

Functional MRI-based lie detection: scientific and societal challenges

Martha J. Farah, J. Benjamin Hutchinson, Elizabeth A. Phelps and Anthony D. Wagner

Abstract | Functional MRI (fMRI)-based lie detection has been marketed as a tool for enhancing personnel selection, strengthening national security and protecting personal reputations, and at least three US courts have been asked to admit the results of lie detection scans as evidence during trials. How well does fMRI-based lie detection perform, and how should the courts, and society more generally, respond? Here, we address various questions — some of which are based on a meta-analysis of published studies — concerning the scientific state of the art in fMRI-based lie detection and its legal status, and discuss broader ethical and societal implications. We close with three general policy recommendations.

Over the centuries, human beings have devised many different methods for the detection of deception (BOX 1). Some are low-tech — for example, the skilled recognition of facial expressions — and some are high-tech, including the polygraph, a device that measures autonomic arousal and is known in popular culture as the ‘lie detector’. At present, no method of lie detection has been proven to perform with high accuracy, and the search for a better method continues¹.

Recent efforts to detect lies have focused on measures of the brain. The appeal of this brain-based lie detection approach is that, in contrast to most previous methods — which detected the emotional arousal resulting from deception — it measures physiological changes associated with cognitive processes during deception and could therefore, in principle, be detecting the process of deception itself. Most functional imaging attempts to discriminate lying from truth telling have used functional MRI (fMRI), although a few early studies used positron emission tomography, and other methods (event-related potentials and functional near-infrared spectroscopy) have been applied to the related problem of detecting concealed knowledge^{2,3}. Scientific and legal interest in fMRI-based lie detection has developed rapidly. The majority of scientific articles on this topic have been published within the past decade, and there have been at least three attempts to have fMRI-based lie detection admitted into US courts since 2010.

In this Perspective article, we assess the current state of the science in fMRI-based lie detection and review some of the legal

and societal issues raised by this technology. Beginning with the science, we address three questions about the current state of the art in fMRI-based lie detection. First, do current findings on lie detection, from different laboratories and using different experimental tasks, identify a consistent set of brain regions and, if so, which areas are they? Second, how confidently can we interpret the results of these studies with respect to the neural substrates of deception per se, and what alternative interpretations have yet to be ruled out? Third, what additional challenges do we face in the effort to use fMRI for the detection of deception in real-world contexts? We then raise a series of issues concerning the ethical, legal and societal impact of attempting to detect lies with fMRI.

The science of fMRI-based lie detection

Although fMRI-based lie detection has been commercialized and is used by some for real-world applications, research on this topic began as a form of basic science with the goal of identifying the neural systems involved in deception^{4–6}. In such studies, blood-oxygen-level dependent (BOLD) activity is measured under conditions in which subjects are instructed or explicitly permitted to make deceptive versus truthful responses. Deception has been operationalized in many different ways in fMRI-based lie-detection research as well as in lie-detection research more generally (BOX 2). The designs of these studies are crucial for understanding the degree to which they successfully isolate the neural correlates of deception, so several examples of research tasks are given here.

In one of the earliest studies, subjects were given two playing cards and were instructed to deny possession of one and acknowledge possession of the other⁵. Subjects underwent fMRI scanning while they viewed a series of cards, including the two critical cards and other cards that they had not been given. A comparison between ‘truth’ trials and ‘lie’ trials revealed brain activity associated with deception. In a similar task design, subjects mentally picked a number between three and eight, and when shown a series of numbers on a screen, denied having picked that number (the critical lie item) and also denied having picked the other numbers (truth items) while undergoing fMRI⁷. Here again, activity on lie trials was compared with that on truth trials to discover which areas were associated with deception in this task.

A small step towards more realistic experimental paradigms to study lying was taken by Kozel and colleagues⁸. They devised a mock crime scenario, in which subjects were given a choice whether to ‘steal’ a ring or a watch and were instructed to place the chosen item in a locker. Subsequently, they answered questions about the ‘crime’ while in the scanner and were instructed to deny having taken either item. In addition to lies (denials regarding the item they took) and truthful statements (denials regarding the item they did not take), subjects were asked yes-or-no general knowledge questions such as ‘is it 2004?’ or ‘do you live in the United States?’ The difference between the lie and truth trials was taken to index the neural activity associated with lying. Finally, researchers have compared subjects’ responses when answering questions about past events or personal information, for which they were cued to respond truthfully or deceptively⁹.

In summary, for most fMRI studies of lie detection, subjects are instructed by the researchers to lie and to tell the truth on specific trials of the experiment, and the activation in lie trials and truth trials is either directly contrasted or compared after contrasting each trial type to a baseline condition. The regions showing significantly greater activation for lies than for truth are taken to be the neural correlates of deception.

Consistency of results across laboratories and tasks. A comparison between the deceptive and truthful response task conditions often results in activation in certain regions, particularly the prefrontal cortex, anterior cingulate cortex and parietal cortex. How reliably are these regions associated with

Box 1 | Lie-detection technology through the ages

Spotting a liar is hard to do. Studies have repeatedly proven this using tests in which observers perform barely better than if they were blindly guessing⁵⁸. Many different methods have been devised to increase the accuracy of detecting deception. Among behavioural methods, some focus on non-verbal behaviours that are indicative of emotion, such as fleeting facial 'microexpressions' (REF. 59), tone of voice⁶⁰ and many other facial, vocal and bodily cues, with modest success at best^{61,62}. Others focus on the verbal content of speech, which can reveal both emotional and cognitive concomitants of lying⁶³.

Functional MRI (fMRI)-based lie detection is only the most recent attempt to exploit changes in nervous system function associated with lying as indicators of deception. The earliest methods focused on the autonomic nervous system, and in particular on the tendency for lies to be accompanied by activation of the sympathetic nervous system (SNS). In ancient China, suspected liars were forced to fill their mouths with dry rice and then spit it out. As sympathetic arousal suppresses salivation, the rice would adhere more to the mouths of liars, who as a result would take longer to spit the grains out⁶⁴. In the early twentieth century, Harvard University psychology student and creator of Wonder Woman comics William Marston explored systolic blood pressure — another sign of SNS activation — and its relation to deception. Wonder Woman's 'Magic Lasso', which forced its captives to tell the truth, was inspired by the blood pressure cuff⁶⁵.

The modern polygraph was based on Marston's invention, with the addition of other measures of SNS activity — namely, heart rate, respiration and perspiration. Uncertainty regarding the accuracy of polygraphy under field conditions and its dependence on examiner skills and attitudes^{1,66} had curtailed its use by the end of the twentieth century⁶⁷. In the United States, it is outlawed for non-governmental pre-employment screening and is rarely used in court⁴⁷. However, to the dismay of many analysts, the polygraph nevertheless remains widely used in US national security employee screening¹.

The past few decades have seen efforts to develop measures of CNS activity to index deception, including scalp-recorded electroencephalography (EEG) measures and functional near-infrared spectroscopy (fNIRS). Among EEG-based methods, so-called 'brain fingerprinting' uses components of the event-related potential to infer whether the subject is familiar with aspects of a crime that only the perpetrator would be expected to know⁶⁸. Brain Electrical Oscillations Signature (BEOS) profiling is similarly intended to detect 'guilty knowledge' and has been widely used in Indian courts⁶⁹, although we are not aware of any peer-reviewed studies of the method. Finally, fNIRS uses NIR light, to which the skull is relatively transparent, to measure blood oxygenation and has been explored as a means of detecting CNS activity associated with lying⁷⁰. Its advantages over fMRI include price and portability; its major disadvantage is its field of view, which is limited to superficial brain regions, as well as a lower signal-to-noise ratio.

deception, and can one identify more specific subregions within these relatively broad anatomical areas that are activated during deception? Earlier reviews of the literature found substantial consistency across studies^{10–12}, but the literature has continued to grow rapidly in recent years. To assess the consistency of deception-related activity across laboratories and tasks, we therefore carried out a new meta-analysis of the fMRI-based lie-detection literature to date.

Like the earlier analyses of Christ *et al.*¹¹ and Wagner¹², our meta-analysis used the activation likelihood estimation (ALE) method¹³ (for details, see [Supplementary information S1](#) (box)). ALE quantifies the degree of anatomical overlap across published neuroimaging studies that are based on peak-voxel coordinate information. This enabled us to quantify the reliability of anatomical overlap of observed activation foci rather than simply analysing commonalities in terms of activations that occur in pre-selected regions of interest. Published lie-detection studies were included if they: used fMRI; reported

results from a whole-brain, group analysis of healthy, young adults; conducted a statistical contrast indexing deception and reported one or more foci in standardized coordinate space; and reported data that had not been reported (in part or in full) in any other study included in the meta-analysis (that is, a given dataset contributed to the analysis only once). Critically, the activations analysed were group-level contrasts of deceptive versus truthful responding ([Supplementary information S2](#) (table)). As detailed in [Supplementary information S1](#) (box), this analysis was performed over 321 foci from 28 independent statistical contrasts between lie versus truth conditions that were reported in 23 different studies.

As shown in [FIG. 1](#) (see also [Supplementary information S3](#) (table)), the meta-analysis revealed a number of regions that were active during deception versus truth conditions across studies at an above-chance rate, including the bilateral dorso-lateral and ventrolateral prefrontal cortex, inferior parietal lobule, anterior insula and

medial superior frontal cortex. As previously noted by others, there was considerable variability from study to study, as no region was active in all (or nearly all) studies. This may be due in part to differences in tasks and stimuli⁹, in data acquisition procedures (for example, magnet field strength and acquired functional data resolution as well as the number of trials per condition and number of subjects, both of which affect statistical power) and in statistical procedures (for example, choice of statistical thresholds). Nevertheless, there is considerable agreement across studies on which brain areas are more active during instructed lying than during truth telling.

Could the patterns documented by this meta-analysis provide the scientific foundation for a useful lie detector? We suggest that several other empirical questions need to be addressed before fMRI-based lie detection can be considered for real-world use, such as: are the observed brain activations due to deception per se or to confounds within the experimental designs? More generally, is the observed activation specific to lying or does it reflect something about the way lies are usually (but not necessarily or invariably) produced? Can fMRI discriminate lies from truth in individual subjects with sufficient accuracy to be useful in at least some circumstances? If so, for what types of subjects? Do laboratory-derived indicators generalize to the real world, in which stakes and hence emotions may be high, base rates may be unknown and subjects may attempt countermeasures?

Experimental confounds and other questions of specificity. The experimental designs used in fMRI studies of deception are most naturally described in terms such as 'lie' and 'truth', and this language encourages a certain presumption of specificity. However, on the basis of simple experimental contrasts between lie trials and truth trials, we cannot know what psychological processes other than deception might evoke the same patterns of activity. What can we conclude from the current literature concerning the specificity of deception-related activation?

In most of the tasks used to study deception with fMRI, a number of experimental factors are confounded with the lie-versus-truth manipulation^{7,12,14,15}. For example, in one study^{16,17}, BOLD activity differed in instructed-lie versus truth trials, but the frequency of the motor response required on truth trials was much lower than that required on lie trials. Because of this confound, rather than reflecting neural

Box 2 | Varieties of deception

Lying is not a single homogeneous category of behaviour^{9,71–73}. There are lies that we tell for personal gain or to spare the feelings of another, lies that we regard as consequential or trivial, lies that we tell to others or to ourselves and lies that deceive by active assertion of a falsehood or by mere omission of information. Each of these may have different neural correlates, depending on factors such as those discussed in the article, including the degree of rehearsal or emotion associated with the lie⁹.

An issue that has occupied lie-detection researchers working with polygraphy is the distinction between lies elicited in different types of task^{74,75}. In the 'control question test', which was developed for use with the polygraph, subjects are asked three types of question: relevant questions (that is, questions that directly elicit the suspected lie), control questions (that is, questions about topics that would cause most people discomfort or shame, such as childhood stealing or betrayal of a friend) and irrelevant questions (that is, questions that elicit innocuous information such as name or birthplace). Truthful responding to the relevant questions is expected to evoke less arousal than responses to the control questions, whereas lying is expected to evoke more arousal. In the 'concealed information test', also known as the 'guilty knowledge test', the subject is presented with correct and incorrect answers to questions about facts that only the perpetrator of a crime would be likely to know, and so truthful subjects are assumed to show equivalent arousal to all of the answers and deceptive subjects are assumed to show more arousal to the correct answers. The functional MRI literature includes tasks that are similar to both types of test, although the exact conditions and comparisons differ from the polygraph literature.

correlates of deception, critical aspects of these data may reflect neural responses associated with selecting a frequent versus an infrequent motor action.

In an experiment with even broader implications for the interpretation of existing fMRI-based lie-detection data, Hakun and colleagues⁷ carried out an experiment in which subjects were scanned while viewing a series of numbers after having chosen one in advance. In one condition, subjects were instructed to lie by responding that they had not chosen the selected number when it was shown during scanning. In a different condition, subjects simply viewed the numbers without responding during scanning. Strikingly, in three out of three subjects, there was greater lateral prefrontal and parietal activation in response to the chosen number relative to other numbers, both when subjects simply passively viewed the numbers and when they were instructed to lie about the chosen number⁷. This finding suggests that in many of the studies in our meta-analysis, the greater activity observed for instructed lie stimuli may not reflect neural processes related to deception but rather may reflect other cognitive differences produced by the task. For example, the mere act of selecting a stimulus at the outset of the experiment (whether it is selecting a specific number from the range of three to eight or selecting a ring rather than a watch from a drawer) may attach particular significance to the stimulus that alters subsequent cognitive responses to the stimulus when it appears during fMRI scanning. The selected stimulus may be more salient relative to the other ('truth')

stimuli, resulting in differential engagement of neural mechanisms of attentional orienting, and the selected stimulus may also be associated with stronger or richer memories, resulting in differential engagement of neural mechanisms of memory. Are the activation differences between 'lie' and 'truth' stimuli due to the act of deception or to such confounding effects of attention and memory?

The presence of memory confounds in fMRI-based lie-detection studies was directly addressed in an important study by Gamer *et al.*¹⁴. In that study, subjects were instructed to encode critical items (a banknote and a playing card) into memory, and were then scanned while simply viewing the critical items and control items (subjects pressed a button to indicate the presentation of each stimulus). Because this stimulus viewing task did not require subjects to respond deceptively or truthfully, differences in the BOLD signal for critical items versus control items must be due to differences in the items' histories; that is, whether the subject does or does not have a memory for the item. Strikingly, the results revealed greater activation in response to the critical items versus control items in the same prefrontal and anterior insular cortical areas that were previously observed in fMRI studies comparing instructed lie versus truth trials. Because memory and attention confounds are present in many fMRI-based lie-detection studies, the data of Gamer *et al.*¹⁴, along with those of Hakun *et al.*⁷, suggest that the consistent findings obtained in our meta-analysis may reflect activation related to processes other than deception per se.

Related to these concerns about experimental confounds is the broader issue of process (or functional) specificity. Even if it were possible to correct the memory and attention confounds inherent in the tasks just described, there would probably remain an association between deception on the one hand and executive function and other cognitive processes on the other hand, because deception generally places greater demands on memory and executive functions than does truthful responding: the liar must generally keep two versions of events in working memory and must inhibit the more natural response of responding in accordance with reality. It has been noted that the regions in which activation is associated with deception in the fMRI-based lie-detection literature are also associated with executive function, attention and memory processes, which is consistent with the association between deception and cognitive load^{11,12}. As many have argued, it is therefore possible that fMRI-based lie detection measures differences in the engagement of these more general-purpose cognitive processes. To the extent that deception typically (but not necessarily) imposes a higher cognitive load than truth telling, it may be possible to dissociate the two using fMRI. Under certain circumstances, true responding could tax these processes to an equivalent or greater extent than deceptive responding, a pattern that has in fact been observed in at least one study¹⁸. Thus, despite the encouraging consistency revealed by the meta-analysis (FIG. 1; Supplementary information S3 (table)), truth telling could be mistakenly interpreted as deception according to current methods of fMRI-based lie detection.

From the laboratory to the real world

Even if all of the uncertainties just described could be overcome, many other issues would have to be addressed before fMRI-based lie detection could be used responsibly in the real world. We turn to these issues here.

Inferences about individual subjects. Real-world uses of lie detection will of course involve inferences about the truthfulness or deceit of individuals. What do the published studies tell us about the accuracy with which deception can be identified at the individual level? Most publications report only group analyses, which makes them poorly suited to answering this question. Of the minority of studies that report statistics that are directly relevant to assessing accuracy at the level of individual subjects or individual events^{2,7,8,16,17,19–27}, only two studies




Figure 1 | Results of the ALE analysis of the functional MRI ‘deception’ literature. Overlay of map of activation likelihood estimation (ALE) values (orange) on the lateral (top) and medial (bottom) inflated PALS surface⁷⁶, revealing regions that are consistently implicated in deception across studies. The detection thresholded was set at $P < 0.05$, and the false-discovery rate was corrected as per the method described in REF. 13. IFG, inferior frontal gyrus; IPL, inferior parietal lobule; MFG, middle frontal gyrus; m/SFG, medial and superior frontal gyrus.

(to our knowledge) report data relevant to detecting deception at the individual-event level. Specifically, Langleben *et al.*¹⁶ and Davatzikos *et al.*¹⁷ focused on whether instructed lie and truth events could be discriminated in the same dataset, using either logistic regression¹⁶ or non-linear machine-learning analyses¹⁷. Although these event-level analyses showed accuracy rates of 78% and 88%, respectively, the impact of these early findings is limited because of the above-noted response-frequency confound in the experimental design.

Other studies focused on examining whether fMRI BOLD activity differs when an individual is lying compared to telling the truth (pooling data across events). Various statistical approaches have been implemented, including single-subject univariate analyses^{2,7,19,20}, univariate analyses combined with the counting of above-threshold voxels in targeted regions of interest^{8,22,25} and machine-learning classification². The reported accuracies in these individual-subject-level analyses have ranged from 69% to 100%, suggesting that these statistical approaches have promise. However, here again, the noted concerns about attention and memory confounds undermine data interpretation^{7,12,26}. Indeed, so long as studies use tasks in which the effects of deception cannot be separated from the effects of attention and memory, the problem of confounds will remain, regardless of

whether the neural correlates of deception are sought using univariate or multivariate analyses and regardless of whether the correlates are discovered by simple regression analysis or machine-learning algorithms. Furthermore, even high accuracy rates may decline precipitously when subjects use countermeasures in an attempt to conceal their ‘deception’ (REF. 2).

The laboratory studies assessing the accuracy of fMRI-based lie detection on the individual-subject level assess the sensitivity and specificity within an individual by differentiating trials on which the individual is deceptive or truthful. However, determining the accuracy of a test in a general population also requires an assessment of the test’s sensitivity and specificity across individuals within that population: that is, what is the likelihood of detecting deception when it is present in a member of the population (sensitivity), and what is the likelihood of correctly indicating when deception is absent (specificity)? To date, fMRI-based lie-detection tests examining accuracy within individuals have generally not assessed the sensitivity or specificity of the test across individuals. One exception is a study by Kozel and colleagues²², which tested participants who had been successfully classified using the researchers’ method as ‘lying’ on a prior mock crime task (25 out of 36 participants). These pre-selected participants were then examined on a secondary mock crime task. On this secondary

task, some of the participants committed the mock crime and others did not, but all were instructed to indicate that they did not. The authors were able to correctly detect deception, using fMRI, in 100% of the participants in the ‘mock crime present’ condition. However, they also mistakenly detected deception in 67% of the participants in the ‘mock crime absent’ condition. In the language of diagnostic testing, the sensitivity of this test was high but the specificity was low.

Determining the real-world accuracy of a detection test also depends on a critical third factor, which is the probability of the event occurring within the population — the base rate. Indeed, the risks associated with a lie-detection test with low specificity will depend on the base rate of lying in the population assessed^{28,29}. Imagine that the test described in REF. 22 was given to 101 people, 100 of them truthful and 1 deceptive. On the basis of the false-positive rates of Kozel and colleagues²², the test would identify 68 participants as ‘lying’ — the 1 participant who lied about the mock crime and 67 who did not. In other words, given a positive result, the probability of the test accurately indicating someone as lying is 1 in 68, or less than 1.5%, and the likelihood of incorrectly indicating deception when it is not present is over 98%. As this example illustrates, even in an ideal circumstance in which a laboratory lie-detection test is developed and used in identical situations, and is sensitive to deception within an individual 100% of the time, its accuracy in a larger population may still be unacceptably low if the specificity of the test is low and the base rate of lying is low.

The real-world validity of fMRI-based lie detection will also depend on the generalizability of the findings obtained in laboratory studies (which typically involve undergraduate students — that is, healthy, educated young adult subjects) to the individuals whose veracity is to be assessed by these methods. Consider the differences between criminal offenders, a group that is likely to be subjected to lie-detection methods, and the undergraduate students on which these methods have so far been tested. A relatively high proportion of criminals meet the criteria for psychopathy, a condition that is associated with frequent acts of deception and with alterations in both structural MRI and fMRI studies²⁷. A study of fMRI-based lie detection in criminal offenders with a diagnosis related to psychopathy — specifically, antisocial personality disorder — found that a large proportion of these participants did not show typical prefrontal BOLD response patterns during instructed deception²⁹.

Cognitive, personality and brain factors associated with a wide range of individual differences may also affect the validity of fMRI-based lie detection. Structural MRI and fMRI changes observed with advancing age, a range of psychiatric conditions (for example, schizophrenia and post-traumatic stress disorder) or individual traits (for example, high anxiety and extraversion) limit the applicability of lie-detection tests that have not been validated in these populations. There may also be important individual differences in the neural systems involved in deception per se that we have yet to characterize. In a clever study in which participants were given the opportunity to gain money by being dishonest, those who tended towards dishonesty showed an increased BOLD signal in regions related to cognitive control when behaving dishonestly and when behaving honestly, whereas participants who tended to be honest did not show this pattern¹⁸.

Differences between lies in the laboratory and in real life. Another potential obstacle to real-world lie detection is that the lies examined in laboratory tasks are generally quite unlike those that we would try to detect outside the laboratory (BOX 2). Although researchers have been concerned with real-world effectiveness, and have presented their studies as “ecologically valid” (REF. 30) or “emulat[ing] as closely as possible a real world situation” (REF. 31), the tasks differ in many important ways from the situations in which lie detection would be used in the real world. In the laboratory studies, subjects lie because they are instructed to, about matters with little personal relevance, in highly constrained and contrived situations. In addition, the familiarity of the information being concealed and the level of emotion associated with it are typically much lower in laboratory studies than in real life.

Consider the situation in which a lie is highly rehearsed and thus familiar, and the truth is abhorrent. In this circumstance, it seems very possible that truth telling is more effortful than lying. Evidence from both fMRI and behavioural studies suggests that practice or rehearsal may alter the neural signature of deception. As indicated in FIG. 1, some of the regions commonly activated in fMRI studies of lie detection include prefrontal regions that have been proposed to underlie cognitive functions required during the more effortful ‘lie’ condition. Studies examining practice effects across a number of cognitive tasks routinely show diminished activation in the prefrontal cortex³². This reduction is thought to reflect

the diminished executive control required as highly practiced tasks become more automatic. An early fMRI study of lie detection found that memorized lies resulted in less BOLD activation when compared with unpractised lies in every deception-related region of interest identified except for one associated with memory retrieval⁹. A behavioural deception study showed that training on deception eliminated pre-training differences in reaction times between deception and truth trials, which is consistent with enhanced automaticity of lying³³. Outside the laboratory, if someone expects to be interrogated about a lie, it is very likely that they will rehearse and memorize the lie, which might eliminate many of the detectable differences in the behavioural and neural expression of deception.

In addition, real-world deception is likely to be highly emotional and personally relevant. Emotion could influence the neural circuitry of lying in two ways that might make it more difficult to distinguish truth from lies. First, truthfully answering questions about highly emotional events may be more effortful or require more (emotional) control and/or inhibition than truthfully answering questions about neutral events. To the extent that a lie-detection test measures non-specific brain signals of effort or inhibition, it may be more likely that a true statement about an emotional event is classified as a lie (as has been found with event-related potentials³⁴). However, it is also possible that the emotional qualities of the event may shift the neural signature of deception, making it less likely that a lie will be detected. In fMRI studies, emotion has been shown to alter the neural circuitry of memory³⁵, inhibition and cognitive control³⁶ and working memory interference³⁷ — all of which are processes thought to underlie the differences between brain activity during deceptive responses versus truthful responses. Indeed, emotional valence has been found to affect the neural localizations of deception-related processing³⁸. If a lie-detection test is developed based on lying about non-emotional events, its applicability to assessing deception concerning emotional or important, personally relevant events may be limited.

Countermeasures. Methods of lie detection inevitably spawn methods designed to evade detection, so the mere possibility of countermeasures should not be grounds for rejection of fMRI-based lie detection. However, the ease and success of countermeasures are relevant to the real-world usefulness of fMRI-based lie detection. In the one study

cited earlier reporting 100% accuracy, the investigators further demonstrated that if subjects adopt a simple countermeasure strategy of making imperceptible finger and toe movements, accuracy fell to 33%. At present, researchers have only begun to explore possible countermeasures for fMRI-based lie detection. Preliminary data from studies in the laboratory of A.D.W. that focused on whether subjects can conceal memory-related patterns of the BOLD signal also indicate that countermeasure strategies can reduce machine learning-based decoding of memory states from well above chance to chance levels^{39,40}.

In sum, even if the challenges facing fMRI studies of deception in the laboratory that were described in the previous section were met, a number of additional challenges await the successful translation of this method into the real world. The accuracy with which individual subjects can be assessed when telling lies or truth and the suitability of these accuracy rates given the base rates of lying and truth telling in the population demand a major new empirical research effort. How these accuracies vary as a function of an individual’s age, health, personality, life history and other variables is also a crucial question and would require an even larger programme of research to be adequately addressed. Just as differences among individuals would be expected to influence the validity of the method, so too would differences in the nature of the lie and its context: whether the subject has lived with the lie for a long time, whether the truth is more emotionally charged than the lie and what is at stake if the lie is discovered. Finally, the susceptibility of fMRI-based lie detection to countermeasures would need to be more fully explored before it is applied in the real world, and this too requires extensive research.

Uses of fMRI-based lie detection

Lie detection using fMRI has moved rapidly from basic research in the laboratory to commercial application in the real world. In 2006, two companies began offering fMRI-based lie detection services: No Lie MRI, based on the method developed by Langleben and colleagues at the University of Pennsylvania (Philadelphia, USA), and Cephus, based on the method of Kozel and colleagues at the University of Texas at Dallas (USA). These companies have suggested a number of uses for fMRI-based lie detection, spanning business, family life, criminal justice and national security contexts. For example, No Lie MRI recommends its services for diverse problems,

ranging from combating insurance fraud (<http://www.noliemri.com/customers/GroupOrCorporate.htm>) and increasing public trust of US and foreign leaders (<http://www.noliemri.com/customers/Government.htm>) to “risk reduction in dating” (<http://www.noliemri.com/customers/Individuals.htm>).

The potential application of fMRI-based lie detection that has received the greatest public scrutiny has been for assessing the truthfulness of legal testimony. In at least three cases, US courts have been asked to admit evidence from fMRI-based lie detection. The courts determine admissibility of scientific evidence by applying one of two standards, depending on jurisdiction. Both standards are designed to keep ‘junk science’ from influencing jury decisions, although they do so in different ways. According to the Frye standard, set forth in *Frye v. United States* in 1923 (REF. 41), admissibility hinges on the general acceptance of the method within its particular scientific field. According to the Daubert standard, set forth in 1993 by the US Supreme Court in *Daubert v. Merrell Dow*⁴², judges in federal cases must be more active gate-keepers of scientific evidence in court rather than simply deferring to general scientific opinion. In doing so, they should take five factors (among others that they may consider relevant) into account when making their decisions: whether the method is testable and has been tested; whether it has been reported in peer-reviewed publications; whether there is a known or potentially knowable error rate; whether there are standards for the way in which the method is used; and finally, as in Frye, whether the method is generally accepted within the relevant scientific community.

The first two attempts to introduce fMRI-based lie detection occurred in 2010. In the first case, *Wilson v. Corestaff Services*⁴³, the plaintiff in an employment discrimination case sought to introduce evidence gathered by Cephos to support the credibility of a witness’s testimony. The judge in this case ruled the evidence inadmissible on the grounds that credibility assessment is the job of the jury but also noted that the method would not meet the Frye criteria, stating: “even a cursory review of the scientific literature demonstrates that the plaintiff is unable to establish that the use of the fMRI test to determine truthfulness or deceit is accepted as reliable in the relevant scientific community” (REF. 43).

In the second case, *United States v. Semrau*⁴⁴, the defendant standing trial for Medicare fraud claimed that he did not

intentionally violate the law and sought to present the results of fMRI-based lie detection as evidence that his testimony was truthful. A hearing was held in Memphis Federal Court, before a Magistrate Judge to whom had been delegated the task of making a recommendation on admissibility — that is, to determine whether fMRI-based lie detection meets the Daubert criteria. The 39-page opinion included the recommendation (later accepted by the presiding district judge) that the fMRI evidence be excluded. The analysis noted: the lack of general acceptance for the method within the scientific community; the substantial differences between laboratory research designs and the real-world use of fMRI-based lie detection in this case; the lack of ‘real-life’ error rates; and the lack of suitably controlling standards for the use of the method. (The latter determination was prompted by Cephos’s discounting of one of three scanning sessions, which had indicated deception, on the grounds that the defendant was fatigued.) However, the opinion also observes that precise validation in real-world contexts, of the sort that scientists might require in their own research, is not always legally necessary: “in the future, should fMRI-based lie detection undergo further testing, development, and peer review, improve upon standards controlling the technique’s operation, and gain acceptance by the scientific community for use in the real world, this methodology may be found to be admissible even if the error rate is not able to be quantified in a real world setting.” (REF. 44).

In the case of *Smith v. State of Maryland*⁴⁵, the defendant was being retried for second-degree murder in 2012 and sought to introduce evidence of the truthfulness of his own testimony from No Lie MRI. The judge refused to admit the evidence on the basis of the Frye standard, after concluding that experts in the field (including A.D.W. and E.A.P) did not agree with the company’s experts, suggesting lack of ‘general acceptance’.

Legal, social and ethical issues

Science can, in principle, tell us the accuracy of fMRI-based lie detection for any particular population of individuals under any particular circumstances that we might specify. It cannot, however, tell us how accurate fMRI-based lie detection should be for any particular use to which it might be put. That decision depends on the needs and values of the people using the method. We have seen that the US legal system has

criteria, which were developed in legal cases that involved scientific evidence other than fMRI-based lie detection, for deciding when a method has sufficient validation: namely, the Frye and Daubert criteria. These criteria exemplify the interdependence and the independence of scientific and societal decision making. Although scientific research provides essential input into decisions regarding admissibility, the decisions themselves are not made by the kinds of conventional scientific criteria applied, for example, by reviewers of research articles. It is fair to say that, in some respects, legal standards may be lower than scientific standards where scientific evidence such as fMRI-based lie detection is concerned. As argued by Schauer⁴⁶, the societal needs served by the law often require a more pragmatic approach to the vetting of scientific evidence. In the absence of better methods of discovering the truth, an imperfect method may be better than nothing: “... the exclusion of substandard science, when measured by scientific standards, may have the perverse effect of lowering the accuracy and rigor of legal fact-finding, because the exclusion of flawed science will only increase the importance of the even more flawed non-science that now dominates legal fact-finding.” (REF. 46). As Langleben and Moriarity have noted⁴⁷, this reasoning can be extended to other scientific methods as well, given that many of the methods of forensic science have been found wanting⁴⁸.

As scientists, we are not accustomed to endorsing methods on the grounds that they are ‘lesser evils’. The difference between the scientific and legal approach to accepting questionable sources of evidence rests in part on the scientist’s choice of hypotheses to test. If no good method is available to test a particular hypothesis, then a scientist will normally simply decline to test it. The legal system cannot make the analogous decision; the question of ‘guilt beyond a reasonable doubt’ must be addressed with the evidence available, even when that evidence is acknowledged to have serious weaknesses.

What about societal decision making regarding potential uses of fMRI-based lie detection outside the courtroom? The world has only begun to engage with the question of how best to use, and limit the use of, fMRI-based lie detection. Questions concerning the necessary degrees of accuracy and validity will undoubtedly require different answers for different tasks in different contexts, according to the potential benefits of correct

lie detection, the costs of wrong calls and the intersection of this technology with moral principles such as the right to privacy.

How accurate is accurate enough? The most immediate ethical and social issues raised by fMRI-based lie detection arise because of the method's lack of demonstrated accuracy and validity. Given the scientific and technical problems reviewed earlier, the most likely harms would result from false determinations — lies wrongly identified as truths and truthful statements wrongly identified as lies. Some commentators have suggested a ban or moratorium on fMRI-based lie detection pending better evidence concerning its accuracy^{49,50}. The history of polygraphy offers tragic reminders of the cost, to national security and human life, of over-reliance on an apparently high-tech but inaccurate method for detecting deception.

In two well-known examples of the polygraph's false-negative results, the American CIA (Central Intelligence Agency) agent Aldrich Ames and the Jordanian CIA informant Humam Khalil Abu-Mulal al-Balawi both passed polygraph testing (twice in the case of Ames), even though Ames spent years selling American secrets to the Soviets and Russians, and al-Balawi killed seven CIA agents as a double agent. False-positive polygraph results have cost honest people job opportunities and even their liberty, such as when they led to false confessions⁵¹.

Marks⁵² has pointed out that the impressive visual appearance of fMRI-based lie detection results may eclipse concerns about the method's technical weaknesses, including the likelihood of false positives. He suggests that government agents interrogating detainees would naturally tend to increase the aggressiveness of their tactics if a detainee's fMRI scan indicates deception. Thus, although lie-detection methods might be thought to reduce the use of harsh interrogation, they might in fact be used to justify abusive treatment of detainees, a particularly deplorable outcome in the case of false-positive results⁵².

As noted earlier, false positives in fMRI-based lie detection will have the greatest negative impact when the method is used to identify relatively rare cases of dishonesty in a population. When applied to large numbers of mostly honest people, even modest false-positive rates will result in many wrongly accused individuals. Applications such as the routine screening of job applicants or travellers are therefore problematic, as they would probably result in large numbers of falsely accused individuals in

relation to the occasional correct identification. As with other questions concerning requirements for accuracy, the question of how many individuals we are willing to falsely accuse of deception for the sake of an occasional correctly identified liar will be determined not by research but by society's needs and values.

Ethical issues beyond validity and accuracy.

Not all ethical issues surrounding fMRI-based lie detection depend on the method's accuracy; even a technically perfect lie detector would raise ethical issues and be subject to societal deliberation and regulation. Indeed, some issues would become more pressing in the event of successful fMRI-based lie detection.

Lie detection using fMRI raises privacy issues that require societal control, much as we place limits on other practices that intrude on privacy, from DNA collection to wire tapping. For example, if everyone's phone conversations and e-mail messages were generally available to their families, employers and the state, and if the DNA of all citizens were on file with law enforcement authorities, much crime and misbehaviour would be discovered or, better still, averted. But societies place limits on the collection and use of such information in order to protect personal privacy. An additional reason to limit access to such information is to increase the benefit it provides to those in possession of it. Societal management of fMRI-based lie detection would presumably be aimed at balancing the cost to individual privacy against the collective benefits of reduced crime and terrorism, enhanced personnel selection and the generally increased honesty between people that might result from the knowledge that the veracity of one's statements could be tested.

In the United States, the legal framework for the protection of privacy includes the Fourth and Fifth Amendments to the US Constitution, which several legal scholars have discussed in relation to fMRI-based lie detection and other forms of brain imaging that provide psychological information^{53–55}. The Fourth Amendment protects against warrantless search, including physical tests such as fMRI. The Fifth Amendment protects against compelled self-incriminating testimony. It remains to be decided whether fMRI-based lie detection should be viewed as physical evidence or testimony⁵⁶.

Considerations of individual autonomy and freedom arise in connection with the process of consenting to undergo fMRI-based lie detection. Consent procedures therefore

constitute another aspect of societal management that arises regardless of the method's known accuracy. At present, when the method's accuracy in the real world is unknown and even laboratory accuracy estimates are unavailable for subjects of different ages and states of health, subjects must understand the questionable accuracy of the method in order to make an informed decision on their own behalf. They must also understand that not all outside parties will interpret the results of the test with appropriate caution. In the future, when the method may become more accurate, subjects should still be informed that the method is not perfect. Consent is especially fraught when testing is not requested by the subject but by another party, such as the government, an employer or a jealous spouse. Additional safeguards may be needed to prevent coercion, including the indirect coercion that results when refusal to take the test is seen as indicative of guilt.

In sum, the question of whether and how to use fMRI-based lie detection cannot be answered solely on the basis of the method's performance. No method will ever be known to have 100% accuracy for any context in which it might be deployed. Deciding what level of uncertainty is acceptable depends on how different kinds of outcomes are valued. Correct and incorrect identifications of lies and of truth may be weighted very differently under different circumstances and in different societies. For example, the priority placed on outcomes related to security relative to the rights of individuals will determine whether it is worse to miss a liar or to falsely accuse an honest person. The strength of a society's commitments to principles including individual privacy, autonomy and freedom will also shape its policies concerning fMRI-based lie detection.

Policy recommendations

How should the development and use of fMRI-based lie detection be managed in light of the scientific, legal, ethical and societal issues just reviewed? We offer three general recommendations.

First, different policies should be considered for different applications of fMRI-based lie detection. We do not join calls to ban fMRI-based lie detection across the board. Despite the enormous shortcomings of the current evidence, reviewed above, we suggest that restrictions should be proportional to the outcomes and principles at stake. Risk reduction in dating calls for different standards of certainty and different protections of individual rights than the interrogation of terrorist suspects.

Second, publicly funded research should be undertaken to explore the potential of fMRI-based lie detection while keeping in mind potential conflicts of interest for researchers associated with companies that offer the service. The two highest research priorities are: first, the removal of, or accounting for, experimental confounds noted earlier; and, second, the validation of the methods under more realistic conditions, including with countermeasures and with more diverse subjects. If fMRI-based lie detection passes these hurdles, then a substantial investment in real-world validation⁴⁷ would be justified.

Third, although we acknowledge that the standards of science with regard to truth and certainty may not always be the appropriate ones in legal and other societal contexts, scientists have a vital part to play in the application of neuroscience to the law⁵⁷. In the case of fMRI-based lie detection, it is our duty to raise questions about its accuracy, validity and specificity, to provide accurate answers to these questions and to communicate their relevance to non-scientist citizens.

Martha J. Farah is at the Center for Cognitive Neuroscience, Center for Neuroscience & Society, Department of Psychology, University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA.

J. Benjamin Hutchinson is at the Department of Psychology, Green Hall, Princeton University, Princeton, New Jersey 08540, USA.

Elizabeth A. Phelps is at the Center for Neural Science, New York University, New York, New York 10003, USA, and Nathan Kline Institute, Orangeburg, New York 10962, USA.

Anthony D. Wagner is at the Department of Psychology and Neurosciences Program, Stanford University, Stanford, California 94305, USA.

Correspondence to M.J.F.

e-mail: mfarah@neuroethics.upenn.edu

doi:10.1038/nrm3665

- National Research Council. *The polygraph and lie detection* (The National Academies, 2003).
- Ganis, G. & Keenan, J. P. Neural correlates of deception. *Soc. Neurosci.* 465–472 (2009).
- Heckman, K. E. & Happel, M. D. in *Educating information. Interrogation: science and art* 63–94 (NDIC, 2006).
- Spence, S. A. *et al.* Behavioural and functional anatomical correlates of deception in humans. *Neuroreport* 12, 2849–2853 (2001).
- Langleben, D. D. *et al.* Brain activity during simulated deception: an event-related functional magnetic resonance study. *Neuroimage* 15, 727–732 (2002).
- Lee, T. M. C. *et al.* Lie detection by functional magnetic resonance imaging. *Hum. Brain Mapp.* 15, 157–164 (2002).
- Hakun, J. G. *et al.* fMRI investigation of the cognitive structure of the Concealed Information Test. *Neurocase* 14, 59–67 (2008).
- Kozel, F. A. *et al.* Detecting deception using functional magnetic resonance imaging. *Biol. Psychiatry* 58, 605–613 (2005).
- Ganis, G., Kosslyn, S. M., Stose, S., Thompson, W. L. & Yurgelun-Todd, D. A. Neural correlates of different types of deception: an fMRI investigation. *Cereb. Cortex* 13, 830–836 (2003).
- Bles, M. & Haynes, J.-D. Detecting concealed information using brain-imaging technology. *Neurocase* 14, 82–92 (2008).
- Christ, S. E., van Essen, D. C., Watson, J. M., Brubaker, L. E. & McDermott, K. B. The contributions of prefrontal cortex and executive control to deception: evidence from activation likelihood estimate meta-analyses. *Cereb. Cortex* 19, 1557–1566 (2009).
- Wagner, A. in *A Judge's Guide to Neuroscience: A Concise Introduction* (eds Gazzaniga, M. S. & Rakoff, J. S.) 13–25 (University of California, 2010).
- Eickhoff, S. B. *et al.* Coordinate-based activation likelihood estimation meta-analysis of neuroimaging data: a random-effects approach based on empirical estimates of spatial uncertainty. *Hum. Brain Mapp.* 30, 2907–2926 (2009).
- Gamer, M., Klimecki, O., Bauermann, T., Stoeter, P. & Vossel, G. fMRI-activation patterns in the detection of concealed information rely on memory-related effects. *Soc. Cogn. Affect. Neurosci.* 7, 506–515 (2012).
- Kanwisher, N. in *Using Imaging to Identify Deceit: Scientific and Ethical Questions* (eds Bizzi, E. *et al.*) 7–15 (American Academy of Arts and Sciences, 2009).
- Langleben, D. D. *et al.* Telling truth from lie in individual subjects with fast event-related fMRI. *Hum. Brain Mapp.* 26, 262–272 (2005).
- Davatzikos, C. *et al.* Classifying spatial patterns of brain activity with machine learning methods: application to lie detection. *Neuroimage* 28, 663–668 (2005).
- Greene, J. D. & Paxton, J. M. Patterns of neural activity associated with honest and dishonest moral decisions. *Proc. Natl Acad. Sci. USA* 106, 12506–12511 (2009).
- Kozel, F. A. *et al.* A pilot study of functional magnetic resonance imaging brain correlates of deception in healthy young men. *J. Neuropsychiatry Clin. Neurosci.* 16, 295–305 (2004).
- Kozel, F. A., Padgett, T. M. & George, M. S. A replication study of the neural correlates of deception. *Behav. Neurosci.* 118, 852–856 (2004).
- Monteleone, G. T. *et al.* Detection of deception using fMRI: better than chance, but well below perfection. *Soc. Neurosci.* 4, 528–538 (2009).
- Kozel, F. A. *et al.* Functional MRI detection of deception after committing a mock sabotage crime. *J. Forensic Sci.* 54, 220–231 (2009).
- Jin, B. *et al.* Feature selection for fMRI-based deception detection. *BMC Bioinformatics* 10 (Suppl. 9), S15 (2009).
- Nose, I., Murai, J. & Taira, M. Disclosing concealed information on the basis of cortical activations. *Neuroimage* 44, 1380–1386 (2009).
- Lee, T. M. C., Leung, M.-K., Lee, T. M. Y., Raine, A. & Chan, C. C. H. I want to lie about not knowing you, but my precuneus refuses to cooperate. *Sci. Rep.* 3, 1636 (2013).
- Ganis, G., Morris, R. R. & Kosslyn, S. M. Neural processes underlying self- and other-related lies: an individual difference approach using fMRI. *Soc. Neurosci.* 4, 539–553 (2009).
- Anderson, N. E. & Kiehl, K. A. The psychopath magnetized: insights from brain imaging. *Trends Cogn. Sci.* 16, 52–60 (2012).
- Wolpe, P. R., Foster, K. R. & Langleben, D. D. Emerging neurotechnologies for lie-detection: promises and perils. *Am. J. Bioeth.* 5, 39–49 (2005).
- Jiang, W. *et al.* A functional MRI study of deception among offenders with antisocial personality disorders. *Neuroscience* 244, 90–98 (2013).
- Mohamed, F. B. *et al.* Brain mapping of deception and truth telling about an ecologically valid situation: functional MR imaging and polygraph investigation—initial experience. *Radiology* 238, 679–688 (2006).
- Abe, N. The neurobiology of deception: evidence from neuroimaging and loss-of-function studies. *Curr. Opin. Neurol.* 22, 594–600 (2009).
- Raichle, M. E. *et al.* Practice-related changes in human brain functional anatomy during nonmotor learning. *Cereb. Cortex* 4, 8–26 (1994).
- Hu, X., Chen, H. & Fu, G. A repeated lie becomes a truth? The effect of intentional control and training on deception. *Front. Psychol.* 3, 488 (2012).
- Proverbio, A. M., Vanutelli, M. E. & Adorni, R. Can you catch a liar? How negative emotions affect brain responses when lying or telling the truth. *PLoS ONE* 8, e59383 (2013).
- Phelps, E. A. Emotion and cognition: insights from studies of the human amygdala. *Annu. Rev. Psychol.* 57, 27–53 (2006).
- Bush, G., Luu, P. & Posner, M. I. Cognitive and emotional influences in anterior cingulate cortex. *Trends Cogn. Sci.* 4, 215–222 (2000).
- Levens, S. M. & Phelps, E. A. Insula and orbital frontal cortex activity underlying emotion interference resolution in working memory. *J. Cogn. Neurosci.* 22, 2790–2803 (2010).
- Lee, T. M. C., Lee, T. M. Y., Raine, A. & Chan, C. C. H. Lying about the valence of affective pictures: an fMRI study. *PLoS ONE* 5, e12291 (2010).
- Uncapher, M. R., Chow, T., Rissman, J., Eberhart, J. & Wagner, A. D. Strategic influences on memory expression: effects of countermeasures and memory strength on the neural decoding of past experience. *Soc. Neurosci. Abstr.* 905.13 (2012).
- Rissman, J., Greeley, H. T. & Wagner, A. D. Detecting individual memories through the neural decoding of memory states and past experience. *Proc. Natl Acad. Sci. USA* 107, 9849–9854 (2010).
- Frye v. United States, 293 F1013 (D. C. Cir. 1923).
- Daubert v. Merrell Dow Pharmaceuticals, 509 U. S. 579 (1993).
- Wilson v. Corestaff Services L. P., 900 N. Y. S.2d 639 (N. Y. Sup. Ct. 2010).
- United States of America v. Lorne Allan Semrau 07–10074 (2010).
- Gary James Smith v. State of Maryland 418 Md 587, 16 A.3d 977 (2011).
- Schauer, F. Can bad science be good evidence — neuroscience, lie detection and beyond. *Cornell Law Rev.* 95, 1191–1220 (2009).
- Langleben, D. D. & Moriarty, J. C. Using brain imaging for lie detection: where science, law, and policy collide. *Psychol. Public Policy Law* 19, 222–234 (2005).
- National Research Council. *Strengthening Forensic Science in the United States: A Path Forward* (The National Academies, 2009).
- Greeley, H. T. & Illes, J. Neuroscience-based lie detection: the urgent need for regulation. *Am. J. Law Med.* 33, 377–431 (2007).
- Tovino, S. A. Imaging body structure and mapping brain function: a historical approach. *Am. J. Law Med.* 33, 193–228 (2007).
- Costanzo, M. & Krauss, D. *Forensic and Legal Psychology* (Macmillan, 2010).
- Marks, J. H. Interrogational neuroimaging in counterterrorism: a 'no-brainer' or a human rights hazard? *Am. J. Law Med.* 33, 483–500 (2007).
- Farahany, N. A. Incriminating thoughts. *Stanford Law Rev.* 64, 351 (2012).
- Pustilnik, A. C. in *The Constitution and the Future of the Criminal Law in America* Ch. 7 (eds Parry, J. T. & Richardson, L. S.) (Cambridge Univ. Press, 2013).
- Shen, F. X. Neuroscience, mental privacy, and the law. *Harvard J. Law Public Policy* 36, 653–713 (2013).
- Stoller, S. E. & Wolpe, P. R. Emerging neurotechnologies for lie detection and the Fifth Amendment. *Am. J. Law Med.* 33, 359–375 (2007).
- Jones, O. D., Wagner, A. D., Faigman, D. L. & Raichle, M. E. Neuroscientists in court. *Nature Rev. Neurosci.* 14, 730–736 (2013).
- Bond, C. F. & DePaulo, B. M. Accuracy of deception judgments. *Pers. Soc. Psychol. Rev.* 10, 214–234 (2006).
- Ekman, P. & Friesen, W. V. Nonverbal leakage and clues to deception. *Psychiatry* 32, 88–106 (1969).
- Scherer, K. R., Feldstein, S., Bond, R. N. & Rosenthal, R. Vocal cues to deception: a comparative channel approach. *J. Psycholinguist. Res.* 14, 409–425 (1985).
- DePaulo, B. M. *et al.* Cues to deception. *Psychol. Bull.* 129, 74–118 (2003).
- Vrij, A. *Detecting Lies and Deceit: Pitfalls and Opportunities* 2nd edn (Wiley, 2008).
- Vrij, A. Nonverbal dominance versus verbal accuracy in lie detection a plea to change police practice. *Crim. Justice Behav.* 35, 1323–1336 (2008).
- Schafer, E. D. in *Forensic Science* (Embar-Seddon, A. & Pass, A. D.) 40 (Salem, 2008).
- Bunn, G. C. The lie detector, wonder woman and liberty: the life and works of William Moulton Marston. *Hist. Hum. Sci.* 10, 91–119 (1997).
- Ben-Shakhar, G., Bar-Hillel, M. & Lieblich, I. Trial by polygraph: scientific and juridical issues in lie detection. *Behav. Sci. Law* 4, 459–479 (1986).
- Alder, K. *The Lie Detectors: The History of an American Obsession* (Simon and Schuster, 2007).
- Farwell, L. A. & Smith, S. S. Using brain MERMER testing to detect knowledge despite efforts to conceal. *J. Forensic Sci.* 46, 135–143 (2001).

69. Aggarwal, N. K. Neuroimaging, culture, and forensic psychiatry. *J. Am. Acad. Psychiat Law* **37**, 239–244 (2009).
70. Tian, F., Sharma, V., Kozel, F. A. & Liu, H. Functional near-infrared spectroscopy to investigate hemodynamic responses to deception in the prefrontal cortex. *Brain Res.* **1303**, 120–130 (2009).
71. Bok, S. *Lying: Moral Choice in Public and Private Life* (Random House Digital, 2011).
72. Barnes, J. A. *A Pack of Lies: Towards a Sociology of Lying* (Cambridge Univ. Press, 1994).
73. DePaulo, B. M., Kashy, D. A., Kirkendol, S. E., Wyer, M. M. & Epstein, J. A. Lying in everyday life. *J. Pers. Soc. Psychol.* **70**, 979–995 (1996).
74. Lykken, D. T. Why (some) Americans believe in the lie detector while others believe in the guilty knowledge test. *Integr. Physiol. Behav. Sci.* **26**, 214–222 (1991).
75. Ben-Shakhar, G., Bar-Hillel, M. & Kremnitzer, M. Trial by polygraph: reconsidering the use of the guilty knowledge technique in court. *Law Hum. Behav.* **26**, 527–541 (2002).
76. Van Essen, D. C. A population-average, landmark- and surface-based (PALS) atlas of the human cerebral cortex. *Neuroimage* **28**, 635–662 (2005).

Acknowledgements

The authors thank O. Jones for guidance on the legal issues discussed herein, and T. Chow for assistance with the meta-analysis. They gratefully acknowledge the support of the Law and Neuroscience Project, which is funded by the John D. and Catherine T. MacArthur Foundation. The writing of this article was partially supported by the US National Institutes of Health grant R01-HD055689. This article reflects the views of the authors and does not necessarily represent the

official views of either the John D. and Catherine T. MacArthur Foundation or the MacArthur Foundation Research Network on Law and Neuroscience.

Competing interests statement

The authors declare no competing interests.

FURTHER INFORMATION

Nature Reviews Neuroscience Series on Neuroscience and the law: <http://www.nature.com/nrn/series/neurosciencelaw/index.html>

Neurolaw.org: <http://neurolaw.org/>

SUPPLEMENTARY INFORMATION

See online article: S1 (box) | S2 (table) | S3 (table)

ALL LINKS ARE ACTIVE IN THE ONLINE PDF