

# Visualization and Measurement of the Cortical Surface

**Brian A. Wandell, Suelika Chial, and Benjamin T. Backus**

Stanford University

## Abstract

■ Much of the human cortical surface is obscured from view by the complex pattern of folds, making the spatial relationship between different surface locations hard to interpret. Methods for viewing large portions of the brain's surface in a single flattened representation are described. The flattened representation preserves several key spatial relationships between regions on the cortical surface. The principles used in the

implementations and evaluations of these implementations using artificial test surfaces are provided. Results of applying the methods to structural magnetic resonance measurements of the human brain are also shown. The implementation details are available in the source code, which is freely available on the Internet. ■

## INTRODUCTION

The rapid growth of functional magnetic resonance imaging (fMRI) has served to motivate the development of new methods for analyzing and visualizing neuroimaging data. In this paper, we describe methods that we have developed, which fall at the intersection of visualization and analysis. Specifically, we consider methods for visualizing and analyzing the spatial distribution of activity measured along the two-dimensional cortical surface.

The motivation for developing new methods of analysis and visualization flows from consideration of the limitations with current methods. Perhaps the most widespread representation of fMRI measurements is a multislice image. In this representation, measurements of the brain volume are shown as a set of two-dimensional images that represent both the anatomical and functional activity in a series of planar slices through the brain volume. Using this format, it is difficult to appreciate the spatial relationships between points in separate images. Furthermore, when seen in separate images, the position of any but the most familiar anatomical features is very hard to infer.

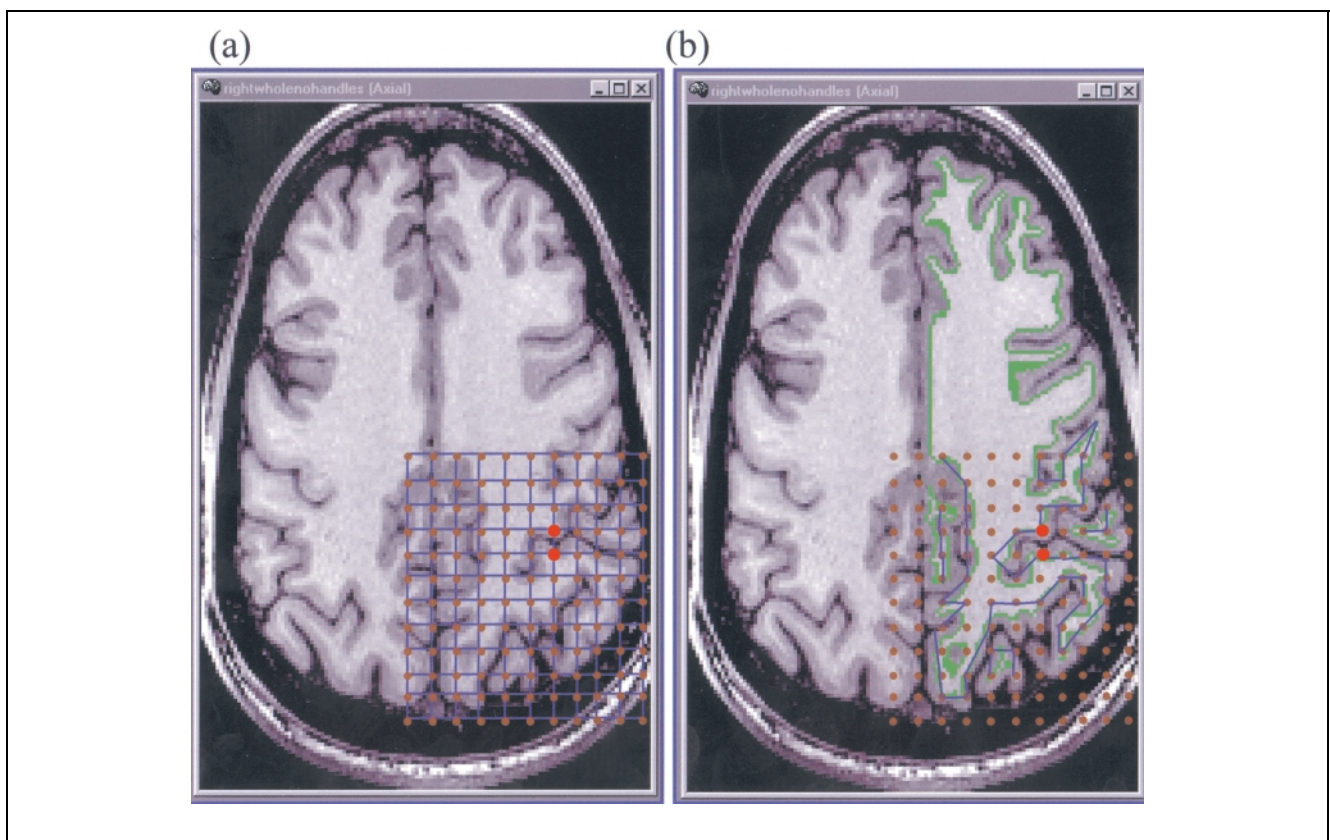
A second widely used format colors the surface of a three-dimensional brain rendering. This format greatly improves one's ability to see the relative spatial relationships between different loci of activity. The difficulty with this method is that much of the cortical gray matter is buried within sulci whose shape varies across individuals and whose form is quite complex. Consequently, this approach offers only a coarse representation of the spatial position of the activity, useful mainly for distinguishing between activity separated by a few centimeters

along the cortical sheet. Such observations can be important in many cases, but this visualization method does not approach the potential resolution of the fMRI signal nor can the method be used for various types of measurements often made in computational neuroimaging (Wandell, 1999).

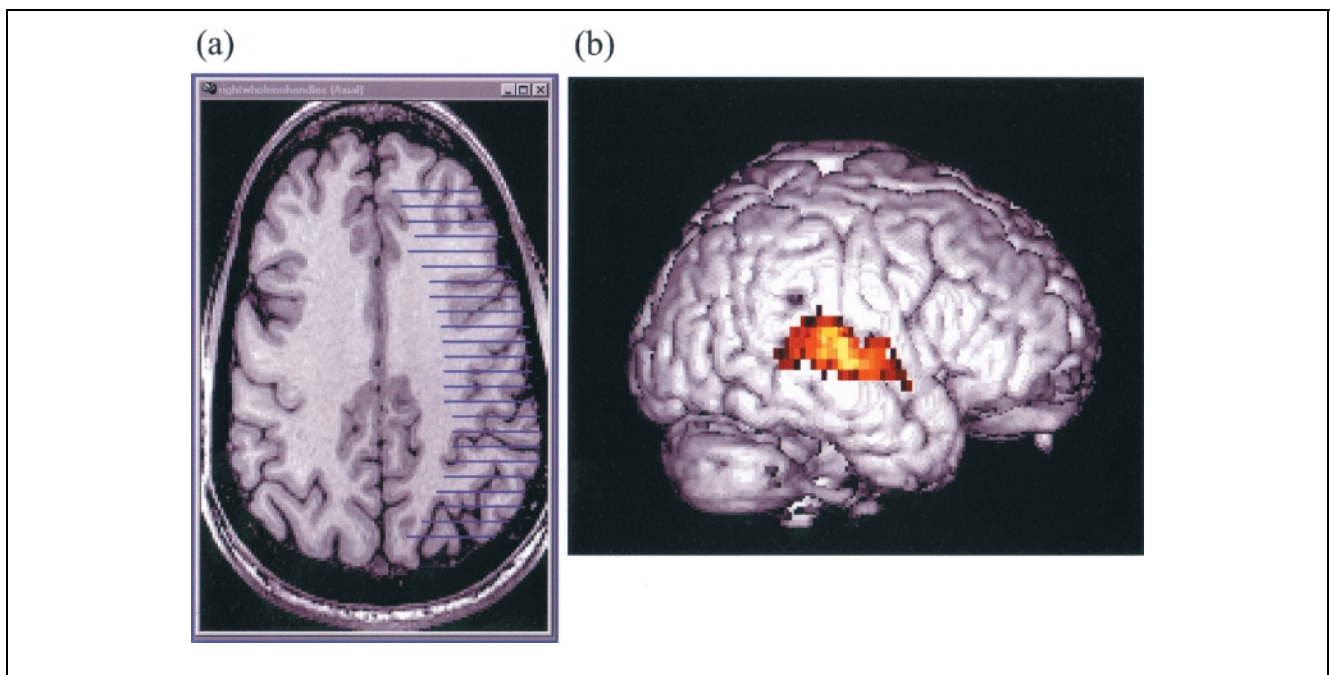
During the last 15 years, a number of authors have described a third approach to visualization methods, and this paper forms part of that series. Schwartz, Dale, Drury, Goebel, and their collaborators define a variety of methods of transforming the folded cortical surface into simpler forms that can be viewed in a single image while preserving some aspects of the spatial relationships (Fischl, Sereno, & Dale, 1999; Drury et al., 1996; Wandell, Engel, & Heil-Or, 1996; Carman, Drury, & Van Essen, 1995; Dale & Sereno, 1993; Dale, Fischl, & Sereno, 1999). The conceptual issues concerning these visualization methods can be divided into two main theoretical topics.

The first important theoretical challenge is to specify how to identify and then transform the spatial relationships inherent in the MR measurements to the relationships inherent in the cortical surface. These two representations have fundamentally different neighborhood relationships that are illustrated in Figure 1. The instrumental sampling of the MR signal falls on a regular grid of points that spans a volume. In conventional analysis, one treats adjacent points as connected for averaging and visualizing.

The rectangular sampling lattice of the MR measurements is illustrated in panel (a). Adjacency relationships inherent in the cortical surface do not follow this pattern of connectivity. Rather, as shown in panel (b), the



**Figure 1.** The sampling grid of the fMRI signal and the spatial structure of cortical gray matter. The nearest neighbors in cortex differ from the usual assignment of nearest neighbors to the fMRI grid so that adjacent points in the fMRI grid are not adjacent in gray matter. Methods that average across the fMRI grid, without respecting the cortical neighborhoods, will provide an inaccurate picture of brain activity. See text for details.



**Figure 2.** Spatial averaging of the fMRI signals used in the “glass-brain” representation. This representation averages across gray matter positions that are widely spaced in the brain. (a) An axial anatomical image with superimposed lines showing the 3-cm trajectory used to create volumetric representations of statistical parameter maps. (b) A glass-brain representation obtained from SPM '96, showing the values of a parameter map on the cortical surface.

connections trace out a two-dimensional sheet along a complicated path within this volume and points that are adjacent in the MR data are not adjacent in the cortical sheet.

For example, the two highlighted points are adjacent in the MR measurements and would be blurred under conventional convolution operations, but they are not adjacent on the cortical sheet. We refer to the neighborhood relationships as defining the configuration or “topology,” and we describe the discrepancy in the identity of neighboring points as a topological inconsistency between the instrumental samples and the cortical surface samples.

In recent years, the instrumental sample spacing has become significantly finer than the size of important brain structures such as the principal sulci and gyri. Hence, to represent and visualize data acquired at the instrumental resolution accurately, it is necessary to develop computational methods to relate the instrumental sampling and the cortical topology. Identifying and using the proper neighbor relationships is an important step in developing algorithms both for visualization and analysis.

One example of why neighborhood relationships matter may be appreciated by considering two different types of brain atlas maps. The widely used Talairach representation of brain coordinates (Talairach & Tournoux, 1988) defines a rectangular grid that is closely coupled to the measurement space, but not to the intrinsic shape of the cortex. In contrast, Van Essen, Drury, Joshi, and Miller (1998) propose a surface-based representation that registers activity with respect to sample positions along the cortex. This representation is used in several laboratories for visualization, spatial averaging, and noise removal (Engel, Glover, & Wandell, 1997; Tootell, Dale, Sereno, & Malach, 1996; Tootell et al., 1997), and it is also used in recent software products (<http://www.brainvoyager.de>).

Figure 2 illustrates a second reason why identifying the proper topology is important in the visualization of MR images. The figure illustrates how inconsistent topology can cause problems in visualizing signals using the “glass-brain” representation. Panel (a) shows an axial anatomical image. In the glass-brain representation used in SPM '96, statistical parameter maps measured along lines 3 cm from the surface of the brain are averaged and shown as activity on a three-dimensional rendering of the whole brain (panel (b)). The lines superimposed on the axial image show the set of integrating lines whose parameter maps are averaged and then displayed. This averaging, with respect to the topology of the instrumental coordinates, does not capture the topology of the brain. Hence, the activity portrayed on the surface can only provide a coarse summary of the local activity. The mismatch between the averaging in instrumental space (straight lines) and the curved shape of the cortical

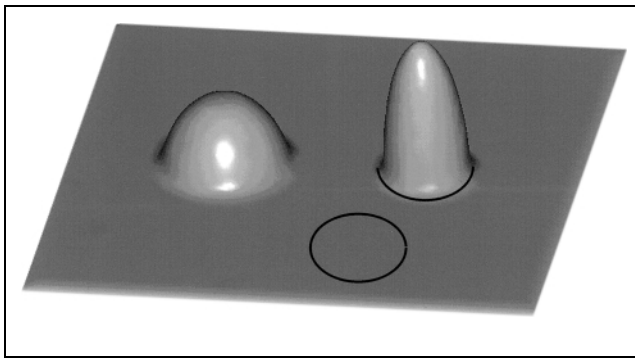
surface degrades the effective spatial resolution of the representation.

Several aspects of visualization and averaging can be improved by using the proper cortical gray matter neighborhoods. First, the neighbor relationships can be used to create flattened representations so that activity from the entire surface can be viewed in a single image. How this can be accomplished is the main focus of this paper. Second, the neighbor relationships serve as the basis for three-dimensional rendering of the brain. Often, renderings of the boundary between the gray and white matter provide a satisfactory method for visualizing most of the surface area in a small number of views. With the widespread use of inexpensive processors for three-dimensional rendering, these visualizations are now practical and satisfactory for interactive data analysis. Third, the neighbor relationship can be used for statistical processing that combines and compares data properly with respect to their locations on the surface of the gray matter. Combining and comparing measurements in a way that respects the location of the data with respect to the gray matter is far superior to methods that are based on the instrumental sampling grid. We believe that the use of cortical neighborhood relationships will become a very important part of data analysis in the next few years.

Elsewhere, we have described and distributed algorithms for segmenting cortical gray matter and producing representations that specify the connectivity (neighborhood relationships) of points on the surface (Teo, Sapiro, & Wandell, 1997). We refer to the connectivity between points as the mesh topology. Here, we describe procedures for transforming the connected representation to a flattened surface. In the first part of the paper, the key principles of the algorithms are described, and a mathematical method for initiating the unfolding process using a representation without any twists in the surface is explained. Quantitative evaluations of the algorithm using analytic test sets to illustrate the method and its limitations are described. The method is then illustrated using anatomical MR images of the human brain. Finally, related work in the literature is discussed, and the limitations of common measures, such as Gaussian curvature, are described. The complete algorithm includes many implementation details, and the interested reader can see these in the source code that is distributed on the Internet (<http://white.stanford.edu>).

## ALGORITHM OVERVIEW

The motivation for the algorithms that we describe can be understood by considering the fundamental flattening problem illustrated in Figure 3. The image shows a plane along with two quadratic surfaces, like bowls extruded through the sheet. These extrusions



**Figure 3.** Why flattening distorts distances. The surface contains three regions. One is planar, a second has a portion of a sphere inserted, and the third has a portion of an ellipsoid inserted. The perimeters of the two circles are the same, but the areas contained within them differ. The flattened representations described here mainly distort the distance relationships between nodes on the spherical and ellipsoidal regions.

are typical of cortex, though their precise shape varies and some extrusions are more similar to ridges than bowls. The entire cortex consists of many pieces similar to this one but folded into a crumpled shape.

Were there no extrusions, or technically were the surface “developable,” unfolding the crumpled representation and preserving the distances between points (as measured along the surface itself) would be possible (Horn, 1986). If all the distance relationships are preserved, other metric quantities (e.g., angles and areas) would be preserved as well. When the cortical surface includes regions like the half-sphere and half-ellipsoid, no flattened representation can precisely preserve the distance relationships.

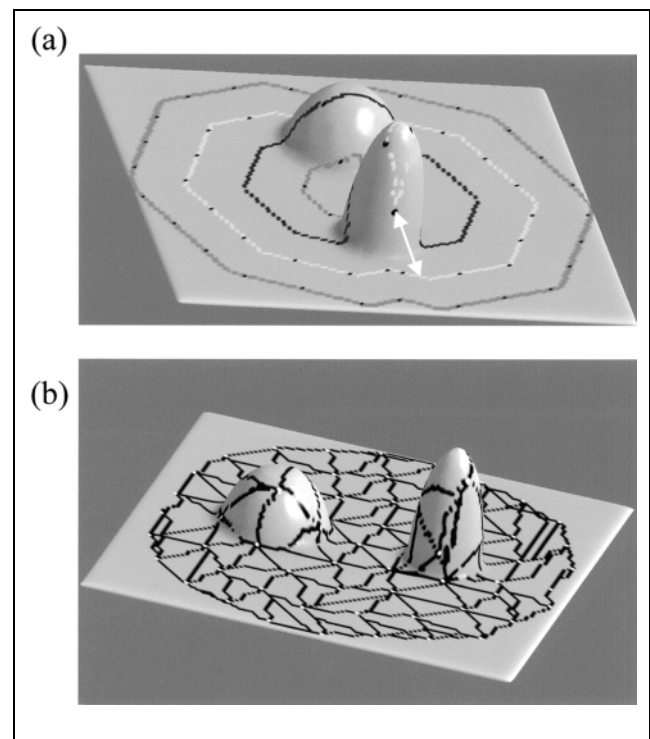
One way to see why a distance-preserving flattened representation is impossible is to consider the area of cortex within the regions marked on the figure. The circle shown on the plane contains much less surface area ( $\pi r^2$ ) than the equal-perimeter circle that bounds the ellipsoidal extrusion ( $> 2\pi r^2$ ). There is no way to make these circles map in the same way to a flattened representation while simultaneously preserving the relationships between the points falling inside the circles. One must choose the most important quantity to preserve.

The quantity we place the greatest value on preserving in the flattened representation is illustrated in Figure 4. Panel (a) shows a set of closed curves on the surface in Figure 3. The curves are equidistant (within the cortical sheet) from a central point between the sphere and ellipsoid. Were the surface a simple plane, the curves would be a series of concentric rings. The conic forms inserted within the surface deform the equidistant curves, and in one case split the distance contour into two unconnected parts, each of which is itself a closed curve. The dark points on the curves show a set of sample points, also called nodes, which are spaced evenly on the contours. These nodes form a coarse sampling grid that spans the surface. Panel

(b) shows the result of applying a mesh generation algorithm to this sampling grid. Adjacent points on the closed curves are connected, and nearby points on adjacent rings are connected by edges. These edges are then pruned to remove any intersections, except if doing so would create a node with fewer than three edges. A graph containing no intersecting edges is a planar graph (Bollobas, 1991).

Like Drury et al. (1996), a main objective in creating the flattened representation is to preserve the arrangement of the nodes and edges. We preserve neighborhood relationships between the sample nodes and prevent large twists from occurring in the flattened representation. A key element of the algorithm, then, is the method used to create the mesh on the cortical surface that guarantees that the flattened representation has no twists or tears. A novel aspect of our design is that nodes can be placed within the plane, preserving the mesh topology, in a single step. No iterative processing is required, resulting in a significant computational advantage.

The initial placement of the nodes can result in a spacing that differs significantly from the cortical spacing. After creating the initial flattened representation, the two-dimensional positions of the nodes are adjusted in order to match cortical and planar graph distances



**Figure 4.** (a) Curves equidistant from the central point are shown. The curves show points at 5, 10, 15 mm, and so forth. The ellipsoidal insertion breaks the 15-mm curve into two disconnected parts (arrow), each of which forms a simple contour. The dark points on the curves are sample nodes placed along the curves at roughly equal steps. (b) The results of the mesh generation applied to the sample nodes in (a).

between sample nodes. To match these distances, we find the shortest distance (geodesic) between all pairs of sample nodes, storing both the path and the total distance. In an iterative loop, sample node positions are adjusted to match the sum of the edge lengths along the stored path with the lengths of the corresponding cortical geodesics. Only adjustments in the two-dimensional node position that preserve the mesh topology (no edge intersections can be created) are considered. Although the planar and cortical distances cannot be matched perfectly, the initial placement can be substantially improved. A quantitative evaluation of the distances is presented in the Results.

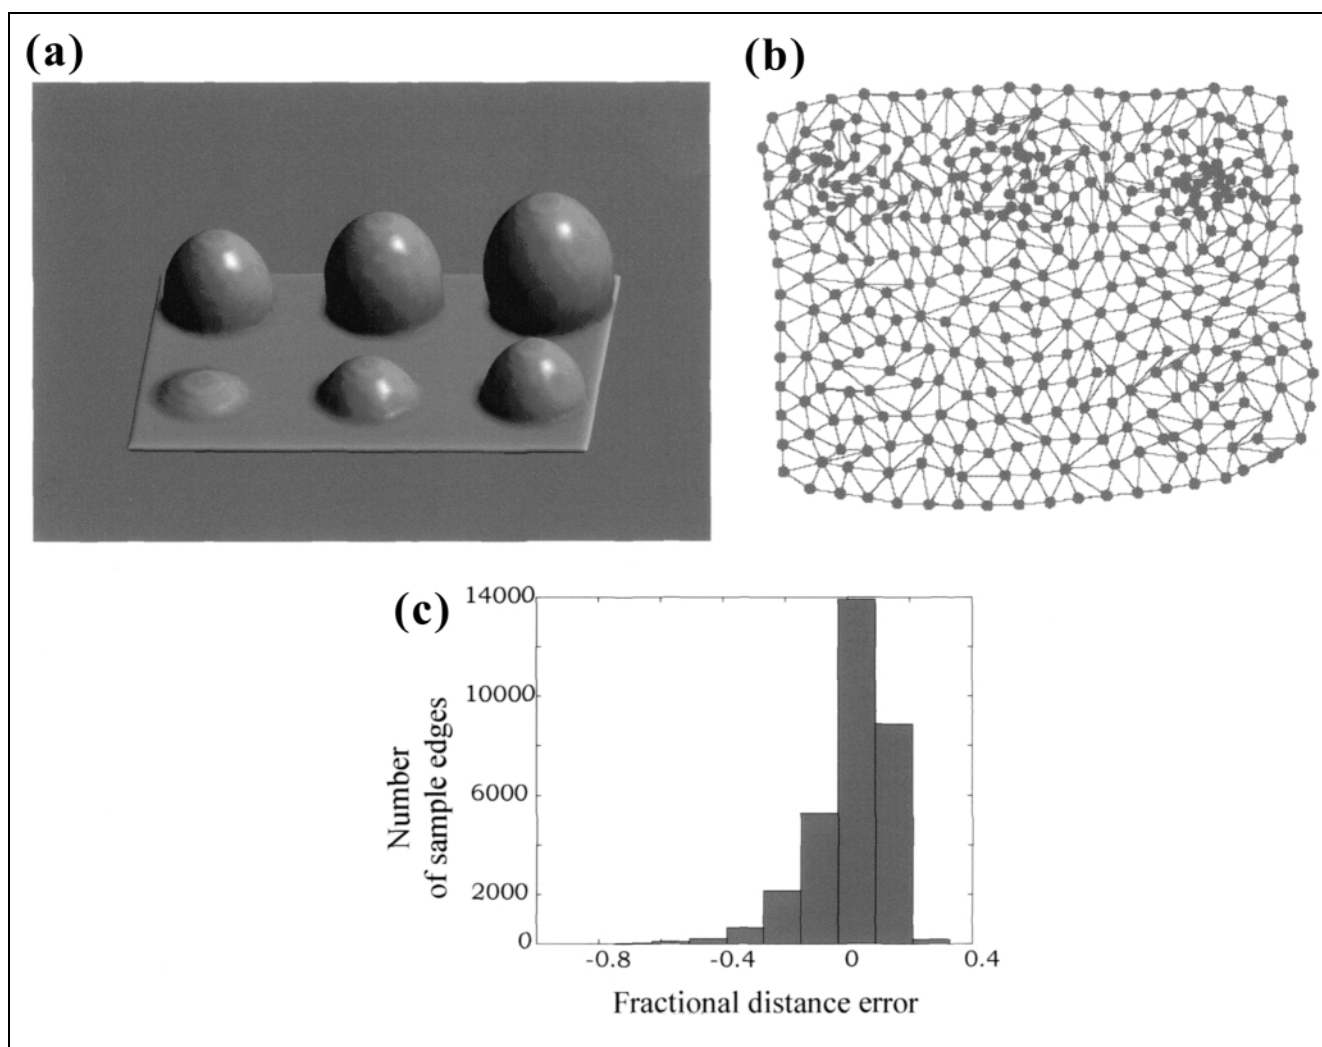
## RESULTS

In this section, we measure the distortions introduced by the unfolding algorithm. We first illustrate some of the properties of the algorithm by making measurements using three test surfaces designed to evaluate the algorithm. Then, we show examples of unfolded

human brains. The distortions in these example brains can be compared to the distortions observed in the test surfaces.

### Curvature Distortions

The distance distortions introduced by the unfolding method are evaluated here by analyzing the results of unfolding several different types of test surfaces. Consider first the surface shown in Figure 5a. The surface consists of a plane that has several spheres inserted into it. The spheres are inserted to different degrees, some just emerging above the surface and others ballooning out significantly. Consider this surface as representing the white matter, so that gray matter would cover the surface. The sample nodes, which are not shown on the surface, were chosen automatically (but with a spacing that was small enough so that the nodes were fairly evenly spaced as measured along the surface). Panel (b) shows the final locations of the sample nodes after unfolding the surface. The sample nodes are relatively compressed at locations



**Figure 5.** (a) A test surface comprising a plane with different portions of a sphere inserted. (b) The positions of the sample nodes after flattening. (c) A histogram summarizing the distance errors (distance in cortex minus distance on the plane) between sample nodes.

corresponding to the large spheres. The amount of the distortion is larger at the larger spheres.

Panel (c) measures the distortion by comparing the distance between points measured along the geodesic path through the sample nodes on the original surface ( $d_m$ ) with the distance measured along the same path in the planar representation ( $d_p$ ). In order to compare distances between nearby and widely separated points, we plot the fractional error,  $(d_p - d_m)/d_m$ , so that positive and negative errors represent an expansion and compression of the planar distances, respectively. The extreme distortions of the planar distances are roughly 40% of the distance on the surface, though for most of the points, the distance distortions are relatively small. The values in this figure can be used to compare with the values presented later for flattened representations of human brains.

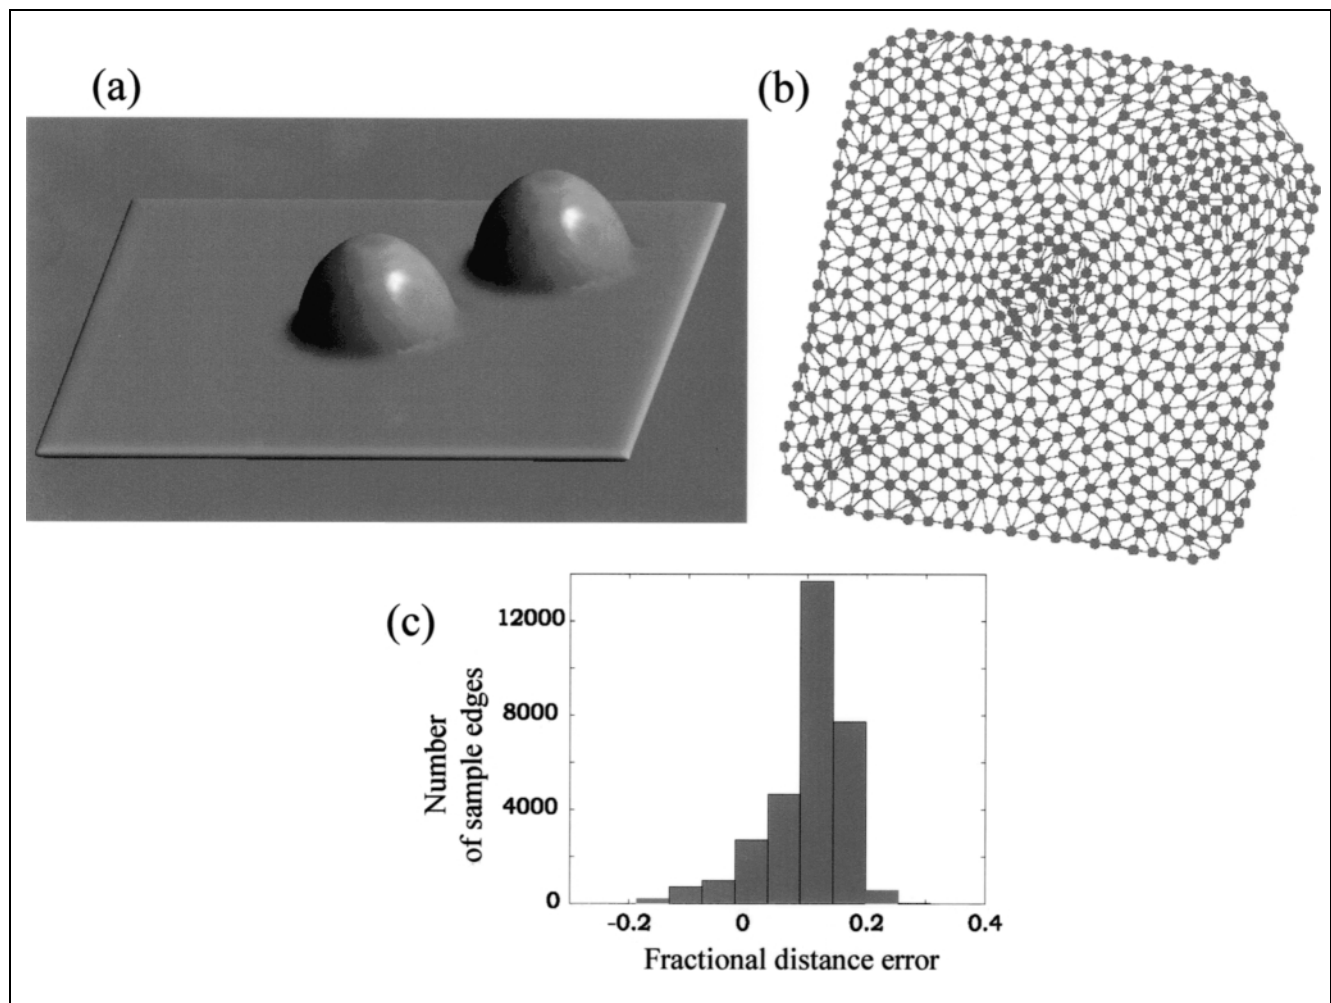
### Perimeter Distortions

The unfolding algorithm introduces distortions in a space-variant manner. Specifically, nodes on the perimeter are placed down separately from the other sample nodes. Consequently, the errors associated with a dis-

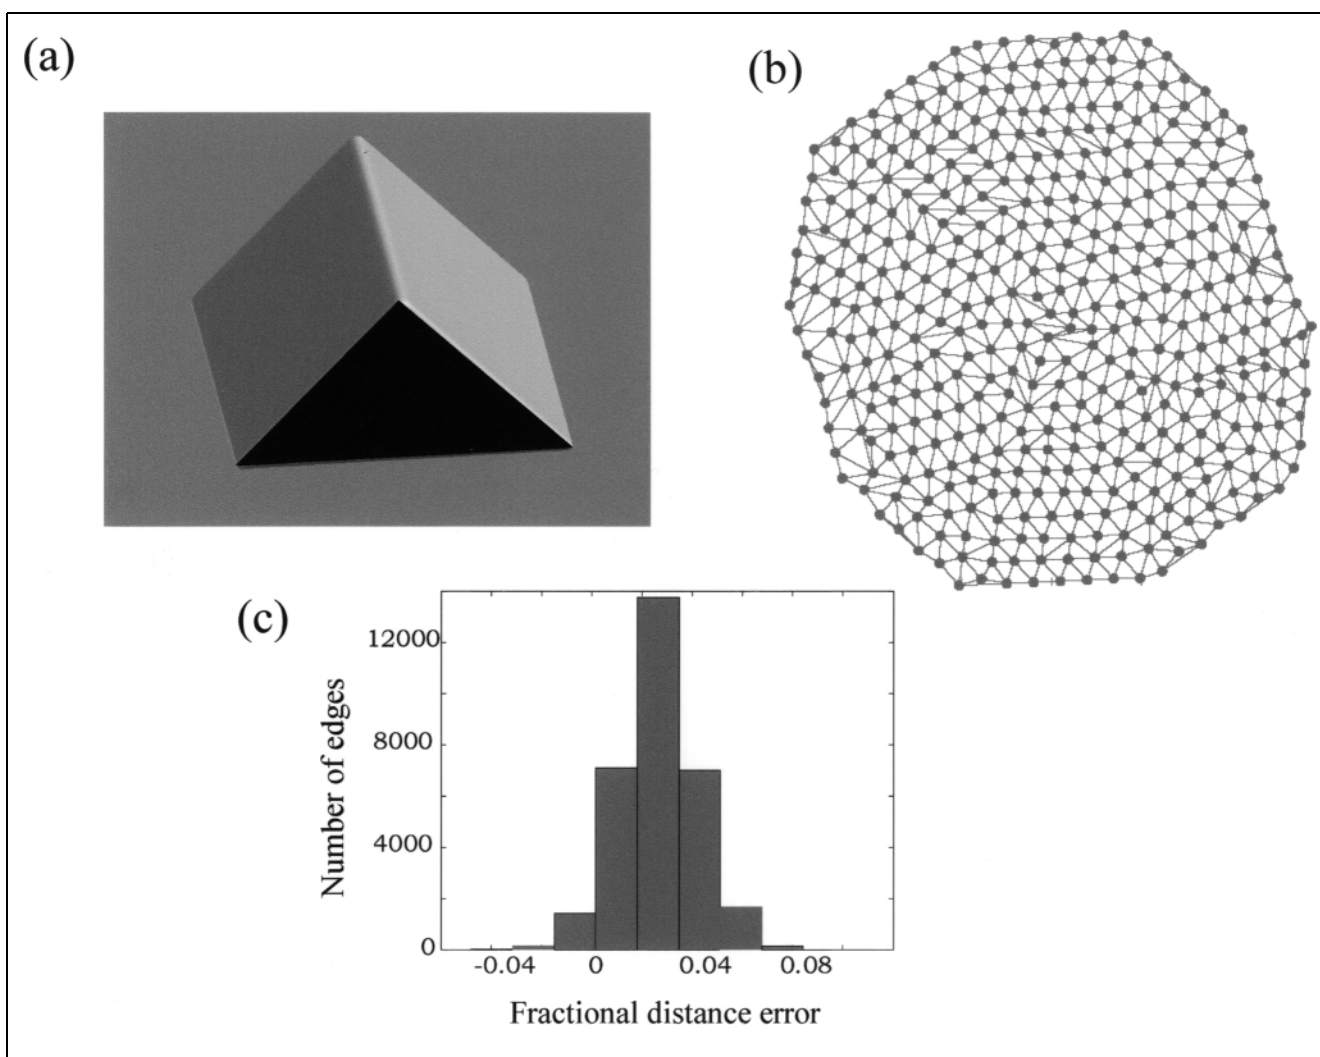
tortion in the center of the unfolded region are slightly different from those near the perimeter. We have observed this variation in normal unfolding, and it is illustrated for the test surface shown in Figure 6a. This figure comprises a single plane with two identical spheres inserted into it, one near the center of the plane and the other near a corner. The representation of the sample nodes shows the usual compression near the spheres, but the nature of the compression is somewhat different for the central and peripheral spheres. Panel (b) shows that the majority of the distortions of the central sample nodes are more localized than those of the peripheral nodes. In practice, perimeter distortions that influence the visualization can be avoided simply by setting the perimeter to be more than 5 mm from the area of interest. Notice that the fractional distance error for this surface is considerably smaller than for the surface in Figure 5.

### Folded Planar Regions

Finally, the discrete sampling of the original MR measurements and numerical errors in the algorithm



**Figure 6.** A test surface comprising a plane with a portion of a sphere near the center and near the perimeter. Other details as in Figure 5.



**Figure 7.** A test surface comprising a folded plane. The flattened representation represents both the mesh properties and the distance properties accurately (typical error of less than 1%). Some error is expected due to quantization noise. Other details as in Figure 5.

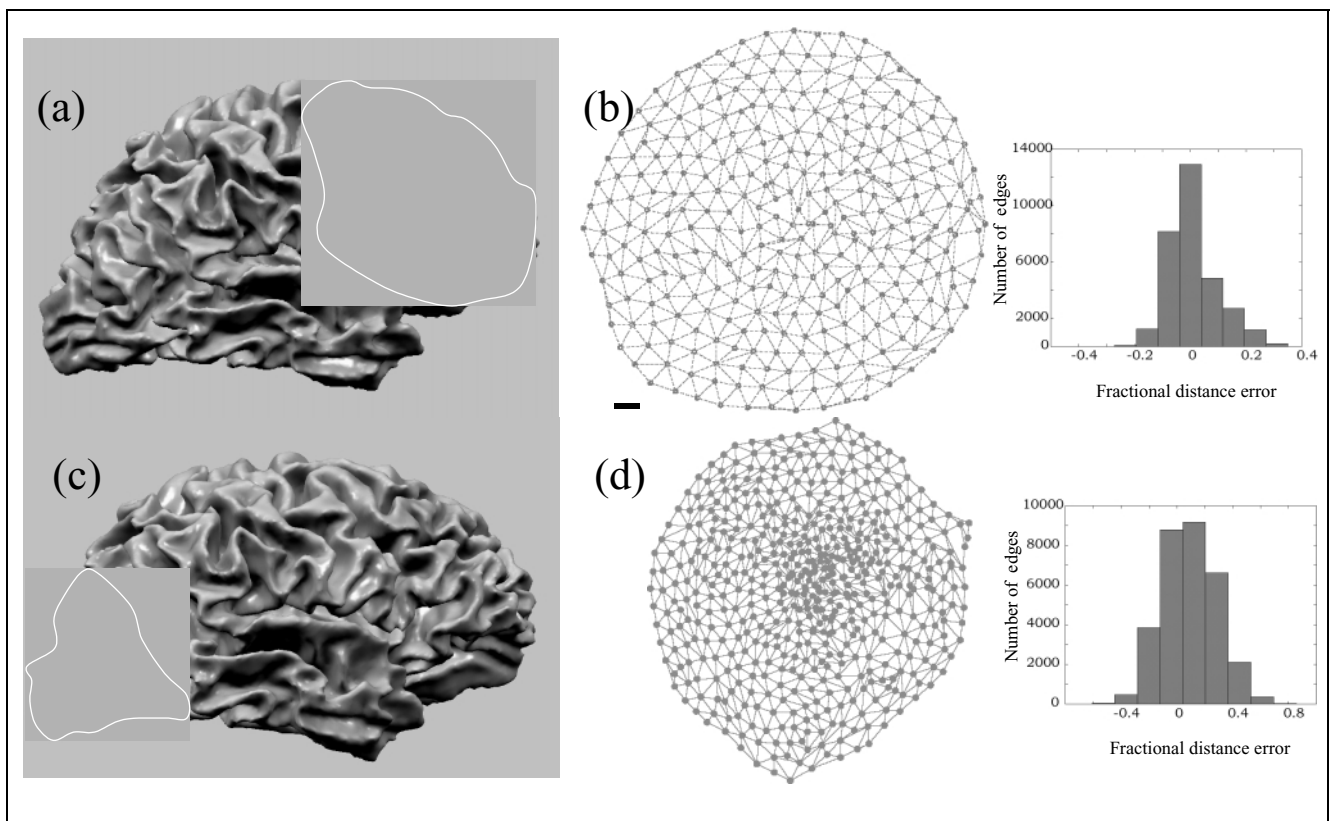
will introduce some amount of error into the unfolding process as well. We can measure this by unfolding a surface with no intrinsic curvature, such as the bent plane shown in Figure 7a. The spatial array of sampling nodes does not have any large regions in which the nodes are compressed (Figure 7b). The size of the fractional distance error is greatly reduced, with large error corresponding to only 0.05, compared to the more extreme values of 0.4 and 0.2 in Figure 5 and Figure 6.

### Human Brain

Planar sample node positions derived from the frontal and occipital lobes of a human brain are shown in Figure 8. Similar number of sample nodes were used, though the area of the frontal lobe is significantly larger. The accuracy achieved in these two portions of the brain is quite different as illustrated by the dense cluster of dots in the occipital lobe. The

frontal lobe can be flattened with a fractional distance error that, mainly, is within 10–15%. The occipital-lobe error is mainly within 30%. These error values are similar to values reported by Drury et al. and Fischl et al. using their procedures. The difference in the size of the errors presumably reflects differences in the intrinsic curvature of cortex: The occipital lobe is more like the extruded spheres in Figure 5, and the forebrain is more like a crumpled sheet.

Finally, Figure 9 is an image of a flattened hemisphere. While our work does not require us to flatten the entire hemisphere at once, we include this image to demonstrate that the procedures described here succeed on large unfoldings. In this image, the shading represents the local curvature of the folded brain: Dark and light regions tend to fall near the sulcal fundi and crowns of gyri, respectively. The positions of the perimeter, cuts, and several principal sulci are shown in panel (a). The flattened representation is labeled with these same sulci, and the boundaries of



**Figure 8.** Flattened representations of the frontal (a, b) and occipital (c, d) lobes of a human brain. The region selected for each of these lobes are shown by the lightly shaded areas in panels (a) and (c). The sample node positions are shown in (b) and (d) along with the fractional errors of distances in the flattened representation. Scale bars indicate 1-cm distances in the flattened representation.

the four lobes are marked. The picture is quite similar to the representation obtained by Drury et al. (1996) in which they applied their methods to data from the visible human data set.

Recall that the algorithm positions points representing the gray matter nodes, and the final representation is created by interpolating across these unevenly positioned points. We do not interpolate curvature beyond a few millimeters, so that in a few places, holes are apparent. These occur when the flattening process stretches a local region more than it should. In practice, we eliminate holes by using a different start point or smaller selection.

## DISCUSSION

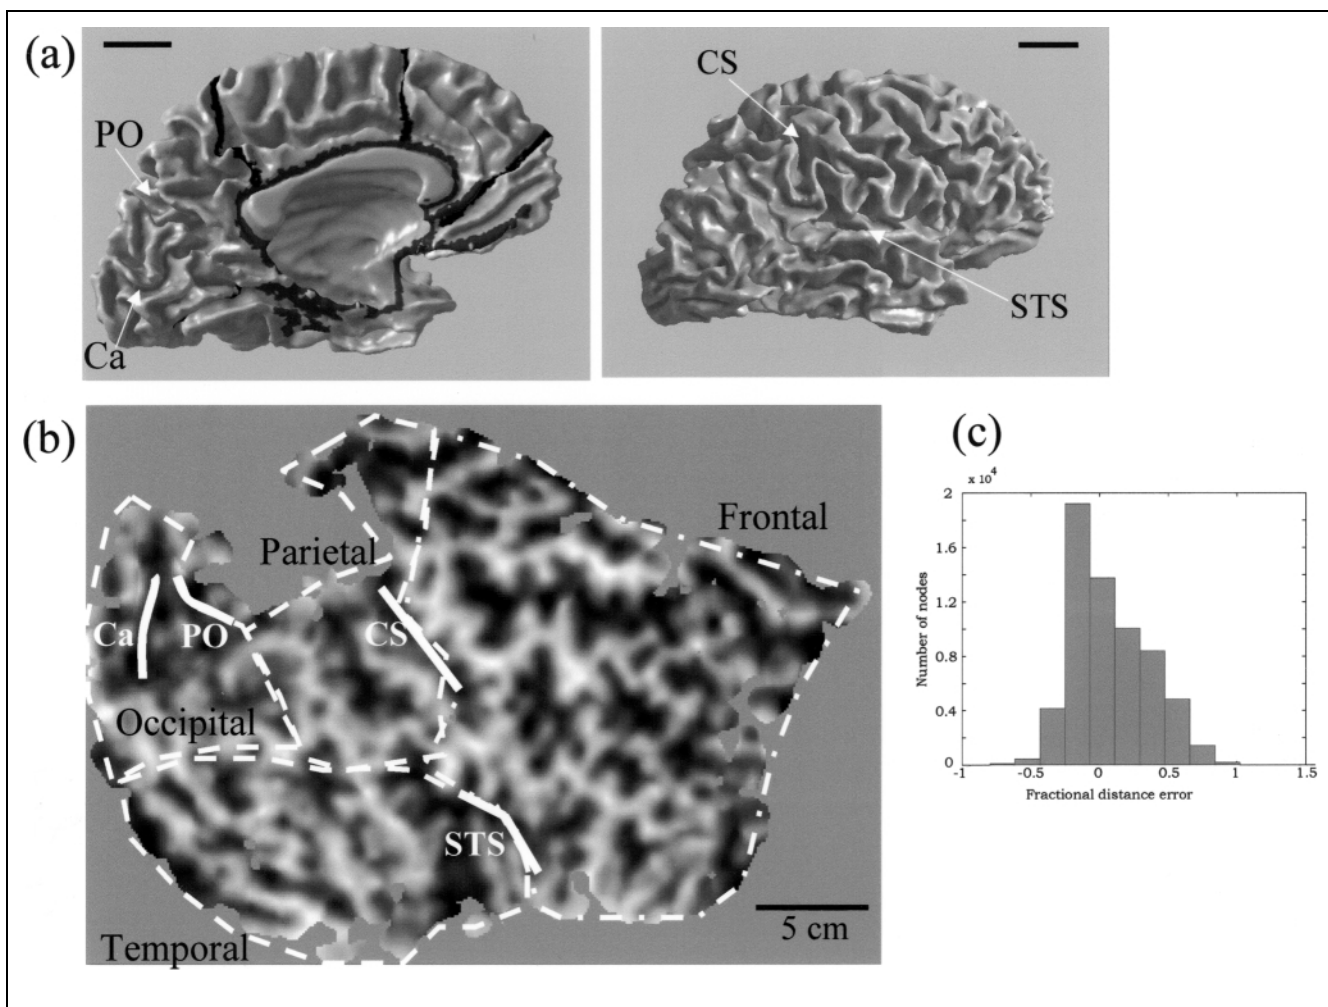
### Curvature

Within the literature on cortical flattening, Gaussian curvature plays a deservedly important role. A nice example of the use of curvature and its application to surface characterization is in the work of Van Essen and Drury (1997), who use it to distinguish the properties of different cortical lobes. Curvature is also often used in the form of a graphical aid, as in Figure 9, to label sulci and gyri in the flattened representation.

This measure has some important limitations, however, if one is interested in understanding how difficult it will be to flatten a surface and preserve metric properties. For example, while surfaces with nonzero Gaussian curvature cannot be flattened without metric distortion, not all surfaces with zero Gaussian curvature can be flattened without distortion. An example is the soap film in a saddle-shaped wire frame. Hence, curvature is not a complete characterization of how well a surface can be flattened.

Another limitation on the use of Gaussian curvature is shown in Figure 10. Two cylinders are shown that differ only in height. The Gaussian curvature along the surfaces is zero, except near the base and top where it is the same for both cylinders. Yet, the two surfaces will not have similar distortions when flattened. Because the area of the tall cylinder is much greater, that surface will have larger metric distortions when flattened.

The locality of the Gaussian curvature fails to capture some of the problems that will be observed in the flattening algorithms. In addition to local measures, such as Gaussian curvature, measures of the ratio of surface area to perimeter capture much of the problem with unfolding. On a truly planar surface, points within a distance,  $r$ , sweep out area ( $\pi r^2$ ), and define a perimeter of length  $2\pi r$ . Hence, the area to perimeter ratio is  $r/2$ .



**Figure 9.** Flattened representation of an entire hemisphere. (a) The locations of the perimeter and cuts are shown as black lines. Principal sulci, also labeled on the flat representation, are indicated by the labels and white arrows. (b) The flattened representation is shaded according to the local convexity of the cortical surface. Sulci and gyri are shaded as black and white, respectively. Several sulci are labeled and marked as solid white lines. Scale bars measure 5 cm. Panel (c) and other details as in Figure 5.

Deviations from this value indicate that there will be metric distortions in the flattened representation.

### Related Work

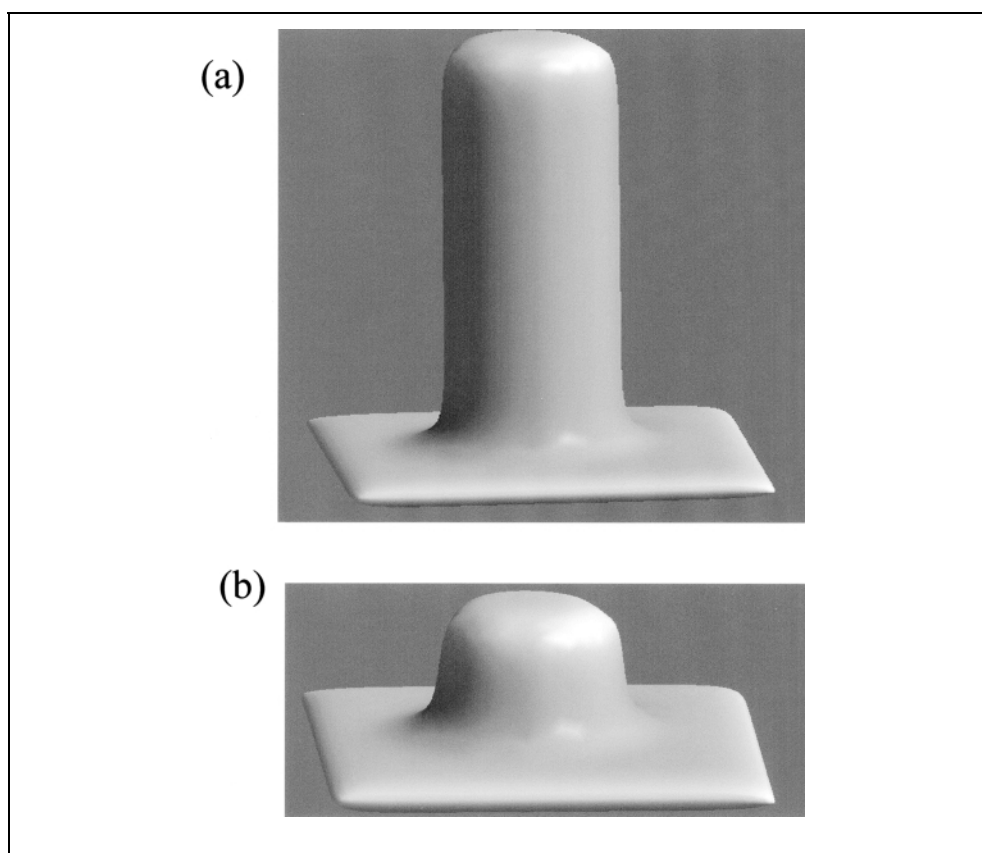
The algorithms described here are most closely related to the contemporary work of Drury, Van Essen, and their colleagues (Van Essen & Drury, 1997; Drury et al., 1996). The methods are also related to the work of Dale and his colleagues (Fischl et al., 1999; Dale & Sereno, 1993; Dale et al., 1999). This work, in turn, draws on ideas from the seminal work of Schwartz (1990) and his colleagues. There are a number of very closely related developments in the graphics literature (Levy & Mallet, 1998; Floater, 1997; Eck et al., 1995), as well as the mathematical literature on surface properties and topological graphs referenced above. In this section, we mainly discuss the general principles and the work directly pertaining to cortical flattening.

Drury et al. (1996) describe two related algorithms for creating flattened representations. In both methods, a

wire-frame reconstruction of the surface is created, such as the one produced by software developed in our laboratory (Teo et al., 1997). In the method most similar to the one described here, a seed point on the brain is chosen and placed on the plane. Then, a rapid flattening is achieved by the method they call “accretion of concentric rings.” They find all the points connected to the seed node and place these points on the plane around the seed node. Then, they place the next ring down around that node, and so forth.

The sample rings that we create and place on the plane are similar in concept, though there are some significant differences. First, we place down only a coarse sampling of the nodes rather than the full complement. Second, Drury et al. did not comment on the fact that sometimes, rings at equal distance from a seed node are split into two sections, as illustrated in Figure 4. This case can be problematic, and we suspect that it is the reason they observe substantial shearing in some instances when using this method (Drury et al., 1996, p. 24). Third, the initial placement of the nodes contains twists and folds,

**Figure 10.** Intrinsic curvature does not predict the metric distortion caused by flattening. Although the intrinsic curvature of these two surfaces is the same, the metric distortions after flattening will differ.



and significant resources are applied in order to eliminate these. The methods used here begin with a topologically correct flattened representation.

Fischl et al. (1999) begin by segmenting gray and white matter and creating a planar graph describing the surface between the gray and white matter using the methods outlined in Dale et al. (1999). The nodes of the graph are projected onto a plane, a procedure that introduces folds. The flattened representation is then improved by an iterative process that minimizes an energy functional containing two terms. One term corrects for the topological folds and twists in the surface. The second term tries to adjust node positions to match planar distances with those measured along the cortical surface. Because the initial planar graph contains many folds, in the early stage of the flattening algorithm, the functional that undoes the folds is given a large weight. As the process continues, the functional that governs the metric properties is given an increasing weight. Hence, the success of the algorithm depends on the properties of the energy functional and the minimization process described in their work. The methods described are smooth and contain well-defined gradients. Even so, the performance of these algorithms in achieving their objectives is notoriously difficult to characterize.

Both the groups include substantial discussions on the use of flattened representations to develop atlases for coregistering different functional measurements.

These papers correctly point out that surface-based atlases have better registration properties than atlases based on three-dimensional coordinates, such as the Talairach–Tournoux scheme (Talairach & Tournoux, 1988 #2009, cf. Figure 1). We note, however, that the use of flattened representations does not solve many of the generic problems with atlases. Brain structures differ significantly and combining data across individuals based on structural registration does not equate functional areas in either three-dimensional brains or their flattened representations.

A fundamental application of connected segmentations will be to integrate cortical neighborhoods with statistical analyses. Analyzing functional data with respect to the gray matter surface, rather than three-dimensional space, can improve the signal-to-noise ratio of the analysis and help identify functional regions. Integrating data across subjects based on functional responses rather than structural form is a profitable future direction.

### Limitations

Within our laboratory, segmentation algorithms are the most important bottleneck for successfully unfolding. Poor unfolding results are usually caused by small errors in the segmentation. Pulse sequences that produce reliable gray/white contrast in MR images greatly

reduce the amount of operator intervention needed in the segmentation process. Checking the topology of the white matter segmentation to make sure it contains no holes or handles, an algorithm developed by Robert Taylor in our laboratory and that we distribute with our code, has been important in helping us to obtain good unfolding results. In addition, our laboratory continues to use tools for identifying regions of the flattened map whose metric properties do not match the brain distances and then help the user verify the segmentation in the corresponding part of the brain.

Measuring distances on a mesh remains a research problem. We expect that the implementation of Dijkstra's algorithm that we use may be replaced in the next few years as the mathematics of computing geodesics on manifolds with less than fully triangulated graphs is solved (Kimmel & Kiryati, 1998).

Our unfolding is usually done to understand the activity in a specific part of the brain, rather than the entire hemisphere. Consequently, algorithm speed has not been an important problem for us, and we have been able to use a simple, high-level programming language (Matlab) that runs across platforms. Our current implementations, however, are not suitable for regularly flattening many hemispheres, as the processing time on a 300 MHz Pentium can be as long as 2 days. There are no extreme memory requirements. For example, the entire hemisphere was flattened using a PC with 256 MB of memory.

## CONCLUSION

Creating flattened representations of the cortical surface and other biological surfaces is a useful visualization tool. In this paper, we have described the challenges in producing flattened representations and the principles that we have adopted to meet these challenges. The implementation details can be seen and further evaluated in the software implementation, as the source code for segmenting and flattening is available on the Internet.

## METHODS

Flattened representations of the cortical gray matter are obtained in six basic steps. First, we identify a region of the cortex that we plan to render in a flattened view. Within this region, gray matter voxels that fall along the boundary of the white matter are identified. These are called "first-layer points." Second, cuts are introduced. Third, a subset of the first-layer points are chosen to serve as sample points. Based on these sample points, we create a triangular mesh that spans all first-layer points. Fourth, the sample points are assigned positions in the plane using a procedure that preserves the mesh topology, as explained above.

Fifth, sample point positions are adjusted within the plane so that distances along geodesic paths defined on the cortical surface match the distances through these same points in the flattened representation. Sixth, the other gray matter points are assigned positions on the plane.

## Gray Matter Selection

The anatomical images used in the examples in this paper were T1-weighted MRI scans (voxel size is  $0.9 \times 0.9 \times 1.2$  mm). The principles of gray matter selection are described in a separate publication (Teo et al., 1997). The current implementation that we use and distribute (mrGray) produces a description of gray matter locations (nodes), as well as their connections (edges). The tool also includes integrated methods for verifying that the topology of the white matter surface is correct, specifically that there are no holes or handles in the surface.

In our work, we commonly flatten regions of interest rather than the entire hemisphere. While it is possible to flatten the entire hemisphere (see Figure 9), flattening a region of interest produces maps with much higher precision, an important concern in the work that we do. To initiate the unfolding process, cuts are introduced and perimeter points are identified. We commonly choose the region of interest by selecting a single gray matter location (start node) in the region of interest and then specifying a distance that we wish to include from the start node. All the nodes within that distance from the start node are selected for the flattened representation. Perimeter points are identified by their distance from the start point and whether their neighbors are selected. It is also possible to select a start node and a perimeter manually.

The measurement of distance along the cortical surface is the most fundamental operation used throughout the code. To measure distances, we use a modified version of Dijkstra's algorithm on the connected representation (Dijkstra, 1959). Our implementation of this algorithm takes into account distances between individual nodes in the gray matter, including differences in the sample spacing in the cardinal directions and differences in distances for connections made between diagonally connected nodes. We have made our implementation of this algorithm publicly available for several years, and it has been independently built and confirmed (H. Yamamoto, personal communication).

## Mesh Construction

Next, a sparse mesh (selected nodes and edges to connect them) is created, which covers the portion of the cortical surface that will be flattened. This mesh generation operation comprises four steps.

First a set of connected isodistance contours are found, with each contour falling at a user-defined sample spacing from the previous contour. The last contour is set to be the points on the perimeter of the unfold area, the perimeter nodes. This ensures that we cover the entire cortical area of interest.

Second, the nodes falling along the cuts are removed. The cuts always extend from a perimeter node. The points falling along the boundary of the cut are added to the perimeter nodes.

Third, sample nodes are selected along each of the isodistance contours. These are selected by finding a geodesic between the start node and a node on the perimeter, i.e., an orthogonal geodesic contour. The locations where this geodesic intersects each of the contours is the first sample node for that contour. The first sample node is the start point. We select additional sample nodes on the remaining contours as follows. For the second contour, we identify nodes within one-half the sample spacing from the contour's first sample point and eliminate these as candidate sample nodes. We choose the second sample node as the unmarked point whose distance from the start point is most similar to the sample distance. Second sample nodes on subsequent contours are chosen to be the sample spacing away from the first node and close to the second sample of the previous contour. Continue selecting sample nodes for each contour. If no points matching the criteria are found, and, yet, there are unmarked points on the contour that are more than a sample spacing away, the contour must be split into two disjoint parts as illustrated in Figure 4a. In that case, continue choosing sample nodes on the second part of the contour.

Fourth, edges connecting the sample nodes are introduced. Edges are placed between all pairs of sample nodes separated by less than twice the sample spacing. The distance requirement is relaxed in the event that a node does not have at least three edges. Next, we remove intersecting edges. Edge intersections are detected by finding the geodesics between connected nodes and detecting common points. The order in which edges are removed is determined by the requirements that (a) each node must have at least three edges and (b) longer edges are removed first. After this step, the mesh covers the region of interest and has roughly uniform sample node spacing.

### Flattening the Mesh

Until this stage of the computation, all operations are performed on the three-dimensional positions. The next step is to assign two-dimensional locations to the sample nodes. This assignment begins by assuming the start node position is (0,0). The perimeter sample nodes are assigned positions based on their distance from the start node and their spacing along the perimeter. The perimeter points

are first placed on a circle whose radius matches the distance from the start node to the perimeter. Then, maintaining the angular position, the perimeter node positions are adjusted individually so that their planar and cortical distances from the start node match.

The interior sample nodes are assigned initial locations on the plane by placing each node at a position equal to the average position of its neighbors (the sample nodes connected by an edge). When the neighbors form a convex shape, placing the node at the average position of the neighbors will not introduce any edge intersections and the mesh will continue to be a planar graph (Levy & Mallet, 1998; Floater, 1997; Tutte, 1960). If the node's neighbors are a concave shape, this is not guaranteed. But in our experience over hundreds of unfolds, the mesh generation algorithm combined with the initial placement procedure introduces no edge intersections among the edges of interior nodes.

The initial placement of the nodes can be implemented in a single efficient linear calculation. The linear equation can be understood from the following considerations. Suppose there are  $n$  interior sample nodes and  $p$  perimeter points. Let  $\mathbf{X}$  be the  $n \times 2$  matrix containing the planar positions of the sample nodes. Let  $\mathbf{N}$  be an  $n \times n$  matrix whose  $(i,j)$ th entry is  $1/m_i$ , where  $m_i$  is the number of edges connected to the  $i$ th node, or 0 if  $(i,j)$  are not connected. The  $p \times 2$  matrix  $\mathbf{X}_0$  contains the two-dimensional locations assigned to the perimeter points. Finally,  $\mathbf{P}$  is an  $n \times p$  matrix whose  $(i,j)$ th entry is  $1/m_i$  when perimeter node  $j$  is connected to sample node  $i$ . Each node is at the average position of its neighbors when the following equation is satisfied.

$$\mathbf{X} = \mathbf{N}\mathbf{X} + \mathbf{P}\mathbf{X}_0$$

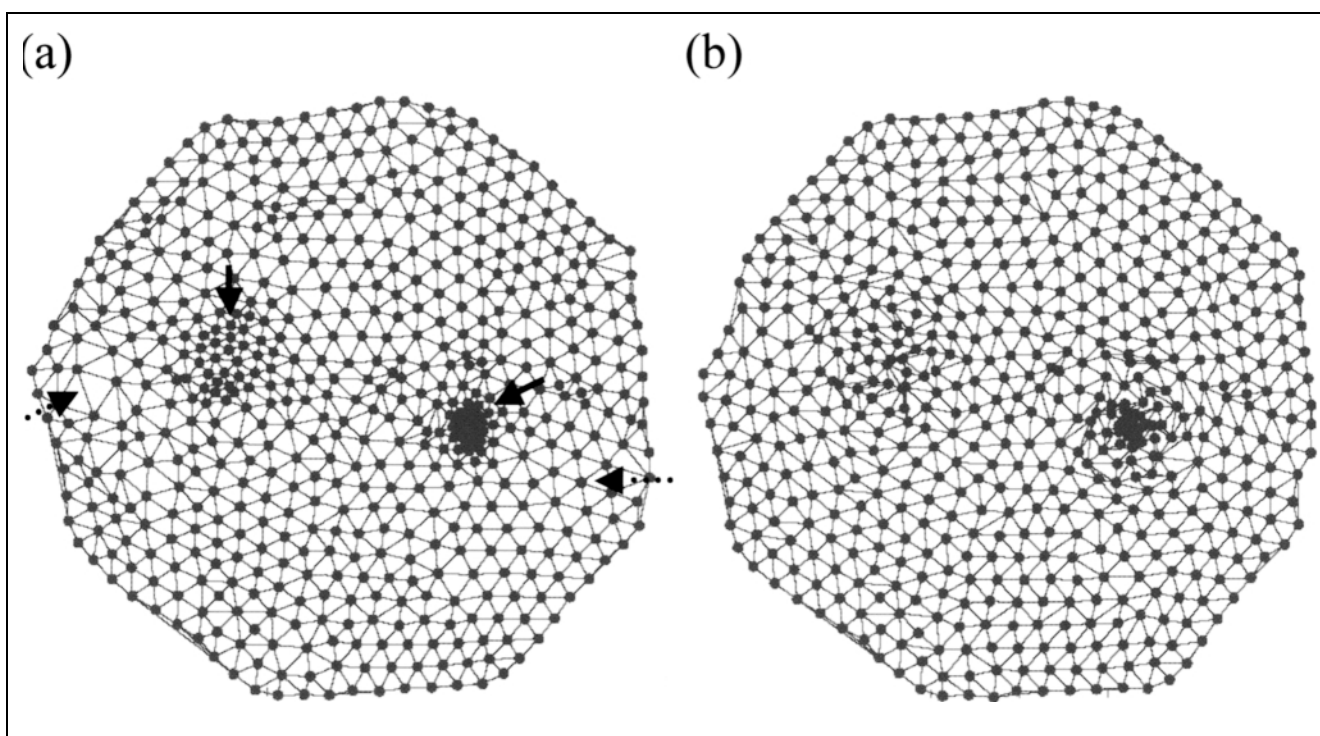
The planar positions are found by solving

$$\mathbf{X} = (\mathbf{I} - \mathbf{N})^{-1} \mathbf{P}\mathbf{X}_0$$

Hence, the initial positions,  $\mathbf{X}$ , are determined by the edges attached to each node (specified in matrices  $\mathbf{N}$  and  $\mathbf{P}$ ) and the initial positions of the perimeter points (specified in  $\mathbf{X}_0$ ). Placing the nodes and edges down in a method that preserves the mesh topology, and in a single step, is a computationally efficient beginning to the algorithm that could be widely applied.

### Distance Matching

The initial placement of the sample nodes meets the topological requirement we have imposed: The interior edges between nodes do not intersect. By preserving the graph planarity, the sample node placements do not introduce any twists in the surface representation. The edges do not comprise a triangulation. Hence, at this point, we replace the original edges with a new planar representation using the Delaunay triangulation



**Figure 11.** (a) The initial planar position of the sample nodes. A Delaunay triangulation of the nodes is also shown. The spacing of points in the high-density regions (solid arrows) and low-density regions (dashed arrows) are poor distance matches to the corresponding distances on the test surface. (b) The triangulated sample node positions after applying adjusting the node positions to improve the planar and cortical distance matches. Even after this procedure, density remains very high in the region near the ellipsoid. Throughout the representation, however, the topology is correct.

of the nodes. This set of edges preserves the planar graph and is efficient. The initial position of the sample nodes and the edges of the Delaunay triangulation are shown in Figure 11a.

While the sample node positions preserve the planar graph, the internode distances do not match the distances along the original surface. In the next phase of the algorithm, the positions of the sample node positions are adjusted so that their planar and cortical distances match. The adjustment is made subject to the constraint that no edge intersections are introduced. The adjustments are made by an iterative procedure that minimizes the difference between the cortical and planar distance between each node and its neighbors. The penalty function for evaluating a position change includes the squared error between the cortical and planar distances and a very large penalty in case a position causes two edges to cross. The positions of the sample nodes after adjusting to match distances are shown in Figure 11b. The picture shows that the overall density of sample nodes is more uniform, corresponding better to the original relatively uniform sampling.

### Position the Remaining Gray Matter Points

At this point, the major goals of the flattening algorithm have been achieved. The sample nodes are positioned

in the plane and the surface topology is preserved. All that remains is to position the other gray matter points within the plane. This is performed in two steps. First, we assign positions of the gray matter points adjacent to white matter. Each of these points falls within one of the triangles identified in the planar representation. We find this triangle and assign a planar location based on the cortical distance between that gray matter point and the three vertices of the triangle. Second, we assign positions of the gray matter points that are not adjacent to white matter. We assign these points the position of the closest gray matter point that is adjacent to the white matter. In this way, the thickness of the gray matter is suppressed and gray matter is represented as a thin sheet. This is not a requirement, of course, and we have experimented with representations that preserve some of the cortical thickness by using thin representations in which some points are allowed to float out of the plane.

### Acknowledgments

We thank Heidi Baseler, Geoffrey Boynton, Jon Demb, Robert Dougherty, David Fleet, David Heeger, William Press, Guillermo Sapiro, and Alex Wade for discussions and comments. Supported by NIH EY30164 and a McKnight Senior Fellowship.

Reprint requests should be sent to Dr. Brian Wandell, Psychology Department, Jordan Hall, Bldg. 420, Stanford University, Stanford, CA 94305. e-mail: wandell@stanford.edu.

## REFERENCES

- Bollobas, B. (1991). *Modern graph theory*. New York: Springer.
- Carman, G. J., Drury, H. A., & Van Essen, D. C. (1995). Computational methods for reconstructing and unfolding the cerebral cortex. *Cerebral Cortex*, 5, 506–517.
- Dale, A. M., Fischl, B., & Sereno, M. I. (1999). Cortical surface-based analysis: I. Segmentation and surface reconstruction. *Neuroimage*, 9, 179–194.
- Dale, A. M., & Sereno, M. I. (1993). Improved localization of cortical activity by combining EEG and MEG with MRI cortical surface reconstruction: A linear approach. *Journal of Cognitive Neuroscience*, 5, 162–176.
- Dijkstra, E. W. (1959). A note on two problems in connection with graphs. *Numerische Mathematik*, 1, 269–271.
- Drury, H. A., Essen, D. C. V., Anderson, C. H., Lee, C. W., Coogan, T. A., & Lewis, J. W. (1996). Computerized mappings of the cerebral cortex: A multiresolution flattening method and a surface-based coordinate system. *Journal of Cognitive Neuroscience*, 8, 1–28.
- Eck, M., DeRose, T., Duchamp, T., Hoppe, H., Lounsbery, M., & Stuetzle, W. (1995). Multiresolution analysis of arbitrary meshes. *Proceedings of SIGGRAPH* 26, 173–182.
- Engel, S. A., Glover, G. H., & Wandell, B. A. (1997). Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cerebral Cortex*, 7, 181–192.
- Fischl, B., Sereno, M. I., & Dale, A. M. (1999). Cortical surface-based analysis: II. Inflation, flattening, and a surface-based coordinate system. *Neuroimage*, 9, 195–207.
- Floater, M. S. (1997). Parameterization and smooth approximation of surface triangulations. *Computer-Aided Geometric Design*, 14, 231–250.
- Horn, B. (1986). *Robot vision*. Cambridge: MIT Press.
- Kimmel, R., & Kiryati, N. (1998). Computing geodesic paths on manifolds. *Proceedings of the National Academy of Science, U.S.A.*, 95, 8431–8435.
- Levy, B., & Mallet, J. L. (1998). Non-distorted texture mapping for sheared triangulated meshes. *SIGGRAPH '98*, 25, 343–352.
- Schwartz, E. L. (1990). Computer aided neuroanatomy of macaque visual cortex. In E. L. Schwartz (Ed.), *Computational neuroscience* (pp. 295–315). Cambridge: MIT Press.
- Talairach, J., & Tournoux, P. (1988). *Co-planar stereotaxic atlas of the human brain*. New York: Thieme.
- Teo, P., Sapiro, G., & Wandell, B. (1997). Creating connected representations of cortical gray matter for functional MRI visualization. *IEEE Transactions on Medical Imaging*, 16, 852–863.
- Tootell, R. B., Dale, A. M., Sereno, M. I., & Malach, R. (1996). New images from human visual cortex. *Trends in Neuroscience*, 19, 481–489.
- Tootell, R. B., Mendola, J. D., Hadjikhani, N. K., Ledden, P. J., Liu, A. K., Reppas, J. B., Sereno, M. I., & Dale, A. M. (1997). Functional analysis of V3A and related areas in human visual cortex. *Journal of Neuroscience*, 17, 7060–7078.
- Tutte, W. T. (1960). Convex representation of graphs. *Proceedings of the London Mathematical Society*, 10, 743–768.
- Van Essen, D. C., & Drury, H. A. (1997). Structural and functional analyses of human cerebral cortex using a surface-based atlas. *Journal of Neuroscience*, 17, 7079–7102.
- Van Essen, D. C., Drury, H. A., Joshi, S., & Miller, M. I. (1998). Functional and structural mapping of human cerebral cortex: Solutions are in the surfaces. *Proceedings of the National Academy Science, U.S.A.*, 95, 788–795.
- Wandell, B. A. (1999). Computational neuroimaging of human visual cortex. *Annual Review of Neuroscience*, 22, 145–173.
- Wandell, B. A., Engel, S. A., & Hel-Or, H. Z. (1996). Creating images of the flattened cortical sheet. *Investigative Ophthalmology and Visual Science*, 37, S1081.