

IEEE Trans. Medical Imaging, to appear

Creating Connected Representations of Cortical Gray Matter for Functional MRI Visualization

Patrick C. Teo

Computer Science
Stanford University
Stanford, CA 94305
teo@white.stanford.edu

Guillermo Sapiro

Electrical and Computer Eng.
University of Minnesota
Minneapolis, MN 55455
guille@ece.umn.edu

Brian A. Wandell

Psychology and Neuroscience
Stanford University
Stanford, CA 94305
wandell@stanford.edu

Abstract

We describe a system that is being used to segment gray matter from MRI and to create connected cortical representations for functional MRI visualization. The method exploits knowledge of the anatomy of the cortex and incorporates structural constraints into the segmentation. First, the white matter and CSF regions in the MR volume are segmented using some novel techniques of posterior anisotropic diffusion. Then, the user selects the cortical white matter component of interest, and its structure is verified by checking for cavities and handles. After this, a connected representation of the gray matter is created by a constrained growing-out from the white matter boundary. Because the connectivity is computed, the segmentation can be used as input to several methods of visualizing the spatial pattern of cortical activity within gray matter. In our case, the connected representation of gray matter is used to create a flattened representation of the cortex. Then, fMRI measurements are overlaid on the flattened representation, yielding a representation of the volumetric data within a single image. The software is freely available to the research community.

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

1 Introduction

Magnetic resonance scanners are 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). The development of 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 in addition to structure, is an important breakthrough, providing us with a new opportunity to study the activity of single human brains at relatively high spatial resolution [15, 14, 36, 12, 28].

The human brain is composed mainly of two types of tissue: gray matter and white matter. Gray matter forms the outer layer (the cortex), encasing the inner white matter almost completely. Gray matter tissue contains a high density of heavily interconnected neurons (approximately $10^5/\text{mm}^3$ [40]). The activity of these neurons is the computational basis of sensation, thought and action. White matter is comprised 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 indirectly measures the activity of neurons in the gray matter.

Among the various parts of the brain, the cerebral cortex is the most prominent, and one of the most intensely studied. The cortex, is divided into two hemispheres that are connected by many nerve fibers (making up cortical white matter). Despite its complex outward appearance, the structure of the gray matter in each hemisphere is quite straightforward and consistent across human brains. Cortical gray matter is highly convoluted, and its topology is that of two crumpled sheets having no holes or self-intersections.

There are various ways to visualize gray matter. One approach is to create a gray matter surface model. In this approach, some form of segmentation is applied, usually a simple classification of gray matter from the surrounding cerebral spinal fluid (CSF, the fluid that fills the cranial cavity), though more elaborate labeling of anatomical structures may also be applied. The gray matter is rendered in three dimensions as a surface, and the user sees mainly those portions on the exterior surface.

Much of cortical gray matter, however, is buried deep within the folds of the brain, called sulci. Visualizing the neural activity recorded by fMRI within these sulci requires novel

¹This 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.

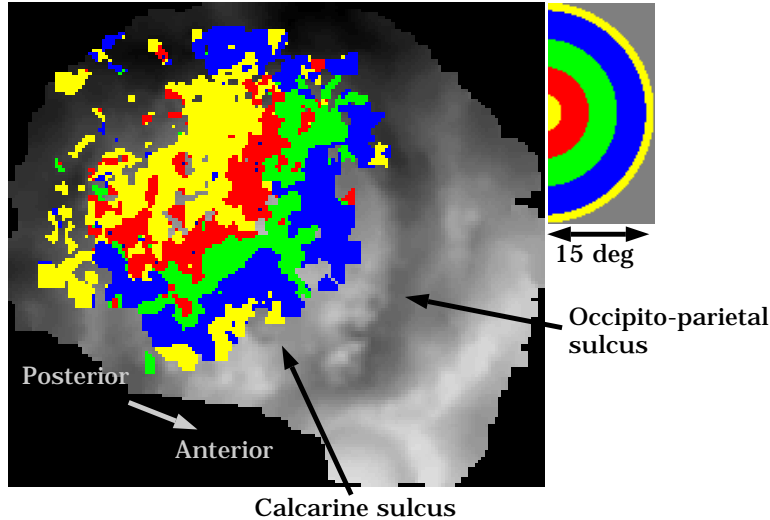


Figure 1: Functional MRI measurements are shown as a (color) overlay upon a (grayscale) representation of the flattened occipital lobe. The color overlay indicates the visual field eccentricity of a stimulus that evoked a response at that cortical position. The visual field eccentricity represented by each color is shown by the legend on the upper right. Increasingly peripheral stimuli cause responses at increasingly anterior cortical locations; also, note the large representation of the foveal region compared to the peripheral. The grayscale value represents the cortical position along the medial (bright) to lateral (dark) direction prior to flattening. Human V1 is located within the calcarine sulcus, which ultimately merges with the occipito-parietal sulcus. While most of calcarine is obscured by the color activity map, its position is indicated in the figure. The grayscale image extends beyond the color overlay because the portion of flattened cortex extends beyond the portion of brain measured during the fMRI experiments. The fMRI data measure one aspect of the retinotopic organization within human visual cortex. The experimental methods are described in [14].

visualization techniques. An increasingly popular way of visualizing such mappings is to superimpose fMRI measurements on flattened representations of the cortical surface [10, 13, 14, 41]. One method of creating a 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 (geodesic) distances within the gray matter.

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 [40]. The figure illustrates how the spatial structure of these traveling waves, represented on the flattened cortical surface, can be used to determine the retinotopic organization within the occipital lobe of individual observers.

Our specific need for a gray matter representation is to create a flat map for representing activity measured using fMRI. To create a flat map, like the one shown in Figure 1, one must be able to measure distances within the segmented gray matter; hence, one must identify both the gray matter and its topological connectivity. For this application, the most important portion of the connected gray matter representation is the first layer that falls along the gray/white matter boundary. With a connected representation of this boundary, one can measure distances, curvature and other important surface features. The flattened representation complements conventional methods of viewing volume data as a series of separate images. Using the flattened representation, one can appreciate the activity across a large region of cortex within a single image.

One difficulty in achieving a topologically connected gray matter segmentation is that MR intensity levels within gray matter significantly overlap with the levels from both white matter and non-brain matter. With the current spatial resolution of MRI, regions of gray matter voxels can be as narrow as one or two voxels so that a large percentage of the gray matter suffers from *partial volume* effects.² Partial voluming limits the effectiveness of intensity-based gray matter segmentation algorithms.

A second difficulty is determining gray matter topological connectivity from gray matter segmentation alone. Gray matter voxels on opposite sides of a sulcus may be adjacent to one another on the sampling grid. Yet, the gray matter voxels on opposite sides of the sulcus are not connected. Hence, gray matter connectivity cannot be discerned from the segmentation alone; information about the nearby white matter is necessary to determine connectivity.

Various gray-matter segmentation techniques have been proposed. Many techniques use image segmentation methods that do not incorporate knowledge of basic features of cortical anatomy [42, 38, 43, 39, 8, 2]. Most importantly, it is difficult to compute accurate gray matter connectivity from these segmentations. There are techniques that do use some information about cortical anatomy, the manner in which such knowledge is employed tends to be local and statistical [21, 37]. The methods that are closest to ours are those proposed by Joliot and Mazoyer [22], Mangin *et al.* [27], and Dale and Sereno [10]. We discuss the relationship between this work and ours in Section 4.

Methods using deformable surfaces (so called snakes or balloons) can produce connected segmentations [6, 7, 10, 23, 26, 33, 35] (see [27] for a detailed discussion of major problems with this technique when used for this task). These methods have several useful features.

²Partial 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.

They incorporate smoothness as part of their segmentation criterion; they are capable of producing sub-pixel classification (of the boundary between white and gray matter, for example); when surfaces are initialized to be topologically equivalent to a sheet, they will be consistent with the topology of the gray/white matter boundary, having no holes or self-intersections. These methods also have a severe problem: the minimization process used to deform the surface is prone to local minima. This frequently occurs near deep sulci with narrow openings that are present in the occipital lobe of human cortex (see [5], Fig. 7, for an illustration of this problem). Hence, deformable surface algorithms must begin with a very good initialization. The segmentation method proposed in this paper can be used to initialize such algorithms [7, 27].

To evaluate the quality of the system described herein, we visually compared the system output with manual segmentations produced by trained users. The comparison included both the gray matter segmentation and the appearance of the fMRI data on the flattened representation. We describe these comparisons in Section 3.

2 Method

The segmentation method is comprised of four steps that we will explain in this section. First, the white matter and CSF regions in the MR volume are segmented.³ Second, the user selects the desired cortical white matter component. Third, the white matter structure is verified by checking for cavities and handles. Fourth, a connected representation of the gray matter is created by growing out from the white matter boundary. The gray-matter growing is subject to two main constraints: (a) new gray-matter cannot grow into voxels that have been already classified as CSF, white matter or gray matter, (b) connectivity of the segmented gray matter must be maintained during the growing process.

2.1 Segmentation of White Matter and CSF

In the first stage, voxels containing white matter tissue and “non-brain” material, principally CSF, are segmented. We begin by segmenting white matter for two main reasons. First, with the T1 weighted scans that we use the intensity levels of white matter have smaller variability than gray matter. Second, beginning with white matter simplifies the computation

³The images are obtained from a T1 weighted gradient echo volumetric acquisition system with TE set to the minimum full, TR set to 33 msec, NEX set to 1, and with 40 degrees flip angle.

of connectivity in gray matter.

At this step, we create three classes: white matter, CSF (non-brain), and unknown. The unknown class contains mainly gray matter, but the segmentation of this class is unreliable and it contains no connectivity information. Hence, this class will not be used as the basis for deriving the connected gray matter representation.

In the first stage of classification, the voxel intensities within each class are modeled as independent random variables with normal distributions. Thus, the likelihood of a particular voxel, V_i , belonging to a certain class, C_i , 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 the index c stands for one of the classes {white, unknown, CSF}. V_i and C_i correspond to the intensity and classification of voxel i respectively. To establish the classification, the user adjusts the parameters μ_c and σ_c in real time to obtain a visually satisfactory segmentation, as judged by examining the segmentation in a few anatomical slices. The values of these parameters typically remain unchanged across different MR data sets collected using the same pulse sequence.

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

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

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, if white matter voxels occur more frequently than gray matter voxels, the prior probability of white matter is larger than that of gray matter. The exact prior depends on the part of the cortex being segmented and can be set in advance by the user. We have found that the segmentation results are robust to variations in the value of the priors.

In the second step, the posterior volumes are smoothed anisotropically in three dimensions, but preserving discontinuities. Figure 2 shows an example of a posterior derived from a homogeneous prior and its smoothed counterpart. The anisotropic smoothing technique




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.

applied is a 3D extension of the original 2D version proposed by Perona *et al.* [29]. This step involves simulating a discretization of the following partial differential equation for a small number of iterations: ⁴

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

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

⁴For the examples in this paper we use the maximal time step that ensures stability of the numerical implementation of this type of equation [29], and normally run 5 iterations. The parameter $\eta_c = 0.5$ for the three classes. Using the recently developed techniques in [3], it should be possible to determine η_c automatically from the *median absolute deviation* [31] of each class, and to run the equation until its steady-state.

function $g(\cdot)$ controls the local amount of diffusion such that diffusion across discontinuities in the volume is suppressed. The reason for applying anisotropic smoothing to the posterior probabilities, rather than to the MR data, is deferred to the discussion section.

Finally, the white matter and non-brain classifications are obtained using the maximum a posteriori probability (MAP) estimate after anisotropic diffusion. That is,

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

where $\Pr^*(C_i = c | V_i = v)$ corresponds to the posterior following anisotropic diffusion. The upper panels of Figure 2 shows the MAP classification prior to anisotropic diffusion. The lower panels of that figure show the segmentation after applying anisotropic diffusion.⁵ The white matter and CSF classification in the lower panels are smooth and connected. The unknown class classification does not correspond to a plausible description of the gray matter. For this reason, we retain only the white matter and CSF segmentations. The accuracy of this white matter and CSF segmentation will be evaluated later when we compare the overall results of our algorithm with manual segmentation.

2.2 Selection of Cortical White Matter

Gray matter segmentation can be obtained from white matter segmentation because gray matter surrounds white matter. Thus, it is important to ensure that we have obtained an accurate white matter segmentation. To verify the quality of the white matter segmentation, we must investigate some of the topological properties of the selected white matter. Because the data are represented on a grid, before beginning the investigation we must decide upon a digital topology to define connectedness for each class.

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

⁵With our measurement protocol, based on a high quality head coil, it was unnecessary to correct for variations in the mean gray matter intensity levels as described in [42].




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

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) [24].

The initial classification generally yields several unconnected components labeled as white matter. Only one of these is the main section of white matter, the others being either parts of the cerebellum, or other non-brain materials. The user identifies a voxel in the cortical white matter component via a graphical user interface, and a flood-filling algorithm automatically identifies the entire connected component [18]. 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 such as skin or 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

Gray matter is a single sheet that encases white matter. To ensure that there are not multiple gray matter sheets or self-intersections of the gray matter, we must eliminate gray matter grown within cavities or through white matter handles. Cavities are non-white matter regions that are completely surrounded by white matter (for example, the inside of a tennis ball). Handles are non-white matter regions that are partially surrounded by white matter (for example, the middle of a doughnut). Figure 3 shows an example of a white matter cavity and a white matter handle.

In our application, handles may arise when the classification inappropriately assigns the white matter label to voxels that cross a sulcus, cutting through two layers of gray matter and CSF. We check for this condition in the previewer. Because the gray matter is several millimeters thick, we rarely encounter handles. Handles are removed by hand-editing or re-adjusting the parameters used to obtain the white matter classification.

It is possible to automatically compute *the number* of handles using the Euler characteristic, χ , which is equal to the sum of the number of connected components and cavities, minus the number of handles. The first two quantities can be computed using flood-fill algorithms. The Euler characteristic can be computed as the sum of the local Euler characteristic over all $2 \times 2 \times 2$ voxel neighborhoods: ⁶

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

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 [25]. Thus, the number of handles can be computed as the sum of the number of connected components and cavities minus the Euler characteristic. The computation is not implemented in our distribution software at this stage because the computation (a) does not yield the *locations* of the handles, and (b) is rarely needed.

A second type of segmentation error is the presence of cavities within the white matter. The problem created by cavities is that gray matter will grow on the boundary of the cavity and form a surface internal to the white matter. It is possible to eliminate white matter cavities in a number of simple ways, for example by repeatedly initiating the flood-filling algorithm from non-white matter voxels on the volume boundary. All non-white matter connected components that are *not* filled must be encased entirely by white matter and thus be cavities.

In practice, we have found it simpler to grow gray matter from all voxels on the white matter surface, potentially creating gray matter components within cavities. Then, the user selects a voxel from the gray matter component that surrounds the white matter. The program identifies all gray matter voxels connected to the selected voxel. All unconnected gray matter voxels are deleted; this removes unconnected gray matter components caused by cavities.

⁶There 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, and thus, the number of handles, is then computed efficiently using table lookups.




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.

2.4 Gray Matter Segmentation and Connectivity

In this section we describe how gray matter voxels are grown from the boundary of the white matter. The gray matter voxels are identified by growing a sequence of layers that begin on the white matter boundary. The maximum number of gray matter layers is a parameter of the program that is set by the user, and is basically determined by the spatial resolution of the MR and the area of interest in the cortex. For example, suppose the MR data has a spatial resolution of 1 mm along each spatial dimension and we are identifying gray matter near calcarine cortex in the occipital pole where the gray matter is roughly 5 mm thick. Then a maximum of 5 layers are grown. There may be fewer than 5 layers at any particular location if CSF is encountered before the 5 mm limit, or if gray matter from the opposite side of a sulcus is encountered. Thus, the thickness of the final classification depends on (a) the




Figure 5: Representation of three dimensional gray and white matter connectivity. Dark shaded cubes represent gray matter voxels; light shaded cubes represent 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.

maximum thickness of gray matter, (b) the CSF classification, and (c) potential collisions with gray matter growing from different portions of the white matter.

Figure 4 shows an example of gray matter classification. A simple case, in which only two layers are grown from the boundary of the white matter component, is illustrated.

For our application, it is very important to compute the connectivity of the gray matter voxels. Each layer of gray matter grows upon the previous layer (or from the boundary of the white matter component for the first layer) in the same fashion. Adding a 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.

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 classified as gray *only if all its parents*

are connected. Connectivity in this case is determined from the 26-neighbor adjacency of white matter voxel parents. For layer N+1, unclassified voxels that are 6-neighbor adjacent to gray matter in layer N are classified as gray matter voxels belonging to layer N+1. The connected gray matter voxels in layer N are known as *parents* of the voxels in layer N+1. Again, a new voxel is classified as gray only if *all its parents are connected.* The reason for requiring connectivity amongst the parents is this: a voxel must not be assigned a gray classification if doing so results in a contention among unconnected voxels in the previous layer. For example, a voxel falling between two gray matter voxels on opposite sides of a sulcus will not be classified as gray matter.

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. This occurs precisely when one gray matter voxel is a parent of the other. 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) have a common parent or have parents that are connected (as computed in the previous connectivity step). Moreover, 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 results of this segmentation processing are contained in two files. One file represents results of the classification step, including the labels of white matter, CSF, and unknown voxels in the MR volume. The classification file contains the results of the automatic segmentation and the hand-edits. The second file contains the connected representation of the gray matter that is derived from the classification file. The connected representation consists of nodes representing the three dimensional coordinates of the gray matter voxels and edges




Figure 6: Representation of distance within the gray matter measured using different metrics. The intensity within the shaded contours in each image represents the distances from the same selected gray matter voxel (at bottom-center). Bright points represent shorter distances and darker areas denote longer distances. Figure (a) plots the Euclidean distances; 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.

representing the connections between gray matter voxels. Subsequent application software measures distances (geodesics) between pairs of gray matter voxels using the shortest paths between pairs of vertices [9].

Figure 6 shows the distance between a collection of gray matter voxels and a single selected gray matter voxel measured in several ways. Figure 6a represents the 3D Euclidean distance between the voxels. Figure 6b represents the shortest distance along the gray matter segmentation within the selected slice. Figure 6c shows the distance measured along the shortest path within the full three dimensional gray matter segmentation. For each voxel in the slice, the distance from the selected gray matter voxel is shortest when the Euclidean distance is used. The distance measured using the full connected representation is no greater than the distance measured within the selected slice. For applications involving flattening cortex, distance within the full three dimensional connected representation is the appropriate measure. To measure these distances, it is essential then to obtain the topological connectivity between gray matter voxels.

3 Results

The segmentation technique described in this paper has been implemented and is being used to identify 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. Prior to the development of the method described here, gray matter was identified manually. Identifying gray matter in a single occipital lobe of one hemisphere, using rudimentary segmentation tools, required about 18 hours for an experienced person. Much of the time was spent visually inspecting connectivity and ensuring topological correctness. This was because directly segmenting gray matter often produces self-intersections in the gray matter. With the present method, the entire procedure takes about half an hour. The time spent in 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 in a simple windowing system that permits the user to select volume regions of interest, apply the methods described in this paper, verify and edit the automatic segmentation. Figure 7 shows an example of the windows in the system.

How might we evaluate the quality of the methods? One measure is the utility of the method in laboratories that use it as an application tool. Apart from our lab, where the method is in use every day, the tool has already been used to create published material by other groups, e.g., [11], and was made public to the research community for testing. A second measure is to compare the gray matter segmentation obtained from the algorithm with post mortem material in which gray matter can be identified more certainly. We have initiated such a project, but the results will not be available for several years. A third measure is to compare this method with those published by other groups with the same goal. We have made efforts to obtain software from other groups, but we have not yet succeeded. What we can do at this stage then is to compare the method described here with the results from manual segmentation obtained by trained users.

Figure 8 shows several comparisons between gray matter segmentation results derived manually and those computed using the current method. In this, and in all other comparisons shown in this paper, no manual editing of the automatic segmentation was carried out. The automatic segmentation results are (qualitatively) similar to those obtained manually despite the large number of deep and narrow folds in this region of the cortex. The current method has difficulties when the white matter is extremely thin, roughly one voxel thick. In this case, the anisotropic smoothing algorithm tends to remove the narrow regions of white matter in

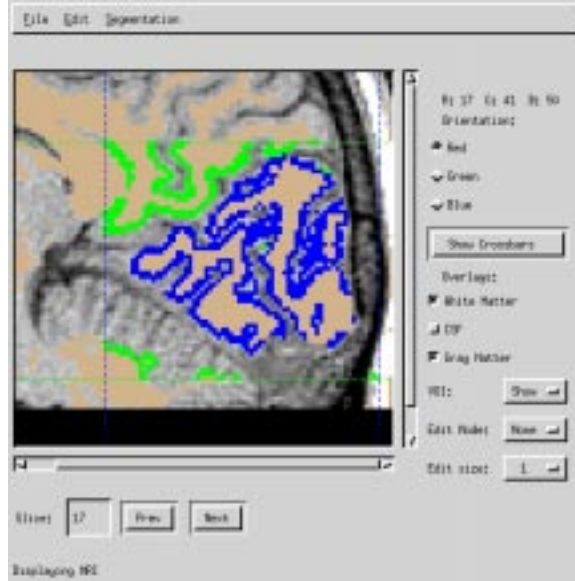


Figure 7: Example of the windows based interactive system used to implement the algorithms described in this paper. The user interactively performs a large number of operations. Through this software automatic segmentation and unfolding are combined with manual corrections.

favor of larger regions of CSF.

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 [41] 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, the two 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 [14]: A rotating wedge (top) and an expanding ring (bottom). The color diagrams are as in Figure 1 for the wedge and similar to it for the ring (green represents the inner ring, moving towards purple on the outside). 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 [14]. 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 retinotopically organized visual areas.

4 Discussion

In this section, we first discuss the segmentation methods proposed in [10, 22, 27]. These methods are closely related to the method introduced here. Then, we review the main decisions made at each stage of our method. Finally, we speculate on alternatives and extensions.

4.1 Related papers

The algorithm for gray matter segmentation proposed by Joliot and Mazoyer [22] shares several common features with the method described here. These authors favor white matter segmentation as a preliminary step, and gray matter is defined from the boundary of the white matter segmentation. There are two main differences. First, the white matter segmentation process we use is based on a novel application of anisotropic smoothing on the posterior probabilities. Second, we compute connectivity relationships of the segmented gray matter voxels. This connectivity is essential for the visualization of cortical activity from fMRI measurements.

Dale and Sereno [10] also begin by classifying white matter, but they do not specify their white matter segmentation methods in enough detail for us to comment upon. Rather than growing gray matter from the white matter boundary, they use the boundary to locate a deformable surface that they then flatten (see [27] for a critique of using deformable models for segmentation of MRI). Dale and Sereno do not segment gray-matter. The white matter segmentation we obtain could also be used to initialize the shape of a deformable surface (see e.g., [6, 7]).




Figure 8: Comparison of manual and automated segmentation procedures. Images are presented in pairs: 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 two pairs in the fourth row are axial and coronal slices of the same region of cortex.




Figure 9: Qualitative test of the automatic segmentation via 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 intensity represents the position in 3D along the medial to lateral dimension. Brighter regions indicate medial positions, while darker regions indicate lateral positions.

The segmentation method proposed by Mangin *et al.* [27] is similar to our method in that white matter segmentation precedes gray matter segmentation. But, the two methods address the difficulties inherent in cortical segmentation quite differently. First, Mangin *et al.* compute white matter segmentation using discrete mathematical morphology. This contrasts with our method of using continuous anisotropic smoothing, which is an approximation to a Markov Random Field formulation with an additional discontinuity field [34]. Second, Mangin *et al.* group the gray matter and CSF together while we segment the data into three groups (gray matter, white matter, and CSF). Because the connectivity relationship of the gray matter is essential to us, their method does not solve our main application problem.

Mangin *et al.* make an important contribution by introducing homotopy constraints that prevent self intersections in the deformable surface. The topology of the segmented gray/white matter boundary is ensured by dilating *inwards* a deformable region which is initialized to a bounding box containing the initial white matter segmentation. The growing process developed here is also designed to prevent self-intersections. Our growing process dilates outwards to identify the gray matter.




Figure 10: Comparison of functional activity on flattened representations created manually (left) and created using the automated method (right). The color overlay represents the temporal phase of the fMRI signal. Images in the top row show the temporal phase of the signal in response to a rotating wedge; images in the bottom row show results obtained using an expanding ring stimulus.

4.2 Design decisions

An important feature of our method is that gray matter is calculated from the white matter segmentation. Because cortical gray matter voxels border either on white matter or on CSF, MR signals from the gray matter often suffer from partial volume effects so that segmentation based on the level of the MR gray matter signal is poor. White matter segmentation does not suffer from these problems to the same extent. The signal to noise ratio (SNR) of the MR data is adequate to achieve good white matter segmentation using only local spatial constraints. Had the SNR been much lower, algorithms that promote global (hierarchical) spatial constraints would probably be required [4].

A second reason for growing gray matter from white matter segmentation is that it simplifies the computation of gray matter connectivity, a main goal for our application. If one begins with a gray matter segmentation, making decisions about the connectivity of gray matter voxels on opposite sides of a sulcus is very difficult, perhaps impossible. By growing gray matter from white matter, we can keep track of which sulcal wall each gray matter voxel is on and thus develop the proper connectivity relationships.

There are several reasons for applying anisotropic smoothing to the posterior probabilities instead of directly on the MR data. First, anisotropic smoothing applied directly to the MR data would not take into consideration that only three classes are being segmented. Second, anisotropic diffusion applied to the raw data is well-motivated only when the noise is additive and class independent. For example, if two classes have the same mean and differ only in variance, anisotropic smoothing of the raw data is ineffective. 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. Applying anisotropic smoothing on the posterior probabilities is feasible even if the classes are described by general probability distribution functions. This novel application is related to (anisotropic) relaxation labeling [16, 20, 30], and is further discussed in [34].

Our implementation of anisotropic diffusion is a 3D extension of Perona and Malik’s method. Other anisotropic smoothing techniques, when applied to the posterior, are likely to be effective as well and should be explored [1, 3, 19, 32].

The MR data we acquire and segment is scalar-valued; the proposed segmentation method could be readily adapted to segment vector-valued (multispectral) measurements [17]. Because anisotropic smoothing is applied to the posterior probabilities, scalar anisotropic smoothing techniques can still be used. The only modification would be to generalize the calculation of class likelihoods.

System performance was evaluated using visual comparisons with results obtained by manual segmentation performed by trained users. For the main goal of this work, visualization of fMRI, the very similar qualitative results obtained between the two approaches are sufficient to adopt the automatic system over the manual one, which takes more than 18 hours of human labor.

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. The anonymous reviewers provided important suggestions, corrections, and related references that have been included in the text. Steve Engel and Hagit Hel-Or provided the initial brain unfolding and segmentation software. This work was partially supported by the Hewlett-Packard Labs Grassroots Basic Research Program, NEI grant ROI EY03164, the McDonnell-Pew Program in Cognitive Neuroscience, ONR Grant N00014-97-1-0509, and NSF Learning and Intelligent Systems. Part of this work was performed while PT and GS were at HP Labs.

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] 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.
- [3] M Black, G Sapiro, D Marimont, and D Heeger. Robust anisotropic diffusion. *IEEE Transactions on Image Processing*, page to appear, March 1997.
- [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] C Davatzikos N Bran. Using a deformable surface model to obtain a shape representation of the cortex. *IEEE Transactions on Medical Imaging*, 15(6):785–795, 1996.
- [6] 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.

- [7] V Caselles, R Kimmel, G Sapiro, and C. Sbert. Minimal surfaces based object segmentation. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 19(4):394–398, 1997.
- [8] 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.
- [9] T Cormen, C Leiserson, and R Rivest. *Introduction to Algorithms*. MIT Press, Cambridge, MA, 1990.
- [10] 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.
- [11] P DeWeerd, A Karni, S Kastner, L G Ungerleider, and P Jezzard. An investigation of fmri resolution in the visual cortex. In A Toga, R Frackowiak, and J C Mazziotta, editors, *Neuroimage*, volume 5. Academic Press, May 1997.
- [12] E Deyoe, G Carman, P Bandettini, S Glickman, J Wieser, and R Cox. Mapping striate and extrastriate visual areas in human cerebral cortex. *Proceedings Of The National Academy Of Sciences Of The United States Of America*, 93(6):2382–2386, 1996.
- [13] 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.
- [14] S Engel, G Glover, and B Wandell. Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cerebral Cortex*, 7:181–192, 1997.
- [15] S. Engel, D. Rumelhart, B. Wandell, A. Lee, M. Shadlen, and G. Glover. fmri of human visual cortex. *Nature*, 369:525, June 1994.
- [16] 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.
- [17] L Fletcher, J Barsotti, and J Hornak. A multispectral analysis of brain tissues. *Magnetic Resonance in Medicine*, 29(5):623–630, 1993.
- [18] J Foley, A van Dam, S Feiner, and J Hughes. *Computer Graphics: Principles and Practice*. Addison-Wesley Publishing Company, Inc., Reading, MA, 1990.
- [19] 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.

- [20] 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.
- [21] 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.
- [22] 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.
- [23] S Kichenassamy, A Kumar, P Olver, A Tannenbaum, and A Yezzi. Gradient flows and geometric active contour models. In *Proc. Int’l. Conf. Comp. Vision*, pages 810–815, Cambridge, MA, 1995.
- [24] T Kong and A Rosenfeld. Digital topology: introduction and survey. *Computer Vision, Graphics, and Image Processing*, 48:357–393, 1989.
- [25] T Lee and R Kashyap. Building skeleton models via 3-D medial surface/axis thinning algorithms. *CVGIP: Graphical Models and Image Processing*, 56(6):462–478, 1994.
- [26] 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.
- [27] J-F Mangin, V Frouin, I Bloch, J Regis, and J Lopez-Krahe. Automatic construction of an attributed relational graph representing the cortex topography using homotopic transformations. In *SPIE Mathematical Methods in Medical Imaging III*, pages 110–121, San Diego, California, 1994.
- [28] R S Menon, S Ogawa, J P Strupp, and K Ugurbil. Ocular dominance in human v1 demonstrated by functional magnetic resonance imaging. *J. Neurophysiology*, 77(5):2780–2787, 1997.
- [29] 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.
- [30] 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.
- [31] P J Rousseeuw and A M Leroy. *Robust Regression and Outlier Detection*. John Wiley & Sons, New York, 1987.
- [32] 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.

- [33] 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.
- [34] P C Teo, G Sapiro, and B Wandell. Anisotropic smoothing of posterior probabilities. *Proc. IEEE Int'l. Conf. Image Processing, Santa Barbara, CA*, October 1997.
- [35] D Terzopoulos, A Witkin, and M Kass. Constraints on deformable models: Recovering 3D shape and nonrigid motions. *Artificial Intelligence*, 36:91–123, 1988.
- [36] R Tootell, A Dale, M Sereno, and R Malach. New images from human visual cortex. *Trends In Neurosciences*, 19(11):481–489, 1996.
- [37] 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.
- [38] 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.
- [39] 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.
- [40] B Wandell. *Foundations of Vision*. Sinauer Associates, Inc., Sunderland, MA, 1995.
- [41] B Wandell, S Engel, and H Hel-Or. Creating images of the flattened cortical sheet. *Invest. Opth. and Vis. Sci.*, 36(S612), 1996.
- [42] 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.
- [43] 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.