

Brain Imaging and Brain Privacy: A Realistic Concern?

Martha J. Farah, M. Elizabeth Smith, Cyrena Gawuga, Dennis Lindsay,
and Dean Foster

Abstract

■ Functional neuroimaging has been used to study a wide array of psychological traits, including aspects of personality and intelligence. Progress in identifying the neural correlates of individual differences in such traits, for the sake of basic science, has moved us closer to the applied science goal of measuring them and thereby raised ethical concerns about privacy. How realistic are such concerns given the current state of the art? In

this article, we describe the statistical basis of the measurement of psychological traits using functional neuroimaging and examine the degree to which current functional neuroimaging protocols could be used for this purpose. By analyzing the published data from 16 studies, we demonstrate that the use of imaging to gather information about an individual's psychological traits is already possible, but to an extremely limited extent. ■

INTRODUCTION

Progress in neuroscience raises a number of ethical concerns. Among the most prominent of these concerns is the potential use of functional neuroimaging to obtain personal information about individuals. Recent studies have reported neuroimaging correlates of personality, attitudes, and intelligence (for reviews, see Gray & Thompson, 2004; Duncan, 2003; Illes, Kirschen, & Gabrieli, 2003; Canli & Amin, 2002). By analogy with the term "genotyping," the use of neuroimaging to determine features of brain function relevant to an individual's traits could be called "brainotyping." Many of the same ethical issues that have arisen with genotyping may also arise with brainotyping. For example, both genotyping and brainotyping, in principle, raise serious concerns about privacy (Illes & Racine, 2005).

Like the tissue samples used for genotyping, brain images can be obtained with consent for one purpose but later analyzed for other purposes. Correlations between brain function and psychological traits are often obtained in the context of tasks that lack an obvious relation to the trait being measured. For example, extraversion and unconscious racial attitudes are both correlated with brain activity evoked by simply viewing pictures of faces (Canli, Sivers, Whitfield, Gotlieb, & Gabrieli, 2002; Phelps et al., 2000). Hence, in these studies, subjects could be told that the purpose of the scan was related to face perception. This feature of neuroimaging-based studies of psychological traits makes it relatively easy to obtain information about a person without their knowledge.

Although a number of writers have commented on the potential threat to brain privacy posed by functional neuroimaging (New York City Bar Association, 2005; Farah & Wolpe, 2004; Hyman, 2004; Kennedy, 2004; Illes, 2003; Canli & Amin, 2002), little attention has been paid to the question of whether this threat is realistic, given the current state of the field. Exceptions include Canli (2006), who argues that brain imaging data may provide better measures of personality than conventional paper-and-pencil tests, and in contrast, Levy (2007, Chap. 4) and Gazzaniga (2006, Chap. 7), who argue that current technologies for brain imaging are unlikely to enable any kind of socially relevant "mind reading."

It is clear that measurement of psychological traits using neuroimaging is, in principle, possible, and that such a use raises ethical concerns. However, it is not clear whether such measurement is a likely development in the near term, a more remote but still realistic goal for the future, or whether it remains effectively a science fiction scenario.

The present article addresses this issue by answering the question: What can the current state of the art in functional neuroimaging of psychological traits tell us about an individual? An explicit answer to this question is not available in the literature because published functional neuroimaging research has not been directed toward the measurement of normal psychological traits for the purpose of characterizing individuals. Rather, the goal of individual differences research with functional neuroimaging has been to understand the brain bases of variation in psychological traits across the population. In other words, current functional neuroimaging research on individual differences is not carried out for the applied science goal of measuring levels of traits in

individuals, but rather the basic science goal of identifying those features of brain function that underlie individual difference in psychological function.

On the basis of this alone, one might conclude that worries about brain privacy are premature. Just as the possibility of genotyping individuals to obtain useful information about psychological traits is, at present, remote, we might conclude that prospects for brainotyping individuals are also unrealistic. Genetic research on individual differences in psychological traits has yet to identify any specific genes with more than extremely weak correlations with such traits (Parens, 2004; Van Gestel & Van Broeckhoven, 2003), and this research has had a few decades' head start on functional neuroimaging research on individual differences.

However, one should not dismiss the prospect of brainotyping on these grounds. Two considerations argue for the possibility of current or near-term measurement of psychological traits using functional neuroimaging. First, the analogy between functional neuroimaging and genotyping is imperfect. Behavioral geneticists have concluded that most complex psychological traits involve multiple genes and are determined by a combination of genes and experience. This would account for the weak relations between individual genes and psychological traits, and predicts only moderate success, at best, in predicting such traits from a complete genetic analysis. In contrast, brain function is one causal step closer to behavior, in that it reflects the joint influence of genome and experience. Although there is no guarantee that the most relevant aspects of brain function are captured by current functional neuroimaging methods, such methods at least target a point on the final common pathway for the full set of genetic and environmental influences on psychology, and are therefore more likely to provide useful correlates.

Second, an empirical basis for brainotyping already exists. Although imaging studies, to date, have indeed focused on relations between psychological traits and brain activity across groups of subjects, the data collected in these studies are, in principle, as applicable to the goal of individual measurement as to the goal of generalizing about the neural bases of these traits. For either goal, one needs the same information, namely, information about the distributions of brain activation measures and psychological measures and the correlation between the two. Here we describe a statistical basis of the measurement of psychological traits using functional neuroimaging, and use it to determine the degree to which current functional neuroimaging protocols could be used for this purpose.

A Statistical Basis for Brainotyping

For a particular neuroimaging study of individual differences in a psychological trait, given the mean and standard deviation of brain activity in the relevant region of

interest (ROI) (\bar{b} and σ_b) and of the trait (\bar{p} and σ_p), and the correlation between them, rho (ρ), we can use the level of brain activity of a hypothetical new subject (b) to predict his or her likely level of the psychological trait (\hat{p}):

$$\frac{\hat{p} - \bar{p}}{\sigma_p} = \rho \frac{(b - \bar{b})}{\sigma_b}$$

$$\hat{p} = \bar{p} + \rho \frac{\sigma_p}{\sigma_b} (b - \bar{b})$$

Of course, for $\rho < 1$, there will be an interval around the predicted value within which the true level of the psychological is likely to fall. This is illustrated in Figure 1. The distribution shown on the y -axis represents the postscan prediction, that is, the range of trait values and associated probabilities for a subject, given a brain activation value. Its spread depends on the length of the prescan prediction interval and the strength of the correlation.

The predictive value of the brain activation can be expressed in terms of the prediction interval, omega (Ω), which is the range of values within which the true value will fall some set percentage of the time. A conventional choice of percentage would be 95%, yielding an interval that extends from roughly 2 standard deviations above the predicted value to 2 standard deviations below it.

Before a scan, the prediction interval for a psychological trait is simply the range of scores that 95% of the relevant normal population would be expected to have; its length is four times the standard deviation of that trait in this population. If a scan is informative about a psychological trait, then after the scan, we will be able to define a shorter interval within which the subject has a




Figure 1.

95% probability of falling, that is, an interval that is some number smaller than 4 standard deviations of the trait in the population.

The length of the postscan prediction interval depends on the prescan interval and the strength of the correlation between the brain activation and psychological trait as follows (Sheskin, 1997, Equation 22.12):

$$\Omega_{\text{postscan}} = \Omega_{\text{prescan}} \sqrt{1-\rho^2}$$

Thus, the reduction in the prediction interval from prescan to postscan is a straightforward measure of a scan's informativeness, which depends on the strength of correlation between the relevant measure of brain activity and the trait.

The foregoing assumes that the correlation between brain activity and trait is known with certainty, which would be a reasonable approximation for studies with large samples of subjects. However, in the current literature, the correlations between brain activation and psychological trait are estimated with small groups of subjects and, therefore, have errors of their own, as illustrated in Figure 2. Although the analyses just presented take into account error in the individual measures (σ_b and σ_p), they do not take into account error in the correlation coefficient (ρ).

In order to take into account the error of estimating the correlation between brain activity and psychological trait, as well as the error in the two measures being correlated, we need to begin with the individual data points, from which the published correlations above were computed, and build a regression model of the relation between the measured activations and traits. The model takes the form of an equation converting a subject's brain activation value into a psychological trait level, along with a confidence interval around that prediction, in other words, the prediction interval, Ω . The proportional reduction in prediction interval resulting from the scan is a




Figure 2.

measure of the scan's informativeness concerning an individual.

METHODS

Search Method

Medline was searched using the following combined keywords: [individual differences OR intelligence OR personality] and fMRI, for articles published in 2006 or before. Additional articles were identified by colleague recommendation. This search procedure yielded a total of 1875 articles. On the basis of the articles' abstracts, a set of 64 articles was identified as being potentially relevant, in that they reported a relationship between individual differences in brain activation and in a psychological trait for a sample of normal human subjects. Four additional criteria were then applied.

First, we required that the psychological trait under study has some social or personal relevance, such that privacy could conceivably be a concern, according to our intuitions. An example of a study excluded by this criterion is Urry et al.'s (2006) study of individual differences in the diurnal slope of salivary cortisol secretion.

Second, we required that the subject sample be unselected for the trait under study (e.g., we excluded extreme group designs). An example of a study excluded by this criterion is Bertolino et al.'s (2005) study of fear responsiveness in which subjects were selected for being prone to phobia.

The third criterion concerned the number of independent tests carried out in each study. We adopted this criterion to limit the number of spurious results in the review, and thus, avoid overestimating current brainotyping capabilities. The most common reason for exclusion was insufficient constraint on ROIs. Only results from a priori ROIs, obtained in studies with just one or two such ROIs, were included. ROIs could be defined anatomically or functionally. The ROI criterion excluded studies that used whole-brain analyses and also excluded studies with several a priori ROIs. Some of the latter found predominantly null results (e.g., Matthews, Ward, & Lawrence, 2004, examined 22 a priori ROIs in a study of attentional control and fearfulness and found a reliable relationship between brain activation and behavioral inhibition in only one ROI), whereas others found predominantly positive results (e.g., Beaver et al., 2006, examined nine ROIs and found reliable relationships between brain activation and reward drive in five of them). In addition to limiting the number of ROIs, we also excluded studies in which multiple traits were examined for correlation with brain activation. An example of an article eliminated by this criterion is Stark et al.'s (2005) study of disgust sensitivity, in which five different subscales of their Questionnaire for Disgust Sensitivity were analyzed separately. A final restriction aimed at limiting the possibility of spurious findings was to confine

the search to studies relating psychological traits to relatively simple and commonly used imaging measures, including percent signal change, number of suprathreshold voxels, and adaptation, rather than to more complex functions derived from these measures such as asymmetry ratios or rate of activation change across conditions. An example of an article eliminated by this criterion is Xue, Chen, Jin, and Dong's (2006) study relating language learning with asymmetry of fusiform activation.

Fourth, we selected those studies for which the published articles included scatterplots showing the relationship between trait scores and brain activations. These scatterplots provided the data for the analyses to be presented here. Specifically, they provided the pairing of a trait value and a brain activation value for each individual subject in each ROI. We also required that the brain activations shown in the scatterplots be from the entirety of each a priori ROI, rather than from voxels within an ROI chosen to illustrate the relationship at its strongest, in order not to bias our results. An example of an article eliminated by this criterion is Singer et al.'s (2004) study of empathy.

Sixteen studies met all four criteria, and these are shown in Table 1. The psychological traits measured in these studies ranged from personality traits and other aspects of affective functioning to intelligence and specific cognitive aptitudes. As shown in Table 1, these were studied in the contexts of tasks that lacked obvious relationships to the traits.

Data Transcoding and Analysis

Published scatterplots were enlarged up to 300% in order to obtain a range of at least 5 cm for the range of psychological trait values. The enlargements were printed onto transparent sheets, which were superimposed on graph paper to facilitate reconstruction of the coordinates of each point. Note that measurement error will have the effect of underestimating the predictive power of brain imaging.

The goal of data analysis was to determine the proportion of reduction in the prediction interval for a hypothetical new subject as a result of scanning that subject. For each study, the brain activations were the independent variables and psychological trait level was the dependent variable in univariate regression analyses using JMP IN 5.1. Regression analysis was used to model the relation between measures of brain activity in a specific ROI and levels of the trait in question, yielding an equation that expresses the predicted level of a psychological trait as a function of brain activation in the ROI. Crucially, these analyses yield the expected error of the predicted level of trait.

The informativeness of a scan was measured as the reduction in the range of possible values of a trait for a hypothetical new individual subjected to the scanning protocol. Before the scan, the 95% confidence interval for

the subject is the mean of the distribution of trait levels, plus or minus 2 standard deviations of that trait. After the scan, we get a new and narrower interval. The ratio of these two is our measure of the scan's informativeness.

RESULTS

Table 2 shows the prediction interval for a hypothetical new subject before and after being scanned using the protocol of each of the 16 studies reviewed here. The reduction of the prediction interval, shown in the rightmost column, indicates by how much the range of possible trait levels would be narrowed as a result of the scan.

The reduction varies across studies, ranging from less than 10% to more than half. For 7 of the 16 studies, a scan would be expected to narrow the range of a person's possible trait levels by at least one third; for 3 of the 16 studies, the range would be reduced by at least half. To put these numbers in perspective, consider what we could expect to learn from some of the scanning protocols reviewed here.

The least informative scanning protocol identified by us concerns reading span (Reichle, Carpenter, & Just, 2000). If a new subject was scanned and the activity in the relevant region of the left inferior frontal cortex was one standard deviation below the average of the sample reported by Reichle et al. (2000), then we would know that his or her reading span score lies between 2.3 and 5.5 with a probability of 95%. Because the range of reading span scores for the sample is approximately the same, namely, from 2.0 to 5.0, the brain imaging protocol would have told us very little about the new subject's likely reading span.

A more informative protocol concerns the personality trait of extraversion as studied in an emotional Stroop task (Canli, Amin, Haas, Omura, & Constable, 2004). A subject whose anterior cingulate activity was one standard deviation below average in this task would be expected, with 95% probability, to have an extraversion score between 18 and 35. Although this remains a fairly wide range, it does exclude a substantial segment of the extraversion scale. For example, Canli's subjects spanned 21 to 42 on the scale. Simply on the basis of this hypothetical subject's scan, one could reasonably infer that the subject is not a particularly extraverted individual.

The most informative protocol identified here concerns navigational ability (Epstein, Higgins, & Thompson-Schill, 2005). If the brain activation of a new subject in this scanning protocol was one standard deviation below normal, then we could infer that his or her navigational ability is 95% certain to lie between 2.5 and 4.5. In comparison, the range for the sample is 3.2 to 6.1. Simply by measuring a person's brain responses when viewing pictures of scenes, without requiring any real or simulated navigation, we can learn a fair amount about his or her navigational ability—for example, in this hypotheti-

Table 1. The 16 Studies Analyzed Here, Grouped According to General Type of Psychological Trait

<i>Trait</i>	<i>Study, Sample Size</i>	<i>Trait Measure</i>	<i>Regions of Interest</i>	<i>Activation Task and Contrast(s)</i>
<i>Intelligence</i>				
Navigational ability	Epstein et al., 2005 (<i>n</i> = 12)	Santa Barbara Sense of Direction Scale	Parahippocampal place area	Place-change relative to viewpoint-change of an environmental scene
Verbal intelligence	Geake & Hansen, 2005 (<i>n</i> = 12)	Verbal Intelligence score from the National Adult Reading Test	Left BA 9 and 45/46	Letter-string analogy task
Visual-spatial skill	Reichle et al., 2000 (<i>n</i> = 12)	Vandenberg Mental Rotation Task	Left and right parietal	Sentence-picture verification task using a visual-imagery strategy
Reading span	Reichle et al., 2000 (<i>n</i> = 12)	Reading Span Test	Left and right inferior frontal	Sentence-picture verification task using a verbal strategy
<i>Negative Affectivity</i>				
Anxiety	Critchley, Wiens, Rotshtein, Ohman, & Dolan, 2004 (<i>n</i> = 11)	Hamilton Anxiety Scale	Right anterior insular/opercular region	Perception of self heartbeat
State anxiety	Bishop, Duncan, & Lawrence, 2004 (<i>n</i> = 27)	Spielberger State-Trait Anxiety Inventory	Left and right amygdala	(I) Subliminal perception of fearful faces relative to neutral faces (II) Attentional modulation of fearful vs. neutral attended faces relative to fearful vs. neutral unattended faces
Threat sensitivity	Cools et al., 2005 (<i>n</i> = 12)	Behavioral Inhibition Scale	Left and right amygdala	Gender categorization of fearful faces relative to happy faces
Pessimism (tendency to believe that negative life events occur for reasons likely to persist, whereas positive life events are exceptions and special cases)	Fischer, Tillfors, Furmark, & Fredrikson, 2001 (<i>n</i> = 13)	Attributional Style Questionnaire	Left and right amygdala	Viewing films of snakes
Harm avoidance	Paulus, Rogalsky, Simmons, Feinstein, & Stein, 2003 (<i>n</i> = 15)	Temperament and character inventory	Left and right insula	Response to punishment in simple gambling game of high-risk gambles
Neuroticism	Paulus et al., 2003 (<i>n</i> = 15)	NEO Five Factor Inventory	Left and right insula	Response to punishment in simple gambling game of high-risk gambles
Negative bias in racial evaluation (the tendency to evaluate other-race strangers more negatively than same race)	Phelps et al., 2000 (<i>n</i> = 12)	Implicit Association Test	Bilateral amygdala	Viewing of facial photographs of unfamiliar black faces, relative to unfamiliar white

Table 1. (continued)

<i>Trait</i>	<i>Study, Sample Size</i>	<i>Trait Measure</i>	<i>Regions of Interest</i>	<i>Activation Task and Contrast(s)</i>
Trait rumination	Ray et al., 2005 (<i>n</i> = 24)	Ruminative Responses Scale; Rumination subscale of the Rumination and Reflection Questionnaire; Anger Rumination Scale	Left and right amygdala	Cognitive reappraisal when viewing a negative or neutral photo
<i>Extraversion and Related Personality Traits</i>				
Extraversion (high energy level, sociability, assertiveness and a tendency to seek enjoyment; a prime dimension in virtually all theories of personality)	Canli et al., 2002 (<i>n</i> = 15)	NEO Five Factor Inventory	Left and right amygdala	Viewing of happy relative to neutral faces
Extraversion	Canli et al., 2004 (<i>n</i> = 12)	NEO Five Factor Inventory	Bilateral anterior cingulate	Emotional Stroop with positive words relative to neutral words
Behavioral Approach (tendency to approach and be activated by rewarding stimuli)	Gray & Braver, 2002 (<i>n</i> = 14)	Behavioral Approach Scale	Caudal anterior cingulate cortex	3-back working memory task

cal case, that it is not very good and, more likely than not, below average.

On the basis of this small sample of studies, certain brain regions appear to be more predictive than others of individual differences, in some cases, predicting individual differences in multiple different traits. For example, the amygdala appears in numerous rows of Table 2, in association with traits as diverse as pessimism, automatic racial evaluations, and extraversion. Of course, these diverse traits correlate with amygdala activations measured in very different tasks. The neurobiological correlate of these traits is not activation of an area per se, but activation of an area by a certain task.

DISCUSSION

The present results indicate that a modest degree of brainotyping capability already exists. The potential use of functional brain imaging to gain knowledge of someone's psychological traits is not science fiction, but rather a realistic possibility, albeit limited in important ways.

One limitation concerns generalizability of the relationship between brain activation and psychological trait to populations other than the typical volunteer subject for functional magnetic resonance imaging (fMRI) ex-

periments. A measure that correlates well in one sample will not necessarily correlate as well with people of other ages and backgrounds. A similar caution applies concerning the generality of fMRI measures across different circumstances. There is no guarantee that a correlation found under one set of circumstances, for example voluntary participation in a cognitive neuroscience experiment, will remain equally predictive under a different set, for example the stress of employment screening. Of course, concerns about generalizability are not unique to brain imaging measures. The specific questions asked on a paper-and-pencil test may be more effective at measuring personality or ability in one population than in another, and differences in the stress or distractions of the testing situation can influence paper-and-pencil test results as well.

Another caution regarding the present results concerns a possible bias in the studies reviewed. If other similar studies were carried out and produced null results that were not published, then a review of published results would overestimate the predictive power of functional neuroimaging. Again, this is a concern that applies equally well to other measures. It will be either born out or dispelled with continuing empirical work.

In some ways, the present review underestimates the predictive power of the current state of the art in

Table 2. The Prescan and Postscan Prediction Intervals for Each Study and Each ROI, Measured on the Scale of the Psychological Trait, and the Percentage Reduction of Prediction Interval from Pre- to Postscan

<i>Trait</i>	<i>Study</i>	<i>Regions of Interest</i>	<i>Prediction Intervals Prescan, Postscan</i>	<i>Reduction of Prediction Interval</i>
<i>Intelligence</i>				
Navigational ability	Epstein et al., 2005	Parahippocampal place area	0.814, 0.496	64%
Verbal intelligence	Geake & Hansen, 2005	BA 9	6.56, 3.54	46%
		BA 45/46	6.56, 3.21	51%
Visual-spatial skill	Reichle et al., 2000	L. Parietal	6.14, 5.51	63%
		R. Parietal	6.21, 5.20	19%
Reading span	Reichle et al., 2000	L. Inferior frontal	0.868, 0.804	7%
		R. Inferior frontal	–	–
<i>Negative Affectivity</i>				
Anxiety	Critchley et al., 2004	R. Anterior insular/opercular region	4.88, 3.55	27%
State anxiety (I)	Bishop et al., 2004	L. Amygdala	9.00, 8.10	11%
		R. Amygdala	–	–
State anxiety (II)	Bishop et al., 2004	L. Amygdala	8.11, 7.04	13%
		R. Amygdala	–	–
Threat sensitivity	Cools et al., 2005	L. Amygdala	–	–
		R. Amygdala	3.77, 2.76	37%
Pessimism	Fischer et al., 2001	L. Amygdala	2.27, 2.03	11%
		R. Amygdala	2.32, 1.94	13%
Harm avoidance	Paulus et al., 2003	L. Insula	–	–
		R. Insula	4.39, 3.88	12%
Neuroticism	Paulus et al., 2003	L. Insula	–	–
		R. Insula	0.412, 0.347	16%
Racial evaluation	Phelps et al., 2000	Bilateral amygdala	71.3, 60.9	15%
Trait rumination	Ray et al., 2005	L. Amygdala	1.44, 1.05	37%
		R. Amygdala	1.42, 1.13	26%
<i>Extraversion and Related Personality Traits</i>				
Extraversion	Canli et al., 2002	L. Amygdala	6.26, 4.58	27%
		R. Amygdala	–	–
Extraversion	Canli et al., 2004	Bilateral anterior cingulate	11.6, 7.14	39%
Behavioral approach	Gray & Braver, 2002	Caudal anterior cingulate cortex	6.65, 4.07	39%

functional neuroimaging. A simple way of boosting predictive power involves examining relationships between traits and higher-order patterns combining brain areas and different task conditions. For example, if the mea-

sure of extraversion obtained from the left amygdala in Canli et al.'s (2002) face-viewing protocol is not entirely redundant with the measure obtained from the anterior cingulate cortex in Canli et al.'s (2004) Stroop protocol,

then one could narrow the extraversion prediction interval further by administering both tasks to a subject.

The results of our review suggest that brainotyping using functional neuroimaging can provide more predictive biological correlates of psychological traits than genotyping. This is not surprising, in principle, because of the difference between the two mentioned earlier, namely, that brain function is a causal step closer to behavior than are genes. It is also consistent with the empirical observation that the percentage of variance in psychological traits accounted for by single genes is typically 5% or less (Parens, 2004; Van Gestel & Van Broeckhoven, 2003), whereas the correlations between imaging measures and such traits are often reported to be .6 and higher, accounting for 36% and higher of the trait variance.

Do our results have implications for current research practice? We believe that they do not, in the sense of requiring any additional protection of research participants' privacy. Rather, they suggest that existing IRB guidelines designed to maintain confidentiality of imaging data, which might once have seemed to be regulatory overkill for cognitive neuroscience researchers, are in fact prudent measures.

The measurement of socially relevant psychological traits with fMRI should be distinguished from the measurement of socially relevant psychological states, defined as more transient characteristics of the mind such as mood, momentary emotional states, and the contents of thought. Brain imaging has also been used in the study of mental states. Examples of states with social relevance include deception, false memory, and consumer interest. For example, the possibility of a brain-based lie detector has captured the imagination of the public and motivated numerous researchers to seek reliable neural correlates of intentional deception (for a review, see Wolpe, Foster, & Langleben, 2005). In contrast to the literature on traits, much of this literature is explicitly concerned with issues of real-world usefulness. Ironically, and in contrast to the trait literature, the prospects for useful measurement of truthfulness seem remote (Stern, 2002).

Bioethical issues have normative ("what should be") and empirical ("fact of the matter") components. The broader context of brain privacy includes several important normative issues. For example, is there a difference between behavioral privacy (i.e., privacy of words and deeds) and brain privacy (i.e., privacy of inner mental or neural processes)? Under what circumstances can brain privacy be breached? What is the proper balance between the rights of individuals to brain privacy and the needs of society, for example, in screening applicants for law enforcement, childcare, and other positions where character is crucial?

An empirical component of the brain privacy debate not discussed here is the public's understanding of brain imaging and its tendency to attribute greater objectivity and certainty to brain images than to other types of

information about the human mind (McCabe & Castel, 2008; Weisberg, Keil, Goodstein, Rawson, & Gray, 2008; Racine, Bar-Ilan, & Illes, 2005; Dumit, 2004). This must be taken into account when considering the social consequences of brain imaging.

Whether the foregoing issues will take on real-world relevance depends on another empirical issue: whether brain imaging, in fact, poses any actual or foreseeable threat to brain privacy, given the current state of the art. The present study is a first attempt to address this issue, and indicates that functional neuroimaging is, indeed, already capable of delivering a modest amount of information about personality, intelligence, and other socially relevant psychological traits.

Acknowledgments

We thank Eric Racine and two anonymous reviewers for their insightful comments on an earlier draft of this article, Pauline Baniqued for assistance with manuscript preparation, and Geena Sankoorik for assistance reviewing the 1875 abstracts. The research reported here was supported by NIH grants R21-DA01586, R01-HD043078, R01-DA14129, and R01-DA18913, and the MacArthur Project on Neuroscience and the Law.

Reprint requests should be sent to Martha J. Farah, Center for Cognitive Neuroscience, University of Pennsylvania, 3720 Walnut St., Philadelphia, PA 19104, or via e-mail: mfarah@psych.upenn.edu.

REFERENCES

- Beaver, J. D., Lawrence, A. D., van Ditzhuijzen, J., Davis, M. H., Woods, A., & Calder, A. J. (2006). Individual differences in reward drive predict neural responses to images of food. *Journal of Neuroscience*, *26*, 5160–5166.
- Bertolino, A., Arciero, G., Rubino, V., Latorre, V., De Candia, M., Mazzola, V., et al. (2005). Variation of human amygdala response during threatening stimuli as a function of 5-HTTLPR genotype and personality style. *Biological Psychiatry*, *57*, 1517–1525.
- Bishop, S. J., Duncan, J., & Lawrence, A. D. (2004). State anxiety modulation of the amygdala response to unattended threat-related stimuli. *Journal of Neuroscience*, *24*, 10364–10368.
- Canli, T. (2006). When genes and brains unite: Ethical implications of genomic neuroimaging. In J. Illes (Ed.), *Neuroethics in the 21st century*. New York: Oxford University Press.
- Canli, T., & Amin, Z. (2002). Neuroimaging of emotion and personality: Scientific evidence and ethical considerations. *Brain and Cognition*, *50*, 414–431.
- Canli, T., Amin, Z., Haas, B., Omura, K., & Constable, R. T. (2004). A double dissociation between mood states and personality traits in the anterior cingulate. *Behavioral Neuroscience*, *118*, 897–904.
- Canli, T., Sivers, H., Whitfield, S., Gotlib, I., & Gabrieli, J. (2002). Amygdala response to happy faces as a function of extraversion. *Science*, *296*, 2191.
- Cools, R., Calder, A. J., Lawrence, A. D., Clark, L., Bullmore, E., & Robbins, T. W. (2005). Individual differences in threat sensitivity predict serotonergic modulation of amygdala response to fearful faces. *Psychopharmacology*, *180*, 670–679.

- Critchley, H. D., Wiens, S., Rotshtein, P., Ohman, A., & Dolan, R. J. (2004). Neural systems supporting interoceptive awareness. *Nature Neuroscience*, 7, 189–195.
- Dumit, J. (2004). *Picturing personhood: Brain scans and biomedical identity*. Princeton, NJ: Princeton University Press.
- Duncan, J. (2003). Intelligence tests predict brain response to demanding task events. *Nature Neuroscience*, 6, 207–208.
- Epstein, R. A., Higgins, J. S., & Thompson-Schill, S. L. (2005). Learning places from views: Variation in scene processing as a function of experience and navigational ability. *Journal of Cognitive Neuroscience*, 17, 73–83.
- Farah, M. J., & Wolpe, P. R. (2004). Monitoring and manipulating brain function. *The Hastings Center Report*, 34, 35–45.
- Fischer, H., Tillfors, M., Furmark, T., & Fredrikson, M. (2001). Dispositional pessimism and amygdala activity: A PET study in healthy volunteers. *NeuroReport*, 12, 1635–1638.
- Gazzaniga, M. S. (2006). *The ethical brain*. New York: The Dana Press.
- Geake, J. G., & Hansen, P. C. (2005). Neural correlates of intelligence as revealed by fMRI of fluid analogies. *NeuroImage*, 26, 555–564.
- Gray, J. R., & Braver, T. S. (2002). Personality predicts working-memory-related activation in the caudal anterior cingulate cortex. *Cognitive, Affective, & Behavioral Neuroscience*, 2, 64–75.
- Gray, J. R., & Thompson, P. M. (2004). Neurobiology of intelligence: Science and ethics. *Nature Reviews Neuroscience*, 5, 471–482.
- Hyman, S. E. (2004). Introduction: The brain's special status. *Cerebrum*, 6, i–iv.
- Illes, J. (2003). Neuroethics in a new era of neuroimaging. *American Journal of Neuroradiology*, 24, 1739–1741.
- Illes, J., Kirschen, M. P., & Gabrieli, J. (2003). From neuroimaging to neuroethics. *Nature Neuroscience*, 6, 250.
- Illes, J., & Racine, E. (2005). Imaging or imaging? A neuroethics challenge informed by genetics. *American Journal of Bioethics*, 5, 5–16.
- Kennedy, D. (2004). Neuroscience & neuroethics. *Science*, 306, 373.
- Levy, N. (2007). *Neuroethics: Challenges for the 21st century*. New York: Cambridge University Press.
- McCabe, D., & Castel, A. (2008). Seeing is believing: The effect of brain images on judgments of scientific reasoning. *Cognition*, 107, 343–352.
- Matthews, A., Ward, J., & Lawrence, A. D. (2004). Individual differences in the modulation of fear-related brain activation by attentional control. *Journal of Cognitive Neuroscience*, 16, 1683–1694.
- New York City Bar Association. (2005). *Are your thoughts your own? Neuroprivacy and the legal implications of brain imaging*. New York: The Committee on Science and Law.
- Parens, E. (2004). Genetic differences and human identities: On why talking about behavioral genetics is important and difficult. *Hastings Center Report*, Special Supplement.
- Paulus, M. P., Rogalsky, C., Simmons, A., Feinstein, J. S., & Stein, M. B. (2003). Increased activation in the right insula during risk-taking decision making is related to harm avoidance and neuroticism. *NeuroImage*, 19, 1439–1448.
- Phelps, E. A., O'Connor, K. J., Cunningham, W. A., Funayama, E. S., Gatenby, J. C., Gore, J. C., et al. (2000). Performance on indirect measures of race evaluation predicts amygdala activation. *Journal of Cognitive Neuroscience*, 12, 729–738.
- Racine, E., Bar-Ilan, O., & Illes, J. (2005). fMRI in the public eye. *Nature Reviews Neuroscience*, 6, 159–164.
- Ray, R. D., Ochsner, K. N., Cooper, J. C., Robertson, E. R., Gabrieli, J. D. E., & Gross, J. J. (2005). Individual differences in trait rumination and the neural systems supporting cognitive reappraisal. *Cognitive, Affective, and Behavioral Neuroscience*, 5, 156–168.
- Reichle, E. D., Carpenter, P. A., & Just, M. A. (2000). The neural bases of strategy and skill in sentence–picture verification. *Cognitive Psychology*, 40, 261–295.
- Sheskin, D. J. (1997). *Parametric and nonparametric statistical procedures*. Boca Raton, FL: CRC Press.
- Singer, T., Seymour, B., O'Doherty, J., Kaube, H., Dolan, R. J., & Frith, C. D. (2004). Empathy for pain involves the affective but not sensory components of pain. *Science*, 303, 1157–1162.
- Stark, R., Schienle, A., Sarlo, M., Palomba, D., Walter, B., & Vaitl, D. (2005). Influences of disgust sensitivity on hemodynamic responses towards a disgust-inducing film clip. *International Journal of Psychophysiology*, 57, 61–67.
- Stern, P. C. (2002). *The polygraph and lie detection: Report of the National Research Council Committee to review the scientific evidence on the polygraph*. Washington, D.C.: The National Academies Press.
- Urry, H. L., van Reekum, C. M., Johnstone, T., Kalin, N. H., Thurow, M. E., Schaefer, H. S., et al. (2006). Amygdala and ventromedial prefrontal cortex are inversely coupled during regulation of negative affect and predict the diurnal pattern of cortisol secretion among older adults. *Journal of Neuroscience*, 26, 4415–4425.
- Van Gestel, S., & Van Broeckhoven, C. V. (2003). Genetics of personality: Are we making progress? *Molecular Psychiatry*, 8, 840–852.
- Weisberg, D. S., Keil, F. C., Goodstein, J., Rawson, E., & Gray, J. R. (2008). The seductive allure of neuroscience. *Journal of Cognitive Neuroscience*, 20, 470–477.
- Wolpe, P. R., Foster, K. R., & Langleben, D. (2005). Emerging neurotechnologies for lie-detection: Promises and perils. *American Journal of Bioethics*, 5, 39–49.
- Xue, G., Chen, C., Jin, Z., & Dong, Q. (2006). Cerebral asymmetry in the fusiform areas predicted the efficiency of learning a new writing system. *Journal of Cognitive Neuroscience*, 18, 923–931.