

Functional MRI and the Study of Human Consciousness

Dan Lloyd

Abstract

■ Functional brain imaging offers new opportunities for the study of that most pervasive of cognitive conditions, human consciousness. Since consciousness is attendant to so much of human cognitive life, its study requires secondary analysis of multiple experimental datasets. Here, four preprocessed datasets from the National fMRI Data Center are considered: Hazeltine et al., Neural activation during response competition; Ishai et al., The representation of objects in the human occipital and temporal cortex; Mechelli et al., The effects of presentation rate during word and pseudoword reading; and Postle et al., Activity in human frontal cortex associated with spatial working memory and saccadic behavior. The study of consciousness also draws from multiple disciplines. In this article, the philosophical subdiscipline of phenomenology provides initial characterization of phenomenal structures conceptually necessary for an analysis of consciousness. These structures include phenomenal intentionality, phenomenal superposition, and experienced temporality. The empirical predictions arising from these structures require new interpretive methods for their confirmation. These methods

begin with single-subject (preprocessed) scan series, and consider the patterns of all voxels as potential multivariate encodings of phenomenal information. Twenty-seven subjects from the four studies were analyzed with multivariate methods, revealing analogues of phenomenal structures, particularly the structures of temporality. In a second interpretive approach, artificial neural networks were used to detect a more explicit prediction from phenomenology, namely, that present experience contains and is inflected by past states of awareness and anticipated events. In all of 21 subjects in this analysis, nets were successfully trained to extract aspects of relative past and future brain states, in comparison with statistically similar controls. This exploratory study thus concludes that the proposed methods for “neurophenomenology” warrant further application, including the exploration of individual differences, multivariate differences between cognitive task conditions, and exploration of specific brain regions possibly contributing to the observations. All of these attractive questions, however, must be reserved for future research. ■

INTRODUCTION

Explaining consciousness is widely regarded as one of a handful of ultimate scientific challenges. For centuries the study of consciousness has been, at best, indirect, speculative, and metaphorical. However, the development of practical methods for observing the dynamic human brain, especially functional brain imaging, offers new prospects for a science of consciousness. Most PET and fMRI research conscious subjects perform tasks with normal awareness. One could argue that all of these studies are already studies of consciousness, and that for a science of consciousness much of the data is already in.

The data may be in, but the hypotheses are not. Contemporary brain imaging informs a developing science of cognition, not consciousness. To employ functional brain imaging in consciousness science, one must first clarify at least some of the distinctive features of consciousness. From that follow empirical hypotheses, which in turn inform the methods of exploration. The

philosophical tradition most immediately concerned with the nature of consciousness is known as “phenomenology,” the study of phenomena, or of things as they appear. It can be divided into two types, “observational phenomenology” and “analytic phenomenology.” Observational phenomenology consists of “field reports” of mental states and their content. As collected by philosophers, observational phenomenological reports are informal first-person descriptions of experience. These are assembled in loose taxonomies, leading to various generalizations. The method of observation is introspection, which must be taken as reliable or even infallible in order to establish the intended phenomenological results. Examples of observational phenomenology abound, appearing in Descartes and especially in the empiricist and associationist tradition of Locke and Hume. It also founds important works in the history of psychology, such as those of William James. The informal distinctions among types of mental states (also known as “folk psychology”) continue to embellish many philosophical books on consciousness. Nonetheless, even with the determined efforts of a generation of introspective psychologists, the observations of observational phenomenology could not

be stabilized from one observer to another. It was a scientific failure, and its failure is often cited as a failure of phenomenology overall.

“Analytic” phenomenology is a different enterprise, however. It does not rely on introspected data. Rather, it concerns itself with “the world as it appears,” a world of which the observing self is a part. And, of the experienced world, analytic phenomenology asks, what must be true of consciousness in order for experience of the world to be possible at all? Analytic phenomenology is thus primarily conceptual rather than empirical. The classical source for analytic phenomenology is Immanuel Kant, but it flowers in the work of Edmund Husserl (1859–1938). Husserl gave the term “phenomenology” its modern philosophical meaning and method, inaugurating a movement that included Heidegger, Sartre, Merleau-Ponty, and many others.

A full survey of analytic phenomenology will not be possible here, but three principal claims delineate the phenomenological anatomy of consciousness, and thus inform empirical study:

(1) *Intentionality is internal to consciousness.* “Intentionality” appears widely in recent philosophy as the general name for the relationship between mind and the external world, or between states of mind and their “intentional objects,” what the mental states are about. But in phenomenology, intentional objects are always regarded exclusively as they present themselves within consciousness. Accordingly, phenomenological intentionality names an always-present duality within subjective states, the regular correlation between a subject pole (how the object is experienced) and an object pole (what the object is experienced as). Intentionality in this sense “moves the world inside the head.” This shift offsets an exaggerated emphasis on self-consciousness and introspection—the subject pole of intentionality—restoring the import of perception, or the exploration of the world as it is experienced—the object pole. But at the same time, it demands that conscious cognition of the objective world be described in its subjective presentation. Phenomenology suggests that conscious percepts are not merely the detection of external conditions, but have subjective characteristics that entail that consciousness is far more than a mechanism for tracking the environment. Within the superstructure of intentionality, analytic phenomenology seeks additional structures of consciousness necessary to constitute the experience of an objective reality. This subjective examination of objectivity leads to two further central claims.

(2) *Both sensory and nonsensory properties are superposed in all perception.* Philosophers and psychologists often pose a variety of distinctions between sensation and perception, but in consciousness, neither appears without the other. The sensory field of properties directly detected by the senses is always completed in consciousness with many elaborations. Husserl dis-

cusses a simple example of a hand placed on a sheet of blank paper (Husserl, 1907, Sec. 22). Rather than seeing a white expanse with a hand-shaped gap in it, consciousness completes the expanse behind the hand, filling in visual qualities that are nonetheless not immediately given in sensation. One also feels the paper beneath the hand, but rather than feeling a hand-shaped patch of smoothness, the texture is also extended in consciousness to the edges of the sheet. These sensory properties fill regions of the object that are not directly sensed, and fill in the hidden backsides of all objects, with various degrees of definiteness. Moreover, both tactile and visual properties coexist in the consciousness of a single, unitary sheet of paper. Finally, the paper is inflected with an indefinite swarm of other meanings and interpretations, which are nonetheless part of the state of consciousness of the object: It may be the page on which you hope to begin a new piece of writing, or the last piece of paper in the house, or a random sheet pulled from a pile in order to illustrate a point in phenomenology. These nonsensory properties are nonetheless present in awareness (Lloyd, 1995, 1996). All of these extensions, fillings, wraps, and bundles go beyond what is immediately sensed of any stimulus, but all are essential in any consciousness of objects “as” objects.

(3) *Temporality is superposed in all states of consciousness.* Time appears in psychology as a stimulus dimension to be detected in judgments of duration and order, or in planning and memory, but in phenomenology, the temporal is a fundamental aspect of all perception. Husserl uses music as his example: Each tone in a melody carries with it its melodic context. If it did not, the perception of the melody could only be a succession of tones with no relationship among them. A single tone is also temporal, as its elapsed duration inflects the present awareness that it is (still) sounding. All objects of perception share in temporality, as perceptions of both change and stability inherently involve comparisons of present perceptions with past and anticipated perception (Husserl, 1893–1917).

The phenomenology of temporality entails two corollaries: (3a.) *All objects, whether stable or changing in themselves, are experienced in continual temporal flux.* Even the most fixed and stable objects are always adding to their duration, increasing their fixity and stability. This nonsensory and always unfolding property is superposed in all perception. You can’t step in the same stream of consciousness twice (Lloyd, 2000a). (3b.) *All states of consciousness encompass in themselves the awareness of the past leading up to the present, and an awareness of future possibilities.* The immediately “given” element of direct perceptual awareness is the primal impression. That is, the phenomenal present is not exclusively occupied with present sensory properties and their occurrent superpositions. Rather, enclosed in the here-and-now is a nonsensory contrail of past nows and a branching sketch

of future possible nows. Husserl called the history built into the present “retention,” and the anticipation “protention.” Only with this structure can the objects of the world appear to us as temporally extended. In adopting special terms, Husserl explicitly distinguished the retentive and protentive superpositions of the experienced present from the phenomenologically distinct processes of recollection and planning.

Husserl offered a dynamic geometric representation of his proposal, updated here as Figure 1. In contrast to a one-dimensional timeline with the experienced present as a moving point, he envisioned the experienced present as a line sweeping across a plane (Husserl, 1893–1917, Sec. 10). At each instant, the line represents past, present, and future, all arrayed within a present state of consciousness. As time passes, the current present shifts into conscious retention, as the current protention slides toward the present. The vertical line represents a temporal “landscape” that both changes in itself and continually shifts in all its content from protention to present to past. But although only the here and now is sensed, the entire temporal landscape (the line) is in conscious awareness

at every moment. Both continuity and change are presented as continuities and discontinuities in the landscape, and thus comprise a facet of the experience of any object.

As conclusions of analytic phenomenology, the three foundations of intentionality, superposition, and temporality are proposed to be essential to consciousness. Establishing this deserves a longer discussion, but, in brief, they can be tested logically by trying to conceive of the awareness of any object without each of the three. Without them, objects collapse into dissociated unrelated instantaneous fragments. The remaining richness is thin, an “awareness” akin to the internal state of a video camera. Complicated as this state may be, it would not be a state of consciousness (for further discussion, see Lloyd, in press.)

Methods for “Neurophenomenology”

Let us suppose that the foundations of analytical phenomenology above do identify distinct and always-present aspects of consciousness. If so, they suggest the general empirical hypothesis that separately or collectively, these aspects of consciousness are implemented in the brain. Applied to functional brain imaging, a preliminary empirical exploration of consciousness would seek neural correlates of these phenomenal structures. Because the phenomenological analysis applies generally to all states of consciousness, its neural expression should appear in all conscious subjects. Accordingly, the exploratory analysis of consciousness should examine many subjects in many task conditions. While it might be possible for an individual researcher to collect the number and variety of datasets required for this analysis, only with a resource like the National fMRI Data Center is it practical for an individual or research community to embark.

Phenomenology also impels a number of methodological choices. Existing interpretive strategies for functional imaging systematically exclude the potential correlates of consciousness. Thus, new interpretive approaches will be needed, methods specifically appropriate to “neurophenomenology” (a term coined by Varela, 1996). The initial method outlined below is exploratory, designed to cast a net widely in search of functional evidence of the structures of consciousness. It will be outlined in this section, detailed in the next section, and revisited in the Discussion. (Other discussions of the interplay of phenomenology and method can be found in Gazzaniga, 2000; Metzinger, 2000; Dennett, 1991.) Initial constraints include the following:

Time Points in a Scan Series Should be Considered Individually. Existing methods partition time points with reference to task or stimulus conditions, and generally average time points within and across subjects. This is inappropriate for neurophenomenology. Temporality implies that even when task or stimulus conditions




Figure 1. Temporality in human consciousness. Husserl proposed that the experience of temporally extended objects required that all awareness, and especially perceptual awareness, must have a temporal structure. In addition to a “primal impression” of the immediately present stimulus, consciousness includes “protention” (anticipation) and “retention” of a stack of former nows stretching back toward a temporal horizon. This diagram is adapted from Husserl’s own illustration of temporality. As each now point passes, it is retained as an ever-changing modification of the next now. As each now subsides, a new primal impression emerges from the future, which is also a constant inflection on present experience. So event a, occurring first in a string of three events, ultimately lies deeper in retention, and the more recent b is stratified nearer the present. Meanwhile, anticipated event c also inflects present experience as it approaches. If time is a point moving along a horizontal line, then experienced time is a vertical structure. It is “constituted” entirely in the present moment.

are static, consciousness at each time point is distinct from points before and after. Temporality, superposition, and intentionality jointly imply that repeating stimuli or tasks are never experienced the same way twice. In general, we cannot assume that external descriptions of task or stimulus describe more than a small component of the progression of consciousness.

Subjects Should be Considered Individually. All three conditions greatly expand the possibilities for individual differences. Intersubject averaging could eliminate individual evidence for the expression of consciousness (see also Aguirre, Zarahn, & D'Esposito, 1997).

Brain States Should be Considered Globally. Initially, we can make no assumptions about the localization of any aspect of consciousness, and phenomenology generally suggests that states of consciousness can incorporate any combination of cognized properties. We seek instead “distributed” states of activation, possibly involving large areas of the brain (Ishai, Ungerleider, Martin, & Haxby, 1999, 2000).

Combining these three methodological constraints, the general strategy for exploring consciousness will regard the individual image volume (time point in a scan series) as the unit for analysis, and the individual subject as the domain. Since each volume is composed of more than 10,000 voxels, considering each state globally entails taking account of more than 10,000 variable dimensions. Since states cannot be partitioned with respect to a reference function or stimulus condition, we must assess them “internally” in relation to each other. Multivariate statistics offers various means for measuring similarity among complex patterns, enabling a schematic description of the high-dimensional relations between and among volumes. Within this framework, how would consciousness manifest itself empirically? The approach taken here considers superposition and temporality separately, and for each seeks indirect evidence of its effects. General predictions include the following:

Temporal Flux

Time is the river that carries all else, so a strong prediction would be that the flux of time would appear as a monotonic increase of intervolumetric multivariate difference as a function of the intervening interval in time, or lag between images. We could think of this as a gradient. From any time point, the next volume will also be most similar in its global pattern, and as lag increases, the intervolumetric difference also increases.

The Tripartite Structure of Temporality

Phenomenologically, each moment of consciousness is a sandwich of past, present, and future. Accordingly, each pattern of activity in the brain will be inflected with past

and future as well. But “past” and “future” can only be understood internally, that is, as past and future states of the brain. To discover tripartite temporality, then, we seek to detect some form of continuous neural encoding of past states, as well as some anticipation of the future. An indirect method for detecting both forms of encoding will be discussed below.

Superposition

From moment to moment, every object carries a shifting set of meanings, which to discern in a brain image would be difficult. Nonetheless, among the many interpretations present in any state of mind, at least a few do track stable conditions in the real environment. So if the multivariate analysis of brain states is valid, among the discernable relations among volumes should be similarities among volumes sharing task or stimulus conditions. This would be a multivariate counterpart to the traditional use of imaging to localize cognitive function.

Superposition and Time

Where sequences are repeated, subjects are also aware of the passing of time relative to repeated and regular events during the scan. That is, if task blocks and intertask intervals are both of fixed length, particular volumes occupying the same relative time points within separate sequences should reflect the awareness of relative duration. This analysis can also be combined with the analysis by task condition.

In sum, where the goal is to isolate the neural implementation of specific cognitive capacities, much of the breadth of conscious cognition is necessarily considered to be noise and “nuisance variables.” In contrast, phenomenology regards isolated cognitive capacities as very thin slices of a much broader continuous consciousness, and thus regards the “noise” as potentially salient data. Accordingly, the established methods for isolating functional brain areas, resting on the intersection of sets of voxel activations (achieved via cognitive conjunction, subtraction, and averaging) are inappropriate for neurophenomenology. Instead, the multivariate approach taken here is one of “cognitive disjunction,” seeking global and continuous properties present in all conscious states.

METHODS AND RESULTS

The Studies

This secondary analysis is based on 27 subjects from five experiments collected in four datasets provided by the National fMRI Data Center (Hazeltine, Poldrack, & Gabrieli, 2000; Ishai et al., 2000; Mechelli, Friston, & Price, 2000; Postle, Berger, Taich, & D'Esposito, 2000). Table 1 lists authors, titles, tasks, and total duration of

Table 1. Scope of the Analysis

<i>Authors, Title, Reference</i>	<i>Task</i>	<i>Conditions</i>	<i>n</i>	<i>Duration Volumes (min)</i>
Hazeltine, Poldrack, and Gabrieli, "Neural Activation During Response Competition," <i>JOCN</i> 12(2): 118–129	Visually cued response	Congruent flanking stimuli, Incongruent flanking stimuli	8	60 (2.8)
Ishai, Ungerleider, Martin, and Haxby, "Representation of Objects in the Human Occipital and Temporal Cortex," <i>JOCN</i> 12(2): 35–51	I. Delayed visual match-to-sample (photographs)	Houses, Faces, Chairs	4	1092 (54.6)
	II. Delayed visual match-to-sample (line drawings)	Houses, Faces, Chairs	5	1092 (54.6)
Mechelli, Friston, and Price, "Effects of Presentation Rate During Word and Pseudoword Reading: A Comparison of PET and fMRI," <i>JOCN</i> 12(2): 145–156	Reading	Words, Pseudowords (varying rates)	6	360 (18.9)
Postle, Berger, Taich, and D'Esposito, "Activity in Human Frontal Cortex Associated with Spatial Working Memory and Saccadic Behavior," <i>JOCN</i> 12(2): 2–14	Spatial memory, tracking	Free horizontal saccade, Guided saccade by moving stimuli, Forward memory of stimuli path, Manipulate memory of stimuli path	4	1632 (54.4)

each experiment. Collectively, the studies fall far short of canvassing all the varieties of conscious experiences, but they nonetheless span diverse aspects of human cognitive life. Perceptual targets included photographs and line drawings of objects from salient perceptual categories (faces, houses, and chairs); spots viewed with and without Stroop-like distractors; moving targets; words and pseudowords presented at differing rates. Tasks included passive viewing, matching to samples, target tracking, spatial memory, spatial memory manipulation, and reading. Motor responses included button presses and saccades. Subjects were omitted from secondary analysis only if data were incomplete.

Preprocessing

The original authors preprocessed their raw scanner data using methods common to many fMRI studies. These include spatial registration and normalization (except in Postle et al., 2000), spatial smoothing (again, except for Postle et al., 2000), and correction for motion artifacts including bulk head motion and spin-echo artifacts. After downloading from the fMRIDC, the preprocessed volumes were thresholded, retaining voxel time series with activations greater than 10% of the mean maximum voxel time series magnitudes for each subject. Multivariate intervolumes distances were then calculated for all volume pairs within each subject, using a squared Euclidean distance metric. That is, "distance" between two volumes is the sum of squared differences between corresponding voxel values in each image.

Geometrically, Euclidean distance as a multivariate measure has the same interpretation as familiar distance measurement in two- or three-dimensional space. By extension, each active voxel can be regarded as an abstract dimension, or axis in a high-dimensional space. Patterns of voxel activity specify vectors, so each pattern is equivalent to a single point in "voxel space," a space with as many dimensions as active voxels. Points in voxel space thus map multivariate relations among patterns of voxel activity.

Accordingly, in this analysis, the squared Euclidean distance is a measure of intervolumes difference or dissimilarity. Since it is the sum of squared differences between corresponding voxels in two images, it represents a global comparison of patterns of voxel activity. Smaller distances imply greater global similarity. Smaller distances also correspond approximately to the inverse correlation between volumes. (If images are normalized, Euclidean distance corresponds exactly to inverse correlation.)

Temporal Flux and Superposition

First, two figures present qualitative observations. Figure 2 is an image based on the numerical distance matrix derived from the subjects in Mechelli et al. (2000). For each volume, the figure represents its multivariate distance from all other volumes, represented in order of the image time series. The organization of the figure parallels Husserl's diagram in Figure 1. Time points are ordered from left to right; from each time point the

volumes to come are above the equator of the present, and those past are below. For example, the first column ($x = 1$) represents distances from the first image volume to all the other volumes in the scan series. The lag between this volume and each of the subsequent volumes is represented on the y -axis. Thus, the distance between the first and last image volumes is represented at the top left of the figure. The magnitude of each volume–volume distance is color-coded. (Alternatively, the graph can be considered to represent, approximately, the inverse of the covariance between volumes, as a function of lag.) As the time series progresses, fewer

volumes remain ahead in the series, while more accumulate behind (relatively past), until the series ends with image number 360.

The figure shows that, generally, as time passes the brain is changing globally, incrementally, and monotonically in the sense that interimage distance is continually increasing with passing time. This is in broad agreement with the phenomenological prediction of continual temporal flux. Figure 3B further abstracts and condenses multisubject data, averaging the intervolumetric distances that comprise Figure 2. Figure 3 shows the mean intervolumetric distance as a function of lag between




Figure 2. Multivariate distance and interscan interval in Mechelli et al. The figure represents the mean of six subjects studied by Mechelli et al. As in Figure 1, time is represented along the x -axis. For each of 360 volumes, the figure represents its multivariate distance to all of the others, with the y dimension representing time intervals (lag) between volumes. At the start of the scan session, all the other volumes are above the timeline (i.e., in the future relative to the current time point). With each subsequent scan, the time series sinks one unit toward the relative past. Color represents multivariate distance between each scan and the current volume. The figure suggests a monotonic correlation between time interval and distance, consistent with the phenomenological expectation of pervasive temporal flux. The smooth continuous variation also suggests that volumes can be individuated “internally,” by virtue of their relations to other volumes, consistent with the phenomenological principle of intentionality (see also Lloyd, 2000).

volumes, in this instance based on four of five subjects in Postle et al. (2000). The figure reflects the monotonic trend shown in the previous figure, suggesting that this trend is strongly linear. (The decline at extreme lag values may reflect the relatively few pairwise distances at large lag values. The mean of these large lag distances will be more sensitive to image noise, reducing sensitivity to temporal flux.)

Figure 3 also suggests either sequence effects or category effects, visible as distinct multivariate waveforms added to the linear effect. The regularity of these profiles reflects specific frequency components introduced by the experimental design.

This effect and the other predicted effects can be quantified. The measurements of temporal flux and superposition proposed here share a common approach. Global multivariate measures begin with calculating all the pairwise distances among image volumes. The result is a distance matrix spanning the entire experiment for each subject. Global effect measures are comparative averages based on intervolumes distances within various subsets of each distance matrix. Within an experiment, various time points may share a property of interest (e.g., temporal proximity, position in sequence, or task condition). If imaging can detect common neural expressions of a property of interest, then volumes at the time points with the property in common should be, on average, more similar to each other than the grand mean of all intervolumes distances. For example, if the brain exhibits a distinct response when detecting a specific stimulus, then all the volumes collected during time points when the stimulus is presented should express the response (along with many other responses). The common neural response should cause the multivariate distances among all the relevant volumes to be smaller, reflecting the similarity of response in the volumes. The null hypothesis in each case will be that volumes sharing the property of interest are, on average, no more similar to one another than the average of all distances in the distance matrix.

The hemodynamic response function presents a significant confound, however. The BOLD effect peaks and attenuates over several seconds following the probable neural event (Bandettini, 1993; Bandettini & Cox, 2000). As a result, for any volume, the immediately subsequent volumes will express the tail of the hemodynamic response, and this would lead to greater similarity (i.e., reduced multivariate distance) between the first volume and its immediate successors. In the context of large-scale multivariate comparisons, serial adjacency contributes only trivially to global measurements (1.6% in the smallest matrix, 0.06% in the largest). Nonetheless, to mitigate this effect, all analyses based on distance matrices excluded a minimum of 10 sec of the scan series immediately preceding and following each volume. This correction is conservative insofar as the removal of the HDR confound also eliminates any contribution of the

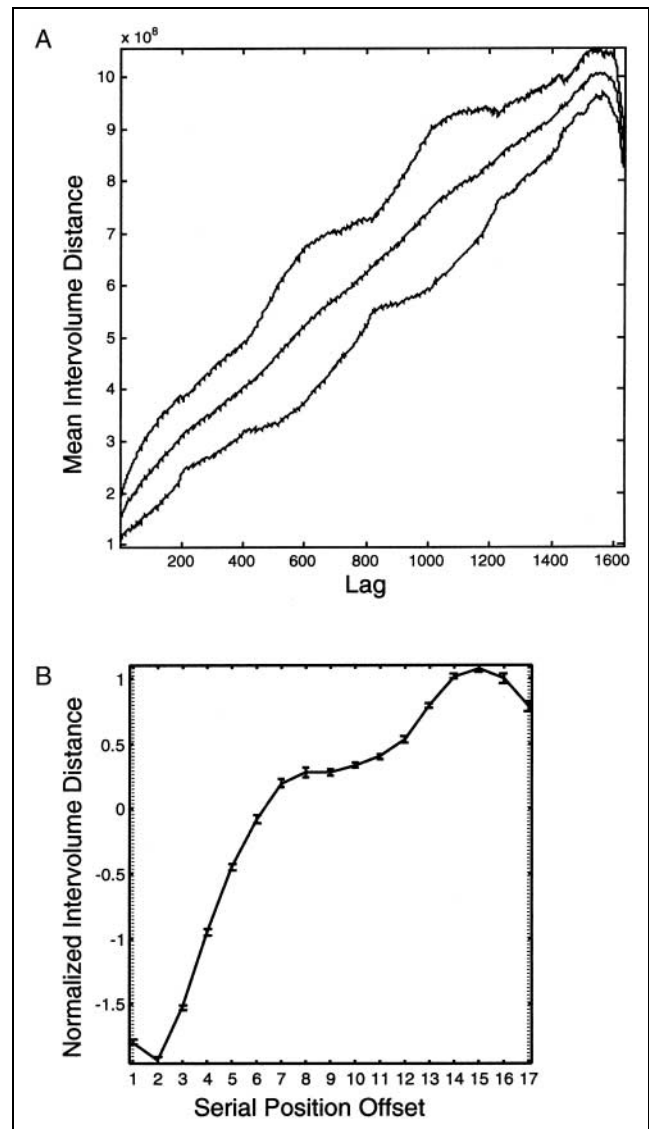


Figure 3. Multivariate distance and interscan interval in Postle et al. (A) The individual volume–volume distances are averaged here, representing the mean relationship between elapsed time interval between scans and multivariate distance. Flanking lines are one standard deviation from the mean. The mean of four subjects from the Postle et al. study reveals a strongly linear time–distance relationship. (The other studies exhibit qualitatively similar linearities.) (B) Mean distance values in A also exhibit a regularity that reflects the structure or duration of the task block in Postle et al. To display this component of the time–distance effect, the mean distance values shown in the top panel were decomposed into subcomponents equal to the task block length of 17 time points. After normalizing to remove the large overall time–distance effect from each subcomponent, the mean result appears in this panel. Error bars represent the standard error of the mean. The small variance for each lag value suggests a strong multivariate effect due to relative serial position. That is, in this instance, the general monotonic and nearly linear temporal flux measured as multivariate intervolumes distance is very sensitive to intervolumes differences at shorter temporal intervals (± 7 images, or 14 sec). This regularity seems unaffected by the distinct tasks used by Postle, including guided and unguided saccades, a sequential spatial memory task, and a spatial memory manipulation task. At lags of around 20 sec, multivariate distance becomes a less precise indicator of elapsed time. This may be an artifact of task block structure in the Postle study. See text for discussion.

excluded volumes to the measurements of the variables of interest here.

The global measurement of temporal flux was calculated as follows: For each volume, the pairwise distances to all the other volumes (excluding volumes within 10 sec) were partitioned into two equal groups according to temporal proximity in the scan series. The mean of intervolumetric distances for all of these temporally more proximate pairings was then compared to the mean for all intervolumetric distances. The first column of Table 2 shows this comparison as the ratio of the two means: (temporally proximate distances)/(all distances). Smaller values indicate larger effects. Significance was assessed with a one-tailed *t* test.

Table 2 shows that 24 of 27 subjects displayed a significant multivariate distance effect correlated with temporal interval, or lag, as measured using all the volumes from the full experiment. Table 2 also shows three other effects conforming in various degrees in different subjects to the qualitative predictions above.

Task (Stimulus) Condition

This was calculated by comparing all intervolumetric distances among volumes sharing a task or stimulus with the mean of all intervolumetric distances. The table reports the number of conditions showing significant ($p < .05$) effects as a fraction of the number of different conditions overall. Twenty-one of 27 subject scan series showed this effect for at least one task condition.

Position in Sequence

Several studies used task blocks and intertask intervals (control blocks) of the same duration. Where this occurred, it was possible to test for an effect reflecting the relative elapsed time in each sequence, independent of task similarity or overall serial position in the scan series. This was calculated similarly to the task comparison. Table 2 shows the count of serial positions within task blocks, with significant similarities as a fraction of the overall number of serial positions within each block, that is, the length (in time points) of each task block. Seventeen out of 21 subjects exhibited at least one instance of this serial position effect. (One study, Mechelli et al., 2000, was excluded due to stimulus sequences mapping onto fractional scanning intervals.)

Tripartite Temporality

Phenomenology's boldest claim is the assertion that the experienced present of an object always carries with it a nonsensory superposition of the object's history and possible future. Empirically, we expect the neural expression of this aspect of consciousness to be complex: Somehow, the present state of the brain represents a nested cascade of prior states and an open-ended laby-

rinth of possible futures. Add to this the complication of this exploratory analysis, namely, that whole brains are subject to multivariate analysis. As a result, the functional specificity of particular regions cannot help here.

Thus, in the absence of specific hypotheses about how tripartite occurrent temporality could be encoded, a more indirect probe was designed to indicate whether temporal information was globally encoded without specifying how this encoding might work. By analogy, the analysis sought to construct a "translator" that uses the present voxel activations to reconstruct immediately prior patterns of activations, as well as predict subsequent patterns. If such a translator can be shown to work, then its success shows indirectly that the past and future information was indeed present to be extracted. Artificial neural networks were used in this capacity here. For each subject in four experiments, a network was created for detecting temporal information in image volumes. To do this, scan series were segmented in serially adjacent pairs. To test for retention, the networks were trained to use the second image of each pair as input to produce as desired output the prior image. To test for protention, another set of networks was trained to map the first volume of each pair onto the second. Successful training in either direction will suggest that images (and underlying patterns of neural activity) contain information about prior or subsequent states of the brains (as expressed in images), consistent with phenomenological predictions.

Artificial neural networks have two other attractive features in this analysis. First, over time, we expect both "fan out" and "fan in" of information in the dynamic brain. That is, at any moment, activity at a particular voxel can be influenced by prior activity in a large and distributed set of voxels (fan in), and likewise, a particular voxel can influence many voxels subsequently (fan out). Neural networks can model this property. Second, neural networks are well suited to model nonlinear dynamics, and this seems appropriate to the brain, considered apart from phenomenology as well. These two capacities of neural networks allow for broader detection of temporal dynamics than, for instance, linear or first-order serial correlations. Overall, the interpretive strategy casts a broad net—spatially and functionally—into a complex dynamic system.

The computational load for this analysis prevented its execution for one of the five experiments (Mechelli et al., 2000), and for the remainder, it required significant preprocessing. Specifically, the subject sets of voxel time series were decomposed into their principal components, orthogonal vectors that capture the largest dimensions of variance common to all the voxel time series. (Principle component analysis was implemented using singular value decomposition in Matlab.) For this study, 50 principal components were retained from each subject series, capturing 80–99% of the variance in the various series. The resulting time series

Table 2. Results of Secondary Analysis of Individual Subject Scan Series

	<i>A. Multivariate Distance/Interscan Time Interval</i>			<i>B. Network Analysis: MSE/(Control MSE) [Trials]</i>	
	<i>Full Experiment [Volumes per Subject]</i>	<i>Within Sequence (<i>p</i> < .05)</i>	<i>Within Condition (<i>p</i> < .05)</i>	<i>Retention</i>	<i>Protention</i>
Hazeltine (Response Competition)	[60]				
	0.97*	1/6	0/2	0.87*** [100]	0.94*** [100]
	0.65***	2/6	0/2	0.86*** [100]	0.78*** [100]
	<i>ns</i>	2/6	0/2	0.93** [100]	0.95** [100]
	0.83***	1/6	1/2	0.83*** [100]	0.86*** [100]
	<i>ns</i>	1/6	1/2	0.88*** [100]	0.87*** [100]
	0.87***	2/6	1/2	0.96** [300]	0.95** [100]
	0.93***	1/6	0/2	0.99* [2000]	0.89*** [100]
	0.76***	3/6	1/2	0.80*** [100]	0.85*** [100]
Ishai Exp. 1 (Object Representation)	[1092]				
	0.88***	3/7	1/2 ; 1/4	0.72*** [10]	0.83*** [10]
	0.67***	5/7	2/2 ; 3/4	0.74*** [10]	0.80*** [10]
	0.59***	0/7	1/2 ; 2/4	0.80*** [10]	0.87*** [10]
	0.52***	0/7	2/2 ; 1/4	0.64*** [10]	0.75*** [10]
Ishai Exp. 2 (Object Representation)	[1092]				
	0.75***	5/7	2/2 ; 0/4	0.76*** [10]	0.86*** [10]
	<i>ns</i>	0/7	1/2 ; 0/4	0.65*** [10]	0.72*** [10]
	0.50***	2/7	2/2 ; 1/4	0.73*** [10]	0.82*** [10]
	0.49***	0/7	2/2 ; 2/4	0.68*** [10]	0.63** [10]
	0.51***	2/7	2/2 ; 1/4	0.79*** [10]	0.92*** [10]
Mechelli (Reading)	[360]				
	0.44***	n.a.	1/2	n.a.	n.a.
	0.76***	n.a.	1/2	n.a.	n.a.
	0.56***	n.a.	1/2	n.a.	n.a.
	0.76***	n.a.	1/2	n.a.	n.a.
	0.57***	n.a.	0/2	n.a.	n.a.
	0.78***	n.a.	1/2	n.a.	n.a.
Postle (Spatial Memory)	[1632]				
	0.58***	8/17	2/4	0.77*** [10]	0.81*** [10]
	0.58***	13/17	1/4	0.75* [10]	0.70*** [10]
	0.45***	3/17	0/4	0.72*** [10]	0.76*** [10]
	0.58***	9/17	2/4	0.65** [10]	0.80*** [10]

of principal component images became the inputs to the neural networks. The time points, mini-volumes of 50 components, are also known as eigen-images. Eigen-images were separated into a training set and a testing (validation) set of the same size, with training and testing pairs alternating in the original sequence. Networks were trained so that each volume would produce its precursor in the test for historical retention, using a standard three-layer backpropagation network. After each training epoch, the network was tested on the novel pairs drawn from the same scan series. This process of training and test continued until performance on the novel test pairs reached its asymptotic best. The residual mean squared error (MSE) measured overall performance at this point. The probe of protention was constructed in the same way, except that the net trained to output subsequent eigen-images in the scan series.

Meanwhile, a companion neural network was created to represent the null hypothesis. This was based on a surrogate dataset that preserved many of the statistical features of the original data. The surrogate data were constructed from the original time series in three steps: First, each voxel time series was decomposed into frequency components using the discrete Fourier transform. Then the phases of the components were re-randomized (“phase shuffling”). Finally, the inverse Fourier transform reassembled each voxel’s surrogate time series. Phase shifting shuffled the peaks and valleys of each voxel time series, thus destroying any systematic relationship based on temporal proximity within the training pairs, but preserved the overall statistical features of the volumes. (The effect is similar to reordering the volumes.) Thus, if any network learning was possible based just on the overall features of the volumes, learned (as it were) by rote association, this should be equally possible in both the original and the phase-shuffled networks and no difference in performance should be observed. This analysis also controls for the confound of the attenuating hemodynamic response. In the original principal component series, the effect of HDR appeared as a large average 1-lag autocorrelation in the subject time series. That is, each component time

series was quite similar to a copy of itself shifted one time point, reflecting the persistence of hemodynamic patterns. The phase-shuffled surrogate time series have the nice feature that autocorrelations within time series are preserved. Thus, the null network also had the same opportunity to exploit autocorrelation as the original. Finally, as an additional precaution, the original voxel time series were detrended using SPM96 prior to the principal components analysis. In short, if the original and the null networks performed equivalently, then temporal information would not be affecting performance, since in the null networks temporal information has been removed. In that case, we could infer that network performance was due to nontemporal information, and the original net would not be detecting any aspect of its inputs that would enable it to extract part of the pattern of its past.

Both the original and its null counterpart were trained and validated the same way, with their best performances (lowest MSE) recorded. Because neural network connection weights are initially random, individual network performance can depend on initial conditions. Repeated trials ameliorate these random effects. Therefore, multiple copies of each were created and tested, and the means of several trials were compared for significance. The same process was used to train and test networks with future volumes as well.

Table 2 shows these results as the ratio of probe network MSE to null network MSE. Since lower errors represent better performance, probe/null ratios less than 1 represent networks that are more successful than their null controls. The table shows that in all 21 subjects analyzed, this interpretive strategy resulted in a significant positive effect. In all subjects, that is, a neural network tuned to the specific image series succeeded in recovering information from images about the immediately prior volume, and a distinct neural network also succeeded in recovering information about the immediately subsequent volume. Controls for each probe suggest that the probe network performance depends on more than simple serial correlation in voxel time series, and on something other than general statistical profiles of the training and test images. This suggests that the

Notes to Table 2:

(A) Measures of relationship between multivariate distance and elapsed time or lag. “Full Experiment” compares all volumes; values are the ratio of the mean of intervolum distances between each volume and all other volumes at <0.5 of maximum lag, compared to the grand mean of all intervolum distances. Lower values indicate multivariate similarity correlated with shorter lags, or a “time–distance effect,” in short. “Within Sequence” compares volumes in the same serial position in task blocks (regardless of task), reporting the number of serial positions with significant time–distance effects, compared to the length of the repeating sequence (in scans). For example, “2/6” means that significant time–distance effects were detected in image volumes from two serial positions in task blocks of six images. “Within Condition” compares volumes with the same task, reporting number of task conditions with significant similarity compared to number of task conditions overall. (With Ishai subjects, two categories of conditions were compared: [passive viewing, match-to-sample] and [houses, faces, chairs, scrambled controls].)

(B) Network analysis. “Retention” reports relative success of neural networks in reconstructing scans prior to current volumes, compared to randomly phase-shuffled control scan series. “Protention” reports success of networks in reconstructing scans following the current time point, compared to controls. “Success” was measured by the MSE in networks, so lower values represent more accurate performance. Accordingly, comparative values of less than one indicate probe networks that perform better than null networks, consistent with phenomenological predictions. The number of trials (distinct networks) examined is given in brackets. See text for details.

* $p < .05$; ** $p < .01$; *** $p < .001$.

brain encompasses a distributed encoding of its own past and future. That past and future brain state information is embedded in present brain states is consistent with the phenomenological claim that retention and protention are superposed in the conscious awareness of the subjective present.

Finally, it is worth noting that probe neural network performance was better in discovering retentional (past) information than protentional (future) information ($p < .01$, paired t test). This asymmetry potentially mirrors a phenomenological asymmetry between the richness of determinate information available about the immediate past, and the comparable scarcity of information available about the immediate, and open, future.

DISCUSSION

The Gradients

The single strongest effect reported here is the pervasive gradient correlated with intervolumetric lag. It is hypothesized that this gradient is at least a partial reflection of an underlying monotonic temporal flow in consciousness. However, it is also possible that it is some form of artifact. Indeed, long baseline “drift,” “1/f noise,” “global signal,” and other longer-term effects have been often noted, usually amidst puzzlement and conjecture (Lowe & Russell, 1999; Turner, Howseman, Rees, Josephs, & Friston, 1998; Lange, 1996; Bandettini, 1993). Note that, in general, the results derive some initial validation from their sources. All the analyses here were based on preprocessed data supplied to the fMRIDC by the original authors. The same data serve to support the original studies of these authors. Had the data been problematic, its problems should have also affected the published results, or should have led to the exclusion of subjects. The common database to both these and the published studies thus offers an initial validation for these results.

However, notwithstanding this innocence by association, the divergent methods employed here might be open to several artifacts that traditional analyses either correct or escape (Lowe & Russell, 1999; Turner et al. 1998; Aguirre et al., 1997; Aguirre, Zarahn, & D’Esposito, 1998; Lange, 1996; Bandettini, 1993). While unknown artifacts cannot be precluded, a few prominent possibilities can be discounted (See also Zarahn, Aguirre, & D’Esposito, 1997.):

Motion Artifact

A motion artifact could have the effect of producing an increase in multivariate distance. However, the expected effect of this artifact would be sporadic. It might occur at different times, and should be flanked by moments of relative immobility. The effect on intervolumetric distance, accordingly, would be to produce a stepwise change.

Moreover, the motion artifact effect would not be expected to be monotonic, as subjects might with equal likelihood move back toward an earlier position. The observed gradient shows none of these effects. It is generally smooth, monotonic, and present in the majority of subjects. In addition, all subject scans were preprocessed to remove potential motion artifacts.

Long Baseline Physiological Cycles

Respiration and heart rate both create physiological artifacts in scan series, so a possible artifact here is another (undetermined) physiological cycle (Lange, 1996). The most telling counter-consideration is, once again, the observation that the basic temporal gradient is monotonic. The duration of the longest study is at least 54 min. (See also Zarahn, Aguirre, & D’Esposito, 1997.)

“Scanner Drift”

Another source of artifact is the scanner itself. Although explicit efforts to detect and eliminate these artifacts are brought to bear in scanner design and data processing, perhaps some elusive component is emerging in the observed gradient. Here it is helpful that this analysis depends on the recordings at four different research sites. Two scanner manufacturers supplied the magnets. A single type of malfunction, even a subtle one, seems unlikely to have affected all four machines. In addition, the diversity of scan series lengths, of voxel size, and of series structures (discrete runs vs. continuous scan sessions) all render this artifact somewhat more unlikely.

These three artifacts would most plausibly result in global changes in voxel magnitudes that would be detectable as global trends or “linear drift” in mean activation intensity (Bandettini, 1993). However, the grand mean of 21 subjects reveals a mean slope of 0.02 ($SD = 0.17$; again, Mechelli et al. (2000) datasets were excluded for computational reasons). The relatively large standard deviation derives mainly from the short (60 scan) series of Hazeltine et al. (2000). Mean slope for all subjects in Ishai et al. (2000) and Postle et al. (2000) is -0.03 , standard deviation is 0.03. From this indirect evidence then, it seems that a global artifact is an unlikely cause of the observed effects. There may also, however, be artifacts with more subtle manifestations:

Neural Processes Unrelated to Consciousness

A philosophical skeptic may claim that any alleged neural expression of consciousness is in fact an unconscious process that is related to, but distinct from, conscious processing. This philosophical skepticism is always available, and always refuted in practice by the general success of science, going about its business unperturbed by skeptical worries. Since this article rests

in part of phenomenology, philosophical skepticism is perhaps a larger worry, but if so, it is a worry for any attempt at a science of consciousness. In any case, with respect to global philosophical skepticism, the burden of proof in the scientific arena is always with the skeptic.

A more plausible skepticism arises within science itself. In this case, the skeptical possibility is that the temporal gradient tracks a neural process, but one unrelated to temporal awareness. The weakest possibility is that this unknown neural effect has no behavioral, cognitive, or phenomenal manifestation whatsoever. This seems implausible. The multivariate effect described here is very large, especially bearing in mind the inherent noisiness of the fMRI signal. That it would be completely without cognitive or phenomenal effect would be surprising. A more plausible alternative is that the gradient reflects another cognitive process, and that the true temporal gradient (if any) is not responsible for the observed effects. One candidate is fatigue. Though time dependent, fatigue is not itself the temporal awareness described above. It might, in some sense, accumulate through a scan. It is surprising, however, to see the effect in the short scan series (Hazeltine et al., 2000), lasting 120 sec. Moreover, the effect seems indifferent to factors that would seemingly affect fatigue, like age of subject and time of day.

Tripartite Temporality

The interpretation of the time–distance correlation as an expression of conscious phenomena is also indirectly supported by the general success of the neural network probes of the subject time series. The network probes, if valid, suggested that encodings of protention and retention were part of the information expressed in voxel activation patterns. If so, this would suggest an underlying temporal structure in the conscious brain that is in continuous flux. This flux in turn could be detected in the gradients.

As discussed above, the network analysis was contrastive, and the baseline case was constructed to replicate potential confounds in the original time series. Global “drift” artifacts should therefore have been controlled, as discussed above. Detrending prior to the network analysis reinforced this control.

The network probe is nonetheless limited by several constraints and assumptions. One of these is the prior processing using principle component analysis. Network outputs are accordingly compared to component patterns, not the original voxel patterns. Open questions of appropriate pre- and postprocessing remain.

These issues are ameliorated by the conservative conclusion drawn. The networks based on the original data do better on the task of reconstructing past and future images than their matched controls. Therefore, there is some information specific to the original data in its original temporal order that is lost when that order is

phase-randomized, and the networks are exploiting that information. The average of MSE for the multiple trials of the network probe reveals that the nets are still only approximating their targets (and in some cases, just barely). The reasons for their marginal performance could be many; certainly, noise in the original series might carry over into the component analysis (and even dominate it), leading to unpredictability of precursors and successors. There is also evidence in the shortest scan series (Hazeltine et al., 2000), consisting of 60 time points, of “overfitting.” The network in this case is training on just 15 input–output pairs, with 15 others reserved for testing. Differences in this sample space were only detectable over many trials: In one of Hazeltine et al.’s (2000) subjects, only after around 2,000 trials did the effect reach significance. The two longer experiments (Ishai et al., 2000; Postle et al., 2000) each afforded several hundred training and testing pairs. Generalization and success with novel tests improved, and robust contrasts with the scrambled controls were quickly established.

Further Questions

Regarding both the temporal gradients and the network probes of tripartite temporal structure, the goal of this article has been to illustrate general effects in accordance with neurophenomenological predictions. The important conclusion is that in each of five diverse experiments, and in 27 different subjects, these effects occur. Subsequent study will turn to the host of specific questions that remain unexplored. In addition to the general questions of validation of the methods, it would be useful to examine the different tasks and many others with respect to the temporal qualities of specific tasks. No less important is consideration of the exceptions, the subjects in which particular effects did not reach significance. Naturally, these subjects were conscious, too, so ultimately explanations for their failing to display the temporal effects will be needed.

Another topic, too large to consider here, is the source of these effects. Do particular brain regions contribute more or less? A localized “temporalizer” is not excluded, although on phenomenological grounds one might look for distributed temporality first. Temporality is continuous and inextricably present in all conscious cognition, quite unlike the isolable and contingent occurrences of specific cognitive capacities.

CONCLUSIONS

Daniel Dennett (1991) has raised what may be the fundamental question for the science of consciousness: Why should any of the information processing in the brain be “conscious”? As we examine each of the many cognitive capacities humans enjoy, Dennett’s question recurs: Considered one by one, why would any cognitive

function necessarily be accompanied by awareness? Traditional fMRI research only sharpens the question by localizing cognitive capacities in small regions of the brain, each of which by itself seems unlikely to be a seat of consciousness.

Phenomenology, I've suggested, offers a new set of questions. It asks, first, what are the features of consciousness at the most general level of description? Then the "why" question can be answered, again at the most general level: The phenomenal structures of consciousness organize information into the composite structure we experience as reality itself. As such, consciousness (i.e., reality as it is experienced) is implicated in all cognitive functions. The global perspectives of phenomenology in turn suggest several global multivariate probes of the conscious brain. The potential for new methods, I hope, has been illustrated in this article. More questions are raised than answered, but at this early moment in the science of consciousness, new questions may be timely.

Acknowledgments

The author thanks, first, the authors of the original studies: Eliot Hazeltine, Alomit Ishai, Andrea Mechelli, Bradley Postle, and their many collaborators: Jeffrey Berger, Mark D'Esposito, Karl Friston, John Gabrieli, James Haxby, Alex Martin, Russell Poldrack, Cathy Price, Alexander Taich, and Leslie Ungerleider. Equally, thanks are due to the fMRIDC and its staff. Vincent Clark has been a generous and wise guide in all of the author's initial explorations. Elizabeth Chua and Gregory Rubin assisted in data preparation. The Trinity College Neuroscience Program assisted with material and intellectual support, and the Trinity College Faculty Research Committee provided essential research leave time. An anonymous reviewer for the "New Perspectives in fMRI Research" award suggested several useful clarifications and a better method for surrogate data construction for the neural network probes. To all, the author extends heartfelt thanks.

Reprint requests should be sent to Dan Lloyd, Department of Philosophy and Program in Neuroscience, Trinity College, Hartford, CT 06106, USA, or via e-mail: dan.lloyd@trincoll.edu.

Data were postprocessed in MatLab version 5, with Neural Networks and Statistics Toolboxes, running on a Macintosh G4 laptop computer. All MatLab functions were written by the author, except: SPM_detrnd, from SPM96 (Friston et al., 1995); PCA functions from the PLS Toolbox (Copyright Eigenvector Research, 1991–1998); and Read_analyze from the fMRI Toolbox (© 2000, Colin J. Humphries).

REFERENCES

National fMRI Data Center Datasets

Hazeltine, E., Poldrack, R., & Gabrieli, J. D. (2000). Accession #: 2-2000-11173. Neural activation during response competition. *Journal of Cognitive Neuroscience*, 12, 118–129.

Ishai, A., Ungerleider, L. G., Martin, A., & Haxby, J. V. (2000). Accession #: 2-2000-1113D. The representation of objects in the human occipital and temporal cortex. *Journal of Cognitive Neuroscience*, 12, 35–51.

Mechelli, A., Friston, K. J., & Price, C. J. (2000). Accession #: 2-2000-1189. The effects of presentation rate during word and pseudoword reading: A comparison of PET and fMRI. *Journal of Cognitive Neuroscience*, 12, 145–156.

Postle, B. R., Berger, J. S., Taich, A. M. & D'Esposito, M. (2000). Accession #: 2-2000-1112R. Activity in human frontal cortex associated with spatial working memory and saccadic behavior. *Journal of Cognitive Neuroscience*, 12, 2–14.

Published References

Aguirre, G., Zarahn, E., & D'Esposito, M. (1997). Empirical analyses of BOLD fMRI statistics. II. Spatially smoothed data collected under null-hypothesis and experimental conditions. *Neuroimage*, 5, 199–212.

Aguirre, G., Zarahn, E., & D'Esposito, M. (1998). The variability of human, BOLD hemodynamic responses. *Neuroimage*, 8, 360–369.

Bandettini, P. (1993). Processing strategies for time-course data sets in functional MRI of the human brain. *Magnetic Resonance in Medicine*, 30, 161–173.

Bandettini, P., & Cox, R. (2000). Event related fMRI contrast when using constant interstimulus interval: Theory and experiment. *Magnetic Resonance in Medicine*, 43, 540–548.

Dennett, D. (1991). *Consciousness explained*. Boston: Little Brown.

Friston, K. J., Holmes, A. P., Poline, J. B., Grasby, P. J., Williams, S. C., Frackowiak, R. S., & Turner, R. (1995). Analysis of fMRI time-series revisited. *Neuroimage*, 2, 45–53.

Gazzaniga, M. (Ed.) (2000). *The new cognitive neurosciences*. Cambridge: MIT Press.

Hazeltine, E., Poldrack, R., & Gabrieli, J. D. (2000). Neural activation during response competition. *Journal of Cognitive Neuroscience*, 12, 118–129.

Husserl, E. (1893–1917). Zur Phänomenologie des inneren Zeitbewusstseins [Phenomenology of inner time consciousness]. In R. Boehm (Ed.), *Husserliana* (vol. 10). The Hague: Martinus Nijhoff, 1966.

Husserl, E. (1907). Ding und raum [Thing and space]. In U. Claesges (Ed.), *Husserliana* (vol. 16). The Hague: Martinus Nijhoff, 1974.

Ishai, A., Ungerleider, L. G., Martin, A., Schouten, J., & Haxby, J. V. (1999). Distributed representation of objects in the human ventral visual pathway. *Proceedings of the National Academy of Sciences, U.S.A.*, 96, 9379–9384.

Ishai, A., Ungerleider, L. G., Martin, A., & Haxby, J. V. (2000). The representation of objects in the human occipital and temporal cortex. *Journal of Cognitive Neuroscience*, 12, 35–51.

Lange, N. (1996). Statistical approaches to human brain mapping by function magnetic resonance imaging. *Statistics in Medicine*, 15, 389–428.

Lloyd, D. (1995). Consciousness: A connectionist manifesto. *Minds and Machines*, 5, 161–185.

Lloyd, D. (1996). Consciousness, connectionism, and cognitive neuroscience: A meeting of the minds. *Philosophical Psychology*, 9, 61–81.

Lloyd, D. (2000). Beyond "The Fringe": A cautionary critique of William James. *Consciousness and Cognition*, 9, 629–637.

Lloyd, D. (in press). *Radiant cool: A novel theory of consciousness*. Cambridge: MIT Press.

Lowe, M., & Russell, D. (1999). Treatment of baseline drifts in fMRI time series analysis. *Journal of Computer Assisted Tomography*, 23, 463–473.

Mechelli, A., Friston, K. J., & Price, C. J. (2000). The effects

- of presentation rate during word and pseudoword reading: A comparison of PET and fMRI. *Journal of Cognitive Neuroscience*, *12*, 145–156.
- Metzinger, T. (Ed.) (2000). *Neural correlates of consciousness: Empirical and conceptual questions*. Cambridge: MIT Press.
- Postle, B. R., Berger, J. S., Taich, A. M., & D'Esposito, M. (2000). Activity in human frontal cortex associated with spatial working memory and saccadic behavior. *Journal of Cognitive Neuroscience*, *12*, 2–14.
- Turner, R., Howseman, A., Rees, G., Josephs, O., & Friston, K. (1998). Functional magnetic resonance imaging of the human brain: Data acquisition and analysis. *Experimental Brain Research*, *123*, 5–12.
- Varela, F. (1996). Neurophenomenology: A methodological remedy for the hard problem. *Journal of Consciousness Studies*, *3*, 330–349.
- Zarahn, E., Aguirre, G., & D'Esposito, M. (1997). Empirical analyses of BOLD fMRI statistics. I. Spatially unsmoothed data collected under null-hypothesis conditions. *Neuroimage*, *5*, 179–197.