CAN BRAIN IMAGING REPLACE INTERROGATION AND TORTURE?

Zen Faulkes
Department of Biology
The University of Texas-Pan American

ABSTRACT

Many techniques have been used to extract reliable information from individuals who are unwilling to divulge it, including interrogation, torture, and "lie detectors," all of which have shortcomings in their ability to get and / or evaluate information. Brain imaging technologies present the possibilities of determining if an individual is lying, concealing information, or has predispositions to particular behaviors. Functional magnetic resonance imaging (fMRI) is the best known brain imaging technique, and can already be used to determine hidden conscious states of an individual, and to determine true and false statements with accuracy greater than chance. Thus, the main empirical question is no longer if brain imaging can be used productively in security situations, but rather how practical it might be and how confident users may be in the information obtained. Ethical questions about appropriate uses of brain imaging technology in security situations are immediate and urgent, but ethical concerns about privacy and similar issues raised by brain imaging appear minor compared to the ethical issues raised by torture. Brain imaging may be able to render some arguments about the use of torture moot by providing a more reliable method of getting and evaluating information from individuals.

Global Virtue Ethics Review
Volume Six, Number 2, pp. 55-78
2011

INTRODUCTION

Concealed information lies at the heart of many security issues. Probably the most common form of concealed information is when one person may know something that another does not. Trust relationships are valuable because they remove reasons and desires for one party to withhold information, but such relationships take time to develop, or, in a worst case scenario, may be effectively impossible to develop. Thus, many security issues would be solved if there were a reliable method to determine if someone was concealing information. To date, no such method has proved reliable, but the development of new brain imaging technologies, particularly functional magnetic resonance imaging (fMRI), has provoked both a flurry of new research on deception, and an almost equal flurry of commentary about the associated ethical implications of the technology.

Because ethical questions turn on empirical facts, the first question to ask about a new technology like fMRI is, "What *is* possible?", followed quickly by, "What *might* be possible?" Answering even the first of these is difficult because of the volume and pace of research. One recent review estimates a new fMRI paper is published about every *three hours* (Logothetis, 2008). Thus, this review is only a single snapshot of a rapidly moving target.

What is fMRI?

Several techniques are used to visualize the brain activity of intact, alert individuals, and the resulting images are colloquially known "brain scans." The most commonly used technique currently is functional magnetic resonance imaging (fMRI). fMRI is used to generate most of the brain scans that are described in the popular press. While fMRI is not as familiar an acronym to the general public as DNA, it is undoubtedly rising in public consciousness.

fMRI images are usually described as showing "brain activity," but what constitutes activity is a complex issue. A few brief reminders of basic neurobiology will suffice for understanding, in broad strokes, what fMRI shows. Brains are made of cells of several types, but the main cells responsible for processing information are neurons. Neurons use two major mechanisms for signaling: electrical and chemical signals. Electrical impulses (action potentials) send signals along the length of a single neuron; chemicals are released to send signals between two different neurons. fMRI measures *neither* of these two things.

There are several ways to get an image using fMRI, but the most commonly used technique is blood oxygen level dependent contrast (BOLD), introduced in the early 1990s (Ogawa; et al., 1990). The magnetic properties of hemoglobin molecules in red blood cells differ depending on whether the blood is carrying oxygen or not. Thus, fMRI is not a direct measure of neuronal activity. Instead, it measures the oxygen consumption of neural tissue, on the hypothesis that highly active neurons consume more oxygen. This assumption is increasingly supported by corroborating evidence (Lauritzen & Gold, 2003; Lee, et al., 2010). Because there is never a time when neurons are not consuming oxygen, however, "brain activity" in fMRI requires comparing a test condition and some baseline condition. Finally, connections between neurons come in at least two varieties: excitatory and inhibitory. fMRI does not show if the neurons in a region are exciting or inhibiting their targets downstream, any more than measuring the volume of sound in a crowded theatre tells you whether the audience is cheering or booing. The nature of signaling is at least as important as the amount of signaling.

Similarly, fMRI's resolution, which seems impressive when viewed in the context of a whole brain, is still low in resolution in terms of the neurons and the connections between them. The shortest time scan fMRI can now resolve is about a tenth of a second (Logothetis, 2008), but neurons can generate hundreds of action potentials per second (Softky & Koch, 1993). The smallest volume of brain tissue that fMRI can now resolve could contain as many as 5.5 million neurons and $2.2–5.5\times10^{10}$ synapses (Logothetis, 2008). For some perspective on the challenge of understanding several million neurons, understanding about 30 neurons has kept many large research labs busy for decades (Harris-Warrick; et al., 1992). Even when the connections between such small numbers of neurons are fixed, the resulting behaviors are wide-ranging, and not locked into a small number of clockwork-like patterns (Marder; et al., 2005). Interactions between neurons are many and subtle, as are the functions they can generate.

Associating cognitive or behavioral functions is further complicated by brain regions being multifunctional (Healy & Rowe, 2007) although this problem may be overcome by imaging the entire brain (2008, p. 19). The function of brain regions can change over time on the order of days and weeks (Kass, et al., Merzenich; et al., 1983a, b) although this is unlikely to cause problems in experiments lasting minutes or hours.

Properly used, fMRI scans do no physical harm to the subjects, although (Kulynych, 2007) notes that fMRI-related injuries have occurred, sometimes due to people being hit by metal objects what were not removed from the area of the scanner.

What is currently possible using fMRI?

In a major review of fMRI, (Logothesis, 2008, p. 869) wrote, "fMRI is not and will never be a mind reader." If "mind reading" is defined as determining someone's private mental state, that statement was incorrect even when it was published. (Haynes & Rees, 2005) used fMRI to study binocular rivalry, in which two different images are presented simultaneously, one to each eye. Because the images do not share common elements, as is typical for images, the viewer perceives these images alternating, one after the other. The timing of the "flip" between these images is not predictable by the viewer. Thus, binocular rivalry is a subjective experience. Using fMRI, Haynes and Rees could correctly predict which image a person was seeing 85% of the time, much better than the 50% expected by chance (Haynes & Rees, 2005). Thus, in this case, fMRI was a mind reader: a crude, imprecise, and limited mind reader, but a mind reader just the same.

More recent studies have performed more complex discriminations. Binocular rivalry only gives people two possible perceptions, but other researchers have used fMRI to distinguish what people are seeing when looking at more complex visual scenes. (Miyawaki, et al., 2008) decoded what shapes and letters individuals were looking at from fMRI signals. (Kay, et al., 2008) were able to identify which naturalistic picture people were viewing from a large set, noting that this is not due to simple mapping of visual stimuli. Preliminary reports indicate that fMRI can be used to predict what movie clips subjects have seen (Nishimoto, et al., 2009).

With regards to detecting deception, a summary of several studies can be found in (Greely & Illes (2007). Briefly, several brain regions show significant increases in fMRI signal during deception compared to truth telling (Ganis, et al., 2003; Luan Phan, et al., 2005). One study using fMRI to detect deception achieved an accuracy rate of 99% within a given tested subject; when those patterns of fMRI signals from one person were used to try predicting deception by other subjects, the accuracy rate dropped to 88% (Davatzikos, et al., 2005). Another study reported accuracy of 93% (Kozel, et al., 2005). In a “mock sabotage” situation, fMRI detected all participants in a mock crime, but wrongly “accused” non-participants of committing sabotage, also known as “false positives”, (Kozel, et al., 2009). A common concern about lie detection is whether intentional deception can be distinguished from an honest mistake. There is some evidence that fMRI may be able to discriminate between these two situations, and discriminate false memories from memories

generated by experience (Abe, et al., 2008; Slotnick & Schacter, 2004) although it appears that fMRI is more likely to be able to decode what someone thinks she or he has seen as opposed to what she or he actually observed (Rissman, et al., 2010). A common pattern described in the literature is that there is no consistent fMRI
signature for true statements; rather, deception appears to place additional demands on cognition, and the concomitant increases in neural activity associated with the additional processing are detectable (Spence, 2008).

(Monteleone, et al., 2009) reanalyzed (Luan Phan, et al., 2005) and concluded that fMRI's record at detecting deception was "well below perfection," with use of one brain region detecting deception in 71% of cases and no false positives, similar to success rates for polygraph testing. Such problems with reliability were also at the heart of one of the first major American legal rulings, which disallowed fMRI evidence (Miller, 2010). The judge wrote that while fMRI did not yet meet the required standards for scientific evidence, he left open the possibility that it might do so in the future. Although precedent is a useful guide both scientifically and legally, it would be premature to think either the re-analysis or the court ruling has "debunked" fMRI's use in detecting deception once and for all, given the furious pace of research using fMRI techniques.

Currently, another problem with using fMRI to detect deception is that fMRI can only be done effectively with an awake and cooperative individual (Heckman & Happel, 2006). The equipment used to take fMRI images is bulky and immobile, and could certainly not be used without a person's knowledge. Likewise, fMRI images are prone to movement artifacts, meaning it would be near impossible to generate an image with someone who is not cooperating. Also, a person cannot be claustrophobic or have any metal in their body, such as metal plates or pins implanted as part of a surgical procedure (Marks, 2007).

What might be possible using fMRI?

Arguably, fMRI has demonstrated "proof of concept" in detecting deception, but the investigation of the neural bases of deception is better described as emerging, basic research rather than as a mature scientific discipline. This seems to be the characterization of the state of the art reached by a recent court decision (Miller, 2010). In the future, given that available computing power tends to double every two years (Moore's Law) and how much of fMRI depends on computation, fMRI will probably become more powerful and accurate than it is now. As basic research progresses, fMRI should also increase in the scope of subjects and mental states it is able to recognize, to the point where it might be possible not only to detect deception, but other unspoken mental states, such as recognition, in a non-cooperative, but otherwise alert subject (Simpson, 2008).

For any imaging technology, the most obvious routes for improvements are to increase spatial or temporal resolution; i.e., to be able to measure smaller areas faster. As data storage and computation power continues to increase, increasing the duration that can be imaged may also be important. There is probably much to be gained in terms of not simply comparing averaged states, but longer and more dynamic processes, which would require more sophisticated algorithms.

Broadening the scope of people and situations where fMRI can be used will also require more basic research on topics such as the amount of variation across individuals and how effective countermeasures might be (Greely & Illes, 2007; Simpson, 2008). The commentaries outlining the challenges of moving fMRI research on deception from controlled lab conditions to real world applications outline a clear research program for the future.

How much accuracy *could* be increased is an open question, but more to the point, few have asked how accurate fMRI *should* be for it to be widely accepted in security settings. For example, Simpson wrote, “Improvements in the technology that would reduce the error rate from 10 percent to something comparable with the billions-to-one accuracy of DNA testing are difficult to conceive of, given the mechanics of the science involved” (Simpson, 2008, p. 497). Putting aside that neither Simpson nor anyone else can accurately predict how much the error rate might be reduced, it is not clear that a technology needs accuracy on the “billions to one level” to be useful. Many procedures routinely used in the legal system also have issues concerning unreliability and repeatability (Saks & Koehler, 2005; Wells & Olson, 2003). For example, fingerprints analysts have sometimes claimed perfect accuracy, but some empirical studies have shown misidentification rates of about 5% (Saks & Koehler, 2005).

Criticisms of fMRI as a lie detector and responses to them

Although fMRI is a new technology, it is not the first attempt to try to use scientific knowledge to assess trustworthiness or detect deception, nor the first new technology to face skepticism regarding its accuracy. fMRI is sometimes compared to phrenology, the practice of associating character traits with the shape of the skull (Van Wyhe, 2002), primarily because both revolve around associating specific functions with specific regions of the brain. In the context of detecting deception, fMRI is compared to polygraph testing. Here, the comparison is more apt, as many of the same important methodological issues, such as the analysis of error rates, are the same in both cases. Polygraph research has not been able to increase accuracy to the level suitable for use in American courtrooms (Committee to Review the Scientific Evidence on the Polygraph, 2003) although the U.S. government still performs polygraph screening (Greely & Illes, 2007). Authors who compare of fRMI to phrenology (Khoshbin & Khoshbin, 2007; Merikangas, 2008) or polygraph testing (Monteleone, et al., 2009) tend to imply that because those two disciplines failed the empirical test, fMRI is also likely to do so. Indeed, one is sometimes left with the impression that many want to dismiss out of hand the possibility that fMRI could *ever* detect deception accurately, as indicated by phrases like, “Could brain scans ever be safe evidence?” (the title to the online version of Geddes, 2008), “fMRI… will never be a mind reader" (Logothetis, 2008, p. 869), or that, "Far from describing the brain and its functions, fMRI… produce(s) models of the brain that reinforce social notions of deception, truth, and deviance" (Littlefield, 2009, p. 365). Crawford argues that many discussions about what fMRI might do are uncritical scientism (Crawford, 2008) or "overconfidence" (Marks, 2007, p. 486).

It seems premature to dismiss the possibility of using brain imaging to detect deception, for multiple reasons. First, there is an obvious and intuitive reason why polygraph testing proved unreliable: poetic assertions to the contrary, people do not think with their hearts, or their blood, or their skin. People *do* think with their brains. This fact alone gives fMRI greater *prima facie* credibility than polygraph testing. Polygraph testing, in essence, tried to create a methodology for detecting deception out of whole cloth. In contrast, fMRI research is tied to a long standing research program in neurobiology. Relating cognitive and behavioral functions to specific brain regions was studied for well over a century before the invention of fMRI, often by studying people who suffered brain injury through strokes or closed head trauma. For example, physician Paul Broca identified a region on the left side of the brain that is closely associated with language processing by studying the behavior of his patients (Broca, 1865). This brain region is now sometimes called "Broca's area," and its role in language processing has been confirmed in broad strokes using more modern techniques like fMRI (Dronkers, et al., 2007; Embick, 2000). That those older medical case studies and newer findings from fMRI experiments usually suggest the same functions for a given brain region increases our confidence in assigning functions to them.

Second, fMRI is less than two decades old. Someone reviewing the capabilities of photography less than 20 years after its invention might scoff at the suggestion that one day photographs would be able to resolve images of planets orbiting other stars – which has now been done (Marois, et al., 2008) – or to resolve images less than a fraction of a second long at micrometer scales.

Third, fMRI is not the only brain imaging technique available (Holstege, et al., 2003). To name just two examples, positron emission tomography (PET) is an older technology than fMRI but is still actively used (Holstege, et al., 2003; Raichle, 2009). Magneto encephalography (MEG) is another functional brain imaging technique that is more recent, and is developing rapidly (Hillebrand, et al., 2005). Thus, even if it turned out that fMRI could never perform some particular imaging task, it does not imply that *no* brain imaging system could ever do the task. More recently, brain *stimulation* techniques might allow experimental manipulations that would shed further light on the neural processing involved in deception (Bolognini & Ro, 2010; Luber, et al., 2009). One experiment using such stimulation improved the lying ability of subjects (Karim, et al., 2010). Another experiment significantly altered subjects' moral judgments such that when decide if some act was morally permissible, stimulated subjects weighed the consequences of an action more than the intent of the action, as compared to controls (Young, et al., 2010).

The prospect of using brain imaging as an alternative to torture

As previously mentioned, there are several practical difficulties in using fRMI in place of interrogation or torture to extract information. Some people with metal in their bodies could not be scanned with fMRI. For people who are claustrophobic or uneducated about what fMRI is, an fMRI scan could be a confusing, or even terrifying, experience that may be as upsetting as non-physical methods used in interrogation to stress the prospective informant. Both of these, however, pale in comparison the most obvious problem: that a subject must be cooperative. If a subject were being cooperative, there would probably be no need for an fMRI. Even if these practical difficulties could be overcome, there are ethical concerns about whether fMRI – or something like it – should be used in such cases.

Most of the ethical concerns raised about fMRI concern its potential use in legal cases in the American judicial system (Greely & Illes, 2007; Khoshbin & Khoshbin, 2007; Luber et al., 2009; Miller, 2010). While important, routine legal cases are generally situations with low urgency compared to the security field (Marks, 2007) which is sometimes tasked with detecting deception and concealed information in urgent and high stakes “ticking time bomb” scenarios; see (Bufacchi & Arrigo 2006) for an analysis of the common elements of such hypothetical scenarios. In these conditions, some have attempted to justify torture on utilitarian grounds (Morgan, 2000; Twining & Twining, 1973). Human rights declarations, however, prohibit torture (e.g., United Nations, 1984). Banning torture, rather than trying to justify it, is regarded as a mark of civilized society discussed in (Bellamy, 2006; Bufacchi & Arrigo, 2006). Under conditions of grave threat, however, utilitarian arguments seem to have more force. This seems to have created situations like those in recent years, when United States agencies engaged in techniques that were previously viewed as torture (i.e., waterboarding) and euphemistically called them "enhanced interrogation techniques" instead (Bellamy, 2006).

Even a strict utilitarian argument hinges on the idea that useful information can be gained from torture. There is little research on the effectiveness of torture to elicit information from unwilling participants (Borum, 2006), and there are no rigorous laboratory studies, for obvious reasons. Nevertheless, neurobiology provides many reasons to believe that information retrieved by torture is unreliable (O'Mara, 2009), given that severe stress typically impairs cognitive functions (McEwen & Sapolsky, 1995; O'Mara, 2009).

A final caveat is that I have assumed that the reasons for lie detection, interrogation or torture are rational attempts to get or verify information. Unfortunately, torture is rarely used in detached, rational ways; claiming that its only purpose is to elicit information is "the interrogation fallacy" of (Twining 1978, p. 147). Torture is surely also used to create fear and terror in groups of people, "the terroristic torture" of (Bufacchi & Arrigo, 2006, p. 360) or simply to satisfy the sadism of those involved in the torturing. Technological solutions, like brain imaging, cannot remove these reasons for torture. Presuppositions dramatically affect how people under investigation are subsequently treatment, and in an environment where torture and extreme interrogation are allowed, fMRI could be used to justify the use of torture (Marks, 2007). The goal must be to create an environment where neither torture nor coercive interrogation is permissible.

Any method that can detect deception in urgent situations that does not do physical or mental harm to the person being questioned must count as a significant force for moral good, even while its use in non-urgent situations is appropriately debated (Kulynych, 2007). If ever there were a situation in which the use of fMRI or other brain scanning techniques could be used legitimately, a “ticking time bomb” scenario would surely qualify. A mature brain imaging technology could make it much more difficult to use the argument that “There is no other alternative” to justify torture. To paraphrase what Isaac Asimov reputedly said about computers, I do not fear brain scans. I fear the lack of them.

References

Abe, N., Okuda, J., Suzuki, M., Sasaki, H., Matsuda, T. and E. Mori. (2008). Neural correlates of true memory, false memory, and deception. *Cerebral Cortex, 18*(12), 2811-2819.

Bellamy, A. J. (2006). No pain, no gain? Torture and ethics in the war on terror. *International Affairs, 82*(1), 121-148.

Bolognini, N. and T. Ro. (2010). Transcranial magnetic stimulation: Disrupting neural activity to alter and assess brain function. *The Journal of Neuroscience, 30*(29), 9647-9650.

Borum, R. (2006). Approaching truth: Behavioral science lessons on educing information from human sources. In R. Swenson (Ed.), *Educing Information: Interrogation: Science and Art-Foundations for the Future*, p. 17-43. Washington, D.C.: National Defense Intelligence College Press.

Broca, P. (1865). Sur le siége de la faculté du langage articulé. *Bulletin dé la Societé d'anthropologie de Paris, 6*(6), 337-393.

Bufacchi, V. and J. M. Arrigo. (2006). Torture, terrorism and the state: A refutation of the ticking-bomb argument. *Journal of Applied Philosophy, 23*(3), 355-373.

Committee to Review the Scientific Evidence on the Polygraph. (2003). *The polygraph and lie detection*. Washington, D.C.: National Academies Press.

Crawford, M. B. (2008). The limits of neuro-talk. *The New Atlantis, 23 (Winter 2008)*, 65-78.
Davatzikos, C., Ruparelb, K., Fana, Y., Shena, D. G., Acharyyaa, M., Lougheadb, J. W., (2005). Classifying spatial patterns of brain activity with machine learning methods: Application to lie detection. *NeuroImage*, *28*(3), 663-668.

Dronkers, N. F., O. Plaisant, T.M. Iba-Zizen and E. A. Cabanis. (2007). Paul Broca's historic cases: high resolution MR imaging of the brains of Leborgne and Lelong. *Brain, 130*(5), 1432-1441.

Embick, D., A. Marantz, Y. Miyashita, W. O'Neil and K. L. Sakai. (2000). A syntactic specialization for Broca's area. *Proceedings of the National Academy of Sciences of the United States of America, 97*(11), 6150-6154.

Ganis, G., Kosslyn, S. M., Stose, L. W. Thompson and A. D. Yurgelun-Todd. (2003). Neural correlates of different types of deception: an fMRI investigation. *Cerebral Cortex, 13*(8), 830-836.

Geddes, L. (2008). Brainscans prepare to enter the courtroom. *New Scientist, 199*, 8-9.
http://www.newscientist.com/channel/being-human/mg20026763.700-could-brain-scans-ever-be-safe-evidence.html

Greely, H. T. and J. Illes. (2007). Neuroscience-based lie detection: The urgent need for regulation. *American Journal of Law and Medicine, 33*(2-3), 377-431.

Harris-Warrick, R. M., E. Marder, I. A. Selverston and M. Moulins. (Eds.). (1992). *Dynamic Biological*

Networks: The Stomatogastric Nervous System. Cambridge: The MIT Press.

Haynes, J. D. (2008). Detecting deception from neuroimaging signals - a data-driven perspective. *Trends in Cognitive Sciences, 12*(4), 126-127.

Haynes, J. D. and G. Rees. (2005). Predicting the stream of consciousness from activity in human visual cortex. *Current Biology, 15*(14), 1301-1307.

Healy, S. D. and C. Rowe. (2007). A critique of comparative studies of brain size. *Proceedings of the Royal Society B: Biological Sciences, 274*(1609), 453-464.

Heckman, K. E. and M. D. Happel. (2006). Mechanical detection of deception: A short review. In R. Swenson (Ed.), *Educing Information: Interrogation: Science and Art-Foundations for the Future* (pp. 63-93). Washington, D.C.: National Defense Intelligence College Press.

Hillebrand, A., Singh, K. D., Holliday, I. E., Furlong, P. L. and G. R. Barnes. (2005). A new approach to neuroimaging with magnetoencephalography. *Human Brain Mapping, 25*(2), 199-211.

Holstege, G., Georgiadis, J. R., Paans, A. M. J., Meiners, L. C., van der Graaf, F. H. C. E. and A. A. T. S. Reinders. (2003). Brain activation during human male ejaculation. *The Journal of Neuroscience*, *23*(27), 9185-9193.

Kaas, J. H., Merzenich, M. M. and H. P. Killackey. (1983). The reorganization of somatosensory cortex following peripheral nerve damage in adult and developing mammals. *Annual Review of Neuroscience, 6*(1), 325-356.

Karim, A. A., Schneider, M., Lotze, M., Veit, R., Sauseng and C. P., Braun. (2010). The truth about lying: Inhibition of the anterior prefrontal cortex improves deceptive behavior. *Cerebral Cortex, 20*(1), 205-213.

Kay, K. N., Naselaris, T., Prenger, R. J. and L. J. Gallant. (2008). Identifying natural images from human brain activity. *Nature, 452*(7185), 352-355.

Khoshbin, L. S. and S. Khoshbin. (2007). Imaging the mind, minding the image: Introduction to brain imaging and the law. *American Journal of Law and Medicine, 33*, 171-192.

Kozel, F. A., Johnson, K. A., Grenesko, E. L., Laken, S. J., Kose, and X. S. Lu. (2009). Functional MRI detection of deception after committing a mock sabotage crime. *Journal of Forensic Sciences, 54*(1), 220-231.

Kozel, F. A., Johnson, K. A., Mu, Q., Grenesko, E. L., Laken, S. J. and S. M. George. (2005). Detecting deception using functional magnetic resonance imaging. *Perception, Empathy, and Reward in Attention-Deficit/Hyperactivity Disorder and Autism, 58*(8), 605-613.

Kulynych, J. A. (2007). The regulation of MR neuroimaging research: Disentangling the Gordian knot. *American Journal of Law and Medicine, 33*, 219-317.

Lauritzen, M. and L. Gold. (2003). Brain function and neurophysiological correlates of signals used in functional neuroimaging. *The Journal of Neuroscience, 23*(10), 3972-3980.

Lee, J. H., Durand, R., Gradinaru, V., Zhang, F., Goshen, I. and D.-S. Kim. (2010). Global and local fMRI signals driven by neurons defined optogenetically by type and wiring. *Nature, 465*(7299), 788-792.

Littlefield, M. (2009). Constructing the organ of deceit: The rhetoric of fMRI and brain fingerprinting in post-9/11 America. *Science Technology Human Values, 34*(3), 365-392.

Logothetis, N. K. (2008). What we can do and what we cannot do with fMRI. *Nature, 453*(7197), 869-878.

Luan Phan, K., Magalhaes, A., Ziemlewicz, T. J., Fitzgerald, D. A., Green, C. and W. Smith. (2005). Neural correlates of telling lies: A functional magnetic resonance imaging study at 4 Tesla1. *Academic Radiology, 12*(2), 164-172.

Luber, B., Fisher, C., Appelbaum, P. S., Ploesser, M. and H. S. Lisanby. H. (2009). Non-invasive brain stimulation in the detection of deception: Scientific challenges and ethical consequences. *Behavioral Sciences and the Law, 17*, 191–208.

Marder, E., Bucher, D., Schulz, D. J. and L. A. Taylor. (2005). Invertebrate central pattern generation moves along. *Current Biology, 15*(17), R685-R699.

Marks, J. H. (2007). Interrogational neuroimaging in counterterrorism: A “No-brainer” or a human rights hazard? *American Journal of Law and Medicine, 33*(2-3), 483-500.

Marois, C., Macintosh, B., Barman, T., Zuckerman, B., Song, I. and J. Patience. (2008). Direct imaging of multiple planets orbiting the star HR 8799. *Science, 322*(5906), 1348-1352.

McEwen, B. S. and M. R. Sapolsky. (1995). Stress and cognitive function. *Current Opinion in Neurobiology, 5*(2), 205-216.

Merikangas, J. R. (2008). Functional MRI lie detection. *Journal of the American Academy of Psychiatry and the Law, 36*(4), 499-501.

Merzenich, M. M., Kaas, J. H., Wall, J., Nelson, R. J., Sur, M. and D. Felleman. (1983). Topographic reorganization of somatosensory cortical areas 3b and 1 in adult monkeys following restricted deafferentation. *Neuroscience, 8*(1), 33-55.

Merzenich, M. M., Kaas, J. H., Wall, J. T., Sur, M., Nelson, R. J. and J. D. Felleman. (1983). Progression of change following median nerve section in the cortical representation of the hand in areas 3b and 1 in adult owl and squirrel monkeys. *Neuroscience, 10*(3), 639-641.

Miller, G. (2010). fMRI lie detection fails a legal test. *Science, 328*(5984), 1336-1337.

Miyawaki, Y., Uchida, H., Yamashita, O., Sato, M.-a., Morito, Y. and C. H. Tanabe. (2008). Visual image reconstruction from human brain activity using a combination of multiscale local image decoders. *Neuron, 60*(5), 915-929.

Monteleone, G. T., Phan, K. L., Nusbaum, H. C., Fitzgerald, D., Irick, J.-S. and E. S. Fienberg. (2009). Detection of deception using fMRI: Better than chance, but well below perfection. *Social Neuroscience, 4*(6), 528-538.

Morgan, R. (2000). The utilitarian justification of torture. *Punishment & Society, 2*(2), 181-196.

Nishimoto, S., Vu, A. and L. J. Gallant . (2009). Decoding human visual cortical activity evoked by continuous time-varying movies. *Journal of Vision, 9*(8), 667.

Nishimoto, S., Vu, A. T., Naselaris, T. and L. J. Gallant. (2009). Decoding visual experiences from brain activity evoked by color natural movies. Program No. 166.15/V12. *2009 Neuroscience Meeting Planner*. Chicago, IL: Society for Neuroscience. Retrieved 29 July 2010, from Society formNeuroscience: http://www.abstractsonline.com/Plan/ViewAbstract.aspx?sKey=420e16ba-b7d1-448e-a3c2-ccfa75c2fdf5&cKey=477fa47c-ef9a-466b-a163-dea5507d95e2

O'Mara, S. (2009). Torturing the brain: On the folk psychology and folk neurobiology motivating 'enhanced and coercive interrogation techniques'. *Trends in Cognitive Sciences, 13*(12), 497-500.

Ogawa, S., Lee, T. M., Kay, A. R. and W. D. Tank. (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences of the United States of America, 87*(24), 9868-9872.

Raichle, M. E. (2009). A brief history of human brain mapping. *Trends in Neurosciences, 32*(2), 118-126.

Rissman, J., Greely, H. T. and D. A. Wagner D. (2010). Detecting individual memories through the neural decoding of memory states and past experience. *Proceedings of the National Academy of Sciences*, *107*(21), 9849-9854.

Saks, M. J. and J. J. Koehler. (2005). The coming paradigm shift in forensic identification science. *Science, 309*(5736), 892-895.

Simpson, J. R. (2008). Functional MRI lie detection: Too good to be true? *Journal of the American Academy of Psychiatry and the Law*, *36*(4), 491-498.

Slotnick, S. D. and L. D. Schacter. (2004). A sensory signature that distinguishes true from false memories. *Nature Neuroscience*, *7*(6), 664-672.

Softky, W. R. and C. Koch. (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *The Journal of Neuroscience, 13*(1), 334-350.

Spence, A. J. (2008). Playing Devil's advocate: The case against fMRI lie detection *Legal and Criminological Psychology, 13*(1), 11-25.

Twining, W. (1978). Torture and philosophy - I. *Proceedings of the Aristotelian Society, Supplementary Volumes, 52*, 143-168.

Twining, W. L. and E. P. Twining. (1973). Bentham on torture. *Northern Ireland Legal Quarterly, 24*(3), 305-347.

United Nations. (1984). *Convention against torture and other cruel, inhuman or degrading treatment or punishment*. Retrieved from http://www.un.org/documents/ga/res/39/a39r046.htm.

Van Wyhe, J. (2002). The Authority of Human Nature: The “Schädellehre” of Franz Joseph Gall. *The British Journal for the History of Science*, *35*(1), 17-42.

Wells, G. L. and A. E. Olson. (2003). Eyewitness testimony. *Annual Review of Psychology*, *54*(1), 277-295.

Young, L., Camprodon, J. A., Hauser, M., Pascual-Leone, A. and R. Saxe. (2010). Disruption of the right temporoparietal junction with transcranial magnetic stimulation reduces the role of beliefs in moral judgments. *Proceedings of the National Academy of Sciences*, *107*(15), 6753-6758.

Biographical Sketch

Zen Faulkes is an Associate Professor of Biology at The University of Texas-Pan American. He publishes research on neurobiology, ethology, and evolution. He is also an active blogger (http://neurodojo.blogspot.com), and has been featured in three editions of the annual anthology of science writing, *The Open Laboratory*.