

## Lie Detection, Mind Reading, and Brain Reading

**Abstract** This chapter briefly looks at the ways that those in the mid-to-late twentieth-century developed lie-detection techniques without neuroimaging – and how various neuroimaging techniques promise more sophisticated types of lie detection. It also very briefly explains how different neuroimaging technologies – such as EEF, fMRI, and fNIR – work, and how they might evolve into more sophisticated – and invasive – techniques in the future, and how law enforcement use of them may thus raise privacy concerns (and do so, even in cases that at first seem free of substantial privacy harms).

**Keywords** Electroencephalography · EEG · Functional magnetic resonance imaging · fMRI · fNIR · Guilty knowledge · Lie detection

### LIE DETECTION BEFORE NEUROIMAGING: POLYGRAPH TESTS

People have long understood that that certain physiological phenomena provide a visible component of certain emotional experiences (such as feeling fear, shock, or surprise). Anxiety, for example, may manifest itself in increased sweating or shaking. Shock may result in pale appearance, and embarrassment may cause blushing. As Francis Shen writes, “humans . . . are natural mind readers,” who have long used various strategies, such as “using facial

expressions and body language to gauge intent” (Shen 2013, 658). In fact, systematic approaches to ferreting out dishonesty have existed far longer than neuroimaging. As Sarah Stoller and Paul Root Wolpe write, “[t]he development of a successful lie detector has been a dream of governments and law enforcement since ancient times” (Stoller & Wolpe 2007, 359).

The use of scientific machinery for lie detection is more recent. In 1921, John Augustus Larsen, a police officer and medical student, invented the polygraph – a machine that uses physiological readings, typically of blood pressure, heart rate, respiratory rate, and electro-dermal activity (electric activity due to skin’s secretion of sweat) (Alder 2007, 6–9). Since that time, moreover, some police and private organizations have sought to improve this biology-based mind reading with technology: They have used polygraph machines to measure changes in blood pressure, galvanic skin response or muscular activity in the hopes such changes could reveal when subjects are lying. In recent decades, some scientists have sought to perfect such physiology-based lie detection. Paul Ekman, for example, has proposed that people display involuntary facial “micro expressions” – “very brief facial expressions, lasting only a fraction of a second” – when they “deliberately or unconsciously concea[1] emotions,” and that these can provide clues (but not proof) that a person may be speaking dishonestly, and trying to conceal their nervousness or discomfort in doing so. (Ekman, *Micro Expression*, at <http://www.paulekman.com/micro-expressions/>).

These familiar methods of lie detection already provide one model for a kind of mind reading: An official can make an inference about someone’s unexpressed mental state – and, particularly, whether she believes what she is stating, or whether she finds a particular word or image more significant than others – by measuring that person’s heartbeat, blood pressure, and other physiological indicators. In fact, when Ronald Allen and M. Kristin Mace asked readers to imagine a mind-reading example – in order to explore if police use of it would violate the self-incrimination clause – the hypothetical they present describes lie-detection methods modeled on those familiar from polygraph machines – one where, even as a suspect refuses to cooperate, and indeed, remains stubbornly, silent, the machine operators “record[s] . . . changes in his heart rate, blood pressure, breathing, and electro-dermal responses (electrical conductance at the skin level),” that occur as he is presented with specific information about a crime (Allen & Mace 2004, 248–249).

To be sure, these methods of mind reading by lie detection have been met with great skepticism. Some observers argue that the most common

lie detection methods don't provide anything close to an accurate measure of honesty. Moreover, the nervousness many people experience in such a test situation may well lead to false positives in spite of the controls, and countermeasures can be taken by which a guilty respondent may be able to control his emotions (Hughes 2014). As Daniel Langleben and Jane Campbell Moriarty note, "most U.S. courts have expressed disapproval of polygraph-based evidence . . . and courts remain largely hostile to its admission into evidence" (Langleben & Moriarty 2013, 223).

Still, it is worth looking more closely at the two most common variants of lie detector tests before we look more closely at neuroimaging techniques – because many uses of neuroimaging as lie detection technology have adapted these methods. Lie detection using a polygraph has typically followed one of two paradigms. In the "control question test," or "comparison question test," individuals are interviewed and then asked a series of questions, each of which falls into one of three types. The questions to which law enforcement typically needs an answer, in the criminal context, are "relevant questions," which are questions about what occurred in a crime (such as, did you kill the victim?). But, if a person subjected to such a question manifests the physiological signs of nervousness or other uncomfortable reactions one would expect to see in someone who would rather not answer, it may be not because she actually committed the crime, but rather because simply being asked a question (in the context of a criminal investigation) is unsettling, and perhaps triggers the fear she is under threat from an authority even if she did nothing criminal. So, the test also asks "comparison" or "control" questions, questions which are also designed to feel threatening, although they are not about the crime. Such questions frequently explore other behavior individuals would be reluctant to admit – for example, whether they have misled friends or spouses, or acted unethically in school or work situations. The premise of asking these two types of question is that the physiological responses for each question will be distinct – so researchers can distinguish questions where individuals merely feel threatened (the control questions) from those where they are lying about their innocence in a crime (the relevant questions). Finally, individuals are also sometimes asked irrelevant questions – questions about mundane issues (Is your name Marc? Are you wearing a white shirt?) that don't threaten the subject and in which the physiological reaction should be at the baseline level (Committee to Review the Scientific Evidence on the Polygraph, National Research Council 2003, 14–15).

While the control question test has been the most widely used method, it has been criticized for uncovering the subject's anxiety when answering in response to questions rather than any reliable proof of whether the subject participated in a crime. David Lykken wrote in 1959, for example, that such lie detection relied on "unreasonable assumptions about the consistency of physiological response patterns" (Lykken 1959). A better alternative, he argued, is to test not for dishonesty or deception, but for "guilty knowledge." As a consequence, some researchers advocate a different kind of method of detecting when someone has participated in criminal or other wrongful activity – the "Guilty Knowledge test" or "Concealed Information Test."

The Guilty Knowledge Test (GKT) presents a subject with multiple choice questions, the answer to which only a participant in a crime could know – for example, a multiple choice question that asks what the murder weapon was, or how criminals in a home invasion and armed robbery made their way into the home, or what kind of vehicle they used as the get-away car. Only someone guilty of the crime should be able to identify the correct answer ("the probe") among the various options, since, as one study notes, the "neutral" or "control" alternatives are carefully chosen so that an innocent subject – lacking knowledge of the crime – would not think they are any less likely to have been involved than the "probe." If an individual consistently shows a greater physiological response when presented with the "probe" than when presented with the controls, it is likely because he has knowledge of the crime (Committee to Review the Scientific Evidence on the Polygraph, National Research Council, *The Polygraph and Lie Detection* 2003, 15). A 2003 study noted that this kind of test has been quite successful in discriminating between those with guilty knowledge and those without it – although it also noted that "almost all attempts to examine the validity of GKT were based on simulations (i.e., mock crime experiments) in which some participants (the guilty) are required to commit a mock crime (e.g., to steal an envelope containing a sum of money and piece of jewelry from a specified office)" (Carmel et al. 2003, 261–262).

Some studies have argued that the guilty knowledge method of lie detection is far superior to the comparison control question method, but critics have still worried that it is susceptible to countermeasures – and also lacks accuracy even when countermeasures aren't used. The use of neuroimaging for lie detection – especially EEG and fMRI – thus strikes some as a possible solution, particularly if such neuroimaging technology improves to

the point of overcoming certain challenges that exist in transferring these methods from laboratory experiments to law enforcement use.

### LIE DETECTION WITH NEUROIMAGING: METHODS OF MONITORING THE BRAIN IN ACTION

Many other books and articles have provided explanations of how brain-based mind reading techniques have worked in the laboratory. I will not discuss them with the same level of detail here. But a little background on the issue is useful for understanding how the law might make use of them – and what constitutional implications they might have (which I will look at more closely in the subsequent chapters).

Each of the 100 billion or so neurons in a human brain participates in a massive and complicated exchange of signals by repeatedly generating an “action potential,” that is an electrico-chemical change that generally begins in branch-like structures called “dendrites,” which receive chemicals released by other neurons. The action potential travels over the membrane of the cell, down a wire-like filament called an “axon,” at the end of which it causes the release of chemicals that travel across a “synaptic gap,” separating the neurons – and, by doing so, either increasing or decreasing the likelihood that another neuron down the chain will “fire.” Somehow, the complex ensemble of action potentials that neurons generate and send to each other makes possible our conscious awareness – and the many aspects of mental life that go with it: Our capacity to have feelings, to remember past events, or imagine what the future would be like (or conjure images and sounds that make up purely fictional scenarios).

While no one knows why or how this neuronal action produces conscious awareness, scientists don’t have to know this to find a pattern in brain activity that seem to consistently arise together with a certain type of mental experience. The question raised for neuroimaging-based lie detection, then, is this: Can neuroimaging identify patterns in brain activity that consistently arise when someone is dishonest, possesses “guilty knowledge,” or has some other mental state that those employing lie detectors wish to uncover? If so, can it do so more reliably than more traditional methods of lie detection (including the methods human beings use when they judge another person’s honesty without the aid of technology) – and reliably *enough* for police or courts to admit them into evidence?

Such questions, however, will not necessarily have the same answer for all methods of neuroimaging. One method of neuroimaging that has been promoted as a basis for lie detection is electroencephalography (EEG). It measures the brain's electrical activity via a set of electrodes placed on the surface of the scalp. Rather than measuring the electrical activity of individual neurons (which is done only with the invasive insertion of tiny microelectrodes – essentially small glass pipettes – beneath the skull), EEG measures the electrical currents that arise in regions of the brain as many millions or billions of neuron fire in sync – in coordinated rhythmic generation of action potentials. This rhythmic activity produces “brain waves” the character of which varies with different kinds of activity. Each electrode detects the electrical signals from the underlying region of the brain, and sends the signal to a machine that amplifies it.

EEGs have been used for decades. The German psychiatrist Hans Berger invented the technique in 1929, while seeking a scientific explanation for how telepathy might transmit ideas from one person's brain to another (Buszaki 2006, 5). While he failed to produce evidence for the possibility of such natural mind reading, his invention has, since the late 1980s, been tested as a possible means of artificial mind reading of a kind – and more specifically, as a substitute for older lie detectors (Langleben & Moriarty 2013, 224).

The most widely discussed form of EEG lie detection has recreated a variant of the “guilty knowledge” or “concealed information” test which, instead of using electrodermal or other traditional measures, uses a kind of brain wave pattern called a “P300 wave.” Scientists have found – in many different types of experiments – that the P300 pattern appears when these subjects view specific kinds of stimuli but not others. Electrical activity generated in the parts of the brain in response to a specific stimulus is called an “event related potential.” Scientists often label certain parts of the wave form generated in such a potential with a “P” if the wave's amplitude is positive and an “N” where the wave's amplitude is negative – followed by the number of millisecond seconds that elapse between the time the subject is shown the stimulus and the time that such a high or low point in the wave appears (Ward 2010). A P300 wave is thus a positive peak that occurs in many individuals roughly 300 milliseconds (and generally within a range of 250 to 400 milliseconds) after a perception of a certain kind of stimulus (Ward 2010; Sur & Sinha 2009 Jan–Jun). Scientists have found that P300 is often elicited by an “oddball” stimulus – that is a stimulus that a subject is instructed to watch for that occurs relatively infrequently and differs

markedly from a more frequently presented stimulus (Wolpov & Wolpov 2012, 216–127).

Researchers have thus hoped that among the stimuli that will trigger a distinctive kind of P300 is an item or scene that strikes them as familiar (and, if it is unlikely to be known by those other than a perpetrator of a crime, would constitute “guilty knowledge”). J. P. Rosenfeld, who initiated the study of EEG-based lie detection in the 1980s, has recently conducted experiments testing the usefulness of P300 measurement to reveal hidden knowledge related to a mock crime. In 2010, he and John Meixner had subjects envision themselves being a part of a terrorism group and write a letter to a fictional terrorist leader about the method, timing, and location of the bombing. When later presented with words from the letter, the subjects showed stronger P300 signals than when reading other words about times, places, and possible weapons (Hughes 2014). More recently, Rosenfeld and Meixner performed another study, which elicited strong P300 waves when they showed subjects words related to activities these subjects had recorded the previous day (while wearing a video camera that recorded for four hours as they followed their normal routine) (Hughes 2014).

The most well-known (and heavily-promoted) use of EEG as a lie-detection method is “brain fingerprinting” – which Lawrence Farwell developed using his own variant of P300-wave measurement: Farwell measures what he calls a “P300-MERMER” signal. This variation uses not only the amplitude peak but also a negative peak that closely follows the positive peak (Farwell 2012). Like the other EEG variants of the guilty knowledge test, brain fingerprinting, as Farwell describes it, is a method of detecting “[w]hen an individual recognizes something as significant” in the “context” of the testing. When an individual does have such recognition, the “Aha!” response he has in this circumstance is accompanied by the P300-MERMER response (Farwell 2012). Farwell has promoted his technology as ready for courts, but other researchers have voiced skepticism, given the proprietary nature of the technology (Brandom 2015).

Since 2000, scientists have also explored whether they might use functional Magnetic Resonance Imaging (fMRI) to detect lies, other dishonest behavior, or concealed information of interest to the justice system. fMRI does not measure the electrical signals of neurons directly. Rather, as Jones, Buckholtz, Schall, and Marois state, “[i]n much the same way that the body delivers more oxygen to muscles that are working harder, the body delivers more oxygen to brain regions that work

harder” (Jones et al. 2009). It does so by delivering more oxygen-rich blood to neurons that are firing more actively than they usually do. fMRI uses a combination of magnetic fields and radio wave bursts to map this flow of oxygenated blood within the brain. In short, it magnetically aligns the spin of hydrogen atom protons throughout the brain, knocks their spin out of sync with radio waves, and then measures the energy they emit as the magnetic field brings the atoms back into alignment: This energy emission is different depending on the oxygen level of the blood, so scientists can use this difference to create a “blood oxygen level dependent” or BOLD signal that varies with the level of oxygen. It cannot create such a measurement for each neuron: Rather, it divides the brain into “voxels” – a three-dimensional equivalent of a “pixel” – each of which consists of about a million neurons (Yuhas 2012).

In recent years, fMRI has received far more attention than EEG methods because of the startling information scientists have been able to infer or reconstruct on a screen – such as pictures and words a person was imagining, or episodes they were remembering or dreaming. As Langleben and Moriarty write, those who think about lie detection have also begun to focus more heavily on fMRI because “recent progress in the ability of fMRI to reliably measure and localize the activity of the central nervous system has created the expectation that an fMRI-based system would be superior to both the polygraph and the EEG for lie detection” (Langleben & Moriarty 2013, 223). fMRI is not superior to EEG in all respects, however: As they note, while fMRI “is greatly superior to EEG in its ability to localize the source of the signal in the brain,” EEG “is significantly less expensive, more mobile, and has a better time resolution than fMRI,” as it is measuring electrical activity in the brain as it changes and not relying on blood flow as a proxy (Langleben & Moriarty 2013, 223). In coming years it is possible that at least some of the EEG’s advantages – such as its portability – can be combined with some of fMRI’s ability to localize brain activity: Functional near infrared imaging (fNIR) makes use of the same premise as fMRI – that is, it measures neuronal activation by measuring the flow of oxygenated blood within the brain. But instead of using a combination of powerful magnetic field and radio waves – requiring a room of heavy equipment – to trace the flow of oxygenated blood, it does so by sending specific wavelengths of “near infrared” light into an individual’s cortex and measuring how the light is absorbed by the brain tissue (Ayaz et al. 2011). This neuroimaging technology is in some respects more limited than fMRI (it can only detect



changes up to 4 cm deep in the brain tissue), but is far more portable and able to measure individuals as they engage in routine activities: measurements require that individuals wear a set of probes and headgear, something they can do as they move around (Ayaz et al. 2011).

In any event, certain fMRI devices – like EEG tests – have had some success in revealing deception or “guilty knowledge” in experimental settings. Neuroscientists have worked since the early 2000s to test fMRI’s success on this front in various scenarios. These studies generally showed that certain activity would increase in certain brain regions when a subject was engaged in deceptive behavior. In 2001, Sean Spence and his colleagues conducted a study in which subjects were instructed to tell the truth or lie, while under a scanner, in response to certain questions about their lives (Spence 2004). When lying, the subject’s brains showed greater activation in certain areas: the ventrolateral prefrontal cortex and the medial prefrontal cortex. In 2002, another experiment conducted by Lee found that certain brain areas showed greater activation when individuals feigned poor memory (Lee et al. 2002). And Langleben conducted another fMRI lie detection experiment in which subjects – who had been given playing cards – were instructed to deny only their possession of one of those playing cards (one which they had received with \$20 in an envelope) when asked if they had particular cards. Thus, they would dishonestly deny having that particular card, but acknowledge having the others they were given. This experiment too found that certain areas of the brain (in this case, the anterior cingulate cortex) showed greater activation during deceptive behavior (Langleben et al. 2002) (and this result may stem not only from the mental acts’ deceptive behavior but also from the distinctive reaction to the card the subject recognizes as the one accompanying the envelope).

As Sarah Stoller and Paul Root Wolpe write, more recent neurotechnological lie detection has moved from focusing on brain regions to “the study of general patterns of brain activation distributed over many regions of the brain,” a technique that “could enable more accurate predictions of cognitive and affective states” (Stoller & Wolpe 2007, 361). Others have done fMRI experiments that closely parallel the kind of guilty knowledge or concealed information tests that, in EEGs elicit a P300 response. For example, one study published in 2012 found that “probes were consistently accompanied by a larger percentage signal change than irrelevant items” in a test asking them to view playing cards, one of which they had selected in an envelope (Gamer et al. 2012, 509).

These lie detection tests, of course, raise significant questions about reliability – and about how jurors will react to presentation of fMRI, EEG or other neuroimaging evidence. As Langleben and Moriarty note, some of the models have been critiqued on the grounds that they measure not dishonesty of the kind one would find in the real world but rather behavior (and accompanying cognition) that occurs when “subjects are explicitly instructed to lie” (Langleben & Moriarty 2013, 224). There have also been questions about whether fMRI lie detection that works in situations where the stakes are low (for example, where students can keep money if they lie successfully) will work in real life situations where someone faces conviction and imprisonment.

As I noted in the [Chapter 2](#), these are among a number of concerns that scientists and other scholars have raised about use of neuroimaging to reveal deception or “guilty knowledge.” Some also raise questions about the use of group data to make an inference about an individual. Barbara Sahakian and Julia Gottwald note that before fMRI can be “used in courts for lie detection, we will need larger studies that determine the levels of accuracy of the technique, especially at the individual level” (Sahakian and Gottwald 2017). They note that experimental methods that work on one set of subjects, may not work on people with different characteristics, and that aspects of real-world lie detection may make techniques that successfully uncover instructed deception in laboratory settings ineffective in other settings. They also take note of a 2011 study by Giorgio Ganis, in which brain scans were able to tell, with 100% accuracy, when an experiment subject was lying about not recognizing the date of their birthday – but also found that these subjects could “disrupt the model” in the experiment, and thus make it difficult to tell false from true answers, by taking simple countermeasures, such as moving a finger or toe (Sahakian and Gottwald 2017).

Other neuroimaging technologies might likewise be used to make inferences about honesty, or about other mental states and activities. Like fMRIs, Positron Emission Tomography (PET) and Single Photon Emission Computer Tomography (SPECT) scanning can measure blood flow within the brain and thus determine what regions become more active in particular tasks. PET and SPECT scanning do so by injecting radioactive isotopes of a molecule (like oxygen) into a test subject’s blood, and then detecting gamma rays produced from the collision of positrons (from the radioactive isotope) with electrons nearby. (SPECT and PET differ in the radioactive isotopes used as “tracers”.) As Francis Shen notes, these methods “have [ ] been used in a variety of criminal and civil cases” (Shen 2016, 501).

Another neuroimaging technology – magnetoencephalography (MEG) – has better temporal resolution than fMRI because, like EEG, it measures neurons’ electrical activity not indirectly, by measuring blood flow, but rather by measuring the magnetic fields generated by this electrical activity (MEG’s spatial resolution, however, is poorer than that of fMRI) (Snead 2007, 1282, Sahakian and Gottwald, 2017).

### NEUROIMAGING BEYOND LIE DETECTION: SCREENING MOVIES FROM THE MIND

My focus here is not on all of the legal questions that might be raised about this technology (such as whether they can produce admissible evidence), but rather on the possible privacy implications of these technologies. On the one hand, use of an fMRI to test for deception – or to see if someone recognizes an object presented by a test administrator – seems far less threatening to privacy than the mind-reading devices of science fiction, that can pull memories out of individuals’ minds. As Fox notes, it makes a constitutional difference that such technologies are “not capable of exposing the content of a subject’s cognitive thoughts” – and instead give government access only to “the less privileged sphere of sensory recall and perceptual recognition about a particular set of facts or the state of past events” (Fox 2008, 2). Moreover, it is the research study itself that will typically supply the content initially: In the comparison question test and any neuroimaging analogues of it, it is the researcher who asks the questions – with the subject giving a simple “yes” or “no” response.

Still, it is conceivable that such technology can – even in something close to its present state – be used in ways that threaten individual privacy. First, even a single yes-no question, or multiple choice question designed to elicit guilty knowledge, might undercut privacy when it concerns an issue that a person regards as highly sensitive (such as that person’s sexual behavior, or whether they have a certain medical condition). Second, if officials have the opportunity to pose *many* yes/no or multiple choice questions – and learn something from a subject’s EEG or fMRI response even where the subject refuses to answer – they might learn a good deal about the subject’s private life. Third, neuroimaging may well be used in conjunction with other investigatory tools, which themselves present a threat to privacy. For example, an investigator may not need an fMRI or EEG test to tell where a subject has been, or what she has done in the past

week, if she has access to a week's worth of location-tracking information, video footage, or text messages and e-mails. In a case like that, fMRI and EEG might simply fill in some blanks – for example, about person's attitude toward people and places in the records authorities already possess. In fact, even where an official has significant doubts about an inference she has drawn from neuroimaging evidence, she might conceivably use such evidence as a starting point (or intermediate step) in an investigation in which she later confirms an educated guess, based on neuroimaging, with evidence from other sources.

Still, future development in neuroimaging technology can substantially change the nature of the privacy threat it poses. Lie detector-style neuroimaging does not, of course, provide individuals with the detailed information about memories or thought that one finds in a personal diary. Such detail would require mind-reading technology akin to that in science fiction where investigators turn brain activity that correlates with internal thoughts into verbal descriptions, videos, or some other representation of a person's memories, beliefs, or dreams.

This is not something that current neuroimaging technology can do. But there have been small experimental steps in that direction. These generally work by using many fMRI readings of individual test subjects to create a dictionary of sorts, each entry of which matches specific fMRI readings with a particular act of cognition or perception. Researchers, for example, might establish what pattern of brain activity (as seen in an fMRI machine) arises when one is thinking about a particular word, such as “house,” or viewing a particular image, or a specific person's face, or hearing a voice with certain characteristics. Then, when they see a previously categorized fMRI reading (or something very like it), then can infer that the person is likely to be experiencing something similar to the matching perception or thought.

The process is, of course, a complex one – and relies not only on the power of fMRI machines, but also on advances in computer technology, as it is computer algorithms that match the complex fMRI readings with specific cognitive tasks.

One research team, headed by Jack Gallant, was able to use fMRI technology to reconstruct simple images – and then, in a later experiment, videos – that someone was watching from fMRI readings of brain states. In the image experiment, a computer was able to guess, with a high level of accuracy, which image (in its library of images) a person was likely viewing. The video translation worked on the same principle – but went further: The research did not simply “guess” what parts of a video clip

someone was watching (based on a match with fMRI readings taken during previous viewings of the video), it also used a computational model to predict what brain activity patterns would arise as individuals watched other movies (not yet viewed by the subject) – and then used this model to reconstruct additional videos individuals were watching, as they watched them, from the sequence of BOLD signals it detected in the fMRI readings (Nishimoto et al. 2010; Smith 2013). Another study was similarly able to use fMRI readings of activity in the brain’s hippocampus to tell where in a virtual-reality environment an experiment subject was, as they navigated through it (Chadwick et al. 2010). Still another study allowed fMRI machines to reconstruct the faces people were viewing. This could conceivably allow witnesses to reconstruct the face of a potential culprit by lying in an fMRI machine and imagining a face instead of describing it to a sketch artist or picking it out of a photographic line-up (Cowen et al. 2014).

These experiments reconstructed images from brain activity that occurred as the individual was viewing an image or watching a video. But researchers have also been able to guess or reconstruct what people are imagining in their “mind’s eye.” Marcel Just and Tom Mitchell were able to use fMRI technology to identify what patterns particular words and concepts generate in a person’s brain activity when someone prompts them with the word or concept, and to develop models that can predict what kind of fMRI pattern would arise for other words (Shinkareva et al. 2008). Another group of researchers using fMRI instruments was able to determine which of several film clips someone had seen. And a group of researchers in Japan was even able to produce “dream recordings,” using fMRI readings to determine – with 60% accuracy – what objects people had reported seeing during dreams. In fact, they were able to reconstruct crude videos of the dream’s imagery based on these measurements.

Science fiction has already envisioned how far more advanced variants of this technology might work in a futuristic legal system. In the movie, *Strange Days* (1995), for example, a key witness to a murder is herself killed, but a recording of her memory remains available (and persuasive, in a world where the technology for creating such recordings is well-established). Obviously, if and when neuroimaging should ever become capable of extracting from our mind such vivid movies of our past experience, or transcripts of our silent thinking, the privacy threat it presents will be far graver than it is with technology of the kind that exists now.

It is also possible that instead of revealing what people are thinking or remembering at a particular time, neuroscience technology will instead

reveal certain enduring features of people's personalities. Other biological research – for example, on DNA features – may already provide scientists with methods for making inferences of this kind. For instance, some studies have correlated shyness and anxiety with having a particular variant of the “serotonin transporter promoter” gene (Battaglia et al. 2005, 85, 91). Thus, for example, while character traits are normally inadmissible under the Federal Rules of Evidence in US courts, where character evidence *is* admissible, it is conceivable one could offer evidence of this kind not only through lay testimony, but also through biological evidence of character (where courts allows experts to present scientifically-informed evidence of character traits).

Neuroimaging evidence may also be a source of such information. Various studies have correlated differences in personality with differences in a person's “connectome.” Just as the “genome” is the entire sequence of nucleotides in your DNA,” a connectome, says Seung, is “the totality of connections between the neurons in a nervous system” (Seung 2012, preface). To the extent this pattern of connections partly makes us who we are, should fMRIs be able to discern information about it, they might be able to uncover features of our personal predispositions. Certain studies have used a variant of fMRI called “resting state fMRI” – which looks at how the brain acts when we are resting as opposed to performing a specific cognitive task – to find correlations between certain personality types with certain brain patterns (Adelstein et al. 2011, 4–5). fMRI evidence revealing information about subjects' responses to particular stimuli might also allow inferences about their personalities: An article by Martha Farah and her colleagues, looking at whether fMRI studies might raise privacy concerns, notes that “functional neuroimaging is, indeed, already capable of delivering a modest amount of information about personality, intelligence, and other socially relevant psychological traits” (Farah et al. 2010, 126). It also notes that it is conceivable that scans that an individual takes for one reason (such as measuring face perception) may also incidentally collect information about other aspects of the subject's psychology (for example, because “extraversion and unconscious racial attitudes are both correlated with brain activity evoked by simply viewing pictures of faces”) (Farah et al. 2010, 110).