Are Neuroimages Like Photographs of the Brain?

Adina L. Roskies†‡

Images come in many varieties, but for evidential purposes, photographs are privileged. Recent advances in neuroimaging provide us with a new type of image that is used as scientific evidence. Brain images are epistemically compelling, in part because they are liable to be viewed as akin to photographs of brain activity. Here I consider features of photography that underlie the evidential status we accord it, and argue that neuroimaging diverges from photography in ways that seriously undermine the photographic analogy. While neuroimaging remains an important source of scientific evidence, proper interpretation of brain images is much more complex than it appears.

1. Introduction. Functional MRI provides a noninvasive window into human brain function, allowing us to associate the performance of particular cognitive tasks with measurements reflecting brain activity. The development of this technique over the last decade and a half has spawned a burgeoning new subfield of neuroscience. Several hundred academic journal articles using fMRI are published each month, and an unprecedented number of these studies are brought to the public's attention through the media and lay publications (Racine et al. 2005). Functional neuroimaging[1] stands at an uneasy crossroads between the scientific world and the public eye, with the compelling and accessible presentation of its data as greyscale images of brains with colorful hot spots of activity. The seeming accessibility of these images, and the grip they have on the scientific and public imagination, makes them important conduits of information about the

†To contact the author, please write to: Dartmouth College, Department of Philosophy, Hanover, NH 03755; e-mail: adina.roskies@dartmouth.edu.

‡This work was supported in part by a grant from the Leslie Humanities Center at Dartmouth College. I thank John Kulvicki for helpful comments, and Kim Sterelny, for making it possible for me to spend some time at the ANU with a grant from the Australian Research Council.

1. In this paper I use functional MRI and neuroimaging interchangeably, although technically speaking, fMRI is one of a number of existing neuroimaging techniques.

Philosophy of Science, 74 (December 2007) pp. 860–872. 0031-8248/2007/7405-0026$10.00

progress of neuroscience. The heightened profile of neuroimaging studies, and their increasing prominence in philosophical as well as everyday reasoning about behavior and cognition, prompt an analysis of the proper epistemological role for neuroimaging data.

My goal here will be to begin the task of such an analysis. I begin by assuming a rough characterization of the commonsensical apprehension of neuroimaging data that explains the epistemic weight such data are accorded: Images are thought of as photographs of brain activity. This characterization suggests a evaluative framework that provides us with a number of criteria against which we can measure neuroimaging. An analysis of whether imaging satisfies these criteria clarifies the evidential status of neuroimages as vehicles of scientific information, but also demonstrates the inaptness of the photographic analogy, suggesting that people approach neuroimages as sources of evidence with caution. Arguments such as these are important to bear in mind when, for instance, philosophical arguments rely upon on data from neuroimaging, and will become crucial in any context in which such data is offered as evidence. While most neuroscientists familiar with neuroimaging techniques are not apt to misapprehend the evidential status of brain images in the ways I suggest, these images are disseminated far beyond the neuroimaging community, and it is here that such arguments become particularly relevant. It is imperative that the dangers inherent in naïve public consumption of brain images become widely recognized (Racine et al. 2005; Illes 2006).

Non-specialists think of neuroimages as analogous to photographs of the brain. This is my starting assumption. Since its validation is fundamentally an empirical matter, I will not defend this claim in any detail, but I take it as eminently plausible and supported by informal inquiry. This commonsensical association between neuroimaging and photography is based on their similarity in several respects, among them that the output of both results in an image, that these images are produced mechanically, and that they purport to tell us something about the way the world is. Photography (and other related technologies, such as video) enjoys a privileged epistemic status (Walton 1984; Cohen and Meskin 2004, 2008): photographs are often used as evidence in courts of law, are important resources for historical inquiry, and provide people with a visual record of the past that avoids the ravages of time and the vagaries of memory.

What makes photographs epistemically special? Do the same characteristics lend neuroimaging a similar epistemic status? In order to answer these questions I will review some suggestions for photography's compelling effects on our beliefs, and will analyze neuroimaging with respect to this framework.

Data from neuroimaging studies are usually rendered as brain images. They need not be, but as visual representations they are highly compelling.

One explanation for the compelling nature of the data thus presented is that they provide the sense that we, through imaging, are seeing the brain in action, much as through photographs we see the things that are photographed.[2] However, the characteristics of brain images differ in important ways from the characteristics of photographs. While these deviations do not undermine neuroimaging's usefulness as a scientific method, they do suggest that the analogy between the two techniques is potentially misleading, and that their differences should afford neuroimaging a much different epistemic status than the public are inclined to attribute it.

Some images carry greater epistemic weight than others. From paintings one cannot reliably infer the existence of the subject or its properties, although paintings can faithfully represent these qualities of their subject. Photographs, on the other hand, seem to be more closely tied to reality. Photographs bear a greater epistemic weight than paintings do, though some paintings appear as realistic as photographs; indeed they can be visually indistinguishable. It is clear that if we were to provide evidence in court to associate a suspect with a crime, we would be better served in producing a photograph of that suspect engaged in the crime than we would be if we producing a visually indistinguishable painting of the suspect performing the very same action. Therefore, there are aspects of photography that contribute to its evidential status, but do not derive purely from the visual properties of the photograph. These, I propose, derive from factors related to the technology and conventions employed.

In what follows, I explore a number of factors that contribute to photography's suitability as an evidential medium: (1) visual properties of photographs instantiate the visual properties of their subjects; (2) they are causally and counterfactually dependent on their subjects; (3) they are belief-independent. My concern here is not to defend these properties as constitutive of the evidential weight ascribed to photographs, but rather to evaluate the aptness of the analogy between photography and neuroimaging by analyzing the extent to which neuroimages share the above characteristics.

2. Seeing Through to the Visual Properties. Both Cohen and Meskin

2. There are at least two types of factors to explore when analyzing the epistemic force of brain images. One involves analyzing the characteristics of the methods or techniques that generate the images, and the other involves analyzing the characteristics of the resulting images themselves that make them seem immediate, accessible, or easily interpretable. Both factors contribute to the power of the image, each in different ways. This paper will deal only with the former, but I recognize at the outset that the visual characteristics of the images themselves—their form of presentation—is also partly responsible for their perceived epistemic weight, and should be discussed in any complete account of the epistemic value of neuroimaging.

(2004, 2008) and Walton (1984) locate a great deal of the evidential import of photographs in the fact that many of the visual properties of the photograph are visual properties of the subject of the photograph. This is presumably the reason underlying Walton's claim that photographs are 'transparent', and that we 'see through the photograph to the object being photographed' (1984). While it is not my concern to defend Walton's transparency claim, that claim enjoys whatever plausibility it has because it seems clear that photographs carry information about many of the visual properties of their objects by instantiating those very properties, so that in seeing the photograph we can imagine we are seeing the object. The two-dimensional projections of object shapes, their luminance properties, and in color photographs their (approximate) colors are thus instantiated in the photograph.[3]

When people look at an image of brain activation, they are inclined to think that they are 'seeing' brain activity: neuronal firing acts as sparks of light to 'light up' brain areas in which activity is present. This is an illusion, for there are no visible properties of brain activity to be instantiated in the image. The processes that comprise the generation of a brain image are distinctly different from those that feed us information about the visual world. Functional MRI is not transparent in the sense that photographs are, for the information we are interested in and that the technique is sensitive to is not visual information at all. What fMRI allows us to visualize directly are magnetic properties of the water in the brain.

Explication of the above cannot be achieved without a brief sketch of the fundamentals of fMRI. Protons (water molecules) are magnetic dipoles; placed in a magnetic field the dipoles tend to line up parallel or anti-parallel to that field. A slightly larger proportion of water molecules line up parallel to the field than anti-parallel, resulting in a net magnetization in the direction of the external field. These dipoles can be tipped out of alignment by a radio-frequency pulse, causing them to precess around the net field vector as they decay back to their equilibrium magnetization. In precessing, they generate a magnetic signal. MR measures the time course of the signal generated by this precession. In functional MR, what is usually measured is the decay of the coherence of the MR signal as water molecules initially in-phase precess at different rates due to inhomogeneities in the local magnetic field, causing them to dephase.

3. It is true, of course, that not all photographs accurately reflect the visual properties of their subjects: filters and film properties can lead to systematic deviations. Black and white photographs do not preserve color properties because they do not record hue, so they only preserve some aspects of color properties. It is also clear that processing can have effects on the colors of a photographic image, and processing differences can lead to more or less accurate color portrayals of the true colors of objects.

In functional MRI, the subject's head is placed into the bore in the MR machine, which contains a large magnet. The subject's head is thus placed in a magnetic field, and protons in the person's blood and tissue generate the MR signal. Oxygenated and deoxygenated hemoglobin in the blood have different magnetic properties, and the local inhomogenities whose effects are measured with functional MR stem from changes in the ratio of oxygenated to deoxygenated blood that results from changes in blood flow and oxygen extraction that in turn is a result of neural activity. The signal thus generated is called the BOLD (Blood Oxygen Level Dependent) response. The complex interaction between neural activity, metabolism, blood flow, capillary volume, and oxygen extraction is understood only in general terms. No single quantitative model relating these various parameters is widely accepted by the scientific community.

The temporal relationship between the BOLD signal and neural activity is also complex. the BOLD response is significantly delayed (3–5 seconds) with respect to the neural activity that causes it, and the timescale of the response does not temporally mirror that of the neural activity. More often than not, the cognitive task of interest is complete long before there is an observable change in the MR signal. Empirical studies indicate that the BOLD response is approximately a linear function of neural activity, enabling us to infer some temporal and magnitude properties of the neural activity that generates the observed response (Buckner 2003). However, although both theory and empirical studies confirm a fundamental relation between brain activity and blood flow, our current understanding permits only broad correlations, not precise quantitative predictions about activity from the MR signal.

To return to the issue at hand, we see that as a technique, MR departs significantly from photography. What is being *directly* measured in functional MR is the timescale of the dephasing of water molecules in the brain, not neural activity. The visual properties of the image are not visual properties of the object being measured. Moreover, fMRI is not sensitive to brain activity, but to some of its causal effects. We infer the presence of brain activity from the change in the dephasing of the signal, and are warranted in doing so because of a general theory about the correlation between neural activity, metabolism, and blood flow. Nonetheless, fMRI allows us to represent nonvisual information as visual, so using standard interpretational methods, we are able to imagine that the colors we see 'painted' on a brain section is a seeing of the relative location and intensity of activity in the brain.

I have argued thus far that functional neuroimaging is unlike photography in that it does not let us directly see visual properties of the brain; it allows us instead to visualize nonvisual properties removed from those that we are seeking to understand, namely magnetic signatures resulting from changes in blood flow. We infer from this blood flow dependent signal that neural activity has changed. However, lest it be concluded that MR is misleading and unhelpful, it is important to recognize that although it does not enable us directly to see visual properties of the brain, it does allow us to see some visual properties indirectly. Shape, spatial extent, and location are properties of objects and events that are often detected visually. To the extent that we are interested in these properties of brain structure and in the relative location of neural activity, indirect access to them through visual representations made possible by imaging can be exceptionally useful. More on this below.

3. Causal and Counterfactual Dependence. It has been argued that photographs are suitable as evidence because of the informational relationship they bear to their subjects. Photographs are causally dependent on their subjects in a different way from the way that paintings and drawings are; they are also counterfactually dependent on their subjects (Walton 1984; Cohen and Meskin 2004). Had the subjects of a photograph been differently arranged, differently illuminated, had they had different visual properties, etc., the resulting photograph would have been correspondingly different. Note that this counterfactual dependence is not thoroughgoing: Certain properties of that which is photographed could be altered without altering the photograph: the hidden aspects of objects could change with no change to the photograph; clever wax duplicates of objects could be substituted without a resulting change in the visual properties of the photograph, and so on. However, the types of changes the subject of the photograph can sustain without a concomitant change in the resulting photograph are limited, and we have a good intuitive grasp of what those limits are.[4] Cohen and Meskin (2004) stress the importance of people's knowledge and background beliefs about photography in accounting for its evidentiary status.

Like photographs, data from neuroimaging are causally dependent upon the tasks the subject performs and the neural activity supporting the behavior. However, unlike the case of photography, we lack a clear understanding of what sorts of counterfactual situations could result in the same data. Understanding this will involve: (1) understanding how the technology works and could fail; and (2) understanding the connection(s) between tasks and neural activity and the MR signal. Since this topic bears both on the issue of whether imaging should be viewed as analogous to photography, and on the important scientific question of

4. Our increasingly sophisticated abilities to digitally manipulate photographs clearly change the character of the counterfactual dependence of photographs on their subjects. This will no doubt bear upon the evidentiary weight accorded to photography.

what scientific information we can expect neuroimaging to reveal, it is worth treating this in greater depth.

In order to determine the information content of a signal, we must know what states of affairs could produce it (Dretske 1981). In order to understand what information is carried by an MR image, we must understand what other tasks, activity, or extraneous factors could result in the same MR image. This is a substantive project in neuroscience in which many researchers are currently engaged, but some points can be made even now. At some level, it is certainly the case that different brain activity could lead to the same observed patterns of blood flow changes. We typically are content with stating that the MR signal reflects neural activity, but that is a rather vague account of the relation between neural activity and BOLD signals. There are two concerns here. One regards what the signal reflects, and the other concerns where the signal is generated.

While it is generally thought that the MR signal corresponds to firing rates of local neurons (the local field potentials measured by physiologists) (Logothetis 2003; Logothetis and Wandell 2004), the signal may reflect a variety of different changes in the brain areas at issue. For instance, it could reflect subthreshold activity, simultaneous excitation and inhibition (Tagamets and Horwitz 2001), or modulatory inputs from other areas (Logothetis 2003; Logothetis and Wandell 2004). It could possibly also reflect changes in neuronal synchrony in the absence of a concomitant change in the mean firing rates of neurons, large changes in the firing rates of small numbers of neurons, or small changes in the firing rates of large numbers of neurons. Importantly, MR cannot even distinguish between excitatory and inhibitory neural activity—both involve an increase in metabolic demand and thus an increase in the BOLD signal (Attwell and Iadecola 2002).

The preceding discussion only concerns the potential for distinguishable changes in brain activity without distinguishable differences in the resulting image, and they highlight the difficulties in figuring out what changes in neural activity can be inferred from the MR signal. Luckily, the degree of uncertainty in interpreting what the MR signal corresponds to in neural terms is not matched by the degree of uncertainty involved in localizing the signal. However, interpretive problems can arise here too: there are a number of ways in which non-neural factors can affect the MR signal, and thus the MR image. A host of artifacts can corrupt imaging data, prompting the judgment that neural activity is present at a location when in fact that is not the case. These range from motion artifacts and scanner artifacts to the erroneous interpretation that a signal from a draining vein distant from the actual site of brain activity is in fact a signal from the local tissue. Although these other factors affect the certainty with which we can ascribe changes in the MR signal to changes in neural activity, with proper analysis many of these sources of error can be ruled out.

It is important to bear in mind possible ways in which neuroimaging results can be equivocal, and to realize that elaboration of the kinds of detailed information about brain function that can be extracted from the MR signal is an ongoing scientific project. However, it is equally crucial to remember that imaging data is, within its limits, a reliable indicator of brain activity. MR is quite good at indicating where in the brain a signal is generated, within the bounds of its spatial resolution—in general, if a statistically significant signal is found in a particular location, we are warranted in inferring that there is a change in blood flow in that region. Numerous studies in nonhuman animals and in humans using PET (positron emission tomography) have demonstrated that increases in brain activity vary parametrically with increases in blood flow in many regions of the brain (Fox and Raichle 1984), licensing inferences about increased brain activity from increased MR signal. Although confirmation of this relation is not possible for all tasks and for all brain regions, we are warranted in expecting such a relation to hold generally. As the least invasive technique that allows us to infer localized increases in neural activity during task performance in humans, functional MRI is an invaluable scientific tool.

Thus, an important similarity exists between the informational relation photographs bear to their subjects and the informational relation holding between brain activity and neuroimages. Images generated by both techniques are causally and counterfactually dependent upon the phenomena they are meant to reflect. However, there is not parity in the two cases. The relevant differences between the nature of the counterfactual dependence in the two stem from differences in our mastery of the underlying theory and phenomena in the case of photography, and our relative ignorance of the analogous theory and phenomena in the case of brain imaging. We have a good intuitive theory about what kinds of information is carried by a photograph and how, and the causal processes in photography are in many ways like those in vision (Walton 1984). Non-specialists lack a correct theory in the case of neuroimaging, and their intuitive theories are massively mistaken, because the causal processes involved in imaging are not analogous to vision at all. Even scientists who understand the theory behind the technique don't know precisely what information about neural activity is carried by the MR signal. While neither photography nor neuroimaging bears a perfect informational relationship to that which it reflects, we have a much better grasp of that relationship in photography than in neuroimaging.

4. Belief-Independence. One of the primary reasons that photographs are taken to be objective forms of evidence is that they are belief-independent (Walton 1984). To be sure, the beliefs or the desires of the photographer may affect what she photographs, how she frames it, what lenses she uses, or how she lights it or exposes it. However, given that she does shoot a photograph, the resulting image is independent of her beliefs about what she photographs; it is dependent solely on the scene the camera is pointed at and the light reaching the film. Once the parameters of the photograph are set, the photographer's beliefs do not influence the resulting negative.[5]

Is neuroimaging belief-independent in the way that photographs are? It may be argued that a neuroimage depends only on the parameters of the scan, and the blood flow changes due to the task at hand. It may also be argued that the image is of precisely the cognitive behavior, neural activity, or blood flow that causes it. In this way, it is belief independent in much the way that photographs are.

However, photographs are also belief-independent in another way. They are interpretable without knowing the beliefs, ideas, or intentions of the photographer. In photography, the image itself suffices for interpretation. Interpreting neuroimages is not belief-independent in this way: proper interpretation of a neuroimage is not derivable from information contained in the MR image. What makes neuroimaging importantly different from photography with respect to belief-independence is the way in which beliefs figure in the interpretation of neuroimaging results (Walton 1984).

Let us call a technique belief-transparent if it is possible to appropriately interpret the image it produces with information contained in the image; belief-opaqueness results when the information needed for the interpretation is not present in the resultant image. Without going in to detail, most photography is belief-transparent: we possess an intuitive understanding of how to interpret pictures; we have an antecedent understanding of what the things represented are like;[6] and many factors which affect the image are more or less readable from the image. In contrast, neuroimaging is belief-opaque. We lack an intuitive understanding of how to interpret neuroimages because proper interpretation is highly theory laden; we have no antecedent knowledge of what the objects of our neu-

5. Again, this characterization of photography downplays the role of the photographer too much. Not only do the photographer's beliefs have important effects on the resulting negative, they also play an important role in the processing and developing stages, and thus of the final print. Nonetheless, the more simplified characterization above captures an important feature of photography, and accounts for why photographs are viewed to be better evidence of the nature of reality than are, for instance, paintings.

6. This was suggested to me by John Kulvicki.

roimaging studies are like (the studies themselves constitute our understanding of the phenomena we investigate); and the image fails to contain information crucial to its correct interpretation. These factors virtually guarantee that neuroimaging is opaque.

I focus here on a few crucial ways in which neuroimages are belief-opaque: task design (decomposition), subtractions, and statistical analysis. In the course of explaining these, we will also encounter other fundamental differences from photography.

Task Design. Neuroimaging is predicated on the well founded assumption that cognition results from the coordination of simpler processing steps: the brain is not a holistic cognitive engine, but one that has subregions with varying degrees of computational specificity. The goal of neuroimaging experiments is to help explain cognition by allowing the attribution of functional roles to different brain regions. Imaging data are collected from people performing cognitive tasks, often complex ones, while in the scanner. Most tasks involve perception of stimuli, a number of intermediate processing stages, and some output behavior. Experiments are designed to elucidate the various processing steps, and to do this requires having some initial ideas about how a particular complex cognitive task might be functionally implemented. Experimental design in neuroimaging rests upon hypotheses about how cognitive tasks are functionally decomposed. Which tasks one chooses is based on an implicit or explicit theory of cognition. Because we lack a theory for reading off task demands or behavior from blood flow changes, tasks performed in the scanner cannot be read off from the neuroimage.

Subtraction. More importantly, neuroimages are almost always comparisons of results from two different cognitive conditions. In subtraction paradigms, the data from a scan of a subject performing one task of interest is subtracted from the data of (usually the same) subject performing a different task. The subtraction is meant to remove the blood flow signal corresponding to ongoing processes common to both tasks, included those imposed by general homeostatic demands, being awake, having one's eyes open, and so on. Often the comparison task is not a neutral baseline task, but one that is chosen to match the task demands of the main task on a number of dimensions, the rationale being that the result of this subtraction will highlight brain activity associated specifically with the differences in the demands of the two tasks. Just as in ordinary subtraction, where the answer to a problem is dependent both on the number subtracted and the number it is subtracted from, brain images produced in neuroimaging studies depend critically on both the main task and the comparison task. Whether the differences seen in the signal are correctly interpreted hinges on: (1) knowledge of the particular tasks involved; (2) understanding of possible confounds; and (3) the degree to which the functional decomposition and matching of the two tasks is accurate (see above).

This exposes several ways in which neuroimaging diverges significantly from photography. First, the subtraction images generated by neuroimaging studies are belief-opaque: one cannot infer the comparison tasks that generated the image from the image alone. The very same raw data from the main task can give rise to a very different image depending on which comparison task is employed in the analysis, and very different tasks can give rise to very similar images if appropriate comparison tasks are chosen. Second, neuroimaging, unlike photography, is essentially contrastive. What one sees when one looks at a brain scan is the result of comparing the activity in one task condition with that in another. Photographs are of a scene, an object, etc., not the comparison of one scene with another.

Statistical Analysis. The MR signal generated from neuroimaging experiments is noisy, and the signal to noise ratio is low. A variety of statistical techniques are always employed in order to analyze the data gathered in an experiment. In order to increase the statistical power of the technique, multiple scans are averaged together, usually from the same subject, but sometimes across subjects. While averaging across blocks or trials within a subject significantly improves the technique's signal to noise ratio, it becomes immediately evident that this manipulation affects that which the resulting image represents. While a photograph captures a moment in time (the duration of which is variable, but typically quite short) of a particular scene, it seems as though an image of brain activity captures a temporally-condensed generalization of information about a cognitive process that occurs over a much longer duration. Changes in the way a task is processed may be obscured by this technique (though they can often be extracted from the original data with proper analysis).

This feature of neuroimaging is further complicated by studies between subjects: in these cases, the brains of each subject must be warped to a common anatomical space before they are averaged, and the resulting image is not only a generalization over time, but also a generalization over individuals. Such an image is visually indistinguishable from an image generated from a single subject. It is rather as if multiple-exposure photographs were undetectable as such because the effects of the multiple exposures are invisible.

Finally even when the dataset from which final images are made is established, statistics must be run on the resulting image data in order to determine what effects are significant. Which statistical tests should be used is a matter of debate in the field. The statistics chosen can have noticeable effects upon the resulting image and reported results (Saad et al. 2003). Results thus depend upon the beliefs of the researchers about how to analyze their results; the image presents the results but does not convey how the results were obtained. Similar considerations attend the choice of representational schemes that contribute to the final image.

Unlike photography, neuroimaging is belief-opaque: neither the functional decomposition assumed by the experimenter, the tasks involved in image generation (and thus what the data represents) nor the statistical methods employed are recoverable in information inherent in the image, yet all are necessary for proper image interpretation.

5. Summary. Many people view brain images as evidential devices analogous to photographs. I have argued that the analogy to photography is in some deep respects inapt. We do not 'see through' the visual properties of neuroimages to the visual properties of their subjects; we do not understand the causal and counterfactual relationships between the images and the data they represent to the same extent that we understand the them with photography. Part of the reason for this is that neuroimaging is belief-opaque in a way photography is not. These deviations from photography do not undermine neuroimaging's ability to provide us with important, sometimes invaluable, information about brain function and cognition, but they do affect the ways in which images should be interpreted. One danger with viewing neuroimages as evidentially on par with photography is that the implicit understanding we have of how photographs represent and of how they causally and counterfactually depend on the world can get carried over into our interpretation of neuroimages, with misleading and potentially harmful results.

REFERENCES

Attwell, David, and Costantino Iadecola (2002), "The Neural Basis of Functional Brain Imaging Signals", *Trend in Neurosciences* 25: 621–625.

Buckner, Randy L. (2003), "The Hemodynamic Inverse Problem: Making Inferences about Neural Activity from Measured MRI Signals", *Proceedings of the National Academy of Sciences USA* 100: 2177–2179.

Cohen, Jonathan, and Aaron Meskin (2004), "On the Epistemic Value of Photographs", *Journal of Aesthetics and Art Criticism* 62: 197–210.

——— (2008), "Photographs as Evidence", in Scott Walden (ed.), *Photography and Philosophy*, New York: Blackwell, 70–90.

Dretske, Fred (1981), *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.

Fox, Peter T., and Marcus E. Raichle (1984), "Stimulus Rate Dependence of Regional Cerebral Blood Flow in Human Striate Cortex, Demonstrated by Positron Emission Tomography", *Journal of Neurophysiology* 51: 1109–1120.

Illes, Judy, ed. (2006), *Neuroethics: Defining the Issues in Theory, Practice, and Policy*. Oxford: Oxford University Press.

Logothetis, Nikos K. (2003), "The Underpinnings of the BOLD Functional Magnetic Resonance Imaging Signal", *Jounal of Neuroscience* 23: 3963–3971.

Logothetis, Nikos K., and Brian A. Wandell (2004), "Interpeting the BOLD Signal", *Annual Review of Physiology* 66: 735–769.

Racine, Éric, Ofek Bar-Ilan, and Judy Illes (2005), "fMRI in the Public Eye", *Nature Reviews Neuroscience* 6: 9–14.
Saad, Ziad S., et al. (2003), "The Spatial Extent of the BOLD Response", *Neuroimage* 19: 132–144.
Tagamets, M.-A., and Barry Horwitz (2001), "Interpreting PET and fMRI Measures of Functional Neural Activity: The Effects of Synaptic Inhibition on Cortical Activation in Human Imaging Studies", *Brain Research Bulletin* 54: 267–273.
Walton, Kendall (1984), "Transparent Pictures: On the Nature of Photographic Realism", *Critical Inquiry* 11: 246–276.