Correspondence

A Proof of the Spherical Homeomorphism Conjecture for Surfaces

Lowell Abrams, Donniell E. Fishkind, and Carey E. Priebe*

Abstract—The human cerebral cortex is topologically equivalent to a sphere when it is viewed as closed at the brain stem. Due to noise and/or resolution issues, magnetic resonance imaging may see "handles" that need to be eliminated to reflect the true spherical topology. Shattuck and Leahy [2] present an algorithm to correct such an image. The basis for their correction strategy is a conjecture, which they call the spherical homeomorphism conjecture, stating that the boundary between the foreground region and the background region is topologically spherical if certain associated foreground and background multigraphs are both graph-theoretic trees. In this paper, we prove the conjecture, and its converse, under the assumption that the foreground/background boundary is a surface.

Index Terms—Magnetic resonance imaging, segmentation, topology, topological correction.

I. INTRODUCTION

The human cerebral cortex is topologically equivalent to a sphere when it is viewed as closed at the brain stem. Due to noise and/or resolution issues, magnetic resonance (MR) imaging may see "handles" that need to be eliminated to reflect the true spherical topology.

Shattuck and Leahy [2] present an algorithm to correct such an image. The basis for their correction strategy is a conjecture, which they call the spherical homeomorphism conjecture, stating that the boundary between the foreground region and the background region is topologically spherical if certain associated foreground and background multigraphs are both graph-theoretic trees.

In this paper, we prove the conjecture, and its converse, under the assumption that the foreground/background boundary is a surface. Note that if the foreground/background boundary is not a surface, then it cannot be topologically spherical.

II. DEFINITIONS

For positive integer N, we consider the subset of Euclidean 3-space $\mathcal{I} := \{(x, y, z) \in \mathbf{R}^3 : 0 \leq x, y, z \leq N\}$, endowed with the usual Euclidean metric and topology. For all N^3 triplets of indexes $(i, j, k) \in \{1, 2, \ldots, N\}^3$ we define the i, j, kth *voxel* to be the subset of $\mathcal{I}$ given by

$$v_{ijk} := \{(x, y, z) \in \mathcal{I} : $$
$$i - 1 \leq x \leq i, j - 1 \leq y \leq j, k - 1 \leq z \leq k\}.$$

The index i is the *height* of the voxel v_{ijk}. The *level* L_i is the union of all voxels of height i. The *sheet* $S_{i, i+1}$ is the intersection $L_i \cap L_{i+1}$.

Fig. 1.　Implications of ∂FG being a surface.

Each voxel is a cube and, thus, its boundary has six faces, called *voxel faces*. If the intersection of two distinct voxel faces is a line segment, we call the intersection a *voxel edge*. If the intersection of two voxel edges is a point, we call the intersection a *voxel vertex*. For any set $A \subset \mathbf{R}^3$, we denote the boundary of A by ∂A. If $A \subset S$ for some sheet S, we denote the boundary of A in S by $\partial_S(A)$. Note that ∂A is two-dimensional and $\partial_S(A)$ is one-dimensional.

Each of the voxels in $\mathcal{I}$ is classified as either *foreground* or *background*. For this MR brain imaging application, the segmentation process classifies a voxel as foreground if the corresponding location in the brain consists of tissue which is interior to the cerebral cortex; otherwise it classifies the voxel as background. Define FG as the union of all foreground voxels and BG as the union of all background voxels. We make the following two important assumptions: First, we assume that ∂FG is topologically equivalent to an orientable surface, which includes the assumption that ∂FG is connected. For this brain imaging application, the surface ∂FG is understood to be an approximate representation of the cerebral cortex itself. Thus, in principle, it should be topologically spherical, but may be of higher genus due to noise and/or resolution issues. The second assumption is that foreground is surrounded by background, i.e., $\partial \mathcal{I} \cap \partial FG = \emptyset$. In other words, the foreground image is strictly contained in $\mathcal{I}$.

For each i, each connected component of $L_i \cap FG$ is called a *foreground vertex* of height i. If two foreground vertices of adjacent heights have a nonempty intersection then each connected component of the intersection (which lies in $S_{i, i+1}$) is called a *foreground edge*. (There may be more than one foreground edge between two foreground vertices.) Denote by V_f the set of all foreground vertices of all heights, and denote by E_f the set of all foreground edges. The *foreground graph* G_f is the (multi)graph (V_f, E_f) in which each foreground edge ϵ in the intersection of foreground vertices ω_1 and ω_2, say, is considered to be a graph-theoretic edge whose two endpoints are ω_1, ω_2. Thus, when we speak of a foreground vertex or a foreground edge we not only think of them as subsets of $\mathcal{I}$, but also as graph-theoretic objects, and the perspective in use at any point of this manuscript will be clear from its context. Analogously, define *background vertices*, *background edges*, V_b (the set of all background vertices), E_b (the set of all background edges), and the *background graph* $G_b = (V_b, E_b)$ by substituting BG for FG in the foreground definitions.

The assumption that ∂FG is a surface yields a number of useful implication: 1) If foreground voxels u and v intersect at precisely voxel edge e [see Fig. 1(a)], then at least one voxel a of the other two voxels

sharing e is foreground. 2) If foreground voxels u and v intersect precisely at voxel vertex x[see Fig. 1(b)], then among the other six voxels sharing x there are foreground voxels a and b such that each of $u \bigcap a$, $a \bigcap b$, and $b \bigcap v$ are voxel faces. (Analogs of these two observations hold, *mutatis mutandis*, for the background as well.) 3) The boundary of each foreground or background vertex is an orientable surface. 4) For each foreground and background edge ϵ in, say, sheet S we have that $\partial_S(\epsilon)$ consists of (pairwise disjoint) simple closed curves.[1] 5) For each i and each foreground or background vertex ω in L_i or L_{i+1}, the region $\omega \cap S_{i,\,i+1}$ is connected and its boundary in $S_{i,\,i+1}$ is a disjoint union of simple closed curves. 6) The graphs G_f and G_b are both connected.

Formally, the foreground and background graphs in [2] are constructed differently than is done here. There, the foreground graph is modified using a "special connectivity rule," and adjacency for background vertices is defined using the "D18" rule (which considers sharing of a voxel edge an adjacency). However, under the assumption that ∂FG is a surface, the observations in the previous paragraph show that the "special connectivity rule" never applies, and that the D18 rule defines an adjacency exactly when the construction here does. Thus, when ∂FG is a surface, the foreground and background graphs in [2] are exactly G_f and G_b as constructed here.

For all i, let $|V_{f_i}|, |V_{b_i}|, |E_{f_i,\,i+1}|$, and $|E_{b_i,\,i+1}|$ denote the number of foreground vertices of height i, background vertices of height i, foreground edges with endpoints of heights i and $i+1$, and background edges with endpoints of heights i and $i+1$, respectively.

III. MAIN RESULT

Shattuck and Leahy [2] conjectured that if both their foreground and background multigraphs are trees, then ∂FG is topologically equivalent to a sphere. Our result regarding their conjecture is the following theorem, which restricts attention to the case in which ∂FG is a surface.[2]

Theorem 1: Suppose that ∂FG is a surface. Then ∂FG is topologically equivalent to a sphere if and only if both G_f and G_b are trees.

The graph G_f (G_b) is a tree if and only if $|E_f| = |V_f| - 1$ ($|E_b| = |V_b| - 1$). By Theorem 1 any extra edges beyond this amount indicates that ∂FG is not topologically spherical. In general, there is not necessarily a way to determine the genus of ∂FG exactly from the number of excess edges, but there is one special case where this can be done:

Theorem 2: Suppose that ∂FG is a surface and that $\forall\,\omega \in V_f$ the genus of $\partial\omega$ is zero. Then the genus of ∂FG is $|E_f| - |V_f| + 1$. Thus, ∂FG is topologically equivalent to a sphere if and only if G_f is a tree.

Note that the assumption that the genus of $\partial\omega$ is zero $\forall\,\omega \in V_f$ is equivalent to the assumption that G_b is precisely a path of length N.

For any two-cell embedding of a graph on an orientable surface of genus g with vertex-set V, edge-set E and face-set F, a classical result of Euler and Poincaré gives the relationship

$$\chi(S) := |V| - |E| + |F| = 2 - 2g \tag{1}$$

(see, e.g., [3, p. 268]). This quantity $\chi(S)$ is called the *Euler characteristic* of the surface. If A is the union of any collection of voxels

then ∂A has a natural two-cell embedding with the voxel vertices, edges, and faces on the boundary ∂A playing the role of graph-theoretic vertices, edges, and faces, respectively. The Euler characteristic $\chi(\partial A) := |V| - |E| + |F|$, where V, E, and F denote the sets of voxel vertices, edges, and faces on ∂A, is a well-defined topological invariant even if A is not a surface.

One approach to determining the genus of ∂FG is to determine its Euler characteristic. This may be done on a local basis, by summing $\chi(\partial\omega)$ over all $\omega \in V_f$ and then adjusting for the voxel vertices, edges, and faces that were overcounted. An overcount occurs when a voxel face is on two different foreground vertex boundaries, hence is in a foreground edge and is not on ∂FG. Such a voxel face needs to be subtracted away twice, once for each of the two times it was (erroneously) counted. Note that if ϵ is a foreground edge in sheet S, then all the voxel edges and vertices that are not in $\partial_S(\epsilon)$ were overcounted twice, and the voxel edges and vertices in $\partial_S(\epsilon)$ were overcounted once since they are in ∂FG (once) but were counted twice.

If ϵ is a foreground edge in sheet S with n_ϵ voxel vertices, e_ϵ voxel edges, and f_ϵ voxel faces, the net amount that must be added to the global Euler characteristic count is

$$\chi'(\epsilon) := -2(n_\epsilon - e_\epsilon + f_\epsilon). \tag{2}$$

This is because the components of $\partial_S(\epsilon)$ are simple, closed curves and, thus, they have exactly as many voxel edges as voxel vertices; since contributions from edges and vertices have opposite sign in (2), the fact that we subtract them one more time (that is, the second time) has no net effect.

Indeed, if we set $\chi'_{i,\,i+1} := \sum_{\epsilon \in E_{f_i,\,i+1}} \chi'(\epsilon)$, then we have

$$\chi(\partial FG) = \sum_{\omega \in V_f} \chi(\partial\omega) + \sum_{\epsilon \in E_f} \chi'(\epsilon)$$

$$= \sum_{\omega \in V_f} \chi(\partial\omega) + \sum_{i=1}^{N-1} \chi'_{i,\,i+1}. \tag{3}$$

The boundary of each foreground vertex is a surface, and the handles of these foreground-vertex-surfaces are in one-to-one correspondence with the set of all background vertices except for the N background vertices that intersect ∂I. Thus, we have

$$\sum_{\omega \in V_f} \chi(\partial\omega) = 2|V_f| - 2|V_b| + 2N \tag{4}$$

which gives us the first summand in right-hand side of (3).

To get the second summand in right-hand side of (3), we need two lemmas. Let $T_{i,\,i+1}$ denote the complement in $S_{i,\,i+1}$ of the union of foreground edges, and let $t_{i,\,i+1}$ denote the number of connected regions of $T_{i,\,i+1}$.

Lemma 3: For all i, $\chi'_{i,\,i+1} = -2 - 2|E_{f_i,\,i+1}| + 2t_{i,\,i+1}$.

Proof of Lemma 3: Consider $S_{i,\,i+1}$ with its embedded voxel vertices, edges, and faces. Just for this proof, disregard all voxel edges and vertices not contained in a foreground edge, and merge each of the $t_{i,\,i+1}$ connected components A of $T_{i,\,i+1}$ to a single face without any embedded voxel vertices, edges, or faces in the interior of A. The remaining foreground-edge voxel vertices, edges, and faces and the regions of $T_{i,\,i+1}$ together constitute a two-cell embedding in $S_{i,\,i+1}$ of a planar graph G'. The faces of the embedding are precisely the foreground-edge voxel faces and the connected components of $T_{i,\,i+1}$, and the number of components of G' is exactly $|E_{f_i,\,i+1}|$. Euler's for-

---

[1]Each voxel vertex y in S intersects (at most) four voxel faces of S. Consideration of cases, together with 1) and 2) above, shows that y intersects exactly zero or two voxel edges of $\partial_S(\epsilon)$. Thus, viewing voxel vertices and voxel edges as graph theoretic vertices and edges, respectively, we see that $\partial_S(\epsilon)$ consists of disjoint cycles.

[2]Notice that in the topological sense ∂FG is obtained from a disjoint union of spheres (i.e., boundaries of voxels) via finitely many direct sums (adjacencies of voxel faces) and point identifications (adjacencies of voxel edges or vertices). Thus, if ∂FG is not a surface, then it is either not connected or, topologically, it arises from a surface via point identifications; in neither case is it topologically equivalent to a sphere.

mula shows that $|\{\text{vertices of } G'\}| - |\{\text{ edges of } G'\}| + |\{\text{faces of } G'\}| = 1 + |E_{f_{i,i+1}}|$. This yields

$$\chi'_{i,i+1} \equiv \sum_{\epsilon \in E_{f_{i,i+1}}} \chi'(\epsilon) \stackrel{(2)}{=} -2 \sum_{\epsilon \in E_{f_{i,i+1}}} (n_\epsilon - e_\epsilon + f_\epsilon) \qquad (5)$$

$$= -2 \left(t_{i,i+1} + \sum_{\epsilon \in E_{f_{i,i+1}}} (n_\epsilon - e_\epsilon + f_\epsilon) \right) + 2t_{i,i+1} \qquad (6)$$

$$= -2 \left(1 + |E_{f_{i,i+1}}| \right) + 2t_{i,i+1} \qquad (7)$$

which was to be shown. $\square$

The following lemma provides a bound for one of the terms in the expression for $\chi'_{i,i+1}$ obtained in Lemma 3.

Lemma 4: For all i, $t_{i,i+1} \geq |V_{b_i}| + |V_{b_{i+1}}| - |E_{b_{i,i+1}}|$.

Proof of Lemma 4: Let H denote the [bipartite] subgraph of G_b induced by $V_{b_i} \cup V_{b_{i+1}}$, and let π denote the orthogonal projection of $L_i \cup L_{i+1}$ on $S_{i,i+1}$. Every path in H between vertices ω_1, ω_2 of H is mapped by π into $T_{i,i+1}$, and every path in $T_{i,i+1}$ from $\pi(\omega_1)$ to $\pi(\omega_2)$ can be lifted to a path in H from ω_1 to ω_2. Thus, H has exactly $t_{i,i+1}$ components. For each of these the number of vertices minus the number of edges is at most 1. Summing over each of the $t_{i,i+1}$ graph-theoretic components of H yields the assertion of the lemma.$\square$

We are now in a position to prove the main results.

Proof of Theorem 1: The characteristic of ∂FG may now be computed and bounded

$$\chi(\partial FG) \stackrel{(3)}{=} \sum_{\omega \in V_f} \chi(\partial\omega) + \sum_{i=1}^{N-1} \chi'_{i,i+1} \qquad (8)$$

$$\stackrel{(4), \mathbf{Lm3}}{=} 2|V_f| - 2|V_b| + 2N$$
$$+ \sum_{i=1}^{N-1} \left(-2 - 2|E_{f_{i,i+1}}| + 2t_{i,i+1} \right) \qquad (9)$$

$$\stackrel{\mathbf{Lm4}}{\geq} 2|V_f| - 2|V_b| + 2 + \sum_{i=1}^{N-1} \left(2|V_{b_i}| + 2|V_{b_{i+1}}| \right.$$
$$\left. - 2|E_{f_{i,i+1}}| - 2|E_{b_{i,i+1}}| \right) \qquad (10)$$

$$\stackrel{|V_{b_1}| = |V_{b_N}| = 1}{=} 2 - 2 \left[(|E_f| - |V_f| + 1) \right.$$
$$\left. + (|E_b| - |V_b| + 1) \right]. \qquad (11)$$

If G_f and G_b are both trees, then $|E_f| - |V_f| + 1 = 0$ and $|E_b| - |V_b| + 1 = 0$ and, thus, by (8)–(11), we have $\chi(\partial FG) \geq 2$. Since the genus of ∂FG is at least 0, it follows that $\chi(\partial FG) = 2$, which means precisely that ∂FG is topologically equivalent to a sphere.

To prove the converse, suppose without loss of generality that G_f contains a cycle C. Let l_f be a closed curve in FG that "realizes" C, i.e., l_f lies in the interior of the union of the vertices of C, and traverses those vertices in the [unoriented] cyclic order given by C. Without loss of generality, we may assume that the intersection of l_f

with any foreground edge from C consists of exactly one point, and also that each intersection of l_f with a sheet is in a foreground edge. Let S be a sheet which intersects at least one edge of C. The intersection $S \bigcap l_f$ is a disjoint union of finitely many points $p_1, \ldots, p_n$ and, because each p_i lies in a distinct foreground edge ϵ_i, any path from p_i to p_j, $i \neq j$, must pass through BG. There is necessarily an m such that the "unbounded" component of the complement of ϵ_m in S contains all other ϵ_i and, thus, there exists a simple closed curve l_b in $S \bigcap BG$ such that one region of $S - l_b$ contains p_m, and all other points p_i lie in the other region. It is clear that l_f and l_b, as closed curves in $\mathbf{R}^3$, are linked and, thus, l_f and l_b are not contractible in FG and BG, respectively. This contradicts the generalized Schoenflies theorem [1], which ensures that both regions of the complement of any topological sphere "nicely" embedded in $\mathbf{R}^3$ are simply connected (all loops can be contracted to a point). It follow that ∂FG is not a sphere. $\square$

Proof of Theorem 2: The assumption on genus implies that for all $\omega \in V_f$, $\partial\omega$ is a topological sphere and, thus, $\chi(\partial\omega) = 2$. If two foreground vertices ω_1 and ω_2 intersecting in sheet S have exactly one foreground edge ϵ between them, then $\partial_S(\epsilon)$ consists of exactly one simple closed curve that is contractible to a point. Hence, $\partial(\omega_1 \bigcup \omega_2)$ is also topologically spherical. Thus, we must have $\chi'(\epsilon) = -2$ in order to balance $\chi(\partial\omega_1) + \chi(\partial\omega_2) + \chi'(\epsilon) = \chi(\partial(\omega_1 \bigcup \omega_2))$. If there are multiple edges, then it follows from the fact that $\chi'(\epsilon)$ is calculated locally that each edge $\epsilon \in E_f$ has $\chi'(\epsilon) = -2$. Thus, by (3), we have

$$\chi(\partial FG) = \sum_{\omega \in V_f} \chi(\partial\omega) + \sum_{\epsilon \in E_f} \chi'(\epsilon) \qquad (12)$$

$$= 2|V_f| - 2|E_f| \qquad (13)$$

$$= 2 - 2(|E_f| - |V_f| + 1). \qquad (14)$$

It is clear from (14) and the Euler–Poincaré relationship that the genus of ∂FG is, thus, $|E_f| - |V_f| + 1$, and this quantity is zero precisely when G_f is a tree, since G_f is connected. $\square$

Referee comments suggested an alternative, more geometric argument for the proof of Theorem 2: The assumption on genus implies that all ω in V_f are topologically three-balls. Since each foreground edge ϵ is contractible, gluing the ω together along the edges ϵ yields a subspace which deformation-retracts to a complex that is graph-isomorphic to G_f.

REFERENCES

[1] M. Brown, "A proof of the generalized Schoenflies theorem," *Bull. Amer. Math. Soc.*, vol. 66, pp. 74–76, 1960.

[2] D. W. Shattuck and R. M. Leahy, "Automated graph-based analysis and correction of cortical volume topology," *IEEE Trans. Med. Imag.*, vol. 20, pp. 1167–1177, Nov. 2001.

[3] D. B. West, *Introduction to Graph Theory*, 2nd ed.   Upper Saddle River, NJ: Prentice-Hall, 2001.