

Automated Manifold Surgery: Constructing Geometrically Accurate and Topologically Correct Models of the Human Cerebral Cortex

Bruce Fischl*, Arthur Liu, and Anders M. Dale

Abstract—Highly accurate surface models of the cerebral cortex are becoming increasingly important as tools in the investigation of the functional organization of the human brain. The construction of such models is difficult using current neuroimaging technology due to the high degree of cortical folding. Even single voxel misclassifications can result in erroneous connections being created between adjacent banks of a sulcus, resulting in a topologically inaccurate model. These topological defects cause the cortical model to no longer be homeomorphic to a sheet, preventing the accurate inflation, flattening, or spherical morphing of the reconstructed cortex. Surface deformation techniques can guarantee the topological correctness of a model, but are time-consuming and may result in geometrically inaccurate models. In order to address this need we have developed a technique for taking a model of the cortex, detecting and fixing the topological defects while leaving that majority of the model intact, resulting in a surface that is both geometrically accurate and topologically correct.

Index Terms—Human cerebral cortex, topology, segmentation.

I. INTRODUCTION

THE cerebral cortex is the largest part of the human brain, and a structure that has been the subject of numerous neuroimaging studies [1]–[26]. Although the cortex is highly folded in many mammalian species, its intrinsic “unfolded” structure is that of a two-dimensional (2-D) sheet, several millimeters thick. It is well-accepted that many functional dimensions (e.g., retinotopy, orientation tuning, ocular dominance, somatotopy, tonotopy, etc.) are mapped on the cortical surface; and that these mapped parameters vary much more rapidly in the two dimensions parallel to the surface than they do through the several millimeters of cortical thickness (i.e., they are “columnar”).

The analysis of these cortical properties in the three-dimensional (3-D) embedding space in which imaging data is typically acquired suffers from serious drawbacks. These drawbacks derive directly from the fact that the intrinsic topology of the cere-

bral cortex is that of a 2-D sheet. For instance, estimates of the amount of “buried” cortex range from 60% to 70% [27], [28]. Thus, Cartesian distances measured in 3-D space between two points on the cortical surface will substantially underestimate the true distance along the cortical sheet, particularly in cases where the points lie on different banks of a sulcus.

For this reason, the analysis of cortical data is greatly facilitated by the use of accurate 2-D models of the cortical sheet [1], [28]–[34]. The construction of these models using current neuroimaging data is a difficult task due to the tradeoff between spatial resolution, field of view and signal-to-noise ratio (SNR). These factors typically constrain the spatial resolution of structural neuroimaging data to be on the order of a cubic millimeter or greater. At this resolution, even single voxel misclassifications can result in erroneous connections being generated across disparate parts of cortex due to the high degree of cortical folding. These erroneous connections can make the analysis of the functional topography of the cortex difficult or impossible.

Methods for the construction of cortical models can be broadly divided into two separate types—those that enforce a given topology [35]–[37] and those that do not [30], [31], [38]–[42]. The topology-enforcing techniques typically start with a surface of known topology (usually a supertessellated icosahedral approximation to a sphere¹) and deform it so that it lies on the cortical surface. These methods have the advantage of allowing the user to specify the proper topology and not allowing it to change. Unfortunately, the energy functionals that drive the deformation are highly nonconvex [43], resulting in a difficult and time-consuming numerical integration. Further, these methods often fail to accurately represent deep sulci [44]. In addition, small errors in the assignment of tissue labels to voxels that is the basis for the initial surface can result in large geometric inaccuracies in the final surface due to the global topology constraint. This can occur for example when an erroneous segmentation results in a bridge connecting two banks of a sulcus. In order to maintain the correct topology, the surface must “drape” over the incorrectly classified region. Finally, a more subtle drawback is that the surface representation is uniform on the initial (i.e., spherical) surface, whereas it is usually desirable to sample the target (i.e., cortical) surface uniformly.

Conversely, methods for generating models of the cortex that do not enforce a given topology can focus on accuracy of segmentation and, therefore, may result in more accurate models

Manuscript received December 15, 1999; revised October 18, 2000. This work was supported in part by the National Institute of Neurological Disorders and Stroke, the National Institute of Mental Health, and the National Cancer Institute under Grant R01-NS39581 and in part by the National Center for Research Resources under Grants P41-RR14075 and R01-RR13609. The Associate Editor responsible for coordinating the review of this paper and recommending its publication was T. Taxt. *Asterisk indicates corresponding author.*

*B. Fischl is with the Nuclear Magnetic Resonance Center, Massachusetts General Hospital, Harvard Medical School, Bldg. 149, 13th Street, Charlestown, MA 02129 USA (e-mail: fischl@nmr.mgh.harvard.edu).

A. Liu and A. M. Dale are the Nuclear Magnetic Resonance Center, Massachusetts General Hospital, Harvard Medical School, Charlestown, MA 02129 USA.

Publisher Item Identifier S 0278-0062(01)00796-0.

¹For a detailed discussion on the use of spherical versus planar topology, see [33].

than the topology-enforcing techniques. That is, the segmentation can be performed using local intensity, prior probabilities, and/or geometric information without regard to topology. Regions such as those described above give rise only to small geometric errors in the surface.² This is, however, at the cost of allowing topological defects in the form of incorrect surface connectivity. Furthermore, because the surface is no longer homeomorphic to a patch of \mathbb{R}^2 , the connectivity errors prevent it from being accurately inflated, flattened, or morphed into a spherical space to generate a surface-based coordinate system, techniques that are increasingly common in the computational neuroscience community [28], [32], [33], [37], [45], [46].

The difficulties encountered in the use of the surface-deformation techniques are due to the fact that in addition to inheriting the proper *topology* from the initial surface representation, these methods also use the initial geometry of this surface. In most cases this geometry is far smoother than the target surface. This problem is further complicated by the fact that much of the cortical surface area is buried deep within cortical sulci that are frequently quite narrow. Thus, large regions of the evolving surface must be driven into these narrow openings, then expand to cover the extensive cortical area within. The nonconvexity occurs because these vertices must be moved away from the cortical surface at the opening of the sulcus in order to eventually settle at the surface near the fundus [43].

Models of the cortical surface, and indeed most solid models in general, are typically instantiated as lists of triangles, edges and vertices in which each point on the surface lies in exactly one triangle.³ Such a structure is called a *polygonal tessellation*. The presence of topological defects in a polygonal tessellation can be easily detected by computing the *Euler number* of the tessellation. The Euler number is a topological invariant of a surface, and is given by $v - e + f$ where v , e , and f are the number of vertices, edges, and faces in the tessellation, respectively [47]. The number of defects in the surface is then given by $1 - (v - e + f)/2$. However, while the Euler number specifies the number of topological defects, it reveals nothing regarding either their location or their spatial scale, and is, therefore, of limited use for correcting the topology of a defective surface.

The defects arise due to incorrect surface connectivity that results in the generation of two (or more) qualitatively different paths between a pair of points on the surface—paths that cannot be continuously deformed into each other.⁴ In order to correct the topology and remove the defect, one would need to identify these paths, then choose to retain one and discard the other. The former issue, that of identifying the presence and location of topologically inconsistent paths on the surface, is a purely topological one. In contrast, the latter issue, the arbitration among alternative paths, must be based solely on geometric information, as the two paths are topologically indistinguishable. Thus, for example, a bridge between two banks of a sulcus is topo-

logically equivalent to a small perforation in a thin white matter sheet. In the first case, we wish to “cut” the bridge, while in the second we wish to “fill” the hole, and it is this decision that must be based on geometric information, as will be discussed in Section II-D.

In order to resolve these problems and combine the advantages of the two classes of surface reconstruction methods, we have developed a technique that takes a representation of the cortex and alters its topology to conform to that of a reference surface. The method leaves the surface geometry essentially unchanged, except in regions containing topological defects. This technique has a number of advantages. Like the deformable surface algorithms, the resulting model is guaranteed to be topologically correct, allowing it to be accurately inflated, flattened, and morphed. However, in separating the specification of topology from the classification of tissue classes, it allows a more accurate segmentation to be performed without regard for topological correctness. Furthermore, since the alternative paths making up a topological defect are explicitly identified, geometric information from the remainder of the surface can be used in the correction of the defect. Finally, most of the initial tessellation of the cortex is retained. This is advantageous, as the cortical tessellation is uniform on the surface of the cortex as opposed to the spherical or planar space in which the surface deformation methods are uniform, a property that is important for accurately representing the functional and structural properties of the cortex.

Further, given reasonable quality data, no manual interaction is required in order to generate surfaces that are suitable for visualization or establishing surface-based coordinate systems using the proposed technique. Manual editing may result in a slight increase in surface accuracy, mainly in noncortical regions.⁵ Finally, the models generated with these techniques have been shown to accurately represent the pial and white matter surfaces to within 1/2 mm [48].

II. METHODS

The topology correction procedure begins with a previously constructed model of the cortical surface. All the results shown in this paper use a cortical surface reconstruction procedure that produces a highly detailed geometric description of the gray–white matter boundary, as well as the pial surface of the human cortex, as described in [30], [31], and [48]. It is important to note that the details of the segmentation and initial surface reconstruction are completely separate from that of ensuring the correct topology. Thus, other segmentation/reconstruction procedures can be employed in a modular fashion in place of those used here.

Briefly, the surface reconstruction procedure we employ begins with a high-resolution T1-weighted image, and removes intensity variations due to low-frequency bias fields. The skull is then removed by shrink-wrapping a deformable ellipsoidal template onto the inner-boundary of the skull. Next, a binary seg-

²Here, we use the term “geometric accuracy” to mean a model that accurately follows the folds of the actual cortex. In these terms, a bridge across adjacent banks of a sulcus is a “small” inaccuracy, although the errors introduced into surface geodesics can be substantial.

³Note that degenerate cases exist in which a point falls precisely on leg of a triangle, thus either lying within two (or more!) triangles. Computationally, this situation never arises.

⁴Assuming the deformation lies within the surface and the endpoints are fixed.

⁵For example, studies of the thickness of the cortex require the surface representations to be placed with subvoxel precision. In this type of study, we typically manually inspect and edit the large defects due to anatomical departures from spherical topology, such as the ventricles and the basal ganglia, as well as a small number of cortical locations.

mentation procedure is applied to the skull-stripped image such that each voxel is assigned to one of two classes: white matter or not white matter. An automated algorithm then finds seed points in each cerebral hemisphere and automatically disconnects the corpus callosum in order to generate two connected masses of voxels, each of which represents a cortical hemisphere. The exterior of each connected mass is covered with vertices, edges and triangles resulting in an initial surface representation of each cortical hemisphere.

The surface deformation methods described earlier seek to find a mapping N from an initial surface S (typically a sphere, although patches of the plane have been used as well), to the cortical surface C

$$N: S \rightarrow C. \quad (1)$$

The necessary and sufficient condition for C to have the proper topology is that the mapping N be a homeomorphism, that is a continuous bijection [49].⁶ The difficulty of finding such an N is that the target surface is significantly less smooth than the initial one. Thus, large geometric deformations must be introduced into the initial surface in order to generate an accurate model of the cortex. In order to avoid this difficulty, we propose to find a mapping M

$$M: C \rightarrow S. \quad (2)$$

This is a significantly easier task as S is much smoother than C . If M is a homeomorphism, then the two surfaces are topologically equivalent and no further action is required. Typically, however, the surface model will not be homeomorphic to the target surface (e.g., a sphere). The topology correction then consists of finding regions in which the inverse of the candidate M is multivalued and correcting it, thus constructing a homeomorphism. This is accomplished by discarding the surface tessellation in these regions and generating a new, topologically correct tessellation, resulting in a surface with the proper topology. More specifically, the algorithm proceeds as follows (the times in parentheses were obtained running the algorithm on a typical dataset on a 500-MHz Pentium III):

- 1) Find a mapping M from the original surface onto a surface with the desired topology (a sphere in this case), such that M is a homeomorphism on as much of the surface as possible. We call such an M a *quasi-homeomorphic map*. (1 h)
- 2) Find all regions on the target surface S in which the inverse mapping M^{-1} is multivalued. These are detectable as regions in which edges in the tessellation intersect. Mark all vertices in a face with an intersecting edge as ambiguous. (4 min)
- 3) Segment the ambiguous vertices into connected regions. (<1 min)
- 4) Discard the tessellation in each defect, and generate a new tessellation so that no edge intersection occurs in the

⁶Note that the while the connectivity is typically specified explicitly, topological correctness in the embedding space requires that the surface deformation be constrained to prevent the surface from intersecting itself. Techniques for preventing self-intersection are typically quite computationally intensive and are reviewed in [43].

target surface. This particular tessellation chosen will be one that optimizes a metric measure in the *source* surface. Thus, the topology is specified by the target surface, but the geometry by the source surface. (15 min)⁷

A. Spherical Inflation

The initial mapping of the cortical surface to that of a sphere could be accomplished by simply projecting each point on the folded surface to the closest point on the sphere. However, doing so would result in large regions of the initial mapping being non-homeomorphic (i.e., substantial folds on the spherical surface), causing the subsequent numerical integration to be time consuming and prone to local minima. Instead, we use a simple procedure to unfold and smooth the folded cortical surface so that it approaches that of a sphere whose origin is the centroid of the initial surface. After applying this procedure, the surface can then be projected with relatively few folds. The algorithm consists of iteratively updating the position of each vertex based on a smoothness force \mathbf{F}_S , and a radial spherical force \mathbf{F}_R

$$\mathbf{x}_k(t+1) = \mathbf{x}_k(t) + \mathbf{F}_S(t) + \lambda_R \mathbf{F}_R(t) \quad (3)$$

where \mathbf{x}_k is the position of the k th vertex at iteration number t and the smoothness force \mathbf{F}_S is given by

$$\mathbf{F}_S = \frac{1}{N_k} \sum_{j \in N_k} (\mathbf{x}_j - \mathbf{x}_k) - \frac{1}{V} \sum_i^V \sum_{j \in N_i} (\mathbf{n}_i \mathbf{n}'_i) \cdot (\mathbf{x}_j - \mathbf{x}_i) \quad (4)$$

where

N_k	set of vertices neighboring the k th vertex;
V	number of vertices in the tessellation;
\mathbf{n}_k and \mathbf{n}'_k	surface normal at location k and its transpose, respectively. ⁸

The smoothness term \mathbf{F}_S moves each vertex in the direction of the centroid of its neighbors, while projecting out the average inwards movement this creates over the entire surface.

The radial term simply drives each vertex toward the surface of a sphere with the desired radius R

$$\mathbf{F}_R = (\mathbf{R}_k - \mathbf{x}_k) \quad (5)$$

where \mathbf{R}_k is the radial projection of \mathbf{x}_k onto the sphere with radius R . Note that we have omitted the functional dependence on the iteration number t in (4) and (5) to avoid notational clutter. We use an R on the order of 100 mm as this results in a sphere with about the same total surface area as an average cortex, and a λ_R of 0.25 to allow sufficient smoothing to take place during the spherical inflation. Once the inflation has converged, the surface is projected so it lies precisely on the surface of a sphere of radius R .

Fig. 1 illustrates this process, showing a sequence of lateral views of the spherical inflation process. Sulcal regions are col-

⁷The computational complexity of step 4) is quadratic in the number of vertices contained within the convex hull of each defect. Minimal manual editing to remove large defects caused by anatomical departures from spherical topology due to structures such as the fornix or the basal ganglia can reduce this to a few minutes.

⁸Each vertex is assigned a surface normal that is the normalized average of the surface normals of all the triangles that it is a member of.




Fig. 1. Example of initial transformation of folded cortical surface into spherical form via simultaneous smoothing and radial inflation. From left to right: initial cortical surface (gray/white interface), and surface after 25, 50, 75, 100 and 300 iterations of (3)–(5). The white arrow indicates a large fold caused by a topological defect.

ored dark gray, and gyral light gray (based on mean curvature). Note that by the end of the inflation process the majority of cortex has been unfolded, with the bulk of the remaining folds due to topological defects. An example of this can be seen in Fig. 1 in the form of the large fold at the anterior end of the sylvian fissure (see white arrow at the right). This fold is due to a hole through the basal ganglia, which creates one of the largest defects on the surface.

B. Quasi-Homeomorphic Mapping


Once the initial spherical configuration has been established, we wish to modify the mapping in order to construct a homeomorphism. This procedure consists of two separate steps. In the first step, a mapping M_1 is found from the folded surface to a sphere that is maximally homeomorphic, which we term a *quasi-homeomorphic* mapping. Next, regions of M_1^{-1} that are multivalued are detected and discarded, and a new mapping is established in these regions, resulting in a global homeomorphism.

In generating the mapping M_1 , we are only concerned with its topological properties, that is, we wish M_1 to be as close to a homeomorphism as possible. A mapping M_1 is a homeomorphism if the determinant of the Jacobian matrix of M_1 is non-singular, and the mapping itself is continuous. This is of course the multidimensional analog of monotonicity. To construct the mapping, we minimize an energy functional that directly penalizes regions in which the determinant becomes zero or negative, thus encouraging positive definiteness.⁹ Note that this is the only term in the energy functional—no preservation of metric properties is needed.


⁹We could equivalently construct a negative definite mapping.

The determinant of the Jacobian yields a measure of what happens to an oriented area element under the mapping. With this in mind, we use an energy functional based on the compression/expansion each face experiences under the mapping. This compression/expansion is transformed by a nonlinearity designed to discount the influence of positive definite regions. More specifically, if the initial area on the folded surface C of the i th face in the tessellation is A_i^0 , and the area on the spherical surface S at time t of the numerical integration is A_i^t , then the energy functional we minimize is given by where F is the number of faces in the tessellation, and k is a positive constant. The logarithmic nonlinearity limits the penalization of compression primarily to negative semi-definite regions, as can be seen in the plot to the right of (6) at the bottom of the page, which shows E over the interval $[-0.1, 0.1]$. Computing the derivative of E with respect to a change in the area of the i th face yields the derivative of E is simply a sigmoid (shown at the right, again over the interval $[-0.1, 0.1]$), the slope of which is defined by k . The extent to which highly compressed positive definite regions are penalized is thus determined by k . In the limit of large k , only negative semi-definite regions are penalized. In practice, a somewhat smaller k (e.g., $k = 100$) results in a mapping with fewer folds, as extending the energy functional slightly into positive regions encourages vertices in these areas to move away from folds, giving vertices in negative definite regions spherical surface area to move into. This sigmoid nonlinearity in the areal term is now used in flattening [32], maximally isometric spherical mapping [32], and spherical morphing of cortical surfaces into an atlas [33]. In all of these contexts, we have found that it provides a strong “regularizing” effect, limiting the solution space to positive definite mappings where possible.

$$E = \sum_{i=1}^F \left(\frac{\log(1 + e^{kR_i})}{k} - R_i \right), R_i = \frac{A_i^t}{A_i^0} \quad (6)$$



$$\frac{\partial E}{\partial A_i^t} = \frac{1}{(1 + e^{-kR_i})A_i^0} \quad (7)$$



In order to complete the definition of the topology term of the energy functional, we consider the i th triangle in the surface tessellation depicted in Fig. 2, with unit normal vector \mathbf{n}_i , and edges \mathbf{a}_i and \mathbf{b}_i connecting the vertex \mathbf{x}_i to two of its neighbors (note that bold-faced symbols denote vector quantities). The unit normal \mathbf{n}_i is given on the original manifold by the normalized cross product of the edges \mathbf{a}_i and \mathbf{b}_i , while the area of the triangle is half the cross product of \mathbf{a}_i and \mathbf{b}_i dotted with the unit normal (i.e., the triple scalar product). In the spherical representation, the normal vector field can be given a consistent orientation on the surface¹⁰ using the embedding space, and A_i becomes an *oriented* area, which may take on negative values indicating folds in the surface. We choose an outward pointing normal vector on the surface of the sphere, given by

$$\mathbf{n}_i = \frac{\mathbf{x}_i}{\|\mathbf{x}_i\|}. \quad (8)$$

In order to derive a rule for modifying the vertex positions, we compute the directional derivative of E with respect to the position of the k th vertex

$$\frac{\partial E}{\partial \mathbf{x}_k} = \frac{\partial E}{\partial A_i^t} \frac{\partial A_i^t}{\partial \mathbf{x}_k}. \quad (9)$$

The first factor is given by (7), shown at the bottom of the previous page. The second is the change in the area of the i th triangle caused by moving the k th vertex, which can be computed from the prior description of the metric properties of the tessellation using the chain rule as

$$\frac{\partial A_i^t}{\partial \mathbf{x}_k} = \frac{\partial A_i^t}{\partial \mathbf{a}_i} \frac{\partial \mathbf{a}_i}{\partial \mathbf{x}_k} + \frac{\partial A_i^t}{\partial \mathbf{b}_i} \frac{\partial \mathbf{b}_i}{\partial \mathbf{x}_k}, \quad \frac{\partial A_i^t}{\partial \mathbf{a}_i} = \mathbf{b}_i \times \mathbf{n}_i, \quad \frac{\partial A_i^t}{\partial \mathbf{b}_i} = \mathbf{n}_i \times \mathbf{a}_i. \quad (10)$$

The partials of the change in the legs with respect to a change in the vertex position are dependent on what position the vertex in question occupies in a given triangle (refer to Fig. 2 for subscript meaning)

$$\frac{\partial \mathbf{a}_i}{\partial \mathbf{x}_k} = \begin{cases} [-1, -1, -1]^T, & k = i \\ [1, 1, 1], & k = l \\ 0, & \text{otherwise} \end{cases} \quad \frac{\partial \mathbf{b}_i}{\partial \mathbf{x}_k} = \begin{cases} [-1, -1, -1]^T, & k = i \\ [1, 1, 1], & k = l \\ 0, & \text{otherwise.} \end{cases} \quad (11)$$

The mapping M is modified by numerically minimizing (6), employing the gradient given by (7)–(10) and the numerical integration scheme described in [32]. After each iteration of the numerical minimization, the surface is projected back onto the sphere, obviating any need for additional terms in the energy functional to keep the surface from leaving the sphere. This results in a new mapping M_1 that is typically positive definite across more than 99% of the cortical surface area. Thus, the majority of the initial surface reconstruction can be retained with only a small percentage of the cortical surface requiring further attention. This is illustrated in Fig. 3(a) which displays the sum-squared error $[E$ in (6)] and Fig. 3(b) the percentage of the

¹⁰This is always possible except in pathological cases such as the Möbius strip that are said to be nonorientable [50].




Fig. 2. Metric properties of the triangular tessellation.




Fig. 3. (a) Plot of sum-squared error and (b) percent negative area during minimization of (6).

surface area that is folded on the spheret, each plotted against iteration number of the numerical minimization of E for a typical surface.

C. Detection of Surface Defects

After the quasi-homeomorphic mapping M_1 has been constructed, the spherical surface is examined for regions of non-invertibility, as these are areas where the current tessellation must be discarded and a new one constructed in order to ensure the proper topology. The noninvertible regions are characterizable as portions of the sphere in which more than one triangle overlap, and can thus be detected by checking for intersecting edges in the tessellation. The overlap detection procedure utilizes a spatial lookup table (LUT) that lists the faces passing through each 1 mm^3 voxel, making the procedure quite rapid (less than 4 min). After the spatial LUT has been constructed, each edge in the tessellation is checked for intersection with all the edges in every face that passes through any of the voxels that




Fig. 4. (a) Multivalued regions of the quasi-homeomorphic mapping painted onto the original folded surface as well as (b) an inflated surface of the same brain as Fig. 1 for better visualization. The arrows indicate a defect across the posterior and anterior banks of the central sulcus shown in more detail in Fig. 5.




Fig. 5. Close-up of defective region indicated by arrows in Fig. 4. White coloring indicates region of ambiguity. The white bar at the bottom of (a) is a scale bar indicating 1 cm.

the candidate edge passes through. Defective vertices are then defined as those vertices that are part of a face in which any of the edges are in the set of intersecting edges. The defective vertices are then partitioned into connected components for further processing, with each connected component representing a defective region that will require retessellation.

The results of this procedure are shown in Fig. 4, with the noninvertible regions painted bright white on a folded [Fig. 4(a)] and inflated [Fig. 4(b)] hemisphere (the inflated representation is useful for visualization purposes as it exposes sulcal regions. Details of this technique are given in [32]. The bulk of these ambiguous regions are the product of small topological defects the causes of which are difficult to find in the segmented volume, as they are frequently the result of incorrect surface connectivity that does not lie in any of the three cardinal planes. In fact, although visual inspection only reveals a handful of defects, the

actual Euler number of this surface is -101 , indicating a total of 51 topological defects in the polygonal representation.

The arrows in Fig. 4 indicate a defect that spans the banks of the central sulcus, a close-up of which is shown in Fig. 5. Note that the region of ambiguity encompasses two different paths between the points on the posterior and anterior banks of the central sulcus, paths that cannot be continuously deformed from one to the other. This erroneous connectivity changes the geodesic distance between the two endpoints from the correct surface distance of 5 or 6 centimeters to a few millimeters—an error of over an order of magnitude, a point to which we will return in Section II-D.

Fig. 6 shows the same defect (as indicated by the red arrow) in (a) coronal, (b) sagitta, and (c) horizontal slices through the volume, with the surface intersection overlaid in yellow. As can be seen here and in the prior figures, the defect arises due to




Fig. 6. Sagittal, coronal, and horizontal views of a T1 volume with the surface intersection overlaid in yellow. The red arrows indicate the location of a defect in the region of the central sulcus.




Fig. 7. Schematic of defect types: (a) a “handle”, and (b) a “hole”.

a small misclassification at a particularly narrow point in the sulcus.

D. Surface Retessellation

The labeled defective regions on the surface include incorrect pieces of surface that give rise to erroneous connectivity and, hence, must be removed, as well as regions of true cortical surface that should be retained. While these regions were found using *topological* criteria, namely edge intersection in the spherical space, the decision of what to keep and what to discard must be made on a *geometric* basis. This decision corresponds to the arbitration between alternative, incompatible paths discussed in Section I. Explicitly finding these discrete alternative paths is a difficult problem. Instead, we solve the problem implicitly, through the manner in which the defective surface regions are retessellated. The retessellation algorithm we use is a greedy one [51] in which an ordered list is created of all possible edges in each defect and its border, with the shortest edge that does

not intersect with the existing tessellation being added at each step.¹¹ After all possible edges have been added; two triangular faces are generated for each edge, completing the polygonal representation of the surface. This then gives a topologically correct tessellation of the patch of \mathbb{R}^2 in the defect [51] and, hence, preserves the spherical topology of the entire closed surface.

The decision as to what to keep and what to discard is made implicitly via the ordering of the edges. That is, if the edges in one path are added to the tessellation prior to an incompatible path, the edges in the second path will cause intersections with the (already added) edges from the first path, and will, therefore, be discarded. Thus, the choice of what metric to use to generate the ordering is a critical one, as it embodies the arbitration process.

In practice, two types of defects occur—which we term “holes” and “handles.” Holes consist of small perforations

¹¹Only the convex hull of the defect and its border must be checked to prohibit edge intersection.




Fig. 8. Close-up view of a surface defect during smoothing: (a) 0, (b) 30, and (c) 60 iterations of surface smoothing. The white bar at the bottom is a scale bar indicating 1 cm.

in planar sheets of white matter, while handles are bridges between nonadjacent points in cortex. A schematic example of each type of defect is given in Fig. 7, which shows a “hole” [Fig. 7(a)] and a “handle” [Fig. 7(b)]. Each defect of either type will reduce the Euler number by two, indicating their topological equivalence. The essential distinguishing characteristic between these two situations is the spatial scale of the hole, or, equivalently, whether the surface along the borders of the hole or the surface on the bridge is more in agreement with the remainder of the surface. That is, the hole is an essentially one-dimensional (1-D) tunnel through a 2-D sheet, while the handle is a 1-D bridge between adjacent 2-D sheets. It is these geometric features of the defect that we exploit in order to generate the proper ordering.

The most obvious choice for a metric to us for ordering the edges is the 3-D Cartesian metric in the space in which the folded surface is embedded. However, in general, the use of a simple 3-D metric to generate the edge ordering will fail. The reason for this can be seen in Fig. 7, in which the distance between points 1 and 2 in the hole shown in Fig. 7(b) may be as short or shorter than the distance between points 1 and 3, as sheets of white matter are frequently only a single millimeter thick (which is what causes perforations in the first place!). A similar caveat applies to the handle defect illustrated in Fig. 7(a): the bridge from point 1 to point 2 may be shorter than the distance from point 1 to point 3. In fact, in general, the defective edges will be short, as the regions in which defects occur are precisely those in which sheets of white matter are very thin, or opposing banks of a sulcus pass in close proximity to one another. These problems with the use of a 3-D metric for generating the edge ordering arise because the 3-D metric incurs the same drawback in this context as it does in others [33]—it does not respect the geometry of the surface. That is, distances in 3-D will in general be poor approximations of surface geodesics. Thus, a retessellation algorithm based on 3-D distances will fail in these situations, causing sulci to be “sewn shut.”

What is needed is a distance metric that is governed by the nondefective parts of the surface—that is, something approximating surface geodesics that are only allowed to pass through the “correct” surface. However, since this is precisely the dis-

tingtion we are hoping to clarify, a somewhat different approach is required. Another characteristic of the topological defects is that they create surface paths which give rise to dramatically different shortest path lengths between points on opposite ends of the defect.¹² In order to detect the defect, we wish to modify the surface in such a way as to make this fact more apparent. Toward that end, we design a simple smoothing procedure similar to the one used in [30] and [32] for visualization purposes, as well as the smoothing term in the spherical inflation described in Section II-A. This technique models the surface tessellation as a network of springs with zero resting length. Allowing the surface to evolve under this model smooths the surface by moving each vertex in the direction of the centroid of its neighbors. In defective regions this results in a pronounced stretching as the bulk of the surface pulls the two endpoints of the defect apart, exploiting the essentially 1-D nature of the defect.

An example of this process is illustrated in Fig. 8, which shows snapshots of the smoothing procedure in a region of the central sulcus defect. Note the stretching of the edges in the defect relative to those in the remainder of the tessellation in Fig. 8(c). Ordering the edges in inverse order of their length on this smoothed surface results in the proper retessellation, shown in Fig. 9, with Fig. 10 depicting the corrected surface overlaid onto orthogonal slices.

More specifically, the retessellation algorithm proceeds as follows. First, the tessellation in defective areas is discarded, including all edges and faces in the original polygonal representation contained within defective areas, as well as all vertices that lie within of regions in which the determinant of the Jacobian of the spherical mapping is nonpositive. The vertices in negative semi-definite regions typically comprise less than 1% of the tessellation, and must be removed to ensure that the final mapping is a homeomorphism. An edge list is then generated containing an entry for every pairwise combination of vertices in the defect and its convex hull. Next, the smoothing algorithm is applied to the surface, and the edge list is ordered by Cartesian distance on the smoothed surface (the smoothed surface is similar to the one shown in Fig. 4). The ordered edges are then

¹²In cases where this is not true, the defect is extremely local and any retessellation will be satisfactory.

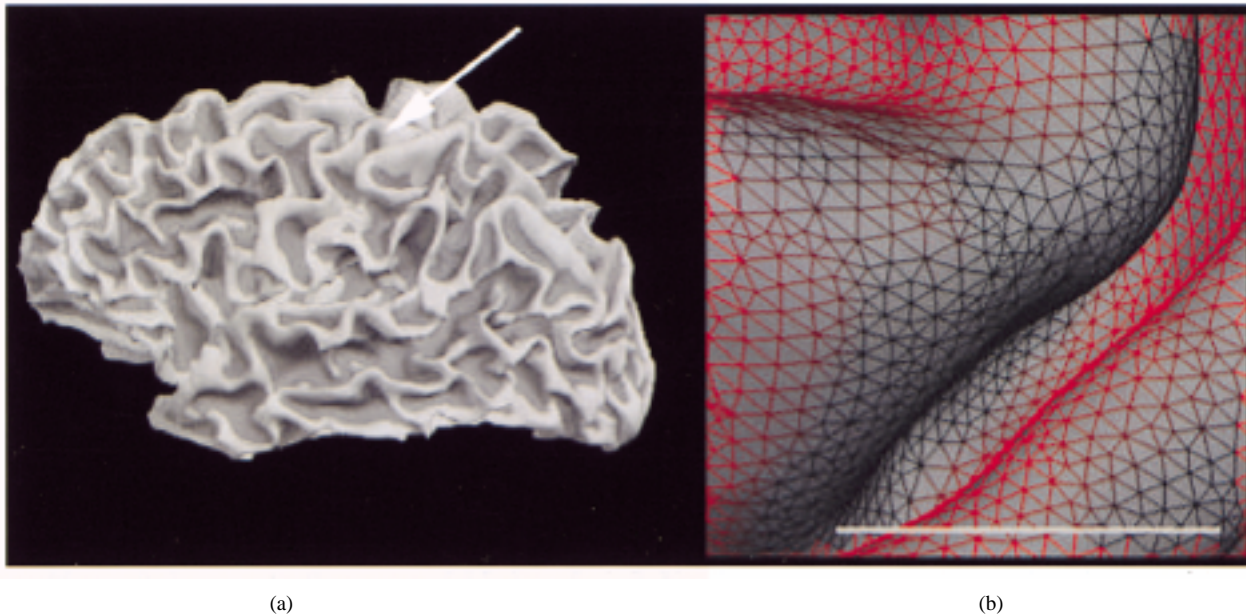


Fig. 9. (a) Topologically correct folded surface with a white arrow indicating the previous location of the defect and (b) close-up of same tessellation in corrected region.




Fig. 10. Topologically correct surface (after fixing) overlaid on T1 volume.

sequentially added, without allowing intersection between each candidate edge and the convex hull of the defect and its border in the spherical space, yielding a topologically correct tessellation of a patch of \mathbb{R}^2 . After all possible edges have been added; each pair of vertices that share an edge is examined for common neighbors, with one triangle being generated for each common neighbor that does not lie within another putative triangle. This results in two triangular faces being generated for each edge, completing the polygonal representation of the surface.

This procedure is illustrated in Figs. 8, 9, and 10, which show a close-up view of the central sulcus surface defect with the original tessellation overlaid in red. Note that the lengths of the edges crossing the sulcus (i.e., the defect) in the original surface shown in Fig. 8(a) are approximately equal to those covering its banks. Thus, a retessellation algorithm based on 3-D distances will fail in this situation, causing the sulcus to

be “sewn shut.” The smoothing process is illustrated in Fig. 8(b) and (c), which shows snapshots of this process in a region of the defect. Note the stretching of the edges in the defect relative to those in the remainder of the tessellation in Fig. 8(c). Ordering the edges in inverse order of their length on this smoothed surface and retessellating in the spherical space results in the proper retessellation, shown in Fig. 9, with Fig. 10 depicting the corrected surface overlaid onto orthogonal slices.

III. CONCLUSION

Models of the human cerebral cortex are important in a variety of contexts, including visualization of functional and structural neuroimaging data, computational modeling of cortical function, as well as statistical analysis of cortical properties. In order for a model of the cortex to be useful in all of these

domains it must be geometrically accurate as well as topologically correct. Unfortunately, many methods for constructing such models sacrifice one of these properties in order to maximize the other. In this paper we have presented a technique that preserves geometric accuracy while at the same time insuring topological correctness.

Topological correctness and geometric accuracy are important for a number of reasons. Incorrect topology manifests itself through spurious connections between potentially disparate parts of cortex. These connections are problematic in that they result in geodesic distances between cortical locations being erroneously by potentially large amounts. Further, such improper connectivity prevents accurate inflation for visualization purposes, or transformation to spherical or flattened formats, as the surface or pieces of it are no longer homeomorphic to a sphere or a sheet, respectively.

Geometric accuracy is necessary for correct and complete representation of the cortex. Techniques that enforce topology can result in surface models that underestimate the true cortical surface area, particularly in narrow deep sulci. These regions are difficult to accurately reconstruct using deformable surface models, as they require the surface to pass through local minima in the deformation energy functional. Using the topology correction procedure outlined in this paper, defects caused by narrow openings in sulci are easily corrected, as they give rise to surface geodesics that typically differ from distances along the majority of the surface by more than an order of magnitude.

Furthermore, the topology correction procedure detailed in this paper can be used as a preprocessing step for a deformable surface algorithm. This may be useful in the context of multiple surface deformations in which the pial and white matter surfaces are deformed simultaneously, each giving clues to the location of the other [35]. Initializing the deformation procedure with surfaces that are correctly positioned over the vast majority of the cortex should result in considerably faster and more robust surface generation.

Another point to note is that while this technique has been used in the context of enforcing spherical topology onto a geometrically accurate model of the cortex, it is applicable to enforcing the topology of any regular orientable manifold onto another such manifold. For example, portions of the cortex should be topologically equivalent to a sheet. This topology could be enforced on cortical patches by applying the topology correction procedure in the plane.

Finally, we have reconstructed over 200 cortical hemispheres using the topology correction technique. Our experiments indicate that the vast majority of surfaces require no manual intervention for use in visualization or establishing surface-based coordinate systems. The software used in the generation of these surfaces is freely available,¹³ and it is our hope that these tools will make surface-based analysis a routine part of neuroimaging studies of functional and structural properties of the human cerebral cortex.

ACKNOWLEDGMENT

The authors would like to thank R. Desikan, E. Busa, M. Glessner and D. Salat for extensive testing of the proposed algorithms.

REFERENCES

- [1] M. I. Sereno, A. M. Dale, J. B. Reppas, K. K. Kwong, J. W. Belliveau, T. J. Brady, B. R. Rosen, and R. B. H. Tootell, "Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging," *Science*, vol. 268, pp. 889–893, May 12, 1995.
- [2] R. B. H. Tootell, K. K. Kwong, J. W. Belliveau, J. R. Baker, C. E. Stern, S. J. Hockfield, H. C. Breiter, R. Born, R. Benson, T. J. Brady, and B. R. Rosen, "Mapping human visual cortex: Evidence from functional MRI and histology," *Investigat. Ophthalmol. Visual Sci.*, vol. 34, no. 4, p. 813, 1993.
- [3] T. D. Griffiths, G. Rees, A. Rees, G. G. R. Green, C. Witton, D. Rowe, C. Buchel, R. Turner, and R. S. J. Frackowiak, "Right parietal cortex is involved in the perception of sound movement in humans," *Nature Neurosci.*, vol. 1, pp. 74–79, 1998.
- [4] G. McCarthy, A. M. Blamire, D. L. Rothman, R. Gruetter, and R. G. Schulman, "Echo-planar magnetic resonance imaging studies of frontal cortex activation during word generation in humans," in *Proc. Nat. Acad. Sci.*, vol. 90, 1993, pp. 4952–4956.
- [5] R. M. Hinkle, X. P. Hu, A. E. Stillman, S. G. Kim, H. Merkle, R. Salmi, and K. Ugurbil, "Functional magnetic resonance imaging of broca area during internal speech," *Neuroreport*, vol. 4, no. 6, pp. 675–678, 1993.
- [6] J. Watson, R. Myers, R. Frackowiak, J. Hajnal, R. Woods, J. Mazziotta, S. Shipp, and S. Zeki, "Area V5 of the human brain: Evidence from a combined study using positron emission tomography and magnetic resonance imaging," *Cerebral Cortex*, vol. 3, pp. 79–94, 1993.
- [7] R. Turner, P. Jezzard, H. Wen, K. Kwong, D. Le Bihan, and R. Balaban, "Functional mapping of the human visual cortex at 4 Tesla using oxygen contrast EPI," presented at the Soc. Magnetic Resonance in Medicine 11th Annu. Meeting, Berlin, Germany, 1992.
- [8] R. B. H. Tootell, N. K. Hadjikhani, W. Vanduffel, A. K. Liu, J. D. Mendola, M. I. Sereno, and A. M. Dale, "Functional analysis of primary visual cortex (V1) in humans," *PNAS*, vol. 95, pp. 811–817, 1998.
- [9] R. Turner, P. Jezzard, H. Wen, K. K. Kwong, D. Le Bihan, T. Zeffiro, and R. Balaban, "Functional mapping of the human visual cortex at 4 and 1.5 Tesla using deoxygenation contrast EPI," *Magn. Reson. Med.*, vol. 29, pp. 277–279, 1993.
- [10] L. G. Ungerleider, "Functional brain imaging studies of cortical mechanisms for memory," *Science*, vol. 270, pp. 769–775, 1995.
- [11] L. G. Ungerleider and J. V. Haxby, "'What' and 'where' in the human brain. [Review] [86 refs]," *Current Opinion Neurobiol.*, vol. 4, no. 2, pp. 157–165, 1994.
- [12] A. Karni, G. Meyer, P. Jezzard, M. M. Adams, R. Turner, and L. G. Ungerleider, "Functional MRI evidence for adult motor cortex plasticity during motor skill learning," *Nature*, vol. 377, no. 6545, pp. 155–158, 1995.
- [13] E. Halgren, A. M. Dale, M. I. Sereno, R. B. H. Tootell, K. Marinkovic, and B. R. Rosen, "Location of human face-selective cortex with respect to retinotopic areas," *Human Brain Mapping*, vol. 7, pp. 29–37, 1999.
- [14] N. K. Hadjikhani, A. K. Liu, A. M. Dale, P. Cavanagh, and R. B. H. Tootell, "Retinotopy and color sensitivity in human visual cortical area V8," *Nature Neurosci.*, vol. 1, pp. 235–241, 1998.
- [15] P. T. Fox, M. A. Mintun, M. E. Raichle, F. M. Miezin, J. M. Allman, and D. C. Van Essen, "Mapping human visual cortex with positron emission tomography," *Nature*, vol. 323, no. 6091, pp. 806–809, 1986.
- [16] P. T. Fox, F. M. Miezin, J. M. Allman, D. C. Van Essen, and M. E. Raichle, "Retinotopic organization of human visual cortex mapped with positron-emission tomography," *J. Neurosci.*, vol. 7, no. 3, pp. 913–922, 1987.
- [17] P. T. Fox, H. Burton, and M. E. Raichle, "Mapping human somatosensory cortex with positron emission tomography," *J. Neurosurg.*, vol. 67, no. 1, pp. 34–43, 1987.
- [18] S. A. Engel, G. H. Glover, and B. A. Wandell, "Retinotopic organization in human visual cortex and the spatial precision of functional MRI," *Cerebral cortex*, vol. 7, no. 2, pp. 181–192, 1997.
- [19] E. A. DeYoe, G. J. Carman, P. Bandettini, S. Glickman, J. Wieser, R. Cox, D. Miller, and J. Neitz, "Mapping striate and extrastriate visual areas in human cerebral cortex," in *Proc. Nat. Acad. Sci. (USA)*, vol. 93, 1996, pp. 2382–2386.

¹³surfer.nmr.mgh.harvard.edu.

- [20] A. Dale, E. Halgren, J. Lewine, R. Buckner, K. Paulson, K. Marinkovic, and B. Rosen, "Spatiotemporal localization of cortical word repetition effects in a size-judgment task using combined fMRI/MEG," *NeuroImage*, vol. 5, no. 4, 1997.
- [21] J. C. Culham, S. A. Brandt, P. Cavanagh, N. G. Kanwisher, A. M. Dale, and R. B. H. Tootell, "Cortical fMRI activation produced by attentive tracking of moving targets," *J. Neurophysiol.*, vol. 80, pp. 2657–2670, 1998.
- [22] J. W. Belliveau, K. K. Kwong, D. N. Kennedy, J. R. Baker, C. E. Stern, R. Benson, D. A. Chesler, R. M. Weisskoff, M. S. Cohen, R. B. H. Tootell, P. T. Fox, T. J. Brady, and B. R. Rosen, "Magnetic resonance imaging mapping of brain function: Human visual cortex," *Invest. Radiol.*, vol. 27, pp. S59–S65, 1992.
- [23] A. K. Liu, J. W. Belliveau, and A. M. Dale, "Spatiotemporal imaging of human brain activity using fMRI constrained MEG data: Monte Carlo simulations," in *Proc. Natl. Acad. Sci. (USA)*, vol. 95, 1998, pp. 8945–8950.
- [24] C. I. Moore, C. E. Stern, S. Corkin, B. Fischl, A. C. Gray, B. R. Rosen, and A. M. Dale, "Segregation of somatosensory activation in the human Rolandic cortex using fMRI," *J. Neurophysiol.*, vol. 84, pp. 558–569, 2000.
- [25] D. L. Schacter, R. L. Buckner, W. Koutstaal, A. M. Dale, and B. R. Rosen, "Late onset of anterior prefrontal activity during true and false recognition: An event-related fMRI study," *Neuroimage*, vol. 6, pp. 258–269, 1997.
- [26] J. D. Mendola, A. M. Dale, B. Fischl, A. K. Liu, and R. B. H. Tootell, "The representation of illusory and real contours in human cortical visual areas revealed by functional magnetic resonance imaging," *J. Neurosci.*, vol. 19, no. 19, pp. 8560–8572, 1999.
- [27] K. Zilles, E. Armstrong, A. Schleicher, and H.-J. Kretschmann, "The human pattern of gyrification in the cerebral cortex," *Anat. Embryol.*, vol. 179, pp. 173–179, 1988.
- [28] D. C. Van Essen and H. A. Drury, "Structural and functional analysis of human cerebral cortex using a surface-based atlas," *J. Neurosci.*, vol. 17, no. 18, pp. 7079–7102, 1997.
- [29] E. L. Schwartz, A. Shaw, and E. Wolfson, "A numerical solution to the generalized mapmaker's problem: Flattening nonconvex polyhedral surfaces," *IEEE Trans. Pattern Anal. Machine Intell.*, pp. 1005–1008, Nov. 1989.
- [30] A. M. Dale and M. I. Sereno, "Improved localization of cortical activity by combining EEG and MEG with MRI cortical surface reconstruction: A linear approach," *J. Cogn. Neurosci.*, vol. 5, no. 2, pp. 162–176, 1993.
- [31] A. M. Dale, B. Fischl, and M. I. Sereno, "Cortical surface-based analysis I: Segmentation and surface reconstruction," *NeuroImage*, vol. 9, pp. 179–194, 1999.
- [32] B. Fischl, M. I. Sereno, and A. M. Dale, "Cortical surface-based analysis II: Inflation, flattening, a surface-based coordinate system," *NeuroImage*, vol. 9, pp. 195–207, 1999.
- [33] B. Fischl, M. I. Sereno, R. B. H. Tootell, and A. M. Dale, "High-resolution inter-subject averaging and a coordinate system for the cortical surface," *Human Brain Mapping*, vol. 8, no. 4, pp. 272–284, 1999.
- [34] D. C. Van Essen, H. A. Drury, S. Joshi, and M. I. Miller, "Functional and structural mapping of human cerebral cortex: Solutions are in the surfaces," in *Proc. Nat. Acad. Sci.*, vol. 95, 1998, pp. 788–795.
- [35] D. MacDonald, N. Kabani, D. Avis, and A. C. Evans, "Automated 3-D extraction of inner and outer surfaces of cerebral cortex from MRI," *NeuroImage*, vol. 12, pp. 340–356, 2000.
- [36] C. Davatzikos and R. N. Bryan, "Using a deformable surface model to obtain a shape representation of the cortex," *IEEE Trans. Med. Imag.*, vol. 15, pp. 785–795, Dec. 1996.
- [37] C. Davatzikos, "Spatial transformation and registration of brain images using elastically deformable models," *Comput. Vision. Image Understanding*, vol. 66, no. 2, pp. 207–222, 1997.
- [38] R. Kikinis, M. E. Shenton, D. V. Iosifescu, R. W. McCarley, P. Saiviroonporn, H. H. Hokama, A. Robotino, D. Metcalf, C. G. Wible, C. M. Portas, R. Donnino, and F. Jolesz, "A digital brain atlas for surgical planning, model-driven segmentation, and teaching," *IEEE Trans. Visual. Comput. Graphics*, vol. 2, pp. 232–241, 1996.
- [39] P. C. Teo, G. Sapiro, and B. A. Wandell, "Creating connected representations of cortical gray matter for functional MRI visualization," *IEEE Trans. Med. Imag.*, vol. 16, pp. 852–863, Dec. 1997.
- [40] W. Wells, W. Grimson, R. Kikinis, and F. Jolesz, "Adaptive segmentation of MRI data," *IEEE Trans. Med. Imag.*, vol. 15, pp. 429–442, Aug. 1996.
- [41] T. Kapur, "Model based segmentation of medical images," Ph.D. dissertation, MIT AI Lab., Massachusetts Inst. Technol., Cambridge, MA, 1999.
- [42] C. Xu, D. L. Pham, M. E. Rettmann, D. N. Yu, and J. L. Prince, "Reconstruct of the human cerebral cortex from magnetic resonance images," *IEEE Trans. Med. Imag.*, vol. 18, p. 467, June 1999.
- [43] D. MacDonald, "A method for identifying geometrically simple surfaces from three-dimensional images," Ph.D. dissertation, Montreal Neurological Inst., McGill Univ., Montreal, P.Q., Canada, 1998.
- [44] A. Manceaux-Demiau, R. N. Bryan, and C. Davatzikos, "A probabilistic ribbon model for shape analysis of the cerebral sulci: Application to the central sulcus," *J. Comput. Assist. Tomogr.*, vol. 22, no. 6, pp. 962–971, 1998.
- [45] P. Thompson, R. Woods, M. Mega, and A. Toga, "Mathematical/computational challenges in creating deformable and probabilistic atlases of the human brain," *Human Brain Mapping*, vol. 9, no. 2, pp. 81–92, 2000.
- [46] H. A. Drury, D. C. Van Essen, C. H. Anderson, C. W. Lee, T. A. Coogan, and J. W. Lewis, "Computerized mappings of the cerebral cortex: A multiresolution flattening method and a surface-based coordinate system," *J. Cogn. Neurosci.*, vol. 8, no. 1, pp. 1–28, 1996.
- [47] B. Bollobás, *Graph Theory, An Introductory Course*. New York: Springer-Verlag, 1979.
- [48] B. Fischl and A. M. Dale, "Measuring the thickness of the human cerebral cortex from magnetic resonance images," in *Proc. Nat. Acad. Sci.*, 2000, pp. 11044–11049.
- [49] J. Dixmier, "General topology," in *Undergraduate Texts in Mathematics*. New York: Springer-Verlag, 1984.
- [50] M. do Carmo, *Differential Geometry of Curves and Surfaces*. Englewood Cliffs, NJ: Prentice-Hall, 1976.
- [51] F. P. Preparata and M. I. Shamos, *Computational Geometry: An Introduction*. Berlin, Germany: Springer-Verlag, 1985.