Brain Warping

Will electronic idiot savants become a doctor's best friends?

By JANET RALOFF

Doctors no longer need a scalpel to probe soft tissues inside the human body. With sophisticated computerized technology, they can obtain high-resolution images of organs or fluid pathways—sometimes in action—without ever opening up a patient.

As remarkable as they are, such noninvasive portraits or videos of the inner body remain only pretty pictures until interpreted by the trained eye and experience of an expert. Yet experts – even physicians and medical researchers – can have trouble seeing, much less understanding, subtle or buried patterns that may connote function, abnormality, or disease. And the likelihood that a lurking pattern will evade notice grows as the volume of images that must be reviewed increases.

Consider, for instance, a magnetic resonance imaging (MRI) scan of the human brain. It collects the data necessary to produce a three-dimensional map of neural territory. But to understand the lay of that land, doctors might have to spend an entire week poring over the 100 or more two-dimensional images that can make up such a map.

Wouldn't it be nice if the doctor could instruct a computer program to analyze those MRI images? Have it look at, say, the cerebellum, or the speech center, or a ventricle and evaluate how any of those structures compares in size or shape or function with some well-characterized average for persons of the same age, size, and sex?

Or imagine a neurosurgeon preparing for some particularly delicate operation by turning on the computer and running through the procedure once, twice—even 10 times before entering the operating room. Moreover, this surgeon would practice not with some representation of the "average" brain but with the patient's brain, as mapped the day before by an MRI scan.

Such computer programs don't exist. But applied mathematicians Michael I. Miller of Washington University in St. Louis and Ulf Grenander of Brown University in Providence, R.I., have made the creation of such programs their goal. A video that Miller's team previewed last month at a meeting of the Council for the Advancement of Science Writing in St. Louis indicates that they may be close to achieving that goal. Very close.

Miller describes his objective as embedding knowledge into a computer program. Today, computers can compare two photos and, through a series of pattern-recognition procedures, identify common elements, even if they vary somewhat in shape. However, because such programs don't know what they're looking at, they're fundamentally ignorant.

In Miller's lab and elsewhere, researchers are "teaching" computers not only how to recognize some structure – a crenellated gray and white mass, for example – but also to understand that this mass is a series of folds in the cortex of the brain. However, even as sophisticated a brain anatomy program as this would be stymied by photos of a peach, turtle, garden, cartoon, or Clipper ship.

Grenander sums it up succinctly: "We're not trying to mimic human intelligence in general. We want to create algorithms that know an awful lot about very little – you know, idiot savants."

Miller and Grenander imbued their first such savant with an eye for mitochondria – those round or rod-shaped structures within a cell that provide energy for cellular functions. After studying countless microscopic images and anatomical descriptions of these subcellular powerhouses, the pair began drafting precise mathematical descriptions of the various shapes and sizes mitochondria take. Creating these descriptions, which the researchers refer to as knowledge representations, "is really the hard part," Grenander says, "and where pattern theory comes in."

Pattern theory, first suggested by Grenander in 1966, offers a means of combining several classical fields (such as algebra, geometry, and probability) in order to describe structures in the natural world – and their inherent variability. The first text on this relatively new field of mathematics will debut later this month: Written by Grenander, *General Pattern Theory* will be published by Oxford University Press.

Miller and Grenander used this theory to model a typical, or average, mitochondrion, describing its structure in terms of several simple geometrical templates, such as a sphere or ellipse. They

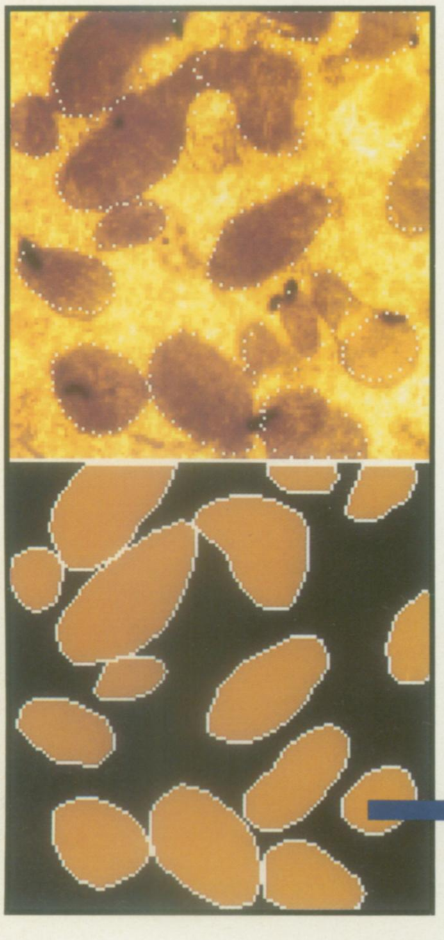

Miller, Saffitz, et al.

Electron micrograph of muscle cell. Below: Outline of structures that the savant recognized as mitochondria.

then described a series of allowable mathematical manipulations for contorting the simple shapes into more complex ones.

"We call these high-dimensional transformations," Miller says. They allow any segment, however small, on the surface of that original geometrical shape to change – to grow or shrink, to twist, even to rotate. Next, the pair assigned specific probabilities to the likelihood that any particular transformation would occur.

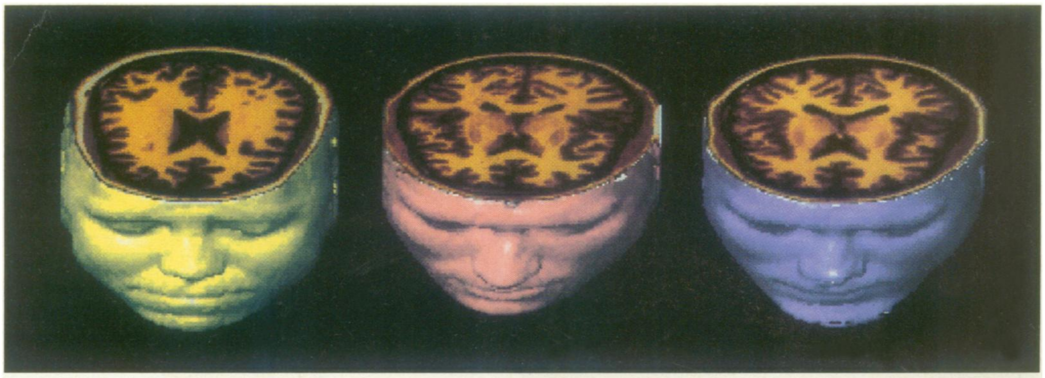

Christensen, Van Essen, *et al.*

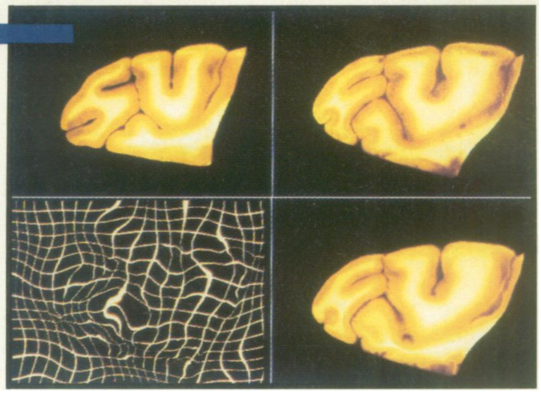

Top: Neuroanatomy savant warps MRI scan of its "textbook" brain (with green face) to match the new patient's brain (pink face). If the match is good, the resulting structure (blue face) will look almost identical to the pink one — and will possess all knowledge initially available for textbook only. Right: Same process for cross sections of macaque visual cortex. Grid (lower left) indicates how much the savant had to warp the textbook structure (upper left) to match the new monkey (upper right).

Miller, Raichle, *et al.*

Finally, the pair melded the mathematical statements into a program to run on an extremely fast, massively parallel computer. This computer breaks problems up in such a way that its 4,000 embedded minicomputers can simultaneously perform separate computations.

When Miller and Grenander finally tested the program on a picture of the interior of a cell, the computer initiated a series of transformations of the program's internal templates. The program attempted to make the templates match structures in the picture without resorting to transformations that had been assigned a very low probability. When a match occurred, the savant designated the structure a probable mitochondrion. A paper describing this mitochondrial recognition program is scheduled to appear early next year in the JOURNAL OF THE ROYAL STATISTICAL SOCIETY.

That initial foray into creating an electronic savant was simple compared to a project in neuroanatomy that Grenander, Miller, and their co-workers are laboring on today.

Several companies offer computerized systems that serve as a neuroanatomical atlas. Each contains a three-dimensional model of the brain based on one individual or a composite of many. The structural features have been labeled, so the user need only designate some region with the cursor and the computer will spit out its name. The user can also retrieve related information, such as text on the designated structure's function, lists of abnormalities that can plague that part of the brain, or citations of research papers on the structure.

Because "the variability in biological shapes is simply awesome," Miller says, no model in an electronic atlas will ever match precisely the brain that a physician is studying. Moreover, if a researcher wants to measure the volume of some part of a patient's brain, he or she must first identify and label the appropriate region on each MRI image for that patient, using the atlas as a guide.

Grenander and Miller decided to teach a computer to identify and label the regions of a brain, much as their earlier savant learned to identify mitochondria. In effect, the program would map features from one atlas – their "textbook" brain – onto the MRI data for any incoming patient. It would do this by matching structures in the textbook brain with those in the patient's brain. Once the match was complete, the computer would then warp the three-dimensional coordinates for structures in the textbook brain until they fit precisely those in the patient's MRI scan.

When this process is successful, all the data available for the textbook brain – anatomical names, functional information, and library references – will be accessible by moving a cursor to corresponding regions in the MRI images of the patient's brain.

The warping of MRI data representing the entire brain can involve 125 million individual deformations and 10 hours of computations on a massively parallel computer. But once that's complete, retrieval of an anatomical label or volume measurement takes only a couple of keystrokes or movements of the cursor.

A pair of papers describing the mathematical underpinnings of this brain-warping program are slated for publication within the next several months – one in the PROCEEDINGS OF THE NATIONAL ACADEMY OF SCIENCES, the other in PHYSICS IN MEDICINE AND BIOLOGY.

The neuroanatomy savant could reap almost instant medical benefits, according to Michael Vannier of the Mallinckrodt Institute of Radiology in St. Louis.

Vannier scouts structural changes in the brain that may reflect – even cause – schizophrenia. MRI scans of schizophrenic patients indicate that, overall, their brains are somewhat smaller than normal, he notes, and that the ventricles (spaces filled with cerebrospinal fluid) are enlarged.

"The question we're trying to answer," he says, "is what has lost volume, because it's not the whole brain." In addition, he says, there's interest in knowing whether the loss occurs in gray matter – the region of the brain where nerve cells reside and information processing occurs – or in white matter – tissue that provides the cable-like connections linking one region of the brain to another.

Preliminary work by Vannier and others (SN: 3/24/90, p.182) indicates that most of the volume reduction occurs in gray matter structures of the limbic system – structures along the bottom surfaces of the brain's temporal lobes.

In attempting to determine whether these changes trace to schizophrenia, to drugs used to treat patients with the disease, or to inherited structural features that may predispose people to develop the disease, Vannier's team is using MRI to measure the volume of several brain structures. They're measuring these not only in schizophrenics who have never received drug treatment, but also in normal individuals, including medical students.

"We can do very high quality MRI scans of the head in less than 10 minutes," Vannier says. "But it takes five working days to extract even a modest amount of information from that scan, using available techniques. So in one week," he continues, "we can collect more MRI data than we can analyze during the rest of the year."

To compute the volume of a structure, its borders must be outlined precisely on every cross-sectional MRI image on which they appear. Labs that do a lot of this work have semiautomated the process already; an operator indicates roughly the structure on each picture to be outlined, then lets a computer painstakingly trace its exact curves.

"The problem," Vannier explains, "is that the computer makes mistakes – obvious mistakes" that the operator must correct. "That's what takes so long," he says, "supervising the machine-based processing and correcting all of its errors." But with the neuroanatomy program Miller's team is completing, Vannier says, structural measurements should be completed "in a few minutes – a big difference."

What's more, he says, the system being developed by Grenander, Miller, and Gary E. Christensen, a doctoral candidate at Washington University, "appears to have a much lower error rate – even with less human input." Although it needs a very powerful computer to operate, Vannier argues that "the quality of the results would certainly seem to justify that technology."

To date, the neuroanatomy savant's accuracy has been tested with only five or six complete-brain MRI scans. Yet even as the system continues to undergo refinement, it's being used to analyze the brains of 10 people involved in Vannier's schizophrenia research.

David Van Essen is also excited about using the neuroanatomy savant. A neuroscientist at Washington University, he studies the cerebral cortex in macaque monkeys. This brain structure possesses many functionally and structurally distinct regions. To find the portion that processes sight information, the focus of his work, he overlays digitized photos of cortical cross sections in which visually responsive areas have been stained.

The first problem confronting anyone in this field is getting past the topographical differences between two individual animals. Though the cortex looks at first much like a smooth, thin pancake, it begins rippling and crumpling up early in development – as an animal grows – in order to fit inside the skull. The number of folds, how deeply they crennelate, and the direction in which they're forced to bend when they confront another part of the brain vary dramatically from animal to animal.

Van Essen has spent almost two decades making unfolded maps of the cortex – essentially warping crenellated structures back into flat pancakes – so he could compare where areas of visual function are located in different macaques. "What Miller and Grenander's team has done – and what really puts them at the head of the pack," he says, "is to come up with a qualitatively different and better strategy for doing this warping."

Unlike previous systems, this one no longer requires researchers to manually insert a series of landmarks on the MRI scans of the brain so the computer can orient itself – a procedure that's tedious and whose accuracy depends on the experience of the individual inserting the landmarks.

As a result, Van Essen says, with the new neuroanatomy savant, "we will get more information and better-quality, higher-resolution, and more meaningful data." It should also greatly improve the quality of comparisons between species, he maintains. "We will be able to take a standard monkey brain and a standard human brain and, using the same warping, or shape transformations, make quantitative comparisons."

Washington University neurologist Marcus Raichle also seeks to correlate structure with function in the brain. Through positron emission tomography (PET), he tracks changes in blood flow that occur as an individual selects particular words or processes specific ideas. But because such PET images of brain function bear so few landmarks, Raichle and others must lay these maps atop MRI scans to identify the particular structures responding to linguistic stimuli.

Comparing data from different people has proved difficult because brain topography varies so greatly. "We have ways of registering one brain on top of another – by trying to make them all the same height, and length, and width," he says. But the precision of these efforts has always "left a lot to be desired," he says.

Of the neuroanatomy savant, Raichle says, its "absolutely elegant way of getting rid of these anatomical differences between individuals" should elevate comparisons of functional PET maps between individuals from a "crude" activity to "a refined exercise."

Kurt Smith believes the brain-warping technique being developed by Grenander and Miller also could speed the ability of firms such as his, Stealth Technologies of Marine, Ill., to develop interactive computer systems that simulate the operating theater, making possible virtual-reality surgery.

Already, hospitals can buy systems that allow sensors on surgical tools to track where the tip of an instrument is. These data feed into a computer program containing three-dimensional MRI or computed tomography (CT) images of the patient taken before surgery. When displayed on a television monitor above the operating table, Smith says, "They show the surgeon where he is while he's operating."

One major drawback of today's systems, he notes, is that they portray brain scan images in gray tones. If a computer had labeled each structure or region, a surgeon could tint the tissue he or she was aiming at, making it easier to see, Smith says. Such labeling also would allow the computer to erase on demand any part of the image, such as structures that might be obscuring or touching the target area.

There's another drawback. Today's systems can display what the body looked like prior to surgery, but once surgery begins, organs may shift, an artery may rupture, or the brain may deform as a fluid-filled sac is punctured. "What you'd like to do is update these [MRI or CT] data sets so that they reflect what's going on during surgery," says Smith. And in a year or two, he predicts, Miller and Grenander's neuroanatomy savant might be capable of warping structures in ways that would represent anticipated responses to surgical actions.

Most of Miller and Grenander's recent work has focused on developing and refining savants for narrow applications in biomedicine: recognizing mitochondria and amoebas, cleaning up low-resolution X-ray videos of the arteries surrounding the heart, and understanding brain structures. But over the past year, they've expanded their efforts into other areas as well.

For instance, in one project they are attempting to distinguish moving objects. The resulting program might search for aircraft against a background of decoys or of electronic noise in a cluttered radar image. Alternatively, it could attempt to identify tanks and other vehicles in a low-resolution night-vision video. In such cases, modeling efforts will focus on developing mathematical statements that describe a range of variability for features other than shape – such as trajectories and accelerations that don't violate the laws of physics, or thermal data that may indicate an engine or its exhaust.

In an even more abstract application of pattern theory, Grenander and Miller are attempting to teach computers to understand long, complicated passages of English text.

"Though the details are different," Grenander says, "we treat all of these [applications] in the same spirit."

Indeed, Miller says, "What is so beautiful, so elegant about pattern theory is that it can be used to extract meaning from so many vastly different types of structures." □