From *SIAM News*, Volume 40, Number 2, March 2007

Invasion of the Shapeshifters

By Barry A. Cipra

I have tried to shew the mathematician a field for his labor—a field which few have entered and no man has explored.

—D'Arcy Thompson, *On Growth and Form*

A topologist, it's often said, is a mathematician who can't tell a donut from a coffee cup. Monica Hurdal, a biomedical mathematician at Florida State University, adds a new wrinkle to the old saw: To her, a brain is just a beach ball. More precisely, she says, "The human brain is topologically equivalent to an orientable 2-manifold."

For Michael Miller, a biomedical engineer at Johns Hopkins University, things go even further. "Anatomy," he explains, "is an orbit under group actions of diffeomorphisms."

Vive la Différence

The study of form may be descriptive merely, or it may become analytical.

—D'Arcy Thompson

Hurdal and Miller are among a morphing cadre of researchers who are applying ideas from topology to problems in biology. In Boston, at last year's SIAM Annual Meeting, Miller gave an invited presentation on computational anatomy and Hurdal spoke on "conformal cortical cartography" in a minisymposium on brain imaging. Other speakers in the minisymposium described computational and geometric methods for analyzing brain structures. One of the goals is to help in the diagnosis of neurological disorders, such as Alzheimer's disease and schizophrenia. Another is simply to understand how the brain goes about its business, normal or otherwise.

Computational anatomy, Miller explains, is the realization of an idea dating back to D'Arcy Thompson's 1917 classic, *On Growth and Form*. Thompson proposed the use of coordinate transformations for comparing the shapes of related creatures. For example, the silver hatchetfish, *Argyropelecus olfersi*, can be transformed into its cousin *Sternoptyx diaphana*, the transparent hatchetfish, with a simple linear shear (Figure 1). (There is also a contraction: *diaphana* is smaller than *olfersi*.) Other comparisons use curvilinear transformations much like the pictures aerodynamic engineers draw when designing airfoils.

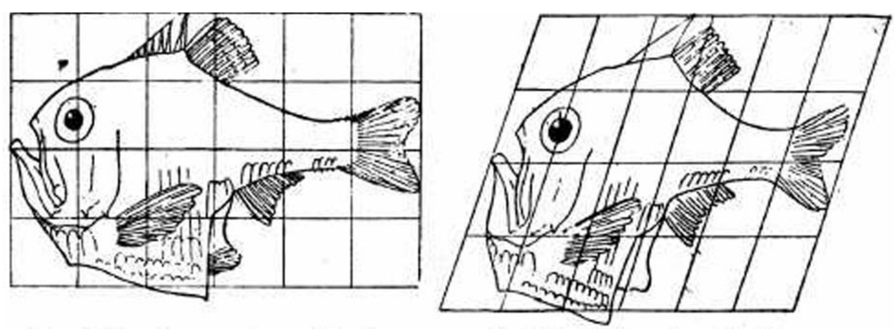

Figure 1. *From D'Arcy Thompson's 1917 classic is his hand-drawn transformation, with a simple linear shear, of* Argyropelecus olfersi *(left) into its cousin* Sternoptyx diaphana *(right).*

A second, more recent and more mathematically advanced inspiration is Ulf Gren-ander's theory of deformable templates. Grenander, an applied mathematician at Brown University, is one of the founding fathers of mathematical pattern recognition, in which researchers aim to understand how it's possible to tell that two things are the same when everything about them is different. A deformable template is an idealization of a two- or three-dimensional object (the theory—and even some of its applications—extend to higher dimensions, but it's best to think in terms of visualizable pictures), which can be stretched, bent, twisted, and otherwise mangled in continuous fashion, to assume a range of "target images" corresponding to the object as it occurs in nature. You can think of it as an airbrushed photo printed on the proverbial rubber sheet of topology.

For anatomical studies, the template should represent the "normal" heart, brain, or other organ, whatever "normal" means. The basic idea is that abnormalities are present in a target image if matching the template to it requires a large deformation—whatever "large" means.

One way to measure the size of a deformation is by defining a function on infinitesimal deformations that assigns a "cost" value at each point, depending on the amount of stretch or shear in its immediate neighborhood. (Rigid translations and rotations are usually cost-free, corresponding as they presumably do to simply lining up the pictures correctly.) As the template is continuously deformed onto the target image, each point racks up a total cost along the path it takes from template to target. The "cost" of the continuous deformation is the sum (actually the integral) of the costs for all the points. The distance between template and target is the smallest such cost among all such continuous deformations.

A continuous deformation that realizes this minimum distance can be thought of as a geodesic in the space of diffeomorphic connections between the template and target images. In fact, this gives, in theory at least, a way to specify a template: If one somehow has a probability distribution on the space of images according to their prevalence in nature, or perhaps according to their desirability, the template can be chosen as the image whose average distance to all other images is the least—a template, in other words, is a kind of infinite-dimensional centroid in a highly warped world.

Knowledge of physiology is crucial in analyzing the infinitesimal deformations, because tissue typically stretches more easily in some directions—along fibers, for example—than in others. This can be modeled mathematically with vector fields. Ideally, the imaging technology will measure the appropriate vectors. Diffusion tensor magnetic resonance imaging (DT-MRI) is a technology that does this. Miller and colleagues at the Center for Imaging Science at Johns Hopkins have developed what they call a "large deformation diffeomorphic metric mapping" (LDDMM) framework for data sets of this type, and applied it in numerical experiments on DT-MRI images of the heart and brain.

In a study of normal canine hearts, Miller and CIS colleagues Yan Cao and Laurent Younes, with Raimond Winslow of the Johns Hopkins Center for Cardiovascular Bioinformatics and Modeling, used the LDDMM software (see http://cis.jhu.edu/software/lddmm/) to match one image, chosen as the template, to the rest. What Thompson did painstakingly by eye and hand the computer now does automatically. By solving a variational problem (for which the researchers prove a general existence theorem), the algorithm zeroes in on the deformation that gives the best match between the vector fields of template and target (see Figure 2).

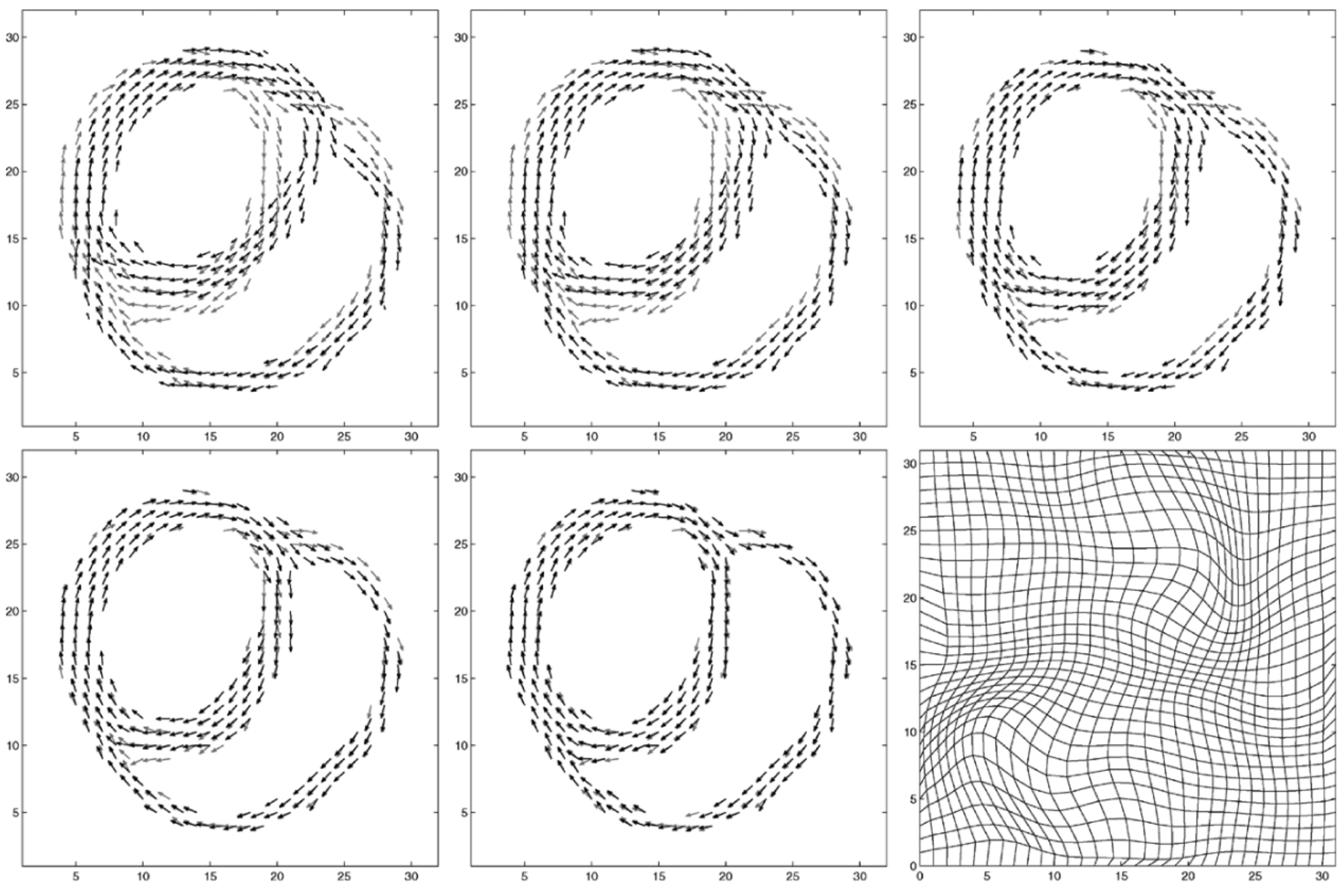

Figure 2. *Deformable dog hearts. Five steps along the way as a template canine heart matches with a target, and the grid that corresponds to the final result. (From Yan Cao, Michael I. Miller, Raimond L. Winslow, and Laurent Younes, "Large Deformation Diffeomorphic Metric Mapping of Vector Fields,"* IEEE Transactions on Medical Imaging, *Vol. 24, September 2005, © 2005 IEEE.)*

Surface Thinking

That I am no skilled mathematician I have had little need to confess.

—D'Arcy Thompson

Hurdal and colleagues, including De Witt Sumners of FSU and Ken Stephenson of the University of Tennessee at Knoxville, are grappling with another deformational aspect of anatomy. They are looking for good ways to flatten human brains in the service of neuroscience. The fundamental problem is familiar from mapmaking: It's impossible to draw any substantial portion of the Earth on a flat piece of paper (or computer screen) without distorting something. The Mercator projection, for example, exaggerates the size of Greenland, making a modestly large island appear to be comparable in size to Africa. One solution, of course, is to carry a globe instead of an atlas, but that can be inconvenient. Mapmakers must choose their poison.

When it comes to mapping the brain, the same is true in spades. It may seem odd even to think about flattening what appears to be a fairly solid lump, but the brain—or at least the gray matter, the part that does most of what passes for thinking—is in fact a thin, flexible sheet, a couple of millimeters thick, that's been crumpled, much like a sheet of paper bearing a failed calculation. The key difference is that crumpled paper is still intrinsically flat, meaning that it can be uncrumpled, whereas the brain's gyri and fundi are intrinsically curved. Topologically speaking, the brain is like a beachball (or beach blanket), but its Riemann metric—the gadget mathematicians use to study curvature—is not only nonzero, it's non-constant.

None of this would much matter if the brain were less complex. The geometry of a head of cauliflower is similarly complex, but one floret is pretty much like another. Even in humans, geometric complexity can often be disregarded. The exact branching pattern of lung tissue, for example, is irrelevant to the pulmonary surgeon. The same goes for the carpet of villi lining the small intestine. For much of medicine, gross anatomy suffices. But not when it comes to the brain. Two chunks of gray matter may look indistinguishable, but one tells the heart to keep beating while the other keeps track of digits of pi. And given the brain's convoluted structure, these disparate chunks can be mere millimeters apart.

It behooves neuroscientists (and neurosurgeons), therefore, to have detailed, accurate maps of the brain that are easy to comprehend and navigate. Hurdal's group proposes a conformal approach. Conformal mapping originates, fittingly enough, in complex analysis. Its definitional virtue is that it preserves angles: Two curves that intersect at, say, 90 degrees on the original surface intersect at the same angle on the map. This means that a conformal map does not create artificial impressions of shear or anisotropic stretching. If you want to know what a tiny, locally almost flat section of the curved surface looks like, all you have to do is magnify the corresponding section of the conformal map.

Another, perhaps even greater virtue, is that conformal maps are essentially unique. This takes a little explaining. Suppose one wants to focus on a particular point and draw a map of some topologically open (and simply connected—a technical condition meaning that there are no holes in the surface) neighborhood of that point. Then, up to rotation, there is a unique conformal map of that neighborhood onto the unit disk, with the focal point at the origin. (One can also make the rotation unique, by specifying a preferred direction at the focal point—say the orientation of a neuron.) This piece of magic is known as the Riemann mapping theorem.

How is such a map actually constructed? Hurdal's group uses a technique called circle packing, an approach to conformal mapping proposed in the mid-1980s by William Thurston, who is now at Cornell University. The central idea is that any simply connected triangular mesh can be converted into a system of mutually tangent circles inside the unit disk in a way that preserves relevant geometric information. As the mesh is refined, the packing gets denser until, in the limit, it fills the entire disk.

Burt Rodin of the University of California at San Diego and Dennis Sullivan of the State University of New York at Stony Brook proved the convergence to full-fledged conformal maps. At any finite state, the result is "quasi-conformal," meaning that there is a bounded amount of angular distortion. Stephenson has done much to develop computer algorithms that implement circle packing. His software, CirclePack (currently for Linux only), is downloadable at http://www.math.utk.edu/~kens/CirclePack/.

For mapping the brain, Hurdal and colleagues start with data from a high-resolution MRI scan. After picking a boundary for the region to be mapped, they triangulate the surface inside—itself a major computational problem. They then apply circle packing to get a quasi-conformal map. In one study with Philip Bowers of FSU, Kelly Rehm and Kirt Schaper of the VA Medical Center in Minneapolis, and David Rottenberg of the University of Minnesota, Hurdal, Stephenson, and Sumners produced a mesh for the left hemisphere with 383,444 triangles, 575,166 edges, and 191,724 vertices. (383,444 – 575,166 + 191,724 = 2, as Euler's famous formula says it must.)

Conformal mapping lends itself, Hurdal notes, to a nice interpretation as a projection into, of all things, the non-euclidean hyperbolic plane. The unit disk provides a model for hyperbolic geometry in which the role of "straight lines" is played by circular arcs that are anchored at right angles to the disk's boundary, which is thus thought of as being "at infinity." Rigid motions in the hyperbolic space correspond to conformal maps of the unit disk to itself, part of a class of analytic functions known as linear fractional transformations. The upshot is that the hyperbolic geometry acts as a kind of natural fisheye lens: If you want a closer look at a region near the periphery of the disk, all you have to do is apply a linear fractional transformation that moves it to the center. Details about all the work of Hurdal and colleagues described here can be found at www.math.fsu.edu/~mhurdal.

Shapeshifting technology itself continues to morph. What it will look like ten years from now is anybody's guess. Researchers expect the subject to offer clinical insights. To invert an old adage, the more things change, the less they stay the same.

Barry A. Cipra is a mathematician and writer based in Northfield, Minnesota.