From *SIAM News*, Volume 37, Number 7, September 2004

***Imaging Science***

# Topologists Take Scalpel to Brain Scans

*By Dana Mackenzie*

In recent years, magnetic resonance imaging has revolutionized the ability of doctors to understand the organization and the functioning of the human brain. But if you look at a brain scan close up—really, really close—you'll see some things that nature never intended. Because of glitches in the digital imaging process, nonexistent connec-tions can appear between adjacent folds of the cerebral cortex, or even between the two hemispheres (see Figure 1).

A few years ago, seeking to banish the unwanted isthmuses and handles, two electrical engineers at the University of Southern California proposed an automated airbrushing method. The corrected image would be a two-dimensional surface with no handles and no holes—or, as a topologist would put it, it would be homeomorphic to a sphere. (If, in other words, you imagine inflating the image like a balloon, it will expand into a sphere, not an inner tube or some more exotic shape.)

There was one problem: David Shattuck and Richard Leahy couldn't prove that their method worked, although in many tests on images with millions of voxels ("leading to numerous combinations and types of handles," Shattuck says), they never encountered a problem. In 2001 they published the algorithm in the *IEEE Transactions on Medical Imaging*, leaving the proof of its validity (which they called the "spherical homeomorphism conjecture") as an open challenge.

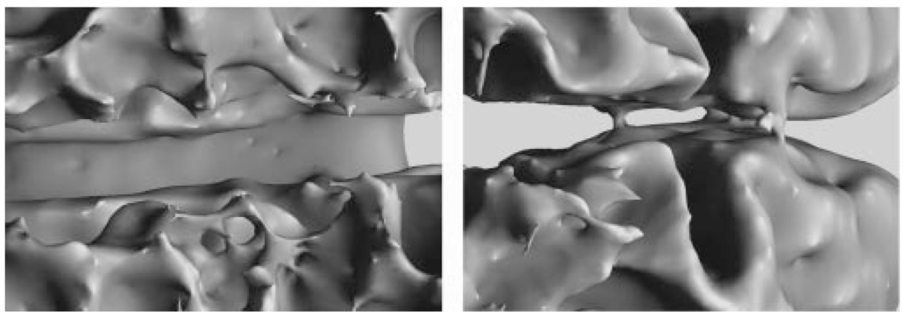

**Figure 1.** *A typical reconstruction of the cortical surface from a three-dimensional MRI scan contains numerous topological defects, which show up as "handles" (left) or tubes (right) connecting the left and right hemispheres. (These are artifacts—such structures do not exist in the brain.) Courtesy of David Shattuck and Richard Leahy, USC.*

Now, a team of three mathematicians have proved themselves equal to the challenge. What started out as a practical problem in brain imaging has turned out to be a beautiful exercise in theoretical mathematics, blending aspects of topology and graph theory. But because the problem and its solution appeared in an engineering journal, it's one that many mathematicians might have missed. "I'm sure they don't read medical imaging journals," says Johns Hopkins statistician Carey Priebe, who headed the group. Priebe's collaborators on the project were graph theorist Donniell Fishkind, also of Johns Hopkins University, and topologist Lowell Abrams of George Washington University.

### The Perils of Three-Dimensionality

Correcting a three-dimensional image is vastly more difficult than fixing a two-dimensional one. For a simple two-dimensional image, consisting only of white "foreground" pixels and black "background" pixels, the foreground topology wll change only when a black speck appears among the white pixels, as in a low-quality fax. According to medical imaging specialist Jerry Prince of Johns Hopkins University, it is "trivial" to write a computer program that will scan a two-dimensional image for specks and white them out.

But the topological errors that crop up in a three-dimensional image are more subtle. A tube of spurious black voxels (the 3-D equivalent of pixels) could burrow through the foreground, or a handle of erroneous white voxels could snake into the background. These defects are hard to detect because the artifacts no longer form black islands in a sea of white, or vice versa. The best fix may also be unclear, although a good rule of thumb is to change as few voxels as possible. Still, this criterion can be ambiguous: Correction of a handle, for example, could be done either by filling in the middle (creating a sort oflozenge-shaped projection) or by cutting it off.

What makes the problem even more difficult is the sheer number of topological defects in a typical MR image. According to Shattuck, a brain scan with a million voxels will usually have several hundred handles. Most of them can be eliminated by changing only one or two voxels, but a few of them may require altering 10 or more. "I've done it by hand, and it took me a few days," Shattuck says. "That definitely was part of my motivation to solve this problem."

Shattuck and Leahy begin their attack on the killer handles by segmenting the MR scan into horizontal slices, each one layer of voxels thick. Each slice contains a certain number of foreground islands; some of the islands may contain lakes, which in turn create islands of background voxels.

Next, they construct two networks, or graphs, that represent the connections between adjacent layers of voxels. One graph corresponds to the foreground, the other to the background. Each contiguous group of voxels in the foreground layer corresponds to one node of the graph. Similarly, each island in the background layer corresponds to one node of the background graph. (In addition, the surrounding sea of background counts as one node.) Whenever an island of voxels in one layer touches an island in the layer immediately below it, Shattuck and Leahy connect the two corresponding nodes in the graph. If the islands abut in more than one place, the two nodes can be joined by more than one line.

A simple example shows how Shattuck and Leahy's method detects handles. The image shown at the upper left in Figure 2 is topologically a torus, or a sphere with one handle. When the torus is aligned vertically, the top layer of the foreground has one island of four voxels. The second layer has two islands of one voxel each, which are illustrated by the two dots in the second row of the foreground graph ($G_f$). Similarly, the third row of the graph has two islands, but in the fourth layer they merge into one again. The edges in the foreground graph form a loop, which indicates the presence of a handle. (The background graph, $G_b$, has no loops.)

When the same image is turned on its side (bottom left in Figure 2), the foreground graph has no loops—in fact, it is merely a single dot. Instead, the handle shows up as a loop in the background graph.

And that, Leahy and Shattuck asserted in their IEEE paper, is the whole story. "Their claim was that handles can't hide," Fishkind says. In other words, any handle in the foreground image must show up as a loop, in either the foreground or the background graph (or possibly in both). This claim served as the foundation for their retouching algorithm. Once they have drawn the foreground and background graphs, they assign a "weight" to each edge. Intuitively, a large weight corresponds to a very fat connection, which is likely to be real; a low weight corresponds to a thin connection, which is more likely to be an artifact. They then erase the lowest-weight edges until all the loops are broken. Finally, they convert the resulting graphs back into a digital image, by editing the corresponding voxels from the original image.

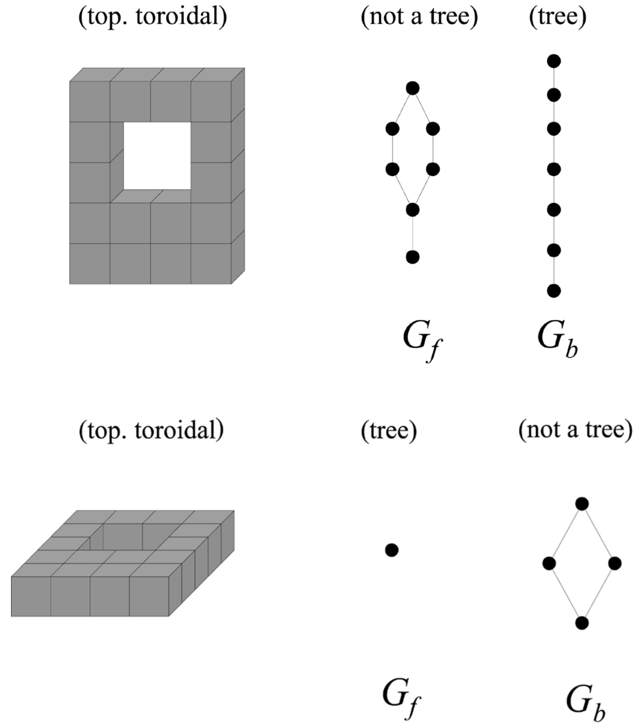


**Figure 2.** *Foreground graph ($G_f$) and background graph ($G_b$) for a toroidal object. Each node in the graphs corresponds to one contiguous island of voxels (either foreground voxels or background voxels) in a given horizontal layer. In the upper image, the handle of the torus is detected by a loop in the foreground graph. In the bottom image, the handle is detected by a loop in the background graph. Courtesy of Donniell Fishkind, Johns Hopkins University.*

The whole procedure is remarkably easy to program on a computer. The foreground and background graphs reduce the complexity of a human brain to the simplicity of a stick figure. "In three-dimensional image reconstruction, you're always looking for things you can do in one dimension," Prince says.

There was just one problem. Shattuck and Leahy had assumed that once they had an object whose foreground and background graphs were loop-free, they were finished: The resulting digital object would have to be homeomorphic to a sphere. Their colleagues in medical imaging thought that they were probably right. "I would have been surprised if it hadn't turned out to be correct," says Bruce Fischl, an imaging expert at Harvard University. Mathematicians, on the other hand, were not so sure, because the history of topology is full of apparently reasonable statements, just like this one, that turned out to be wrong. "It seems too nice to be true," opines Paul Seymour, a graph theorist at Princeton University.

## Strokes of Gen(i)us?

The proof that Priebe, Fishkind, and Abrams devised calculates the genus (i.e., the number of handles) of the foreground object by using one of the oldest and most remarkable formulas in topology, called Euler's formula:

$$V - E + F = 2 - 2g.$$

In other words, if you start with a closed polyhedral surface, add the number of vertices ($V$) and faces ($F$) of the surface, and subtract the number of edges ($E$), the result depends only on the genus ($g$) of the surface. If the polyhedral surface is spherelike, the Euler characteristic ($V - E + F$) will always be 2, regardless of the number or shape of the faces. If the surface is a torus, the Euler characteristic will be 0; if it is a two-holed torus, it will be –2, and so on.

Using Euler's formula, the three mathematicians could easily compute the genus of each layer of the foreground. The hard part was combining the Euler characteristics of the different layers. "If I take a triple torus on one level and a double torus on another, how are they going to connect up?" Fishkind says. "When you join the two, the genus changes. But the way it changes is exactly dictated by which faces, edges, and vertices are being covered over and are no longer on the outside of the union. It turns out that the edges of the foreground and background graphs provide just enough information about how things hook up locally."

In the final analysis, Fishkind and his colleagues showed that the number of handles is bounded below by the number of loops in just one of the graphs, and bounded above by the total number of loops in the two graphs. When neither graph has any loops—the case that Leahy and Shattuck cared about—that means

$$0 \leq g \leq 0 + 0.$$

In other words, the genus of the image is zero, and it is therefore homeomorphic to a sphere.

The original argument, which Priebe, Fishkind, and Abrams presented in a 2002 paper, excluded cases in which two voxels touch only at a vertex or along an edge. Last year, they completed the proof for these defects as well, and that proof appeared in the May 2004 issue of the *IEEE Transactions on Medical Imaging*. The main idea—which actually differs from Shattuck and Leahy's proposal—is to "thicken" the foreground voxels slightly. This breaks the symmetry between foreground and background: Foreground islands are considered contiguous even if some of the voxels touch only along an edge or at a vertex; background voxels have to make full face-to-face contact to be considered contiguous.

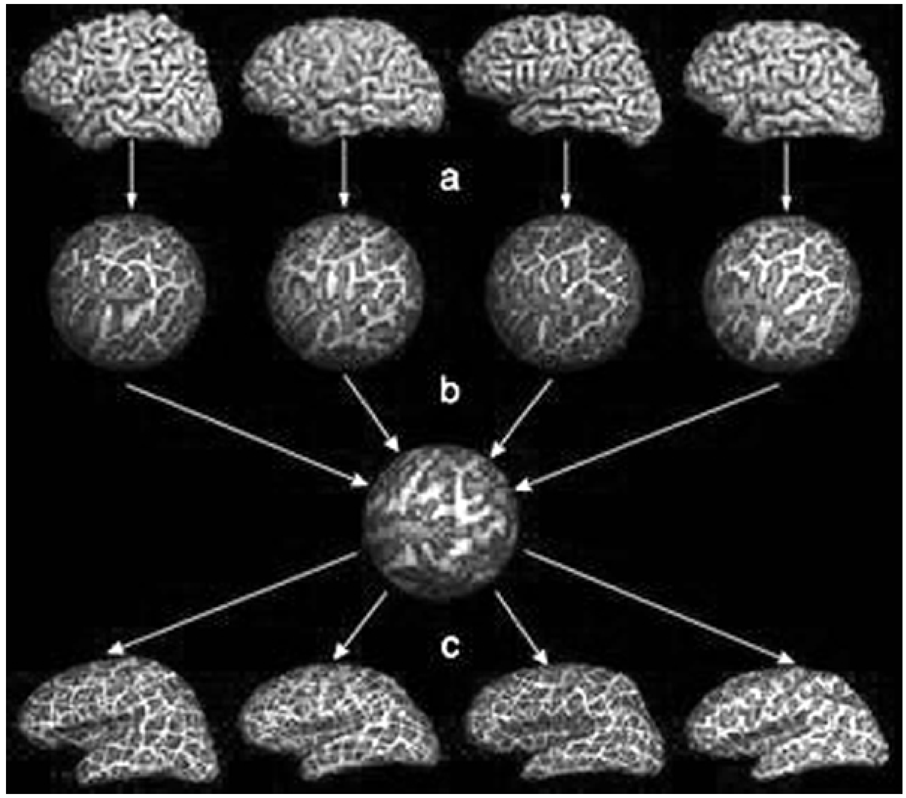


**Figure 3.** *Once topological defects have been corrected, mappings from individual brains (a) to a reference sphere (b) can be used to set up a biologically meaningful coordinate system on the brain cortex (white lines). When morphed back onto the patient's brain (c), the coordinate lines help researchers identify various parts of the brain's anatomy. Courtesy of Bruce Fischl, Harvard University.*

### What It All Means

While Fishkind was working on the spherical homeomorphism theorem, a near tragedy in his family brought home the reality of MR scans. Shortly after giving birth to their first son, his wife started to have slurred speech and weakening in her hands. An MR scan of her brain showed what was wrong: She had had a stroke. Although she eventually recovered, "It was beyond scary," Fishkind says. "For most of those two weeks, math wasn't on my mind, but now and then I'd lie on my back and think of the irony of working on these things at the same time."

In reality, Shattuck and Leahy's retouching algorithm will have little effect on clinical diagnosis. A blocked artery affects far more than one or two stray voxels in an MR image. But the ability to fix the topology is crucial for basic research. Brain scientists like Paul Thompson of UCLA (where Shattuck is now based), for example, deform three-dimensional brain scans until they match either a sphere or a known template. In this way they can place a coordinate grid on the brain, identifying the "latitude and longitude" of various functional regions. Using these maps, Thompson has tracked the effects of conditions like schizophrenia and Alzheimer's disease on different parts of the cortex. Similarly, Fischl uses brain maps to study the visual cortex (see Figure 3). None of this would work if the image had the wrong topology. You can't put spherical coordinates on a torus, and you can't map a brain with 164 handles to a brain with 217 handles.

It is too early to tell if Leahy and Shattuck's method will be the best one for practical use. Several other topological retouching algorithms have been devised that do not require a high-powered mathematical proof of their validity. Fischl, for instance, represents the cortex with mesh surfaces rather than voxels, and therefore cannot use Leahy and Shattuck's voxel-based algorithm. What does seem certain is that Leahy and Shattuck's method will inspire more mathematical work. Fishkind, for example, believes that instead of foreground and background graphs, two foreground graphs can be used alone, formed by slicing in different directions. He and Abrams also hope to discover a continuous, rather than digital, version of the theorem. In the end, Leahy and Shattuck's brainwave may prove to be just as fruitful for mathematics as for brain science.

*Dana Mackenzie writes from Santa Cruz, California.*