

Growing Pains for fMRI

As the use of functional magnetic resonance imaging has exploded, some researchers say the field could use a dose of rigor. Will new experimental approaches come to the rescue?

Last November, the op-ed page of the *New York Times*, which typically airs political controversies, managed to create one of its own. It published a column describing a study in which 20 undecided voters had their brain activity scanned by functional magnetic resonance imaging (fMRI) while viewing photographs and videos of the major candidates in the upcoming U.S. presidential election. The findings revealed “some voter impressions on which this election may well turn,” according to the authors, who included a political scientist, a neuroscientist, and several people affiliated with FKF Applied Research, a company based in

Washington, D.C., that sells fMRI-based marketing studies.

The column infuriated some neuroscientists and ignited an animated discussion in the imaging field. “It was really closer to astrology than it was to real science,” says Russell Poldrack of the University of California, Los Angeles (UCLA), who drafted a letter to the newspaper that was signed by 16 other cognitive neuroscientists and published 3 days later. “It epitomized everything that a lot of us feel is wrong about where certain parts of the field are going, which is throw someone in a scanner and tell a story about it.”

Since its introduction in the early 1990s,

fMRI has transformed neuroscience. Now in its teenage years, the fMRI field is still experiencing growing pains. Some cognitive neuroscientists say they’re frustrated that many studies—including some of those that garner the most attention in the popular press—reveal little about the neural mechanisms of human cognition. “The problem right now with imaging is that doing experiments right is really, really hard, but getting pictures out is really, really easy,” says Steven Petersen, a veteran brain-imaging researcher at Washington University in St. Louis, Missouri.

At the same time, there are signs that the field is maturing, as researchers confront the limitations of fMRI. Such efforts include painstaking experiments that match human fMRI data with analogous fMRI data and electrophysiological recordings of neural activity in monkeys, as well as new analytical methods capable of revealing information processing in the brain that would be impossible to detect with standard methods. “I think [these methods] are really going to revolutionize how we think about our data,” says Poldrack. They also have the potential to introduce more rigor into fMRI research—something that’s badly needed, Poldrack says, otherwise, “people will start to see fMRI as neophrenology, just telling stories and not giving explanations.”

Neuroimagers gone wild


What irked Poldrack and others most about the *Times*’s op-ed was the way the authors inferred particular mental states from the activation of particular brain regions: Activity in the anterior cingulate cortex indicated mixed feelings about Hillary Clinton, for example, whereas amygdala activation indicated “voter anxiety” about Republican candidate Mitt Romney.

The basic problem, the objectors wrote in their letter, is that it’s not possible to infer a particular mental state (such as anxiety) from the activation of a particular brain region (such as the amygdala). Although it’s true that anxiety engages the amygdala, says co-signer Elizabeth Phelps, a cognitive neuroscientist at New York University, so do intense smells, sexually arousing images, and many other things. To conclude that Romney makes voters

anxious based on amygdala activation alone is unjustified, Phelps says.

The neuroscientist co-author on the op-ed piece, Marco Iacoboni of UCLA, stands by the column's conclusions as reasonable and says he's been surprised and stung by what he views as an overly harsh and hypocritical rebuke. After all, he points out, most of his critics use similar "reverse inferences" themselves.

That's true, says Poldrack, and it's a problem the field needs to confront. He and others argue that reverse inferences are particularly common in newer fields such as social cognitive neuroscience and neuroeconomics (not to mention neuropolitics), fields in which researchers are still trying to identify the cognitive processes underlying the behaviors they study. As an example, Poldrack points to a widely cited paper that used fMRI to investigate brain activity in subjects pondering moral dilemmas (*Science*, 14 September 2001, p. 2105); some of the brain regions that lit up had been linked in previous studies to emotional and "rational" cognitive processes, and the authors concluded that these two types of processes are active, to different degrees, in different types of moral judgments. But the strength of such arguments hinges on how specifically a given brain area is linked to a given mental process. Poldrack points out, for example, that some of the "emotional" brain regions in the morality study have also been connected to memory and language—a caveat that is



Political blunder? The *New York Times* used this graphic, showing that U.S. presidential candidates Barack Obama and John McCain stimulated relatively little activity in the brains of undecided voters, to illustrate online a brain-imaging study published as an op-ed column last November.

rarely mentioned in media coverage of the work (*Science*, 9 May, p. 734).

Monkeying around

The general public may be easily seduced by pretty images generated by fMRI (see sidebar, below), but even neuroscientists sometimes seem to fall under the spell and overlook the method's limitations. One constraint is the narrow sliver of the human experience that can be captured when a person has to keep his or her head still for long periods inside an fMRI scanner. Another is the resolution. Using fMRI to spy on neurons is something like using Cold War-era satellites to spy on people: Only large-scale activity is visible. With standard fMRI equipment, the smallest cube of brain

tissue that can be imaged is generally a few millimeters on a side. Each such "voxel" (a mashup of volume and pixel) contains millions of neurons. And although neurons can fire hundreds of impulses per second, the fMRI signal—which indicates an increase in oxygenated blood bringing energy to active neurons—develops sluggishly, over several seconds. This makes fMRI a crude tool for investigating how circuits of intricately connected neurons do the computational work of cognition and behavior, says Roger Tootell, a neuroscientist at Harvard University. "fMRI is really good for telling you where to look," he says, "but I don't think you can ever really come up with mechanisms."

Tootell is one of a handful of researchers trying to circumvent such obstacles by combining human fMRI with monkey experiments. The general idea, he explains, is to follow up on the human findings by using fMRI to identify analogous regions of the monkey brain and then record the activity of individual neurons there with microelectrodes.

In some cases, single neuron recordings have confirmed fMRI findings. In 2006, Tootell and colleagues reported microelectrode data showing that 97% of neurons in the monkey equivalent of the fusiform face area—a region of the temporal cortex that appears in human fMRI studies to respond selectively to images of faces—do indeed respond preferentially to faces (*Science*, 3 February 2006, p. 670). But Tootell says that more recent human fMRI experiments his group has done suggest that neurons in an adjacent "place" region in the temporal cortex respond preferentially to edges, not places per se. The researchers are planning monkey experiments to investigate the preferences of neurons in this region in greater detail.

Such studies, he says, can also begin to reveal mechanisms of visual object processing in the brain, such as how "face" or "place" neu-

DON'T BE SEDUCED BY THE BRAIN

Few advances in neuroscience have generated as much public interest as the ability to see the human brain in action. The enthusiasm isn't hard to understand. Methods such as functional magnetic resonance imaging (fMRI) have enabled researchers to bring distinctly human attributes—love, faith, morality—under scientific scrutiny.

But the images generated by such methods may have a power to captivate that reaches beyond their power to explain. Psychologists David McCabe of Colorado State University in Fort Collins and Alan Castel of the University of California, Los Angeles, recently asked 156 undergraduate students to evaluate several mock news articles describing brain-imaging studies. But the research each described was bogus. One study, for instance, reached the dubious conclusion that because watching television and doing arithmetic problems both activate the temporal lobes of the brain, watching television improves arithmetic abilities.

Students saw one of three versions of each article: the text alone, the text plus an fMRI image depicting activity in part of the brain, or the text plus a bar chart summarizing the fMRI result. Those who saw the brain image rated the scientific reasoning in the article as more compelling than did the others even though the images themselves added no relevant information, McCabe and Castel reported in the April issue of *Cognition*.

People seem to believe that images of brain activity make a behavioral observation more real, says bioethicist Éric Racine of the Institut de Recherches Cliniques de Montréal in Canada. Racine calls this effect "neurorealism" and says it's often amplified by media coverage that oversimplifies research findings and glosses over caveats. In other words, don't let the pretty colors fool you. You don't need an fMRI scan to know that candy tastes good, pain feels bad, and television won't turn you into a genius at math.

—G.M.



rons acquire their selectivity by combining inputs from low-level neurons that respond to simpler features such as texture, curvature, and the orientation of lines. “It’s a beautiful paradigm when you can bring it to bear,” Petersen says of the parallel human-monkey work. The drawback, he says, aside from the incredibly time-consuming experiments, is that it can’t be applied to study many types of cognition—language, for example.

There’s a pattern here

A very different approach to overcoming some of fMRI’s constraints comes from new analysis tools borrowed from machine-learning research. In a standard fMRI study, neuroscientists average together the fMRI activation for neighboring voxels. This averaging makes it easier to detect differences

identified statistically distinct activity patterns elicited by each type of object.

In 2005, two research teams published papers in *Nature Neuroscience* showing that similar methods made it possible to determine the orientation of lines a subject was viewing based on fMRI activation in the primary visual cortex, a feat previously thought impossible because neurons that share a preference for lines of a particular orientation pack into columns narrower than a voxel. That got even more people interested, says Rajeev Raizada of Dartmouth, who organized a session on these methods at an April meeting of the Cognitive Neuroscience Society in San Francisco, California.

Raizada and others at the session presented a variety of new findings illustrating how this new analysis of fMRI data can reveal informa-

Other researchers are taking note of such findings. “This is an exciting new direction,” says Adam Aron, a cognitive neuroscientist at the University of California, San Diego. “Instead of looking at whether this or that brain region is activated, now you’re talking about whether the activity in many different voxels can predict what people are seeing or hearing.” Poldrack predicts that classifiers will help rescue researchers from the logical perils of reverse inference. Instead of inferring that a photo of Mitt Romney induces anxiety, for example, researchers could collect patterns of brain activity evoked by known anxiety inducers (photos of spiders, snakes, and hypodermic needles, perhaps) and see whether the pattern Romney elicits is a statistical match.

An expanding toolbox

Yet even with the promise of these new tools, fMRI remains limited to revealing correlations between cognitive processes and activity in the brain. “The way to use it well is as one tool in a toolbox, as a way of testing hypotheses where you have converging techniques and evidence,” says Aron.

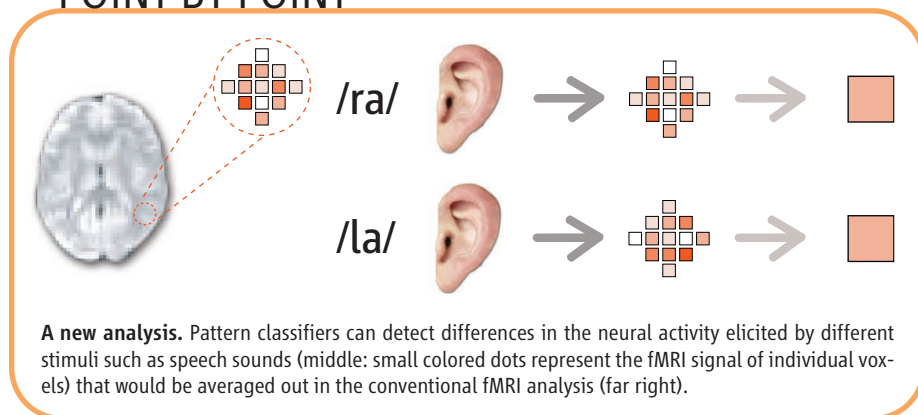
To that end, growing numbers of neuroscientists are using fMRI and related methods to investigate the connectivity between different brain regions involved in cognitive functions such as language and memory. One fMRI approach is to identify brain regions whose activity is synchronized when subjects perform a given task. In some cases, researchers are probing further to determine if those areas that fire together are physically wired together, using a relatively new MRI method called diffusion tensor imaging that can visualize the axon tracts that connect regions in the living human brain.

Others are trying to establish causal links between brain and behavior. Having linked a brain region to a particular behavior using fMRI, for example, some researchers are following up with transcranial magnetic stimulation experiments in which focused magnetic fields noninvasively and temporarily disrupt neural activity in that region. If the behavior is then altered, the region must play a role in controlling it.

With such a convergence of methods and other advances, perhaps one day it will even be possible to divine the intentions of undecided voters. But that day does not seem near at hand. In the *Times* op-ed piece, the authors reported that their scans indicated that voters were “unengaged” with two candidates in particular, Barack Obama and John McCain, ironically, the two men now battling for the U.S. presidency.

—GREG MILLER

POINT BY POINT



between experimental conditions—viewing photos of faces versus places, for example—but it assumes that neurons from different voxels in the region of interest all behave the same way. That’s almost certainly not the case, says Nikolaus Kriegeskorte, a neuroscientist at the National Institute of Mental Health in Bethesda, Maryland.

To sidestep this issue, Kriegeskorte and others have been working with statistical tools called multivariate pattern classifiers to take a finer grained look at brain activity that considers patterns of activation across many individual voxels without averaging. These methods shift the focus from trying to identify specific brain regions that are activated during a particular task to trying to identify how the relevant information is processed in the brain.

The first demonstration of this approach was a study by cognitive neuroscientist James Haxby, now at Dartmouth College (*Science*, 28 September 2001, p. 2425). He and colleagues monitored brain activity elicited by hundreds of images of various types of objects, including faces, cats, houses, and scissors, and

tion processing in the brain that would be overlooked by conventional analyses. Raizada, for example, presented a study in which he and colleagues investigated fMRI responses to the sounds /ra/ and /la/ in the brains of 10 native English and 10 native Japanese speakers. The Japanese language does not distinguish between these sounds, and most native speakers can’t hear the difference.

Inside the scanner, each subject listened to six variations of each /ra/ and /la/ while the researchers collected fMRI data for each variation. Using a pattern classifier, Raizada determined that English—but not Japanese—speakers exhibited distinct activity patterns in the right primary auditory cortex for /ra/ and /la/. In fact, subjects who were best able to distinguish the sounds had the most distinct activity patterns. Each sound is apparently represented by different patterns—but similar overall levels—of neural firing in the auditory cortex of English speakers, Raizada says, which explains why the conventional fMRI analysis can’t pick up this distinction.