

THE BATTLE TO STOP BIRD FLU • GOOGLE VS. CLICK FRAUD

# WIRED

JANUARY 2006 GET SMART

RAISE THE ROOF!  
YOUR D.I.Y.  
GUIDE TO THE  
**WIRED  
HOME**

## 50 BEST BOTS EVER

Einstein Bot, Mars Rovers, R2-D2, Predator Drones, Astroboy ...

**Brain Scans**  
and the New  
Science of  
Lie Detection

**The Worst Movie**  
of 2006 (and  
Why It's Already  
a Cult Hit!)

**PLUS:**  
The Robot Car  
That Won the  
Great Desert  
Road Race

The Albert Hubo robot:  
Facial expression meets  
mobility. Genius!



0 75110 08459 2

## Don't Even Think About Lying

How brain scans are reinventing the science of lie detection.

By Steve Silberman

I'm flat on my back in a very loud machine, trying to keep my mind quiet. It's not easy. The inside of an fMRI scanner is narrow and dark, with only a sliver of the world visible in a tilted mirror above my eyes. Despite a set of earplugs, I'm bathed in a dull roar punctuated by a racket like a dryer full of sneakers.

Functional magnetic resonance imaging - fMRI for short - enables researchers to create maps of the brain's networks in action as they process thoughts, sensations, memories, and motor commands. Since its debut in experimental medicine 10 years ago, functional imaging has opened a window onto the cognitive operations behind such complex and subtle behavior as feeling transported by a piece of music or recognizing the face of a loved one in a crowd. As it migrates into clinical practice, fMRI is making it possible for neurologists to detect early signs of Alzheimer's disease and other disorders, evaluate drug treatments, and pinpoint tissue housing critical abilities like speech before venturing into a patient's brain with a scalpel.

Now fMRI is also poised to transform the security industry, the judicial system, and our fundamental notions of privacy. I'm in a lab at Columbia University, where scientists are using the technology to analyze the cognitive differences between truth and lies. By mapping the neural circuits behind deception, researchers are turning fMRI into a new kind of lie detector that's more probing and accurate than the polygraph, the standard lie-detection tool employed by law enforcement and intelligence agencies for nearly a century.

The polygraph is widely considered unreliable in scientific circles, partly because its effectiveness depends heavily on the intimidation skills of the interrogator. What a polygraph actually measures is the stress of telling a lie, as reflected in accelerated heart rate, rapid breathing, rising blood pressure, and increased sweating. Sociopaths who don't feel guilt and people who learn to inhibit their reactions to stress can slip through a polygrapher's net. Gary Ridgway, known as the Green River Killer, and CIA double agent Aldrich Ames passed polygraph tests and resumed their criminal activities. While evidence based on polygraph tests is barred from most US trials, the device is being used more frequently in parole and child-custody hearings and as a counterintelligence tool in the war on terrorism. Researchers believe that fMRI should be tougher to outwit because it detects something much harder to suppress: neurological evidence of the decision to lie.

My host for the morning's experiment is Joy Hirsch, a neuroscientist and founder of Columbia's fMRI Research Center, who has offered me time in the scanner as a preview of the near future. Later this year, two startups will launch commercial fMRI lie-detection services, marketed initially to individuals who believe they've been unjustly charged with a crime. The first phase of today's procedure is a baseline interval that maps the activity of my brain at rest. Then the "truth" phase begins. Prompted by a signal in the mirror, I launch into an internal monologue about the intimate details of my personal life. I don't speak aloud, because even little movements of my head would disrupt the scan. I focus instead on forming the words clearly and calmly in my mind, as if to a telepathic inquisitor.

Then, after another signal, I start to lie: I've never been married. I had a girlfriend named Linda in high school back in Texas. I remember standing at the door of her parents' house the night she broke up with me. In fact, I grew up in New Jersey, didn't have my first relationship until I went to college, and have been happily married since 2003. I plunge deeper and deeper into confabulation, recalling incidents that never happened, while trying to make the events seem utterly plausible.

I'm relieved when the experiment is over and I'm alone again in the privacy of my thoughts. After an hour of data crunching, Hirsch announces, "I've got a brain for you." She lays out two sets of images, one labeled truth and the other deception, and gives me a guided tour of my own neural networks, complete with circles and Post-it arrows.

"This is a very, very clear single-case experiment," she says. In both sets of images, the areas of my cortex devoted to language lit up during my inner monologues. But there is more activity on the deception scans, as if my mind had to work harder to generate the fictitious narrative. Crucially, the areas of my brain associated with emotion, conflict, and cognitive control - the amygdala, rostral cingulate, caudate, and thalamus - were "hot" when I was lying but "cold" when I was telling the truth.

"The caudate is your inner editor, helping you manage the conflict between telling the truth and creating the lie," Hirsch explains. "Look here - when you're telling the truth, this area is asleep. But when you're trying to deceive, the signals are loud and clear."

I not only failed to fool the invisible inquisitor, I managed to incriminate myself without even opening my mouth.

The science behind fMRI lie detection has matured with astonishing speed. The notion of mapping regions of the brain that become active during deception first appeared in obscure radiology journals less than five years ago. The purpose of these studies was not to create a better lie detector but simply to understand how the brain works.

One of the pioneers in the field is Daniel Langleben, a psychiatrist at the University of Pennsylvania. Back in 1999, he was at Stanford, examining the effects of a drug on the brains of boys diagnosed with attention deficit hyperactivity disorder. He had read a paper theorizing that kids with ADHD have difficulty lying. In Langleben's experience, however, they were fully capable of lying. But they would often make socially awkward statements because "they had a problem inhibiting the truth," he says. "They would just blurt things out."

Langleben developed a hypothesis that in order to formulate a lie, the brain first had to stop itself from telling the truth, then generate the deception - a process that could be mapped with a scanner. Functional imaging makes cognitive operations visible by using a powerful magnetic field to track fluctuations in blood flow to groups of neurons as they fire. It reveals the pathways that thoughts have taken through the brain, like footprints in wet sand.

When Langleben ran an online search for studies of deception using fMRI, however, he found nothing. He was surprised to find "such a low-hanging fruit," as he puts it, still untouched in the hothouse of researchers hungry to find applications for functional imaging.

After taking a job at the University of Pennsylvania School of Medicine later that year, he mapped the brains of undergraduates who had been instructed to lie about whether a playing card displayed on a computer screen was the same one they'd been given in an envelope along with \$20. The volunteers - who responded by pressing a button on a handheld device so they wouldn't have to speak - were told that if they "fooled" the computer, they could keep the money. Langleben concluded in 2002 in a journal called *NeuroImage* that there is "a neurophysiological difference between deception and truth" that can be detected with fMRI.

As it turned out, other researchers in labs across the globe were already reaching for the same fruit. Around the same time, a UK psychiatrist named Sean Spence reported that areas of the prefrontal cortex lit up on fMRI when his subjects lied in response to questions about what they had done that day. Researchers from the University of Hong Kong provided additional confirmation of a distinctive set of neurocircuits involved in deception.

For fMRI early adopters, these breakthroughs validated the practical value of functional imaging itself. "I felt this was one of the first fMRI applications with real value and global interest," Langleben says. "It had implications in crime and society at large, in defense, and even for the insurance industry."

The subject took on a new urgency after 9/11 as security shot to the top of the national agenda. Despite questions about reliability, the use of polygraph machines grew rapidly, both domestically - where the device is employed to evaluate government workers for security clearances - and in places like Iraq and Afghanistan, where Defense Department polygraphers are deployed to extract confessions, check claims about weapons of mass destruction, confirm the loyalty of coalition officers, and grill spies.

The need for a better way to assess credibility was underscored by a 2002 report, *The Polygraph and Lie Detection*, by the National Research Council. After analyzing decades of polygraph use by the Pentagon and the FBI, the council concluded that the device was still too unreliable to be used for personnel screening at national labs. Stephen Fienberg, the scientist who led the evaluation committee, warned: "Either too many loyal employees may be falsely judged as deceptive, or too many major security threats could go undetected. National security is too important to be left to such a blunt instrument." The committee recommended the vigorous pursuit of other methods of lie detection, including fMRI.

"The whole area of research around deception and credibility assessment had been minimal, to say the least, over the last half-century," says Andrew Ryan, head of research at the Department of Defense Polygraph Institute. DoDPI put out a call for funding requests to scientists investigating lie detection, noting that "central nervous system activity related to deception may prove to be a viable area of research." Grants from DoDPI, the Department of Homeland Security, Darpa, and other agencies triggered a wave of research into new lie-detection technologies. "When I took this job in 1999, we could count the labs dedicated to the detection of deception on one hand," Ryan says. "Post-2001, there are 50 labs in the US alone doing this kind of work."

Through their grants, federal agencies began to influence the direction of the research. The early studies focused on discovering "underlying principles," as Columbia's Hirsch puts it - the basic neuromechanisms shared by all acts of deception - by averaging data obtained from scanning many subjects. But once government agencies like DoDPI started looking into fMRI, what began as an exploration of the brain became a race to build a better lie detector.

Paul Root Wolpe, a senior fellow at the Center for Bioethics at the University of Pennsylvania, tracks the development of lie-detection technologies. He calls the accelerated advances in fMRI "a textbook example of how something can be pushed forward by the convergence of basic science, the government directing research through funding, and special interests who desire a particular technology."

Langleben's team, whose work was funded partially by Darpa, began focusing more on detecting individual liars and less on broader psychological issues raised by the discovery of deception networks in the brain. "I wanted to take the research in that direction, but I was hell-bent on building a lie detector, because that's where our funders wanted us to go," he says.

To eliminate one major source of polygraph error - the subjectivity of the human examiner - Langleben and his colleagues developed pattern-recognition algorithms that identify deception in individual subjects by comparing their brain scans with those in a database of known liars. In 2005, both Langleben's lab and a DoDPI-funded team led by Andrew Kozel at the Medical University of South Carolina announced that their algorithms had been able to reliably identify lies.

By the end of 2006, two companies, No Lie MRI and Cephos, will bring fMRI's ability to detect deception to market. Both startups originated in the world of medical diagnostics. Cephos founder Steven Laken helped develop the first commercial DNA test for colorectal cancer. "fMRI lie detection is where DNA diagnostics were 10 or 15 years ago," he says. "The biggest challenge is that this is new to a lot of different groups of people. You have to get lawyers and district attorneys to understand this isn't a polygraph. I view it as no different than developing a diagnostic test."

Laken got interested in marketing a new technology for lie detection when he heard about the number of prisoners being held without charges at the US base in Guantánamo Bay, Cuba. "If these detainees have information we haven't been able to extract that could prevent another 9/11, I think most Americans would agree that we should be doing whatever it takes to extract it," he says. "On the other hand, if they have no information, detaining them is a gross violation of human rights. My idea was that there has to be a better way of determining whether someone has useful information than torture or the polygraph."

Cephos' lie-detection technology will employ the patents and algorithms developed by Kozel's team in South Carolina. Laken and Kozel recently launched another DoDPI-funded study designed to mimic as closely as possible the emotions experienced while committing a crime. In the spring, after this research is complete, Laken will start looking for Cephos' first clients - ideally "people who are trying to show that they're being truthful and who want to use our technology to help support their cases."

No Lie MRI will debut its services this July in Philadelphia, where it will demonstrate the technology to be used in a planned network of facilities the company is calling VeraCenters. Each facility will house a scanner connected to a central computer in California. As the client responds to questions using a handheld device, the imaging data will be fed to the computer, which will classify each answer as truthful or deceptive using software developed by Langleben's team. For No Lie MRI founder Joel Huizenga, scanner-based lie detection represents a significant upgrade in "the arms race between truth-tellers and deceivers."

Both Laken and Huizenga play up the potential power of their technologies to exonerate the innocent and downplay the potential for aiding prosecution of the guilty. "What this is really all about is individuals who come forward willingly and pay their own money to declare that they're telling the truth," Huizenga says. (Neither company has set a price yet.) Still, No Lie MRI plans to market its services to law enforcement and immigration agencies, the military, counterintelligence groups, foreign governments, and even big companies that want to give prospective CEOs the ultimate vetting. "We're really pushing the positive side of this," Huizenga says. "But this is a company - we're here to make money."

Scott Faro, a radiologist at Temple University Hospital who conducted experiments using fMRI in tandem with the polygraph, predicts that the invention of a more accurate lie detector "is going to change the entire judicial system. First it will be used for high-profile crimes like terrorism and Enron. You could have centers across the country built close to airports, staffed with cognitive neuroscientists, MRI physicists, and interrogation experts. Eventually you could have 20 centers in each major city, and the process will start to become more streamlined and cost-effective.

"People say fMRI is expensive," Faro continues, "but what's the cost of a six-month jury trial? And what's the cost to America for missing a terrorist? If this is a more accurate test, I don't see any moral issues at all. People who can afford it and believe they are telling the truth are going to love this test."

The guardians of another Philadelphia innovation that changed the judicial system - the US Constitution - are already sounding the alarm. In September, the Cornell Law Review weighed the legal implications of the use of brain imaging in courtrooms and federal detention centers, calling fMRI "one of the few technologies to which the now clichéd moniker of 'Orwellian' legitimately applies."

When lawyers representing Cephus' and No Lie MRI's clients come to court, the first legal obstacles they'll have to overcome are the precedents barring so-called junk science. Polygraph evidence was excluded from most US courtrooms by a 1923 circuit court decision that became known as the Frye test. The ruling set a high bar for the admission of new types of scientific evidence, requiring that a technology have "general acceptance" and "scientific recognition among physiological and psychological authorities" to be considered. When the polygraph first came before the courts, it had almost no paper trail of independent verification.

fMRI lie detection, however, has evolved in the open, with each new advance subjected to peer review. The Supreme Court has already demonstrated that it is inclined to look favorably on brain imaging: A landmark 2005 decision outlawing the execution of those who commit capital crimes as juveniles was influenced by fMRI studies showing that adolescent brains are wired differently than those of adults. The acceptance of DNA profiling may be another bellwether.

Highly controversial when introduced in the 1980s, it had the support of the scientific community and is now widely accepted in the courts.

The introduction of fMRI evidence at trial may have to be vetted against legal precedents designed to prevent what's called invading the province of the jury, says Carter Snead, former general counsel for the President's Council on Bioethics. In 1973, a federal appeals court ruled that "the jury is the lie detector" and that scientific evidence and expert testimony can be introduced only to help the jury reach a more informed judgment, not to be the final arbiter of truth. "The criminal justice system is not designed simply to ensure accurate truth finding," Snead says. "The human dimension of being subjected to the assessment of your peers has profound social and civic significance. If you supplant that with a biological metric, you're losing something extraordinarily important, even if you gain an incremental value in accuracy."

No Lie MRI's plans to market its services to corporations will likely run afoul of the 1988 Employee Polygraph Protection Act, which bars the use of lie-detection tests by most private companies for personnel screening. Government employers, however, are exempt from this law, which leaves a huge potential market for fMRI in local, state, and federal agencies, as well as in the military.

It is in these sectors that fMRI and other new lie-detection technologies are likely to take root, as the polygraph did. The legality of fMRI use by government agencies will probably focus on issues of consent, predicts Jim Dempsey, executive director of the Center for Democracy & Technology, a Washington, DC-based think tank. "From a constitutional standpoint, consent covers a lot of sins," he explains. "Most applications of the polygraph in the US have been in consensual circumstances, even if the consent is prompted by a statement like 'If you want this job, you must submit to a polygraph.' The police can say, 'Would you blow into this Breathalyzer? Technically you're free to say no, but if you don't consent, we're going to make life hard for you.'"

Today's fMRI scanners are bulky, cost up to \$3 million each, and in effect require consent because of their sensitivity to head movement. Once Cephos and No Lie MRI make their technology commercially available, however, these limitations will seem like glitches that merely need to be fixed. If advances make it possible to perform brain scans on unwilling or even unwitting subjects, it will raise a thicket of legal issues regarding privacy, constitutional protections against self-incrimination, and the prohibitions against unlawful search and seizure.

The technological innovations that produce sweeping changes often evolve beyond their designers' original intentions - the Internet, the cloud chamber, a 19th-century doctor's cuff for measuring blood pressure that, when incorporated into the polygraph, became the unsteady foundation of the modern counterintelligence industry.

So what began as a neurological inquiry into why kids with ADHD blurt out embarrassing truths may end up forcing the legal system to define more clearly the inviolable boundaries of the self.

"My concern is precisely with the civil and commercial uses of fMRI lie detection," says ethicist Paul Root Wolpe. "When this technology is available on the market, it will be in places like Guantánamo Bay and Abu Ghraib in a heartbeat."

"Once people begin to think that police can look right into their brains and tell whether they're lying," he adds, "it's going to be 1984 in their minds, and there could be a significant backlash. The goal of detecting deception requires far more public scrutiny than it has had up until now. As a society, we need to have a very serious conversation about this."

## The Cortex Cop

Your flight is now boarding. Please walk through the "mental detector."

For all the promise of fMRI lie detection, some practical obstacles stand in the way of its widespread use: The scanners are huge and therefore not portable, and a slight shake of the head - let alone outright refusal to be scanned - can disrupt the procedure. Britton Chance, a professor emeritus of biophysics at the University of Pennsylvania, has developed an instrument that records much of the same brain activity as fMRI lie detection - but fits in a briefcase and can be deployed on an unwilling subject.

Chance has spent his life chasing and quantifying elusive signals - electromagnetic, optical, chemical, and biological. During the Second World War, he led the team at the MIT Radiation Lab that helped develop military radar and incorporated analog computers into the ranging system of bombers. In the 1970s, long before the invention of fMRI, Chance began using a related technique called magnetic-resonance spectroscopy to study living tissue. The first functionally imaged brain was that of a hedgehog in one of his experiments. Now 92, Chance still rides his bike to the university six days a week to teach and work in his lab. His mind is as acute as ever. After glancing through a book to confirm a data point, he resumes the conversation by saying, "I'm back online."

He explains that his goal is to create a wearable device "that lets me know what you're thinking without you telling me. If I ask you a question, I'd like to know before you answer whether you're going to be truthful."

To map neural activity without fMRI, Chance uses beams of near-infrared light that pass harmlessly through the forehead and skull, penetrating the first few centimeters of cortical tissue. There the light bounces off the same changes in blood flow tracked by fMRI. When it reemerges from the cranium, this light can be captured by optical sensors, filtered for the "noise" of light in the room, and used to generate scans.

Though near-infrared light doesn't penetrate the brain as deeply as magnetic resonance, some of the key signatures of deception mapped by fMRI researchers occur in the prefrontal cortex, just behind the forehead. The first iteration of Chance's lie detector consisted of a Velcro headband studded with LEDs and silicon diode sensors. Strapping these headbands on 21 subjects in a card-bluffing experiment in 2004, a neuroscientist at Drexel named Scott Bunce was able to accurately detect lying 95 percent of the time. The next step, Chance says, is to develop a system that can be used discreetly in airports and security checkpoints for "remote sensing" of brain activity. This technology could be deployed to check for deception during standard question-and-answer exchanges (for example, "Has anyone else handled your luggage?") with passengers before boarding a plane, or during interviews with those who have been singled out for individual searches.



With funding from the Office of Naval Research, Chance and his colleagues are working to replace the LED headband with an invisible laser and a hypersensitive photon collector to create a system that can pick up the neural signals of deception from across a room.

Before undertaking this project, Chance consulted with Arthur Caplan, director of Penn's Center for Bioethics. "Dr. Chance was a little uneasy about it," Caplan recalls. "But there are certain public places where we lose the right to privacy as a condition of entering the building. Airport security staff is allowed to search your bag, your possessions, and even your body. In my view, there's no blanket rule that says it's always wrong to scan someone without their consent. What we need is a set of policies to determine when you have to have consent."

Chance believes the virtues of what he calls "a network to detect malevolence" outweigh the impact on personal liberties. "It would certainly represent an invasion of privacy," he says. "I'm sure there may be people who, for very good reasons, would not want to come near this device - and they're the interesting ones. But we'll all feel a bit safer if this kind of technology is used in places like airports. If you don't want to take the test, you can turn around and fly another day." Then he smiles. "Of course, that's the biggest selector of guilt you could want." - S.S.