stage, voxels containing white matter tissue and cerebral spinal fluid (CSF) are segmented. White matter is segmented before gray matter because (a) the variability of white matter voxels is lesser than the variability of gray matter voxels, and (b) white matter regions are also much larger in size than gray matter regions. Although only white matter and CSF are being classified, three classes are considered. The third class roughly corresponds to gray matter but, as shall be seen later, the segmentation of this class is poor, and will consequently be rejected in favor of the results of a secondary gray matter segmentation algorithm. The intensity of a voxel belonging to each class is modeled as an independent random variable with a normal distribution. Thus, the likelihood of a particular voxel belonging to a certain class is:

$$\Pr(V_i = v | C_i = c) = \frac{1}{\sqrt{2\pi}\sigma_c} \exp\left(-\frac{1}{2}\frac{(v-\mu_c)^2}{\sigma_c^2}\right) \quad (1)$$

where i is a spatial index ranging over all voxels in the MR volume and c is one of {white, gray, CSF}. V_i and C_i correspond to the intensity and classification of voxel i respectively. Currently, the user interactively adjusts the parameters μ_c and σ_c until a satisfactory segmentation is produced. Typically, these parameters remain the same over all MR measurements that were obtained by a given scanning protocol.

Next, the posterior probabilities of each voxel belonging to each class is computed using Bayes' Rule and anisotropic smoothing. The posterior probability is computed for each voxel independently using Bayes' Rule together with a homogeneous prior:

$$\Pr(C_i = c|V_i = v) = \frac{1}{K} \Pr(V_i = v|C_i = c) \Pr(C_i = c) \tag{2}$$

where K is a normalizing constant independent of c. Adopting a homogeneous prior implies that $\Pr(C_i = c)$ is the same over all spatial indices i. The prior probability typically reflects the relative frequency of each class. For example, since white matter voxels occur more frequently than gray matter voxels, the prior probability of white matter is larger than that of gray matter.

In the second step, the posterior volumes are smoothed anisotropically in three dimensions while preserving discontinuities. Figure 2 shows an example of a posterior derived from a homogeneous prior along with its smoothed counterpart. The anisotropic smoothing technique applied is a 3D extension of the original 2D version proposed by Perona and Malik [22]. This step involves simulating a discretization of the following partial differential equation for a small number of iterations:

$$\frac{\partial P_c}{\partial t} = \mathrm{div}(g(||\nabla P_c||)\nabla P_c) \tag{3}$$

where $P_c = \Pr(C = c|V)$ represents the volume of posterior probabilities for class c. $g(||\nabla P_c||) = \exp(-(||\nabla P_c||/\eta_c)^2)$ and η_c represents the rate of diffusion for class c. The function $g(\cdot)$ controls the local amount of diffusion such that diffusion across discontinuities in the volume is suppressed. These parameters typically remain unchanged across different MR data sets. The intuition for applying anisotropic smoothing on the posterior probabilities is deferred to the discussion section.

Finally, segmentation is carried out by labeling each voxel with the class that yields the maximum posterior probability estimate. That is,

$$C_i^* = \underset{c \in \{\text{white, gray, CSF}\}}{\arg\max} \Pr^*(C_i = c|V_i = v) \tag{4}$$

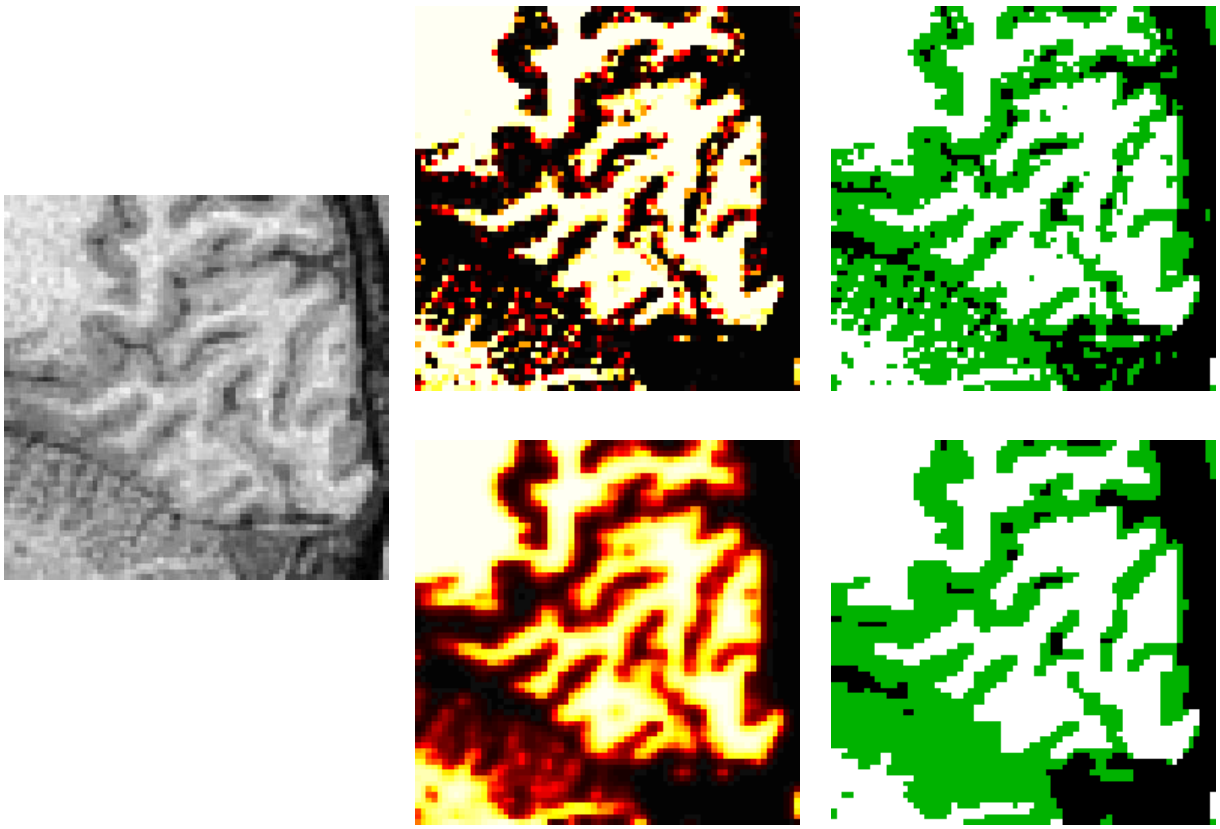

Figure 2: (Top row) Left: Intensity image of MR data. Middle: Image of posterior probabilities corresponding to white matter class. Right: Image of corresponding MAP classification. Brighter regions in the posterior image correspond to areas with higher probability. White regions in the classification image correspond to areas classified as white matter; black regions correspond to areas classified as CSF. (Bottom row) Left: Image of white matter posterior probabilities after being anisotropically smoothed. Right: Image of MAP classification computed with smoothed posteriors.

where $\Pr^*(C_i = c|V_i = v)$ corresponds to the anisotropically smoothed posterior. Figure 2 shows the MAP segmentation results obtained first using a posterior with a homogeneous prior, and subsequently with its smoothed counterpart.[3] As can be seen from the figure, the segmentation results for white matter and CSF are quite good after anisotropic smoothing is applied to the posterior volumes; the gray matter segmentation, on the other hand, remains fairly poor. As a result, only the white matter and CSF segmentations are retained.

[3]Since we were using a head coil and working on small areas, we didn't find any need to use the technique described in [33], which can be easily incorporated in our system if necessary.

2.2 Selection of Cortical White Matter

As gray matter surrounds white matter almost completely, accurate white matter segmentation can be used to obtain accurate gray matter segmentation. Thus, attention is next focused on the white matter segmentation that was computed in the previous stage. Since the white matter segmentation is represented on a grid, its digital topology needs to be defined. White matter connectivity is defined using 26-neighbor adjacency; that is, two distinct white matter voxels are adjacent to each other if their spatial coordinates differ by no more than one. Two white matter voxels are connected to each other if there is a path of white matter voxels connecting the two such that all neighboring pairs of white matter voxels along the path are 26-neighbor adjacent. Gray matter connectivity is also defined using 26-neighbor adjacency. CSF connectivity, on the other hand, is defined using 6-neighbor adjacency; that is, two distinct voxels classified as CSF are adjacent to each other if exactly one of their spatial coordinates differ by one. The reason for defining the connectivity of CSF differently is to prevent intersections between regions of CSF and white matter (or gray matter) [19].

At this stage, the user selects a voxel in the cortical white matter component, via a graphical user interface, and a flood-filling algorithm employing 26-neighbor adjacency is used to extract the corresponding white matter connected component [13]. The flood-filling algorithm begins by marking the user's selection and then proceeds iteratively, marking all unmarked voxels adjacent to existing marked voxels until there are no more unmarked voxels adjacent to marked ones. The purpose of this stage is, primarily, to remove extra-cortical components like the cerebellum. If the MR volume has been cropped to an appropriate region of interest within the cortex, the cortical white matter component typically corresponds to the largest white matter component.

2.3 Verification of White Matter Topology

After the cortical white matter component has been selected, its consistency with cortical anatomy is verified. Since the gray matter is a sheet with no holes or self-intersections, and it encases white matter almost completely, it is sufficient to ensure that the selected white matter component is free of cavities and handles. Cavities are non-white matter regions in the classification that are completely surrounded by white matter while handles are also holes but are not completely surrounded by white matter. An everyday example of a cavity is the inside of a tennis ball; likewise, an example of a handle is the hole of a doughnut. Figure 3 shows an example of a white matter cavity and a white matter handle. Handles

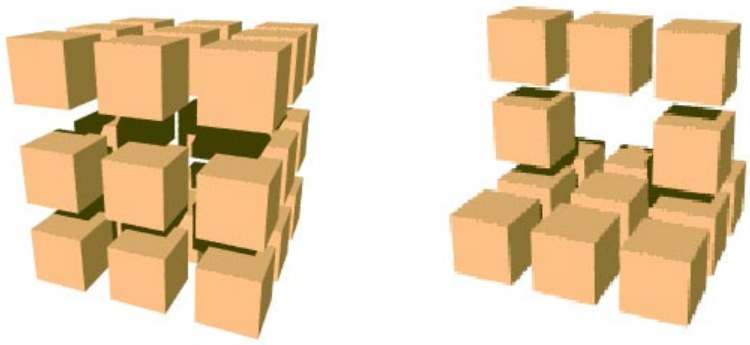

Figure 3: Left: example of a white matter cavity. Note that the middle cube is missing. Right: example of a white matter handle.

and cavities are detected at this stage and brought to the attention of the user who can then edit the segmentation results to remove them.

Cavities are first identified using a flood-filling algorithm. The algorithm identifies non-white matter connected components that are connected to the boundary of the bounding volume containing the cortical white matter component. This is accomplished by repeatedly initiating the flood-filling algorithm from non-white matter voxels on the boundary. After that, all non-white matter connected components that have *not* yet been filled must be encased entirely by white matter; i.e., they must be cavities. The user can then choose to fill these cavities or to edit the segmentation so as to remove the cavities. In practice, these cavities are typically small, and are usually filled.

Identifying handles is more complicated; in the current implementation, the number of handles is computed instead. The number of handles is computed from the Euler characteristic χ which is equal to the sum of the number of connected components and cavities, minus the number of handles. Since the number of connected components is one (as a result of second stage), and the number of cavities is zero (as a result of the previous step), the number of handles is simply one minus the Euler characteristic. The Euler characteristic is, in turn, computed as the sum of the local Euler characteristic over all $2 \times 2 \times 2$ voxel neighborhoods:

$$\chi^{\text{local}} = \sum_i \frac{v_i}{8} - \frac{e_i}{4} + \frac{f_i}{2} - o_i \tag{5}$$

where i ranges over all $2 \times 2 \times 2$ voxel neighborhoods and v_i, e_i, f_i, o_i represent the number of vertices, edges, faces and octants in the i-th neighborhood respectively [20]. Since there are only 256 possible $2 \times 2 \times 2$ neighborhood configurations, the local Euler characteristic of each possible configuration is precomputed and stored in a table. The Euler characteristic,

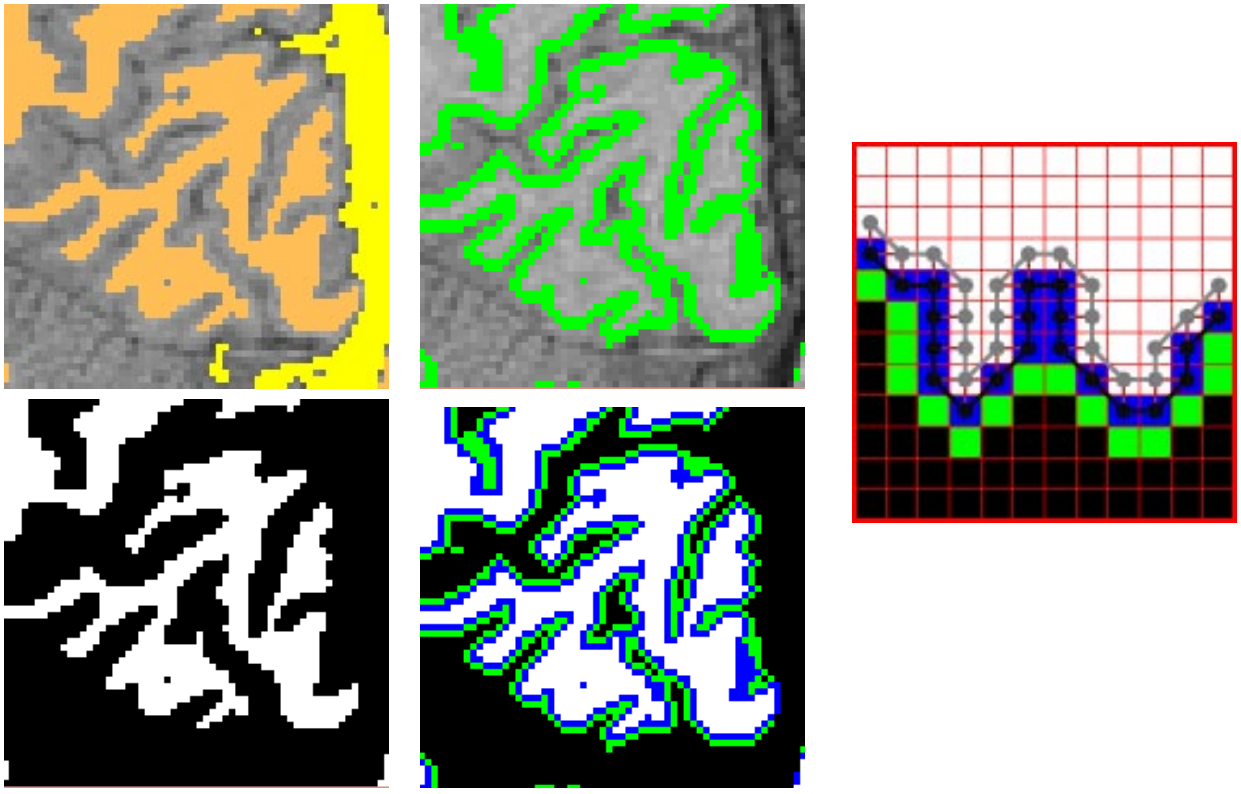

(a) (b)

Figure 4: (a) (Top row) Left: MRI image with white matter and CSF classification overlaid. Right: MRI image with gray matter classification overlaid. (Bottom row) Left: white matter classification. Right: two layers of gray matter classification grown out from white matter classification. (b) Schematic showing two layers of gray matter grown out from the white matter boundary. The connectivity of the white matter boundary and the first layer of gray matter is represented by the links between adjacent filled circles.

and thus, the number of handles, is then computed efficiently using table lookups.

2.4 Gray Matter Segmentation and Connectivity

After the topology of the cortical white matter component has been verified, gray matter segmentation is computed by growing layers from the white matter boundary. Since the white matter component is void of cavities and handles, the gray matter grown out from the white matter boundary, when correctly computed, is topologically the same as a sheet. The number of gray matter layers that are grown is determined by the spatial resolution of the MR data. If the MR data has a spatial resolution of 1 mm along each spatial dimension, then a maximum of 3 layers are grown since gray matter, in the anatomy, is generally 3 mm thick (approximately) perpendicular to white matter. Figure 4 shows an example of gray

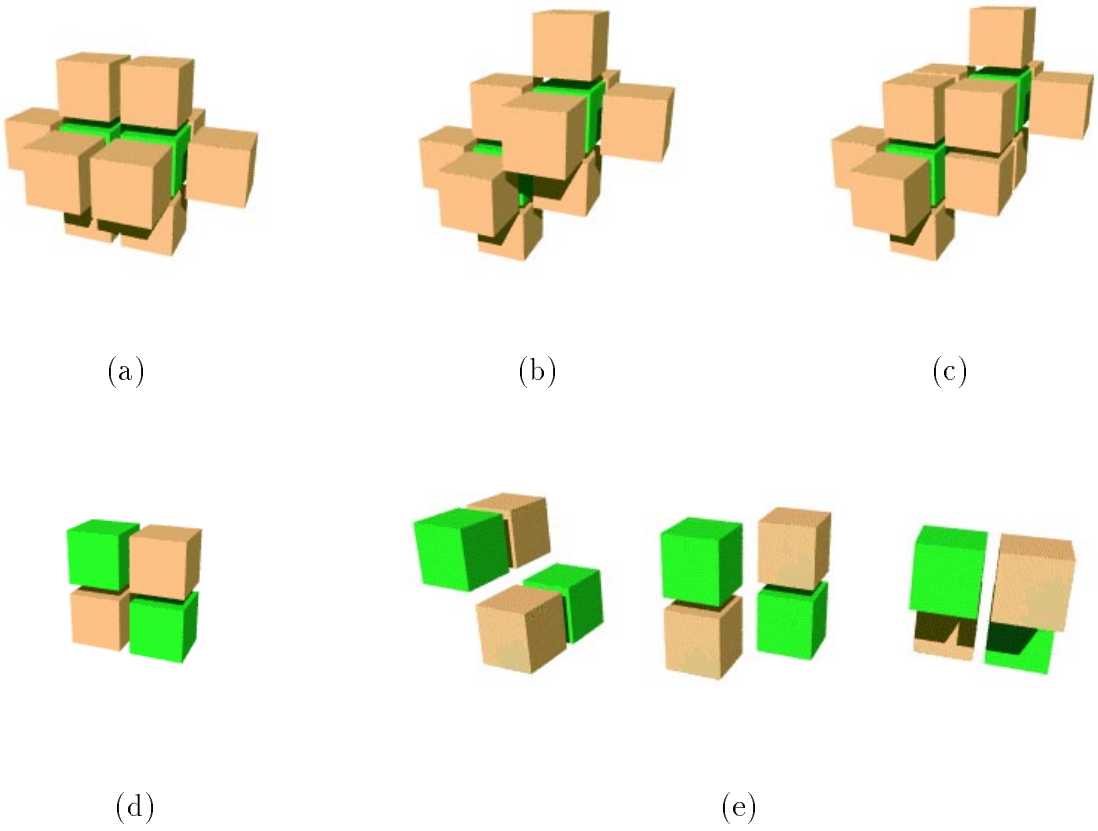

Figure 5: Green shaded cubes represent gray matter voxels; brown shaded cubes represent possible white matter voxels from which these gray matter voxels could have been grown. (Top row) (a)-(c): different configuration of pairs of gray matter voxels and their white matter parents. (Bottom row) (d) 2D exception to the connectivity rule; (e) 3D exceptions to the connectivity rule. See text for details.

matter segmentation produced by growing two layers from the boundary of the white matter component.

Each additional layer of gray matter is grown out from the previous layer (or from the boundary of the white matter component for the first layer) in the same fashion. The process of growing a single layer of gray matter voxels is carried out in two steps: first, new gray matter voxels are identified and labeled; second, connectivity of the new gray matter voxels is determined.

During the first step, unclassified voxels (voxels which have not been classified as white or CSF) that are 6-neighbor adjacent to some gray matter voxel in the previous layer are classified as gray matter voxels belonging to the current layer. The gray matter voxels in the previous layer are known as *parents* of these newly classified voxels. Each new voxel is classified as gray provided all of its parents are connected (as determined by the previous

connectivity step). In other words, the algorithm will not grow a new gray voxel if doing so results in a contention for the new voxel among previous-layer voxels which are unconnected. For example, it will not grow from two gray matter voxels on opposite sides of a sulcus if they are separated by only one voxel. For the first layer, unclassified voxels that are 6-neighbor adjacent to some white matter boundary voxel are classified as gray. Each new voxel is also classified as gray provided all of its parents are connected. Connectivity, in this case, is determined from the 26-neighbor adjacency of white matter voxel parents.

During the second step, connectivity of the newly classified gray matter voxels is computed. Connectivity of gray matter voxels is divided into two categories: inter-layer and intra-layer. Gray matter voxels between different layers are considered connected if they are 6-neighbor adjacent. Ascertaining the connectivity of gray matter voxels within the same layer is a little more involved as it requires examining the connectivity of the voxels' parents. Figures 5 (a)-(c) show the parents of pairs of gray matter voxels in all possible configurations. Two gray matter voxels within the same layer are considered connected if they are (1) 26-neighbor adjacent, and (2) either share a common parent or have parents that are connected (as computed in the previous connectivity step). In addition, the connectivity so determined cannot result in intersecting regions. Figures 5 (d) and (e) show different configurations of voxels that result in intersecting regions. In Figure 5 (d), for example, if the green shaded cubes (gray matter) were labeled as connected, then the digital region formed by these two cubes would intersect the digital region formed by the two brown shaded cubes (white matter parents or gray matter parents from the previous layer). Figure 5 (e) shows all the other remaining cases. For the first layer, since connectivity of white matter voxels is determined using 26-neighbor adjacency, two 26-neighbor adjacent gray matter voxels are considered connected if they either share a common white matter parent or have white matter parents that are 26-neighbor adjacent. Despite the complexity of the connectivity algorithm, it can be efficiently implemented with tables.

The output of this stage, and thus, of the entire algorithm, is (1) a segmentation of the gray matter voxels in the MR volume, and (2) a connectivity graph where the vertices of the graph denote the segmented gray matter voxels and the edges represent the connectivity of adjacent gray matter voxels. Distances (geodesics) between pairs of segmented gray matter voxels in the MR volume are then measured by computing the shortest paths between corresponding vertices within this connectivity graph [7]. Figure 6 plots the distances of gray matter voxels from a selected gray matter voxel for different distance measures. In Figure 6 (a), the 2D Euclidean distance is used. In Figure 6 (b), the distance used is the shortest distance on the connectivity graph restricted to the current slice. Finally, in Figure 6 (c),

(a) (b) (c)

Figure 6: The colored areas in each of the above images represent the same gray matter segmentation of a region in the cortex. The intensities of these areas correspond to distances from the same selected gray matter voxel (at bottom-center) such that bright areas represent shorter distances and darker areas denote longer distances. Figure (a) plots the Euclidean distances on the plane; i.e., the distances between two gray matter voxels in the image is the length of the straight line segment between them. Figure (b) plots distance as the shortest distance within the gray matter connectivity graph restricted to this plane (1D manifold). Figure (c) plots distance as the shortest distance within the original gray matter connectivity graph (2D manifold). Thus, in the latter, the shortest distance between two gray matter voxels may be a path that is partially outside of the plane. The respective distances are ordered such that the Euclidean distance is necessarily the shortest, followed by the 2D manifold distance, and finally, by the 1D manifold distance which is the longest of the three.

the distance plotted is the shortest distance on the full connectivity graph. The different distances between two gray matter voxels is shortest when the Euclidean distance is used. Also, distances measured using the full connectivity graph are no greater than distances measured using the connectivity graph restricted to the current slice. For the purpose of visualizing spatial patterns of cortical activity measured using fMRI, the full connectivity graph provides the appropriate distance relations. Thus, it is important not only to determine the gray matter segmentation, but also the connectivity between gray matter voxels.

3 Results

The segmentation technique described in this paper has been implemented and is being used to segment gray matter voxels in MR data. The segmented gray matter voxels and their connectivity are used together with functional MR data to visualize the spatial pattern of neural activity within the gray matter layer. The portion of the cortex currently being studied are the parietal and occipital lobes, which are located at the posterior end of the cortex. These regions are important for vision, and they are also the most convoluted parts of the

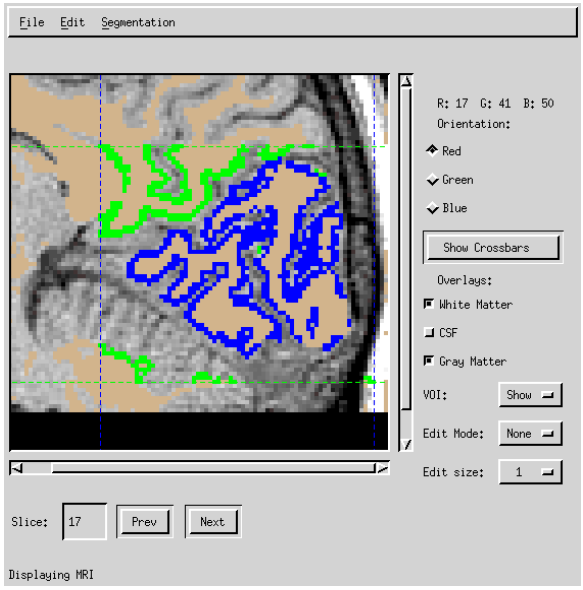

Figure 7: Example of the windows based interactive system where the algorithms described in this paper are implemented. The user can interactively perform a large number of operations, for example combine automatic segmentation and unfolding with manual corrections.

cortex. Manual gray matter segmentation of just the occipital lobe of one hemisphere with rudimentary segmentation tools requires about 18 hours for an experienced person. Much of the time is typically spent on visually inspecting connectivity and ensuring topological correctness. With the present method, the entire procedure takes about half an hour. The total time required by the segmentation algorithm is about 2 minutes; the rest of the time is spent manually verifying the segmentation on each slice.

The segmentation and visualization schemes have been implemented as part of a user friendly interactive windows system, which permits the user to select regions of interest, perform the different algorithms described in this paper, verify and correct the automatic segmentation, unfold the cortex, etc. Figure 7 shows an example of the windows in the system.

Figure 8 shows several comparisons between gray matter segmentation results derived manually and those computed using the proposed method. In this, and in all other comparisons shown in this paper, no manual editing of the segmentation was carried out. The automatic segmentation results are similar to those obtained manually despite the large number of deep and narrow folds in this region of the cortex. The proposed method has difficulties when the folds are extremely thin; i.e., when the white matter is about one voxel thick. This is due to the anisotropic smoothing algorithm; in the algorithm, thin regions of white matter that have small posterior probabilities tend to be smoothed out in favor of

larger regions of CSF. An improvement in the anisotropic smoothing technique should ameliorate this problem. It is also not clear that manual segmentations provide a gold standard for assessing the problem as manual segmentations produced by different people are often different.

Figure 9 shows flattened representations of a portion of the same occipital lobe of the cortex. Figure 9 (a) shows the results computed from a manual segmentation of the gray matter while Figure 9 (b) shows the results computed from gray matter that was automatically segmented. In each case, once the gray matter voxels have been segmented and connectivity determined, a flattening algorithm [32] is then applied to compute the best possible flattened representation of the gray matter layer such that distances between pairs of gray matter voxels within the gray matter layer are as similar as possible to their (Euclidean) distances in the flattened representation. The different intensities in the figures represent different Euclidean distances in 3D of the corresponding gray matter voxel from a fiducial plane; in this case, it is the distance from the leftmost sagittal plane. Brighter regions indicate larger distances while darker regions indicate shorter distances. Although the flattened representation derived from manual segmentation is smoother, both flattened representations are qualitatively very similar in shape as well as in size. The two bright regions in both figures correspond to the lips of the sulcus around which they border. The sulcus itself (known as the calcarine sulcus) is represented by the dark region in the middle.

In Figure 10, fMRI measurements from two different experiments are overlaid on the flattened representations. The images in the left and right columns correspond to overlays on flattened representations of gray matter that have been segmented manually and automatically respectively. The top and bottom rows show results obtained using different visual stimuli. The overlay on each flattened representation shows the temporal phase of the neural activity caused by a periodic, moving visual stimulus that induces a traveling wave within several different cortical regions. The figure shows that the results obtained using the automatic segmentation technique is visually similar to that obtained with manual segmentation. The spatial pattern of these phase maps are used to determine the locations of several different retinoptically organized visual areas.

4 Discussion

In this section, we first discuss the segmentation method proposed in [17] which is closely related to ours. Following that, we review some of the decisions that motivated the design

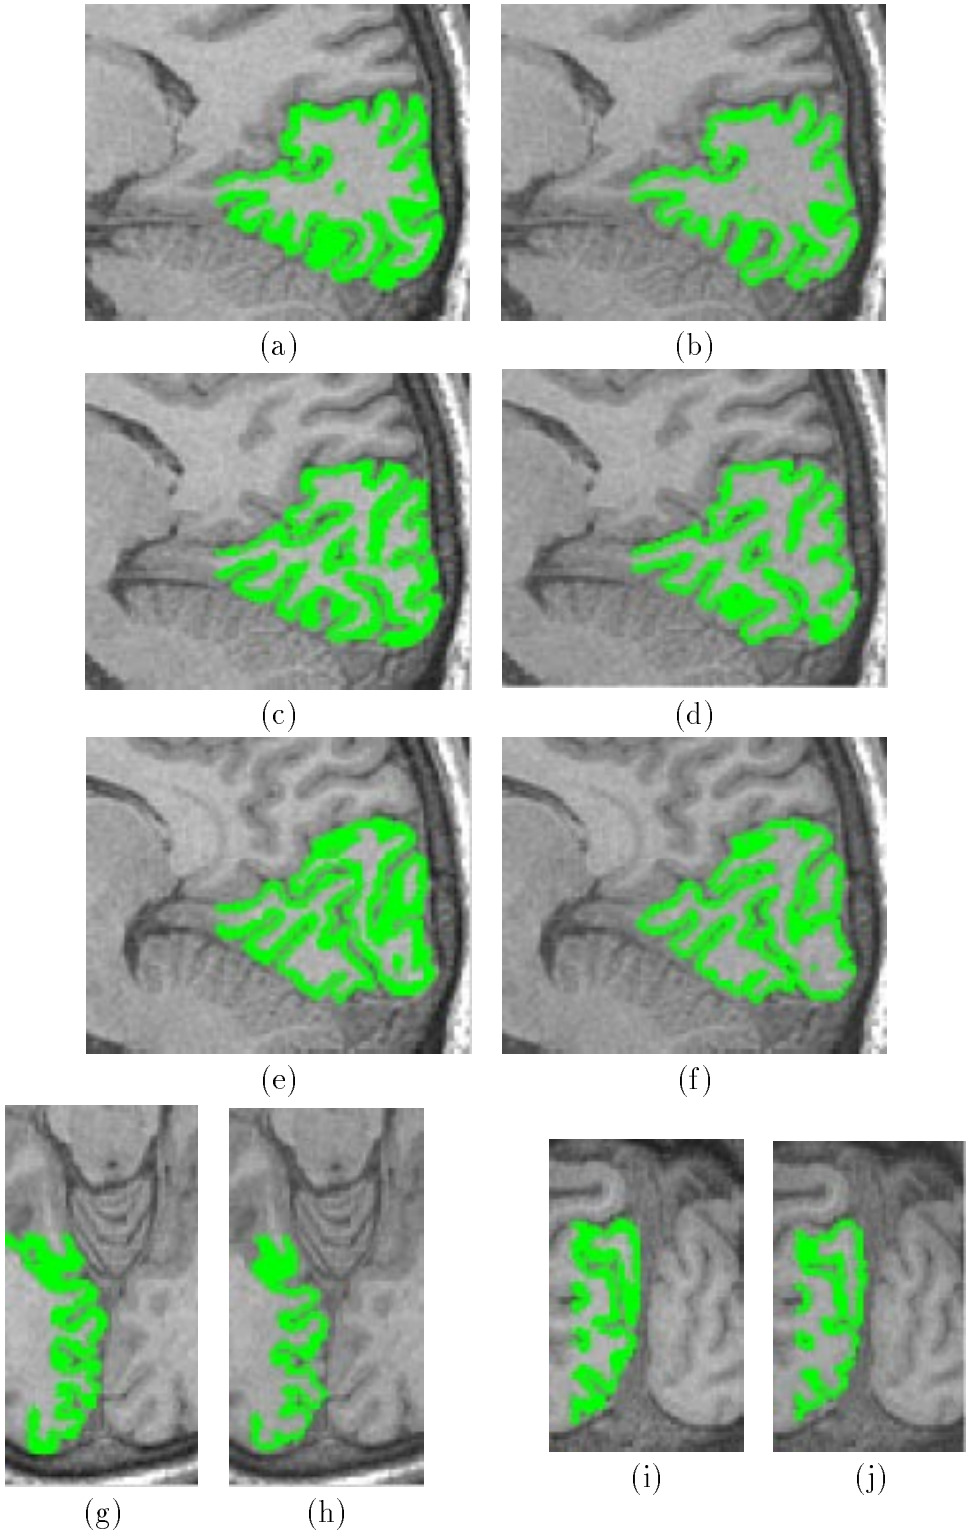

Figure 8: The left column of images show manual gray matter segmentation results; the right column of images show the automatically computed gray matter segmentation. The first three rows show sagittal slices of the occipital lobe. The figures in the fourth row are axial and coronal slices of the same region of cortex.

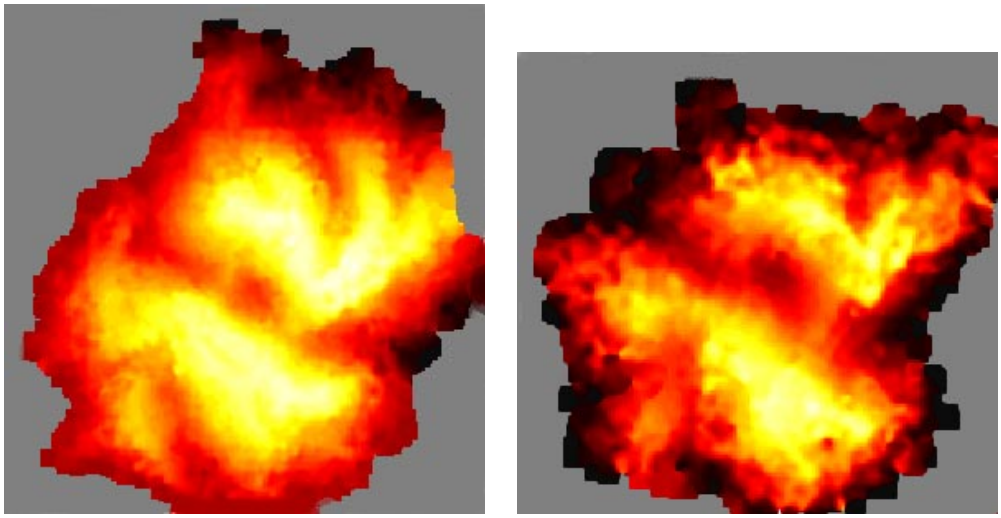

(a) (b)

Figure 9: The left and right images show flattened representations of a portion of the same occipital lobe. On the left is the flattened representation computed from a manual segmentation of the gray matter; on the right is the flattened representation computed using the automatic segmentation technique proposed in this paper. The different intensities represent different distances in 3D from a fiducial plane; in this case, it is the distance from the leftmost sagittal plane. Brighter regions indicate larger distances while darker regions indicate shorter distances.

of the various stages in our method and speculate on feasible alternatives and possible extensions.

Among the many proposed algorithms for gray matter segmentation, the one proposed by Joliot and Mazoyer [17] is conceptually the most similar to ours. In common with our approach, these authors favor white matter segmentation as a preliminary step. Also, gray matter is defined from the boundary of the white matter segmentation. There are also several differences between the two approaches. First, the white matter segmentation used here incorporates structural considerations through the novel application of anisotropic smoothing on the posterior probabilities. Second, anatomical consistency of the gray matter segmentation is enforced by imposing topological constraints on the white matter. Third, connectivity information between segmented gray matter voxels is computed. These last two computations are essential for the visualization of cortical activity from fMRI measurements.

The motivation for using white matter segmentation to guide gray matter segmentation is quite straightforward. In the cortex, the volume of white matter exceeds that of gray matter. Gray matter, on the other hand, is confined to a thin region on the boundary of the cortex. In cortical MR data, this implies that regions of gray matter voxels may be as narrow as one or two pixels. Since gray matter voxels border either on white matter or on CSF, such

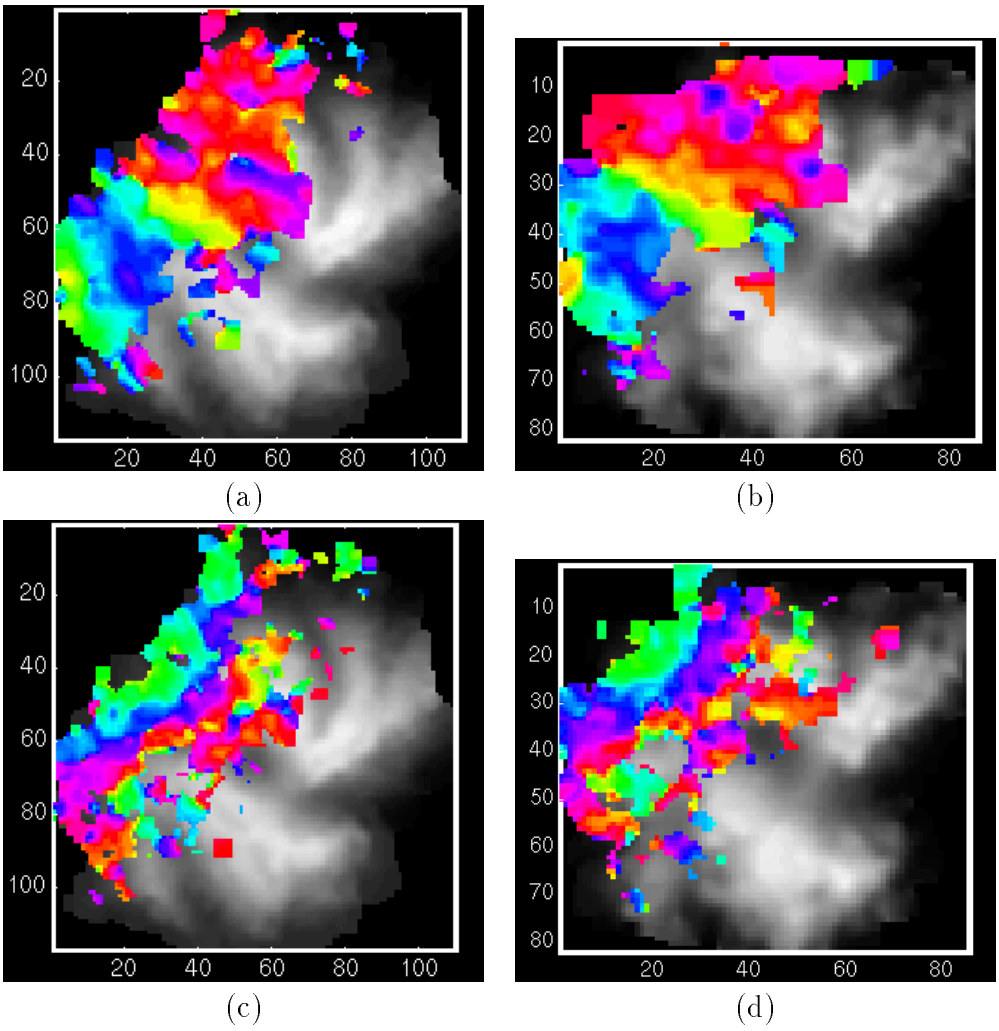

Figure 10: The images show flattened representations of a portion of the same occipital lobe. In the left column are flattened representations computed from a manual segmentation of the gray matter; in the right column are flattened representations computed using the automatic segmentation technique proposed in this paper. Overlaid on each of the flattened representations are phase measurements computed from fMRI data (displayed with different colors). These phase measurements represent the temporal phase of neural activity relative to a temporally periodic visual stimulus. The top and bottom rows show results obtained using different stimuli.

narrow regions suffer from partial volume effects, which if not explicit modeled will reduce the accuracy of the segmentation. White matter segmentation is much easier because the large regions of white matter voxels facilitate the effective use of spatial (structural) constraints to improve segmentation. Since the signal to noise ratio (SNR) of the MR data is relatively high, only fairly local spatial constraints are required. If the SNR is much lower, efficient algorithms that promote more global spatial constraints would probably be required [4].

Using anisotropic diffusion on the posterior probabilities to capture local spatial constraints was motivated by the intuition that posteriors with piecewise uniform regions result in segmentations with piecewise uniform regions. There are several reasons for applying anisotropic smoothing on the posterior probabilities instead of directly on the MR data. First, applying anisotropic smoothing on the MR data does not take into consideration that there are only three classes being segmented. It merely produces a result that is smoother while preserving the main discontinuities in the original image. Second, and in a more general context, anisotropic diffusion applied to the raw data is only applicable when the noise is additive and class independent. For example, if the class means were identical and the classes differed only in their variances, anisotropic smoothing of the raw data would not be effective. On the other hand, applying anisotropic smoothing on the posterior probabilities is still feasible even when the classes are described by general probability distribution functions. This novel technique is related to (anisotropic) relaxation labeling [11, 15, 23], which is further discussed in [26].

The current implementation employs a 3D extension of Perona and Malik's anisotropic smoothing technique, though other anisotropic smoothing techniques, when applied to the posterior, are likely to be equally, if not more, effective [1, 14]. The method proposed by Saint-Marc et. al. [24] is also computationally efficient and a viable alternative.

While the MR data that is currently being acquired and segmented is scalar-valued, the proposed segmentation method could be readily adapted to segment multispectral vector-valued measurements [12]. Since the anisotropic smoothing is performed on the posterior probabilities, scalar anisotropic smoothing techniques, like those mentioned above, can still be used. The only modification would be in the calculation of class likelihoods.

The sole purpose of segmenting white matter first is to detect the white-gray matter boundary accurately. While the current segmentation technique is often sufficiently accurate, more accurate results could be obtained by using deformable membranes initialized with the current segmentation. Likewise, deformable membrane techniques could also be used to determine the gray matter-CSF boundary. The deformation should, similarly, be subject to

the anatomical constraint that the distance to the white-gray matter boundary be no greater than about 3 mm.

Finally, we should note that although the main motivation of the work described in this paper is visualization of fMRI, a number of techniques and concepts here introduced can be useful for other image processing applications in general and medical imaging in particular. In [26] we further study the posterior anisotropic smoothing scheme and use it for other image segmentation tasks. It will be interesting also to test the segmentation ideas here described in problems such as 3D registration [2].

Acknowledgments

We thank Heidi Baseler, Geoff Boynton, Steve Engel, Jon Demb, David Heeger, Hagit Hel-Or, and Tom Malzbender for interesting discussions during the progress of this work. Steve Engel and Hagit Hel-Or provided initial brain unfolding and segmentation software. PT was partially supported by the Hewlett-Packard Labs Grassroots Basic Research Program, and NEI grant ROI EY03164 to BW, while performing this work.

References

[1] L Alvarez, P Lions, and J Morel. Image selective smoothing and edge detection by nonlinear diffusion. *SIAM J. Numerical Analysis*, 29:845–866, 1992.

[2] N Ayache. Medical computer vision, virtual reality and robotics. *Journal of Image and Vision Computing*, 1995.

[3] T Bartlett, M Vannier, D McKeel Jr, M Gado, C Hildebolt, and R Walkup. Interactive segmentation of cerebral gray matter, white matter, and CSF: Photographic and MR images. *Computerized Medical Imaging and Graphics*, 18(6):449–460, 1994.

[4] C Bouman and M Shapiro. A multiscale random field model for Bayesian image segmentation. *IEEE Transactions on Image Processing*, 3(2):162–177, 1994.

[5] V Caselles, R Kimmel, G Sapiro, and C. Sbert. Three dimensional object modeling via minimal surfaces. In *Proc. European Conf. Comp. Vision*, pages 97–106, Cambridge, UK, 1996.

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

[7] T Cormen, C Leiserson, and R Rivest. *Introduction to Algorithms.* MIT Press, Cambridge, MA, 1990.

[8] A Dale and M Sereno. Improved localization of cortical activity by combining EEG and MEG with MRI cortical surface reconstruction: a linear approach. *Journal of Cognitive Neuroscience*, 5(2):162–176, 1993.

[9] H Drury, D Van Essen, C Anderson, C Lee, T Coogan, and J Lewis. Computerized mappings of the cerebral cortex: a multiresolution flattening method and a surface-based coordinate system. *Journal of Cognitive Neuroscience*, 8(1):1–28, 1996.

[10] S Engel, G Glover, and B Wandell. Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *To appear in Cerebral Cortex*, 1997.

[11] O Faugeras and M Berthod. Improving consistency and reducing ambiguity in stochastic labeling. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 3(4):412–424, 1981.

[12] L Fletcher, J Barsotti, and J Hornak. A multispectral analysis of brain tissues. *Magnetic Resonance in Medicine*, 29(5):623–630, 1993.

[13] J Foley, A van Dam, S Feiner, and J Hughes. *Computer Graphics: Principles and Practice.* Addison-Wesley Publishing Company, Inc., Reading, MA, 1990.

[14] G Gerig, O Kubler, R Kikinis, and F Jolesz. Nonlinear anisotropic filtering of MRI data. *IEEE Trans. on Medical Imaging*, 11(2):221–232, 1992.

[15] R A Hummel and S W Zucker. On the foundations of relaxation labeling processes. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 5(3):267–287, 1983.

[16] B Johnston, M Atkins, and K Booth. Three-dimensional partial volume segmentation of multispectral magnetic resonance images using stochastic relaxation. In *Proc. SPIE vol. 2180 Nonlinear Image Processing V*, pages 268–279, 1994.

[17] M Joliot and B Mazoyer. Three-dimensional segmentation and interpolation of magnetic resonance brain images. *IEEE Trans. on Medical Imaging*, 12(2):269–277, 1993.

[18] S Kichenassamy, A Kumar, P Olver, A Tannenbaum, and A Yezzi. Gradient flows and geometric active contour models. In *Proc. Int. Conf. Comp. Vision*, pages 810–815, Cambridge, MA, 1995.

[19] T Kong and A Rosenfeld. Digital topology: introduction and survey. *Computer Vision, Graphics, and Image Processing*, 48:357–393, 1989.

[20] T Lee and R Kashyap. Building skeleton models via 3-D medial surface/axis thining algorithms. *CVGIP: Graphical Models and Image Processing*, 56(6):462–478, 1994.

[21] R Malladi, J A Sethian, and B C Vemuri. Shape modeling with front propagation: A level set approach. *IEEE Trans. Pattern Anal. Machine Intell.*, 17:158–175, 1995.

[22] P Perona, T Shiota, and J Malik. Anisotropic diffusion. In B Romeny, editor, *Geometry-driven diffusion in computer vision*, pages 73–92. Kluwer Academic Publishers, 1994.

[23] A Rosenfeld, R A Hummel, and S W Zucker. Scene labeling by relaxation operations. *IEEE Trans. on Systems, Man, and Cybernetics*, 6(6):420–453, 1976.

[24] P Saint-Marc, J Chen, and G Medioni. Adaptive smoothing: a general tool for early vision. In *Proc. Conf. on Computer Vision and Pattern Recognition*, pages 618–624, 1989.

[25] G Sapiro, R Kimmel, and V Caselles. Object detection and measurements in medical images via geodesic deformable contours. In *Proc. SPIE vol. 2573 Vision Geometry IV*, pages 366–378, San Diego, CA, 1995.

[26] P C Teo, G Sapiro, and B Wandell. Posterior anisotropic diffusion. *In preparation*, 1997.

[27] D Terzopoulos, A Witkin, and M. Kass. Constraints on deformable models: Recovering 3d shape and nonrigid motions. *Artificial Intelligence*, 36:91–123, 1988.

[28] D Vandermeulen, R Verbeeck, L Berben, D Delaere, P Suetens, and G Marchal. Continuous voxel classification by stochastic relaxation: theory and application to mr imaging and mr angiography. *Image and Vision Computing*, 12(9):559–572, 1994.

[29] B Vemuri, S Rahman, and J Li. Multiresolution adaptive K-means algorithm for segmentation of brain MRI. In *Proc. Int'l Computer Science Conf. on Image Analysis and Computer Graphics*, pages 347–354, Hong Kong, 1995.

[30] K Vincken, A Koster, and M Viergever. Probabilistic hyperstack segmentation of MR brain data. In *Proc. First Int'l Conf. on Computer Vision, Virtual Reality and Robotics in Medicine*, pages 351–357, Nice, France, 1995.

[31] B Wandell. *Foundations of Vision.* Sinauer Associates, Inc., Sunderland, MA, 1995.

[32] B Wandell, S Engel, and H Hel-Or. Creating images of the flattened cortical sheet. *Invest. Opth. and Vis. Sci.*, 36(S612), 1996.

[33] W M Wells, W Grimson, R Kikinis, and F Jolesz. Adaptive segmentation of MRI data. *IEEE Trans. on Medical Imaging*, 15(4):429–442, 1996.

[34] J Xuan, Tülay Adali, and Y Wang. Segmentation of magnetic resonance brain image: integrating region growing and edge detection. In *Proc. Int'l Conf. on Image Processing*, pages 544–547, Washington, DC, 1995.