NEUROSCIENCE

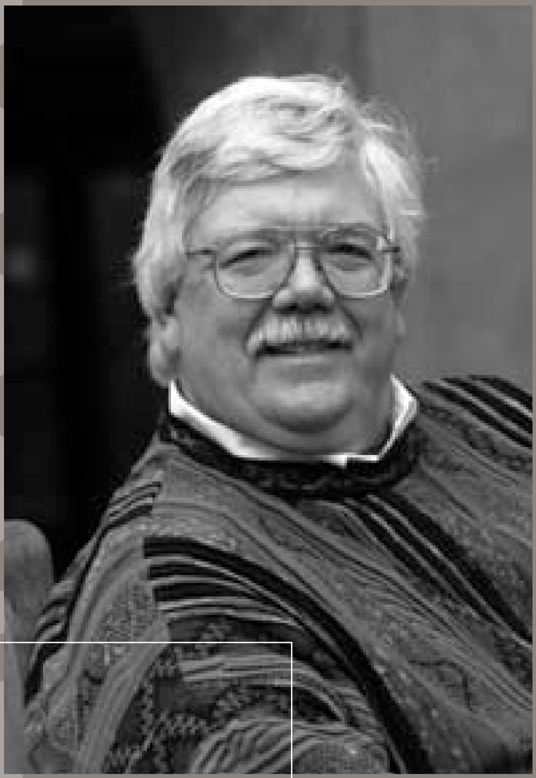

HENRY T. “HANK” GREELY

Deane F. and Kate Edelman Johnson Professor of Law

pp. 82–87

Neuroscience-Based Lie Detection: The Urgent Need for Regulation

HENRY T. "HANK" GREELY (BA '74)

Deane F. and Kate Edelman Johnson Professor of Law;
Professor (by courtesy) of Genetics, Stanford School of Medicine;
Director, Center for Law and the Biosciences; and
Director, Stanford Program on Stem Cells in Society

Neuroscience-Based Lie Detection: The Urgent Need for Regulation

A leading expert and author on the legal, ethical, and social issues surrounding health law and the biosciences, Hank Greely specializes in the legal implications of new biomedical technologies, particularly those related to genetics, neuroscience, and stem cells. He serves as an advisor on California, national, and international policy issues and also chairs the California Advisory Committee on Human Stem Cell Research. Active in university leadership, Professor Greely chairs the steering committee for the Stanford Center for Biomedical Ethics, directs the law school's Center for Law and the Biosciences, directs the Stanford Center for Biomedical Ethics' Program in Neuroethics, and serves on the leadership council for the university's interdisciplinary Bio-X Program. He is one of seven co-directors of a nationwide law and neuroscience project, which recently received $10 million in funding from MacArthur Foundation to examine how breakthroughs in neuroscience might affect the U.S. legal system. He is also a professor (by courtesy) of genetics at Stanford.

The following is an excerpt from Professor Greely's article "Neuroscience-Based Lie Detection: The Urgent Need for Regulation" (*American Journal of Law & Medicine,* 2007), co-authored by Judy Illes, until this year an associate professor of pediatrics (medical genetics) at Stanford and now a professor of neurology and Canada Research Chair in Neuroethics at the University of British Columbia.

D. FMRI-BASED LIE DETECTION—EVALUATING THE RESEARCH

These 12 peer-reviewed articles establish fMRI-based lie detection as a promising technology. They do not prove that it is currently effective as a lie detector in the real world, at *any* accuracy level, let alone the 80 to 90 percent levels being claimed. At least six different issues raise concern about these results: the small number of studies with individual effects, the lack of replication, the small and nondiverse groups of subjects, the inconsistency of reported regions of activity, the artificiality of the deceptive tasks, and the lack of attempted countermeasures.

Of these 12 studies, only three deal at all with determining whether or not individuals are lying. Information that, on average, a group of 12 people showed significant activation in a particular region does not tell us how many of the individual subjects showed activation in that region. It could have been all, many, or only a few. This does not mean these group studies were bad experiments; the researchers did not claim to be testing individual lie detection but instead were looking for broad similarities that might indicate some localization in the brain of lying.

The second problem is the lack of replication of the results by any other laboratories. As is common with fMRI research today, each laboratory tried its own experiments and used its own analytical methods. The only experiments that can be said to have been replicated are the two Langleben experiments (even there with some differences in the test) and the first two Kozel experiments (the third Kozel experiment used a very different model). And only one of those—the second Langleben study—dealt with individual results. A good rule of thumb is to never believe a result until at least one investigator from outside the original group confirms it. Lie detection through fMRI does not pass this test.

A third concern is the number and diversity of subjects. The experiments used healthy young adults, almost all right-handed, with little gender or ethnic

The federal government—or, barring that, state governments—should ban any non-research use of new methods of lie detection, including fMRI-based lie detection, unless or until the method has been proved safe and effective to the satisfaction of a regulatory agency and has been vetted through the peer-reviewed scientific literature.

diversity. No one tested children, the middle-aged or elderly, those with physical or mental illnesses, or those taking drugs, either as medication or illicitly.

Fourth, all of the relevant experiments report finding activation of various regions of the brain (sometimes defined narrowly, sometimes broadly). Together they find activation in many different areas of the brain without strong consistency among the experiments, except when brain regions are very broadly defined. The number of cortical areas activating in these lying and deception tests include anterior prefrontal area, ventromedial prefrontal area, dorsolateral prefrontal area, parahippocampal areas, anterior cingulate, left posterior cingulate, temporal and subcortical caudate, right precuneous, left cerebellum, insula, putamen, caudate, thalamus, and regions of temporal cortex. The activation of many of these regions is known to be correlated with a wide range of cognitive behaviors including memory, self-monitoring, conscious self-awareness, planning and executive function, and emotion. This diversity casts some doubt on the accuracy of any particular method of lie detection.

A fifth problem, and perhaps the greatest, is the artificiality of the deceptive tasks. Most of the experiments involved subjects lying about something unimportant—what card they held or whether they could remember a three-digit number. Only the Kozel paper, involving an instructed "theft" of a ring or a watch from a room, and the Mohamed paper, involving the gun firing, seemed close to more typical real-world lie detection situations. Of course, in those cases, as in every other experiment, the subject telling a "lie" was following an instruction to tell a lie (and in the Kozel experiment the subjects "stealing" the objects were following instructions to "steal" them, knowing it was part of an experiment and not a real theft). Sometimes the subjects were told which lie to tell; other times they got to choose which of two conditions to lie about—but always they were acting not just with permission to lie but under a *command* to do so. It is not clear how this difference from the more usual lie detection settings would affect the results. Although the researchers often told the subjects, falsely, to believe that their success in lying would earn them more money (in fact, researchers with that design paid all the subjects the "extra" money), it is also not clear that this apparent monetary incentive would affect the subjects the same way as the more common—and more powerful—incentives for lying, such as avoiding arrest.

The context points to a deeper problem with the artificiality of the situation—the researchers assume that whatever kind of "lie" they are having the subjects tell is relevant to the kinds of lies people tell in real life. But those lies vary tremendously. Are lies about participation in a crime the same as lies about the quality of a meal or the existence of a "prior engagement"? Do lies about sex activate the same regions of the brain as lies about money, lies to avoid embarrassment, or lies about the five of clubs? Do lies of omission look the same under fMRI as lies of commission? We do not know the answers to these, or many other questions—and neither do the researchers who published these papers. This is not a criticism of the researchers, as scientists have to start somewhere and a well-defined situation is essential for analysis. It is likely to be difficult, and perhaps even impossible, to create good tests of real-world lies. This is a criticism of any attempt to apply this research to the real world without a great deal more work.

All of the concerns discussed so far are reasons to doubt that these experiments did, in fact, prove that one can detect real world lies through fMRI. The last concern is slightly different. Even if the studies had proven that proposition, they did not begin to prove that the method could actually be effective because they did not exclude the very real possibilities that subjects could use countermeasures against fMRI-based lie detection.

The use of countermeasures to polygraphy has been discussed substantially in the past and has even been the subject of some limited research. The National Academy panel on the polygraph spent 10 pages on countermeasures. It concluded as follows:

> If these measures are effective, they could seriously undermine any value of polygraph security screening. Basic physiological theory suggests that training methods might allow individuals to succeed in employing effective countermeasures. Moreover, the empirical research literature suggests that polygraph test results can be affected by the use of countermeasures.

Countermeasures to fMRI-based lie detection could use a wide range of methods. At one extreme, we know a subject can make an fMRI scan useless. Simple movements of the tongue or jaw will make fMRI scans unreadable. Movements of other muscles will introduce new areas of brain activation, muddying the underlying picture. Even less visibly, simply thinking about other things during a task may activate other brain regions in ways that interfere with the lie detection paradigm.

If, as some think, lying is detectable because it is harder than telling the truth and thus requires the activation of more or different areas of the brain, a subject could try doing mental arithmetic or memory tests while giving true answers, thus, perhaps, making true answers harder to distinguish from false ones. Similarly, a well-memorized lie may not activate those additional regions and may look like a truth. The Ganis paper, discussed above, actually reported differences between memorized and improvised lies, though it reported that both were distinguishable on average from the truth.

This issue of countermeasures is both filled with unknowns and vital. If, in fact, countermeasures turn out to be effective, the people we may most want to catch may well be the ones best trained—by criminal gangs, by foreign intelligence agencies, by terrorists, or others—in countermeasures. Of course, if the countermeasures are easy enough, “training” may be as simple as a quick search of the Internet. A quick Google search of “polygraph countermeasures” already turns up many sites offering information on beating the polygraph, some free and some for payment, including one former polygrapher who charges $59.95 (plus shipping) for his manual plus DVD. If fMRI-based lie detection becomes common, efforts to beat fMRI-based lie detection will, no doubt, also become common.

* * * *

Regulation similar to that imposed by the federal and state governments on polygraphs would need to be considered for any somewhat-valid method of lie detection. This would require the federal government plus all state governments (as well as non-American governments) to consider amending many of their statutes, as well as a new surge of litigation in the courts about the admissibility of resulting evidence. At this stage in the technology’s development, however, we believe a simpler solution is preferable. The federal government—or, barring that, state governments—should ban any non-research use of new methods of lie detection, including specifically fMRI-based lie detection, unless or until the method has been proved safe and effective to the satisfaction of a regulatory agency and has been vetted through the peer-reviewed scientific literature. The rest of this section outlines our proposed regulatory scheme in general and raises many additional questions that need to be answered.

* * * *

3. WHY ADOPT THIS REGULATORY PLAN?

The questions raised above are both substantial and difficult. Additionally, they only scratch the surface of

HENRY T. "HANK" GREELY (BA '74)

Neuroscience-Based Lie Detection: The Urgent Need for Regulation

Companies are marketing the age-old dream of lie detection coupled with the high-tech mystique (and beautiful color graphics) of brain scanning. The combination may prove irresistible to many, but with so little evidence that the method is accurate the result may be tragic.

the problems that new methods of lie detection raise. Society will still have to decide a host of puzzling questions, through an analysis of constitutional constraints as well as parsing existing statutes or adopting new ones. The proposal sketched above speaks to whether a lie detection method is safe and effective enough to be used. It does not determine what such a method could be used for, by whom, or under what circumstances. Our rights to and expectations of privacy would have to be weighed against the benefits of lie detection. This is true not only in our dealings with governments, whose actions are limited by the federal Constitution, but in the private spheres of life. Society would need to consider whether laws like the Employee Polygraph Protection Act need to be amended or extended beyond employers to others who may want to use lie detectors—insurers, lenders, contractors, schools, or parents, among others. Completely new licensing schemes might be needed for those who operate these new lie detectors.

Even the preliminary regulation we propose would not pass itself. Congress, preferably, or state legislatures, would have to be convinced to adopt a new, complex, and possibly expensive statutory plan, one dealing with technologies that just barely exist and that may be more widely viewed as blessings than as threats. Why then would—or *should*—Congress act?

Because it is important. Companies are marketing the age-old dream of lie detection coupled with the high-tech mystique (and beautiful color graphics) of brain scanning. The combination may prove irresistible to many, but with so little evidence that the method is accurate the result may be tragic. Honest people may be treated unfairly based on negative tests; dishonest people may go free.

Even in the judicial context, where the *Daubert* and *Frye* tests provide some check on inaccurate evidence, the check is only partial. Each trial judge is empowered to make her own decision, based on the evidence presented in her court. A good lawyer, with a good expert, pushing admissibility of the technology, and a bad lawyer, with a bad (or nonexistent) expert opposing it, could tip the balance in any given court. So could an overly impressionable or scientifically naïve judge. A favorable decision by any single judge anywhere in the country will be trumpeted by the companies selling the technology, in the same way Larry Farwell, the developer of "brain fingerprinting," has publicized his view of the *Harrington* case.

As a result, lives may be ruined. We have seen lives shattered before, with and without these technologies. Wen Ho Lee is one example of a victim of the polygraph. Recent news provides an even clearer example of the costs of investigative mistakes, although not in a case involving (as far as we know) lie detection. In September 2002, Maher Arar, a Canadian citizen who was born in Syria, was returning to Canada with his family from a vacation in Tunisia. While changing planes at Kennedy Airport in New York, he was detained by U.S. officials. After 13 days of questioning—but no formal charges or court action—he was flown to the Middle East, where his American captors delivered him to Syrian security agents. After a year of imprisonment and torture, he was released through Canadian intervention. After a two-year study by a prestigious commission, the Canadian government recently apologized to Arar and agreed to pay him nearly nine million U.S. dollars because of the assistance it gave to the United States in connection with his ordeal. The United States government has not apologized, has sought to have Arar's suit against it dismissed as a result of a "state secrets" privilege, and, indeed, still keeps Arar on its "no fly" list.

Mistakes happen. The inappropriate increase in confidence provided by inaccurate lie detection could make those mistakes worse. Should that happen, and be discovered, the aftermath of the consequent scandal may be discrediting of a tool that, properly verified and controlled, could help prevent such mistakes. Both individuals and the field of lie detection have much to gain from a careful, prudent approach to these new technologies. Currently, nothing enforces such an approach. Requiring proof of safety and efficacy before allowing the use of lie detection technologies is a careful step toward assuring that these technologies are used wisely.

CONCLUSION

We have come a long way, from discussions of the concept of *tu,* mapping, and illustration to the use of individualized, rapidly changing maps of blood flow in the brain to try to detect lies. We need to remember, though, that the map is *never* the territory; the fMRI scan is not the same as the brain it scans. Neuroscience lie detection, if it proves feasible at all, will not be perfect. We need to prevent the use of unreliable technologies and to develop fully detailed information about the limits of accuracy of even reliable lie detection. Government regulation appears to be the only way to accomplish this goal and, by so doing, we take a first step toward maximizing the benefits of these new technologies while minimizing their harms.