

MIND READERS

BRAIN-SCANNING MACHINES MAY SOON BE CAPABLE OF DISCERNING RUDIMENTARY THOUGHTS AND SEPARATING FACT FROM FICTION **BY PHILIP ROSS**

IMAGINE A WORLD YOU COULD TRUST—REALLY TRUST— where truth was transparent and juries, police, locksmiths and gossip columnists were largely overthrown. Human society would be orderly, boring and as alien as an anthill.

This is the promise and the threat of a machine that could read minds. The hoary polygraph has never filled the bill. It measures not thoughts but only the indirect physiological consequences of thoughts—blood pressure and respiration, among others—that hint that a subject may be lying. The result, critics charge, is false positives—an honest answer misjudged as a lie—and false negatives—a lie misjudged as the truth. The courts have long ruled polygraph findings inadmissible as evidence. Just last October the National Research Council damned the device as a “blunt instrument,” of little use in ferreting out criminals, spies and terrorists.

Greek philosopher Diogenes walked with a lamp, in search of an honest man. Yet why shine your lamp into someone’s face when you can look at the very brain? There you might do better than merely tell truth from lies. You might also con-

verse with minds trapped inside paralyzed bodies, expose to analysis the suppressed fears and desires of the stormy unconscious, even observe the insights and errors by which a student moves toward the solution of a math problem.

The idea of looking directly at brain activity to tell truth from falsehood dates back roughly 20 years, to when J. Peter Rosenfeld of Northwestern University observed an interesting feature in the electroencephalograph, or EEG, a chart of the brain’s electrical signals as detected on the surface of the skull. The P300 wave had already been known to be evoked by oddball cues, such as hearing one’s name mentioned in a list of other words. Rosenfeld found that lying elicited it, too. He is now mapping the P300 wave across the scalp to get enough spatial resolution to improve the sensitivity of the test.

The next step appears to have been articulated for the first time by the often prophetic science columnist David Jones, a.k.a. Daedalus, who wrote in 1996 that “a modern magnetic-resonance brain scanner should be a perfect lie detector....




Telling the truth should activate just one site in the brain.... Telling a lie should activate two sites: one holding the lie and the other holding the truth that it is masking.”

Five years later Daniel Langleben of the University of Pennsylvania and his colleagues used functional magnetic resonance imaging (fMRI) to scrutinize the brains of subjects acting out a question-and-answer series. Under certain conditions, the subject would tell a string of falsehoods in such a manner as to mimic lying; in other conditions, the subject would utter a string of truthful statements.

The two brain images from each category were averaged and compared.

It turned out that all areas activated during truth telling were also triggered during lying but that a number of areas were active particularly during lying. “That suggests that the default position is truth, and deception is some sort of process you perform on truth,” Langleben remarks. He notes that several areas activated more during lying—including the anterior cingulate cortex and part of the left prefrontal cortex—are associated with suppression of response, as when the brain decides to go with one of



A graphical rendition of a brain slice highlights two areas that are more active when a subject falsely denies possessing the five of clubs than when the person tells the truth.



two conflicting responses and must therefore inhibit the other one [see illustration above].

According to this theory of “cognitive load,” actor Sean Connery, when asked his name during the filming of a movie, cannot help but flash to the words “Sean Connery”; it is only with a modicum of effort that he chokes off that response to say instead, “Bond. James Bond.” So far two other fMRI groups have published similar research; more have written papers now wending their way to publication.

None of these groups have yet claimed much power in catching a particular hostile witness in a particular lie. “As a practical method, this thing is not even in the proof-of-concept stage,” Langleben admits. “In April [2004] we will take the next step and try to determine the size of the truth-versus-lie effect at a given spot in the brain.” He expects to use a larger sample, 60 to 90 subjects, and to create sit-

polygraphs have almost everything to do with it. Indeed, polygraphs are often used to instill fear as much as to detect it (like the “fear-o-sensor” that a dog waves over a stranger in one of Gary Larson’s cartoons). Second, brain imaging follows a phenomenon that is much closer to thought, in the train of events, than are the pulse, skin conductance, respiratory rate and so on—“output that is 10 times removed from what’s happening in the brain,” Langleben says.

Even fMRI does not sample the neurons themselves, though, but just the oxygen in the nearby bloodstream. More precisely, it measures the ratio of oxygenated to deoxygenated blood. The machine can pinpoint metabolic activity at good resolution, of about four millimeters in diameter, yet it is relatively slow, tracking activity occurring for two seconds or so. That’s not really fast enough to catch a thought.

To capture that level of complexity would require recording a signal lasting for mere milliseconds, providing a snapshot of, say, calcium ions in the neurons themselves. To detect it, however, would require magnets several times as powerful as even Langleben’s four-tesla unit. No such magnets big enough for humans exist, and, for safety reasons, none are likely to be approved for that purpose. “I can tell you there won’t be human studies in 20-tesla machines,” asserts Marcus E. Raichle, an fMRI researcher at Washington University. “It can stimulate the vestibular system, making you feel dizzy; it can heat up the brain, manipulating the very thing you’re supposed to study.”

Another approach to get good resolution in both space and time might come by combining fMRI with EEG. One might measure both things at the same time, or correlate a lie-detecting component of fMRI with a given aspect of EEG. “If we did that, we could

BRAIN IMAGING follows processes that are **MUCH CLOSER TO THOUGHT** than the pulse, skin conductance and respiratory rate measured by **A POLYGRAPH**.

uations closer to real-life deception—perhaps a poker game. (It might be a little hard to simulate, though, inside a churning, claustrophobia-inducing MRI machine.)

In principle, brain imaging is better than a polygraph, he argues, for two reasons. First, it seems to have nothing to do with general anxiety, whereas

discard fMRI and use the EEG signal, and it would be 10 times cheaper,” Langleben says.

Although it may be too hard for today’s brain scanners to trap a hostile witness or a cheating spouse, they may well suffice to divine certain, simple thoughts of a willing communicant, leading to a more general form of mind reading. Unlike lie detection, the task, of course, is made simpler when the subject cooperates with the testing. Already monkeys with electrodes implanted in motor-control areas of their brains have been taught, through biofeedback techniques, to convey neural impulses over an Internet connection to manipulate a robotic arm [see “Controlling Robots with the Mind,” by Miguel A. L. Nicolelis and John K. Chapin; *SCIENTIFIC AMERICAN*, October 2002]. Niels Birbaumer of the University of Tübingen in Germany has reported a degree of success in using biofeedback to train patients immobilized by nerve damage to vary their brain waves and so to spell out sentences on a computer screen.

But true mind reading must do better, by catching a word or concept exactly as it forms itself in the brain. Marcel A. Just of Carnegie Mellon University claims he has done just that with fMRI, by limiting the concepts to a small number and keeping them very simple—carpentry tools, for instance, or kinds of dwellings. “We have 12 categories and can determine which of the 12 the subjects are thinking of with 80 to 90 percent accuracy,” he explains. He is even better at distinguishing brains reading a clear sentence from those reading an ambiguous one or imagining a verb as opposed to a noun.

Just’s colleague Tom Mitchell, a computer scientist, has devised a means to classify the complex brain images that their experiments produce. He analyzes them with neural networks, a type of software that can tune itself to improve its ability to distinguish patterns. “If isolated words can be identified with some degree of accuracy, it ought to be possible to do even better with entire sentences,” Mitchell says. That is because sentence structure constrains the possibilities that the neural network must consider. “If you know that a sentence has two words, then one must be a verb, the other a noun.

“One experiment I would love to do is to find words that produce the most distinguishable brain activity,” he adds. Such words might serve as the building blocks for a neural interface, much as particularly discriminable English words were favored in the early, limited-vocabulary protocols of voice-processing software.

Should this concept-recognition system work with even minimal reliability, it might be coupled with lie-detecting fMRI software to produce a much

more sophisticated tool. In principle, law-enforcement officers might use the combination technology to tell not only that a bank robber is lying but that the loot is stashed in the garage.

A brain decoder that worked on all brains still might not allow for telepathy on the order of a Vulcan mind meld in the *Star Trek* series, which enabled universal translation. An English sentence, beamed into the mind of a non-English speaker, might seem gibberish. Even if the receiving (or eavesdropping) person spoke the same language, he might be puzzled by the idiomatic dialect in which a mind converses with itself, with all its coded entries, abbreviations and emotional associations.

Concocting near-perfect lie detection may, none-

fMRI can determine which
of 12 SIMPLE CATEGORIES
a subject is contemplating with
80 TO 90 PERCENT accuracy.

theless, be much easier than making a sophisticated thought reader—and almost as dangerous to mental privacy. Indeed, it would not be necessary to employ such a machine—the threat of its use would exercise a powerful deterrent force.

As Daedalus concluded, “Like the atom bomb, it is best reserved as a sort of ultimate social weapon. If widely deployed outside the courtroom, it would make social life quite impossible.”

Philip Ross writes on science and technology from New York City. His work has also appeared in Acumen Journal of Sciences, IEEE Spectrum, Forbes, and the New York Times.

MORE TO EXPLORE

Event-Related Potentials in the Detection of Deception, Malingering, and False Memories. J. Peter Rosenfeld in *Handbook of Polygraph Testing*. Edited by Murray Kleiner. Academic Press, 2001. Preprint available at www.psych.northwestern.edu/psych/people/faculty/rosenfeld/NewFiles/P300%20and%20ERP%207-99.pdf

Brain Activity during Simulated Deception: An Event-Related Functional Magnetic Resonance Study. D. D. Langleben, L. Schroeder, J. A. Maldjian, R. C. Gur, S. McDonald, J. D. Ragland, C. P. O’Brien and A. R. Childress in *Neuroimage*, Vol. 15, No. 3, pages 727–732; March 2002. Available at www.uphs.upenn.edu/trc/conditioning/neuroimage15_2002.pdf

The Polygraph and Lie Detection. Board on Behavioral, Cognitive, and Sensory Sciences and Education [BCSSE], Committee on National Statistics [CNSTAT]. National Academies Press, 2003. Available at www.nap.edu/books/0309084369/html

Marcel Just’s Laboratory Center for Cognitive Brain Imaging of CMU Web site is available at http://coglab.psy.cmu.edu/index_main.html