

The Neural-Cognitive Basis of the Jamesian Stream of Thought

Russell Epstein

Medical Research Council Cognition and Brain Sciences Unit, 15 Chaucer Rd., Cambridge CB2 2EF, United Kingdom

E-mail: russell.epstein@mrc-cbu.cam.ac.uk.

William James described the stream of thought as having two components: (1) a *nucleus* of highly conscious, often perceptual material; and (2) a *fringe* of dimly felt contextual information that controls the entry of information into the nucleus and guides the progression of internally directed thought. Here I examine the neural and cognitive correlates of this phenomenology. A survey of the cognitive neuroscience literature suggests that the nucleus corresponds to a dynamic global buffer formed by interactions between different regions of the brain, while the fringe corresponds to a set of mechanisms in the frontal and medial temporal lobes that control the contents of this global buffer. A consequence of this account is that there might be conscious imagistic representations that are not part of the nucleus. I argue that phenomenology can be linked to psychology and neuroscience and a meaningful way that illuminates both. © 2000 Academic Press

To date, the most prominent phenomenological description of the structure of human consciousness is that put forth by William James in the famous stream of thought chapter of the *Principles of Psychology*. In this chapter, James made two observations about the phenomenology of consciousness that are critical to the current article. First, he observed that consciousness is dynamic: It continually moves from idea to idea, thought to thought. Second, he observed that consciousness has a composite structure: Accompanying every explicit thought is a vague *fringe* or *penumbra* of sensed context. James linked these two observations together by proposing that this “fringe” plays an important role in controlling the orderly advance of consciousness from one substantive thought to another. In a famous metaphor, James compared the stream of thought to the life of a bird:

As we take, in fact, a general view of the wonderful stream of our consciousness, what strikes us first is this different pace of its parts. Like a bird’s life, it seems to be made of an alternation of flights and perchings. The rhythm of language expresses this, where every thought is expressed in a sentence, and every sentence closed by a period. The resting-places are usually occupied by sensorial imaginations of some sort, whose peculiarity is that they can be held before the mind for an indefinite time, and contemplated without changing; the places of flight are filled with thoughts of relations, static or dynamic, that for the most part obtain between the matters contemplated in the periods of comparative rest.

Let us call the resting-places the ‘substantive parts’, and the places of flight the ‘transitive parts’, of the stream of thought. (James, 1890, p. 243)

One of James’s key insights was that the sometimes elusive feelings that fill the “transitive parts” of the stream of thought are just as experientially real as the substantive thoughts, and he emphasized the critical functional role they played in guiding the transitions between the “perches.”

After a long period of neglect, James's ideas about the structure of consciousness have been revived by a small number of contemporary authors, including Mangan (1991, 1993a, 1993b), Galin (1994), and Chafe (1994). The goal of this article is to expand this previous work on the fringe by linking James's description of the phenomenology of the stream of thought to recent results in cognitive psychology and cognitive neuroscience. After a brief summary of this phenomenology (section I), I survey the cognitive neuroscience literature and propose separate neural bases for the substantive thoughts (or nucleus) and transitive thoughts (or fringe)¹ (sections II and III). I then briefly discuss how this neural-cognitive framework can accommodate a realm of phenomenology not discussed by James: that of "background" mental representations (section IV).

This project differs from certain other attempts to explicate the "neural correlates of consciousness" (Crick & Koch, 1998) in that I do not attempt to explain what consciousness is or how it arises from neural tissue. I am not interested in the metaphysics or ontology of consciousness, but in its structure. One of my working assumptions is that the level of structure at which James described the stream of thought corresponds roughly to the level of organization at which cognitive neuroscience is attempting to explain the mind/brain. Consequently, the structure of consciousness can be mapped onto the modular structure of mind revealed by cognitive neuroscience. My project is to find correspondences between these two ways of describing mental life.

Given the dynamic state of cognitive neuroscience and our ever-expanding corpus of knowledge about the brain, my conclusions must necessarily be tentative. However, even a preliminary outline can be a fruitful guide for future research—not only in consciousness studies, but in cognitive neuroscience as well. Just as neuroscience can illuminate phenomenology, so can phenomenology illuminate neuroscience by providing an account of the matter that needs to be explained. Ultimately, phenomenology and experiment should converge into a coherent description of mental life that includes both subjective and objective aspects.

I. JAMES' PHENOMENOLOGY

I begin with the brief summary of James' phenomenology. (More detailed expositions can be found in Mangan, 1991, 1993a, 1993b and Galin, 1994). I first consider the "substantive parts" of the stream of thought and then the "transitive parts," or fringe. In the spirit of Mangan (1991, 1993a, 1993b), my goal is to link phenomenology to function whenever possible.

¹ James used the terms *nucleus* (or *image*) and *fringe* when he wished to emphasize the different varieties of experience present in a particular moment of awareness (i.e., "sensorial imaginings" vs "thoughts of relations") and the terms *substantive thoughts* and *transitive thoughts* when he wished to emphasize the dynamic flow of information through consciousness (see Mangan, 1993a; Galin, 1994 for discussion). Although I tend to conflate the two terminologies here, they are not entirely equivalent. Strictly speaking, each substantive thought includes *both* a nucleus of highly conscious, often imagistic features and a fringe that conveys the conceptual significance of those features. During the substantive thought, the nucleus is the most prominent representation in awareness, while the fringe becomes more apparent during the transitions from one substantive thought to another, when no nucleus is present.

The Nucleus

James spent relatively little time discussing substantive thoughts, perhaps because he assumed that their nature would be obvious to his readers. How can they be characterized? As James states, one of their primary features is their relative *stability*: They can be “held before the mind for an indefinite time” (James, p. 243) and contemplated without changing. James used the metaphor of a kaleidoscope to illustrate the relationship between the always-changing stream of thought and the relatively stable substantive thoughts that are its most salient components. Although a kaleidoscope is always rotating, it presents us with a succession of patterns that are themselves momentarily stable (James, p. 246). Similarly, although consciousness is always changing, it presents us with a series of substantive thoughts that are themselves momentarily stable.

Another key quality of substantive thoughts is their *memorability*. Just as one is more likely to remember the stable patterns of the kaleidoscope than the precise way in which the kaleidoscope changes, so one is more likely to remember the substantive than the transitive parts of the stream of thought. James likened substantive thoughts to conclusions that summarize the line of thought that brought them about. Afterward, one does not remember the thought process, all one remembers is the conclusion: “For the important thing about the train of thought is its conclusion.... That is what abides when all its other members have faded from memory” (James, p. 260). We do not usually remember the flight of the bird, only the perches.

A third characteristic of substantive thoughts is *multimodality*: They can potentially involve many different representational modalities (cf. Jackendoff, 1987). James described the substantive “conclusions” of the stream of thought as usually “a word or phrase or particular image, or practical attitude or resolve.” In other words, substantive thoughts can include verbal, visual, or even attitudinal representations. Despite this catholicity, James recognized that substantive thoughts often include a strong imagistic component (in the form of visual images or inner speech). He describes the resting places of thought as “usually occupied by sensorial imaginings of some sort” (p. 243)—the term “usually” implying that most substantive thoughts will be connected with sensory representations, but some will not. Later in the same chapter, James suggests that the substantive parts of thought can be sensations, images, percepts, concepts, or thoughts (p. 247), not all of which are imagistic.² Expanding on this, Galin (1994) has argued that the nucleus contains the few specific features that maximally discriminate the currently conscious representation (i.e., object or idea) from other currently nonconscious representations and that these features are determined by whatever aspect of the representation we are most interested in at the moment. In this formulation, the features can be either perceptual or nonpercep-

² Further evidence that the James did not mean to equate substantive thoughts with the imagistic representations associated with them comes from his discussion of our memory for them. Although James believed that “conclusions” of trains of thought are quite memorable, he wrote that it is not the exact wording of the sentence that is our conclusion that matters; rather, it is the meaning and significance of the sentence. A substantive thought is the entirety of the situation in the brain at the time—not just the image, but the understanding of the image as well.

tual depending on our current interests: For example, when we look at a car, we can either focus on sensory features such as its color (if we are thinking about getting a new paint job) or on nonsensory features such as its cost (if we are thinking about buying) (Galín, 1994, personal communication).

To these three characteristics emphasized by James, we can add two observations made by Bruce Mangan (1991, 1993a). The first is that substantive thoughts appear *sequentially*: we can only have one at a time. This was not explicitly claimed by James, but it is implicit in his metaphors: The bird may only perch on one branch at a time, and the kaleidoscope can only show one pattern at a time. The second is that substantive thoughts are of *limited capacity*: only a small proportion of the total amount of information processed by the brain can be present in any one substantive thought. This last quality does not follow directly from James, but is a consequence of our modern understanding of information processing systems.

To summarize, we can think of substantive thoughts as representations that have *stability* and *memorability*, can be potentially *multimodal*, appear *sequentially*, and have a *limited capacity* to convey information. Note that all of these characteristics are simultaneously phenomenological and functional.

The Fringe

James strongly believed that thought was temporally continuous; thus his metaphor of a “stream” of thought. The formulation that consciousness proceeds from one definite thought to another was a potential problem for this view, since it suggested that experience consisted of temporally discrete elements—a “chain” rather than a “stream.” James resolved this apparent contradiction by arguing that we are aware not only of the substantive thoughts, but also of transitive thoughts, which provide a sense of context that bridges the temporal gaps between substantive thoughts. This sense of context has three components: a faint memory of preceding thoughts, an understanding of the relationship between the current thought and others that might be relevant (“feelings of relation”), and a feeling for where our thoughts are heading (“feelings of tendency”). This somewhat heterogeneous set of feelings—which I am here grouping all under the rubric of “fringe”—acts as a sort of glue that holds the stream of thought together in time:

Every definite image in the mind is steeped and dyed in the free water that flows round it. With it goes the sense of its relations, near and remote, the dying echo of whence it came to us, the dawning sense of whither it is to lead. (p. 255)

What exactly is the fringe? James is not always clear on this point (although at one point he speculated that it involves a sort of direct awareness of the “waxing and waning excitements of the brain”; p. 250.) The fringe is perhaps best approached by examining a set of specific examples that James outlined in which we become aware of it. I refer to these as “fringe experiences.” They include (1) the feeling of expectation we have when our attention has been drawn to something and we have some sense of what it might be, but we have not yet determined exactly what it is; (2) the feeling of having a word “on the tip of one’s tongue”; (3) the sense we have when we know that something is familiar; (4) the sense of connection provided by

words that connote the logical structure of a thought, such as “but” or “nevertheless”; (5) the feeling of intending to say something, when we have a perspective of a scheme of thought that is not yet articulate; (6) the sense we have for the overall scheme or form of a book, a work of art, or a scientific system; (7) the sense we have that there is more to a thought, even when we do not have the words to complete it; (8) the sense of anticipation that occasionally causes us to mistakenly use an anticipated word instead of the correct one; and (9) the sense of being “on the right track” to a conclusion. Obviously, this is a rather heterogeneous set of experiences. However, one commonality is that many of them involve feelings that guide consciousness from one substantive thought to another.

Let us examine the tip-of-the-tongue (TOT) state, which is often taken to be a prototypical fringe experience (Mangan 1991, 1993a; Baars, 1993; but see Galin, 1994, who considers it to be a compound). This state occurs when we try to remember a word or name and know that we know the answer but cannot produce it—instead, we feel that the answer is “on the tip of our tongue.” Despite this retrieval failure, we often have a sense for the sound of the word, or the rhythm of its syllables, and we know its meaning (Brown & McNeill, 1966; Brown, 1991). If it is a name we seek, we may be able to picture the person’s face and recall many facts about him or her. Often, we get “stuck” on a word or name that is similar, though we know it to be incorrect—we know it does not, as James says, “fit the mould.” In TOT states, it is as if we know everything there is to know about a thing without having the word that signifies it. Instead, we have a very vivid sense that we can almost grasp it, but not quite. This sense of grasping is telling us something important about the state of our stream of thought: it informs us that we were *meant* to go somewhere (i.e., the answer word), but we have not arrived.

In addition to conveying information about the state of the retrieval process during the TOT state, the fringe also gives us information about the content of the memory structures that have been activated. For example, suppose we are trying to remember an actor’s name. We may be able to recall the plot of a movie he or she was in, as well as the name of other actors in the movie. In other words, we can activate the part of our memory network associated with the actor. This, when combined with the goal of wanting to retrieve the name, creates a context that normally would lead us immediately to the answer—at which point, the contextual information is no longer needed. The TOT state provides an unusual situation in which the expected word does not arrive and we are left with an unusually vivid sense of guiding context, which in part consists of a sense for memories that are associated with the target.

Such contexts play a role in directing thought even when we are not as painfully aware of them as we are in the TOT state. James writes:

In all our voluntary thinking there is some topic or subject about which all the members of the thought revolve. Half the time this topic is a problem, a gap we cannot yet fill with a definite picture, word, or phrase, but which . . . influences us in an intensely active and determinate psychic way. Whatever may be the images and phrases that pass before us, we feel their relation to this aching gap. To fill it up is our thought’s destiny. Some bring us nearer to that consummation. Some the gap negates as quite irrelevant. Each swings in a felt fringe of relations of which the aforesaid gap is the term . . . Relation, then, to our topic or interest is constantly felt in the fringe, and particular the relation of harmony and discord, of furtherance or hindrance of the topic. (James, p. 259)

In other words, during directed thought, we know the topic and have a sense for the conclusion we are looking for. Knowing the topic may be equated with activation of a certain subset of our associative memory network. Thought progresses by creating different instantiations of the currently active nodes of the memory network in the nucleus (as “images and phrases”). The relationship between these substantive instantiations and the goal is felt in the fringe as a feeling of “on-the-right-trackness,” and this feeling determines which of the various associated ideas and memories get activated next (Mangan, 1991, 1993). Eventually, one arrives at a representation in the nucleus which is satisfactory, and the train of thought has reached its destiny in the form of “a word or phrase or particular image, or practical attitude or resolve” (James, p. 260). This process can also occur in the absence of a specific goal (as when we daydream) in which case the direction of thought is entirely determined by the associative network. Thus, James’s description of the process by which thought advances assumes the existence of (1) an associative memory network and (2) mechanisms that monitor and control the activation of this network.

James does not make it entirely clear whether he believed that we can perceive the structure of the network directly or whether we can only sense it indirectly (as when we feel a sense of harmony when a thought fills the “aching gap”). When James discusses our sense for the form of a work of art or the conceptual structure of a scientific theory and writes that “Great thinkers have vast premonitory glimpses of schemes of relation between terms” (James, p. 255), he makes it sound as if it were possible (at least, for Mozart) to glimpse the network directly. More often, though, he emphasizes the role of evaluative feelings—such as feelings of familiarity—in conveying information about the relationship between the currently active substantive thought and inactive but related thoughts in the network. More recent writers such as Mangan (1991, 1993a) and Galin (1994) have emphasized these evaluative feelings in their descriptions of the fringe. For example, Mangan emphasizes the importance of “cognitive emotions” such as the feeling of on-the-right-trackness in directing the stream of thought, going so far as to suggest that our sense of beauty is a cognitive emotion that reflects an evaluation of the relationship between the different parts of a work of art. Galin suggests that the fringe experiences discussed by James involve only a subset of a wider set of evaluative feelings, including feelings of agency, intention, and emotion. I will not attempt to determine whether it is possible to experience the network directly. However, I do believe that the distinction between the network itself and feelings that evaluate the network is an important one, and I propose different anatomical bases for these two functions in section III.

To summarize, the fringe can be thought of as relating to two distinct functional information processing components. First, there must be an associative memory network that determines relationships between thoughts that could potentially get instantiated in the nucleus. Second, there must be a mechanism that can monitor the state of activation of the network and guide it according to current behavioral goals.³

³ Note that here and in the subsequent discussion, I make the vastly simplifying assumption that thoughts that are not driven directly by incoming perceptual information must be driven by memory associations—both episodic (memory for how events are related to each other in time) and semantic (memory for how things are conceptually related to each other). Of course, it is also possible that one may have a thought that is *creative*, i.e., is not a reinstantiation of a previously encoded representation.

II. NEURAL-COGNITIVE BASIS OF THE NUCLEUS

In the previous section, I listed five distinct characteristics of the substantive parts of thought. In this section, I argue that these thoughts might be formed by dynamic interactions between different parts of the brain. First, I review evidence from the perception literature that suggests that awareness of a percept requires information in different brain regions to be bound together. Then I discuss one possible mechanism for such binding: synchronous firing at a 40-Hz frequency. Unlike other authors, I do not claim that such synchronous firing is the neural correlate of consciousness itself. Rather, I claim that such firing might be the neural level correlate of the phenomenological level nucleus, which is only one aspect of consciousness.

A number of visual perception experiments provide evidence for a distinction between *processing* of a stimulus and *awareness* of a stimulus. Even when many different visual stimuli are presented in very rapid sequence, they can be identified and their meaning extracted (Thorpe et al., 1996, Potter 1975, 1993). However, entry of this information into awareness appears to be a much slower process that can easily be overloaded (Potter, 1976). This serial, capacity-limited process can be dissociated from the more parallel processes that subserves recognition (Marcel, 1983). A number of behavioral paradigms have explored this dissociation (see Potter, 1999 for review). Two of the most prominent are *repetition blindness* and the *attentional blink*. As these phenomena give some insight into how representations reach awareness, I review them briefly here.

In repetition blindness, observers fail to see the second occurrence of a repeated item in a rapidly presented sequence of items. For example, if the sentence “When she spilled the milk there was milk all over” is viewed one word at a time at a rate of eight words per second, observers will often report seeing “When she spilled the milk there was all over,” even though this is ungrammatical (Kanwisher, 1987). Clearly, the second occurrence of the word milk is processed—the brain treats this repeated word differently from all the unpeated words in the sentence—but it is not experienced. Kanwisher (1987, 1991) describes repetition blindness as a failure to bind information about the identity of the stimulus (its visual “type”) with information about where and when it occurred (its episodic “token”).⁴ In other words, the observer processes the stimulus (unconsciously) but does not experience it as a distinct event occurring at a particular place and time.

The attentional blink is another perceptual phenomenon that can be described in terms of a failure to bind type and token (Chun, 1997). In this phenomenon, observers fail to see the second of two targets when both are presented within a rapidly pre-

Even if one argues that most “creative” thoughts are really just novel combinations of elements that have been previously established, it is clear that the existence of creative thoughts indicates that the description of thought offered here—in which one simply recalls previously encoded material that is associated to the current substantive thought—is only a simplified version of the reality.

⁴ Note that this use of the terms “type” and “token” is not equivalent to the logician’s use of the terms, in which “token” refers to a concept that is a specific example of a given “type.” Under the logician’s usage, “Rover” is a token of the type “dog.” In contrast, in the present usage, both “Rover” and “dog” are types—but the spatiotemporal information that distinguishes a specific instance of seeing Rover at a specific time and place is represented by a token.

sented sequence of irrelevant distractor items and the second target appears 200–500 ms after the first (Raymond et al., 1992). For example, when observers are told to report the identities of two letters appearing among a rapidly presented sequence of digits, they will often fail to experience (and hence, report) the second letter (Chun & Potter, 1995). As with repetition blindness, there is evidence that the “blinked” item is processed to the level of identity. For example, a blinked word can semantically prime a subsequent word (Shapiro et al., 1997), and a blinked word that is inconsistent with a previously established semantic context leads to an ERP signal that is characteristic of semantic mismatch (Luck et al., 1996).

Several lines of evidence suggests that the “types” and “tokens” that get disconnected in RB and AB are processed in different parts of the brain, with the occipital-temporal region being particularly important for type processing and the parietal lobes being particularly important for token processing. This division corresponds roughly to the well-known division of the visual system between ventral “what” and dorsal “where” pathways (Ungerleider & Mishkin, 1982). Patients with damage to the parietal lobe often exhibit deficits reminiscent of RB and AB. For example, neglect patients will exhibit a preference for one picture over another based on emotionally salient details that appear in the contralesional (i.e., neglected) hemifield even though they do not explicitly report these details (Marshall & Halligan, 1988). Similarly, extinction patients (who can report a stimulus in the contralesional field if it appears alone but are impaired when a competing stimulus appears in the ipsilesional field) are much worse at reporting contralesional stimuli when they are similar or identical to the competing stimulus in the ipsilesional field (Baylis, Driver, & Rafal, 1993). This effect is similar to repetition blindness insofar as it involves an inability to establish individual tokens for repeated events of the same type. Patients with parietal damage have also been reported to exhibit a prolonged and abnormally severe attentional blink (Husain et al., 1997). In contrast, patients with occipital-temporal damage frequently exhibit *agnosia*: an inability to identify stimuli that can easily be detected (Farah, 1990).

Thus, conscious perception of a visual stimulus requires information from the ventral and dorsal visual streams to be bound together. We can extend this to a general principle: Awareness of a stimulus requires the binding together of information from several cortical regions. This principle is consistent with of a number of theories of attention, including feature integration theory (Treisman & Gelade, 1980) and integrated competition theory (Desimone & Duncan, 1995), both of which assume that attention is involved in the integration of information from many brain regions into a coherent, attended percept. Recently, several writers have suggested that this binding is the mechanism of consciousness itself. For example, Tononi and Edelman (1998) propose that the neural process underlying consciousness is a “dynamic core” or cluster of neurons in the thalamocortical system that fire synchronously together over a period of several hundred milliseconds. This “core” is roughly equivalent to Baars’ global workspace (Baars, 1988). However, while Baars’ global workspace is described as a fixed structure that can contain different representations, the “core” is a dynamic and temporary set of interactions between various representations.

A number of results from animal physiology support the hypothesis that this binding is achieved by temporal synchrony in the firing patterns of neurons in different

brain regions. Neurons that are widely separated in the cortex will fire in synchrony when both respond to the same object (Gray et al., 1989; Roelfsema et al., 1997; Singer et al., 1997), but not when they respond to different objects (Engel et al., 1991a). If the neurons are in different hemispheres, the synchronous firing can be interrupted by cutting the corpus callosum, indicating that the synchrony is achieved by cortico-cortical connections rather than by common subcortical inputs (Engel et al., 1991b). In humans, the mass effects of this synchronous firing is believed to be observable as a 40-Hz oscillation in the EEG signal. Activity in this frequency band has been associated with awareness and attention: For example, it is found when subjects are awake or dreaming, but not when they are in non-REM sleep, when they are presumably not conscious (Llinas & Ribary, 1993). Based on these and other results, some authors have proposed that oscillatory activity at this frequency is the essential neural correlate of consciousness (Engel et al., 1999; but see Zeki & Bartels, 1999).

A recent study by Rodriguez et al. (1999) is particularly relevant to this claim. In this study, subjects viewed upright and inverted Mooney faces and reported whether they saw a face by pressing a button. Mooney faces are two-tone images that can be easily interpreted as a face when upright but appear as nonsense patterns when inverted. The EEG signal when subjects perceived a face exhibited a tripartite pattern: a period of synchronization 200 ms after stimulus onset, followed by a period of desynchronization and then by a second period of synchronization that roughly corresponded with subject reaction time. In contrast, the EEG signal when subjects did not perceive a face only showed the later period of synchronization. Rodriguez et al. conclude that the two periods of synchronization in the perception case correspond to perception of the face and preparation of a motor response, while the period of desynchronization corresponds to the transition between these two distinct cognitive acts.

I propose that the temporary global synchronies observed in the animal and human experiments are not the neural correlates of consciousness itself, but of James's substantive thoughts. For example, in the experiment of Rodriguez et al., I interpret the period of desynchronization as corresponding to the movement of the stream of consciousness from one substantive thought to another. This equation can possibly explain a number of the phenomenological characteristics of the substantive thoughts: (1) *limited capacity*: each region of the brain can independently process an immense amount of information, but only a limited amount of information can take part in the global ensemble; (2) *stability*: a pattern of synchronous firing among disparate brain regions takes several hundred milliseconds to develop—consistent with the 200–500 ms it takes to individuate distinct events (Potter, 1975; Duncan, Ward, & Shapiro, 1994)—and must be maintained for a minimal amount of time in order to be conscious (Tononi & Edelman, 1998); (3) *multimodality*: information from any sensory modality can be a dominant part of any particular ensemble; (4) *sequential processing*: only one such ensemble can exist at any one time in the brain; and (5) *memorability*: the contents of each ensemble are automatically encoded by the medial temporal memory system for subsequent retrieval.

Note that all of the above “explanations” are really hypotheses at this point. For example, we do not really know that only one global ensemble can exist at any one

time, nor do we really understand why only a limited amount of information can be represented in a single ensemble. In order to turn these hypotheses into true neurophysiological explanations of the phenomenological Jamesian stream of thought, we need to understand more about how synchronous firing arises in the brain and how it represents information. The advantage of the above formulation is that it tells us which phenomenological features correspond to which neurophysiological features, which potentially allows phenomenology to constrain experiment and vice versa. For example, because it is commonly assumed that we think only one thought at a time, we suspect that each synchronous ensemble represents only one integrated thought. One could imagine that a deeper understanding of the phenomenon of neural synchrony may suggest that certain ensembles may actually incorporate two distinct thoughts, which would force us to reexamine our phenomenology (the reverse may also obtain; see section IV). Similarly, the observation that substantive thoughts (i.e., nucleus) are easily remembered and recalled suggests that it will be profitable to examine the ways in which the hippocampal memory system encodes this neural pattern and the way in which the frontal lobes and hippocampus interact to reconstitute such patterns (e.g., McIntosh et al., 1999).

III. NEURAL-COGNITIVE BASIS OF THE FRINGE

In this section, I turn to the neural-cognitive basis of the mechanisms that guide consciousness from one substantive thought to another. As discussed in section I, the phenomenology suggests that there are two such mechanisms, which I argue are mediated by different structures in the brain. First, there is an associative memory network that allows one to retrieve information about objects, places, ideas, facts, and events that are related to the current nucleus. A large body of evidence suggests that the medial temporal lobes (including the hippocampus and underlying cortices in the parahippocampal gyrus) play a critical role in supporting this network (Squire & Zola Morgan, 1991). Second, there are mechanisms that control which particular associations get instantiated by evaluating their relevance to the current behavioral context. The evidence suggests that these mechanisms are supported by the frontal lobes.

To offer an analogy: Thinking is like driving a car. The medial temporal regions provide a road of possible associations that one may follow. The frontal lobes are the driver, making decisions about which path to take. In fact, as is shown, this analogy is more than just a metaphor, as many of the areas of the brain involved in memory retrieval are also involved in spatial navigation.

Medial Temporal Lobes

When considering the role that medial temporal brain regions playing in mediating the stream of thought, it is worthwhile to consider their role in both spatial navigation and memory retrieval. On the surface, these may seem like unrelated functions; in fact, there is a long history of competition between spatial (O'Keefe & Nadel, 1978) and declarative memory (Squire & Zola-Morgan, 1991; Squire 1992) accounts of the medial temporal lobes. However, the connection between them becomes clear when one considers the role that memory retrieval plays in internally directed thought.

Declarative memory is the critical associative mechanism that allows us to go from one thought to another without each thought being immediately dependent on incoming sensory information. As such, the mechanisms that support declarative memory mediate a kind of “internal” navigation through associative memory “space” (Eichenbaum et al., 1999). There are a number of functional similarities between this “internal” navigation and external navigation.⁵ Both involve “moving” (i.e., transitive parts of thought) between different “places” (i.e., substantive thoughts). Both involve making decisions about the direction one wants to go on the basis of incomplete information: One cannot fully judge the value of a place without going there or the value of a substantive thought without having it. In this section, I argue that the medial temporal lobes of the human brain are the locus of an associative memory network that originally supported external spatial navigation but later evolved to support internal memory navigation.

The canonical finding linking the medial temporal lobes to spatial navigation is O’Keefe and Dostrovsky’s (1971) observation of “place cells” in the rat hippocampus that fire only when the animal is in a particular location in a given environment. When the animal is put into a new environment, a completely different set of place cells will fire, demonstrating that the hippocampus distinguishes between the two environments and instantiates different representations (or “maps”) for each one. From this result, O’Keefe and Nadel (1978) argued that the hippocampus represents the animal’s location in allocentric (environment-centered) coordinates. A large body of research has generally supported this claim.

The hippocampus is able to identify the animal’s location because it sits at the endpoint of a perceptual processing stream that begins in sensory areas and passes through several intermediate processing stages. The hippocampus can perform two distinct operations in response to this incoming information. If the incoming cues are similar to those that the animal has encountered before, then the hippocampus can use these cues to reconstitute the map of the earlier environment and identify its location on that map. If the incoming cues are novel, the hippocampus can set

⁵ There are also a number of phenomenological similarities between the experience of internal navigation and the experience of external navigation. Consider, for example, the following experience. Think of a friend who lives in another part of the city. Now, imagine going to his or her house. What does the process of imagining this journey involve? Certainly, you do not visualize the whole journey when you think about making it. Perhaps you imagine little snippets of the journey—for example, if you need to take the subway, you might see a brief image of the subway—but certainly you do not visualize the entire journey in advance. However, your inability to do so is in no way troubling to you. You have perfect confidence in your ability to get from here to there. If asked, you will aver that you “know where the house is.” Even if you could not give explicit directions to your destination, you still might have perfect confidence that you could get there—confidence that you will make the right decisions along the way. Now think of what it is like to think about an abstract argument or theory—Kant’s notion of the categorical imperative, for example. Though we neither see the whole abstract idea stretched out before us at once (as Mozart claimed he could do with his symphonies) nor have an ability to cite Kant’s writings word for word, we can develop a confidence in our ability to explicate Kant’s argument in our own words, to explain it to other people, to use it as a tool to reach our own conclusions, and to discuss it reasonably in argument. When we say that we “understand” Kant’s argument, it is because the argument has become an entity that we know well, a structure that we are confident that we can get around in, just as we feel confident that we can get around the city and get to our friend’s house.

up a new map for this new environment. These two “modes” have been referred to as *recall* and *storage* or *pattern completion* and *pattern separation* (Redish, 1999; O’Reilly & McClelland, 1994). A third mode (sometimes called *replay*; Redish, 1999) can occur when there is no incoming sensory information (e.g., during sleep). In this case, the firing of the cells is determined by the intrinsic connections within the hippocampus, which will tend to cause the reinstatement of firing sequences that correspond to routes that have been traveled in the past (Wilson & McNaughton, 1994; Skaggs & McNaughton, 1996). In both recall and replay, signals are sent through back projection pathways to the cortex, causing reinstatement of the *cortical* firing pattern that originally triggered the now recovered *hippocampal* firing pattern.⁶

Thus, the rat hippocampus has two primary roles. The first is to encode a multimodal representation of the surrounding environment. The second is to act as an association machine: when partial cues are encountered that are similar to cues from an environment visited from the past, the hippocampus can complete the pattern and recover the whole set of inputs from the originally encoded environment. If sensory input is temporarily suppressed, the hippocampal representation will drift “along” the route originally traveled by the animal, and the experiences previously encountered along this route will be reinstated in cortex. Although this system may have originally evolved to support the storage, recall, and replay of routes, it is easy to see how it could have subsequently evolved to perform the same operations on experiences that are not explicitly topographical. In fact, cells have been discovered in the rat hippocampus that fire in response to nonspatial aspects of an episode (Wood et al., 1999; Hampson et al., 1999). In the monkey hippocampus, “spatial view cells” have been discovered whose firing is correlated with where the monkey is looking rather than where the monkey is (Rolls et al., 1997). This suggests that the monkey hippocampus is less involved in encoding the animal’s own location than in encoding episodes that are visually observed in other locations. In the human, famous neurological cases like HM provide convincing evidence that the hippocampus and other medial temporal structures play a critical role in supporting nontopographical memory (Squire & Zola Morgan, 1991; Vargha-Khadem et al., 1997).

Given the small number of neurons in the hippocampus, it is not likely that all the details of an episode are encoded within the hippocampus itself. Rather, the hippocampus is believed to encode a summary sketch of an episode that represents the critical *relationships* between the important elements as well as pointers to where those elements are represented in cortex (Moscovitch, 1995). Recent evidence suggests that memory for the *details* of an event does not depend on the hippocampus

⁶ The original role of these back projection pathways in navigation is not entirely clear. In the case of replay, repeated reinstatement of former routes in the cortex is believed to allow the slow modification of cortico-cortical connections through Hebbian learning processes (Marr, 1971; McClelland et al., 1995; Hasselmo, 1999), which eventually allows the routes to be represented in the cortex as well as the hippocampus (Teng & Squire, 1999). Because the learning process is so much slower in the cortex than in the hippocampus, this cortical representation will differ from the hippocampal representation in that it will tend to only include aspects of the environment that are very regular (i.e., have been encountered many times; McClelland et al., 1995). It is also possible that during this process the data are reanalyzed and consolidated to reflect these regularities—for example, creating a representation of an environment in terms of an allocentric map rather than in terms of the original routes.

per se, but on underlying cortices in the parahippocampal region (Murray & Mishkin, 1998; Suzuki et al., 1993), with the perirhinal cortex specialized for encoding object representations (Buckley & Gaffan, 1998), the parahippocampal cortex specialized for spatial representations (Epstein & Kanwisher, 1998; Suzuki & Amaral, 1994), and the entorhinal cortex acting to combine input from these two regions to form representations of objects within locations (Suzuki, Miller, & Desimone, 1997).

Eichenbaum and his collaborators have performed a series of experiments to highlight the different roles of the hippocampus and parahippocampal region in memory. When rats with hippocampal lesions were presented with pairs of odors, they could learn which of the two odors was rewarded (i.e., A is preferred to B; B is preferred to C) but—unlike normal rats—they could not flexibly use this information about preferences between stimuli to create a transitive chain of inference when presented with a novel pairing (i.e., *therefore* A is preferred to C; Bunsey & Eichenbaum, 1996; Dusek & Eichenbaum, 1997; but see Burton et al., 2000). In other words, they could remember which stimuli were rewarded within individual events, but could not abstract this information to encode relationships between stimuli that existed across events. Interestingly, these hippocampally lesioned rats actually learned individual pairings *faster* than normal rats (Eichenbaum & Bunsey, 1995). From these results, Eichenbaum argues that the parahippocampal region and the hippocampus actually work at cross purposes: The parahippocampal region creates a unitary representation of an event or situation, while the hippocampus attempts to distinguish between different representations and encode the relationships between them.

Neuroimaging studies of humans also support the idea that the hippocampus is critical for storing the relationships that allow one to go from one distinct place to another or one distinct representation to another, while the parahippocampal gyrus represents the individual places or representation. As one might expect, the human medial temporal lobe is involved not only in declarative memory but spatial navigation as well (from which declarative memory functions presumably evolved). A number of studies have reported neural activation in the medial temporal region when subjects perform navigational tasks (Aguirre et al., 1996; Maguire et al., 1996, 1997; Ghaem et al., 1997) or view navigationally relevant visual stimuli (Epstein & Kanwisher, 1998; Epstein et al., 1999). In these studies, different roles were observed for the hippocampus and for the underlying cortices in the parahippocampal gyrus. For example, Epstein and Kanwisher (1998) found that parahippocampal cortex (but not hippocampus), responded strongly when subjects viewed photographs of places, but not when they viewed photographs of objects or faces. Complementary to this, Maguire et al. (1997) observed activation in the right hippocampus when subjects planned or imagined complicated itineraries, but not when they imagined the appearance of individual landmarks. Together, these results suggest that parahippocampal cortex is involved in the representation of individual places and scenes, while the hippocampus represents relationships between these places (see also Aguirre & D'Esposito, 1997). Similarly, Gabrieli et al. (1997) found parahippocampal activation when subjects encoded single pictures into memory, but hippocampal (subicular) activation when subjects had to recall a picture in response to a cue (which involves moving the focus of consciousness from the target to the cue). Thus, parahippocampal activity is associated with the encoding of representations of single pictures or places,

while hippocampal activity is associated with movement of the focus of consciousness from one thought to another or one imagined location to another.

It is possible that this individual representation vs relationships distinction that appears to exist between the parahippocampal region and hippocampus can be mapped onto the semantic vs episodic memory distinction proposed by Tulving (1983). One of the most compelling pieces of evidence for such a division between hippocampal and parahippocampal function has been provided by Vargha-Khadem et al. (1997), who reported on three patients who suffered hippocampal damage early in life. Despite a profound inability to recall the episodes of their daily lives, these subjects attended normal schools and were able to amass a normal store of semantic knowledge. They could answer questions like, "Who was Martin Luther King?", but not questions like, "What did you eat for breakfast today?" In all three of these patients, the cortices underlying the hippocampus appeared to be intact. Vargha-Khadem et al. concluded that these neocortical regions are sufficient to support semantic memory (i.e., memories for individual facts independent of any context) but the hippocampus is necessary for episodic memory (i.e., memories for events organized within a rich episodic context).

From these results, it is reasonable to hypothesize that the associative network underlying James's description of the stream of consciousness is mediated by the medial temporal lobes, with the parahippocampal gyrus supporting representations individual places, objects, facts, and events and the hippocampus supporting a representation of the relationships between these individual items.⁷ This associative network guides memory recall and hence the stream of thought. This description is consistent with the position of the hippocampus at the endpoint of the perceptual processing stream: Recognition of a stimulus often results in recall of associated stimuli. In internally directed thought, however, these associations are not driven automatically by incoming perceptual information. In this case, the car must do more than just go down the road: it must follow the will of the driver. In the next section, I make the case that the driver is in the frontal lobes.

The Frontal Lobes

If associative memory is thought of as a space in which the focus of consciousness can move, then "movement" in this space corresponds to changing from one currently active thought or memory to another. James's description of the fringe empha-

⁷ Note that the strictest version of this hypothesis cannot be true. Patient HM is missing much of his hippocampus and a good deal of underlying cortex (Corkin et al., 1997), but no one would claim that he does not have a stream of thought. Although he cannot encode new episodic memories, HM can recall some very old memories and thus can maintain at least a limited stream of associations. His ability to recall old memories indicates that at some point memories are stored outside the medial temporal lobes and can be recalled without benefit of these structures (Teng & Squire, 1999). However, it is interesting to note that the memories that HM does recall tend to have a very fixed structure, which suggests that intact medial temporal lobes are necessary for the formation new associations between old memories. Additional evidence for eventual cortical consolidation of the associative memory network comes from the fact that semantic dementia patients generally have damage to the temporal poles rather than the parahippocampal region (Graham & Hodges, 1997; Mummery et al., 1999).

sizes the role of feelings that monitor and guide these changes—for example, the feeling of harmony one gets when the current thought fits well with the current topic of interest. Although these guiding feelings often convey information about the associative memory network, they do not involve actual activation of that network (cf. the TOT state, which does not involve actual recall). This suggests that these guiding feelings are mediated in a different part of the brain from the associative memory network itself. The frontal lobes have been strongly implicated in the control and monitoring of memory processing (Shimamura, 1995a, 1995b). This inevitably suggests that the topographic locus of the memory navigation mechanisms of the fringe is in the frontal lobes, a proposal originally made by Mangan (1991, 1993a).

Of course, the frontal lobes represent a large and heterogeneous section of cortical territory. The precise functions of each of these subregions has not yet been determined, and many of the claims that have been made are highly controversial (e.g., compare Wilson et al., 1993 and Owen, 1997). It is not my intent to summarize this research. Rather, I discuss some illustrative examples that support the basic idea that the frontal lobes are involved in control of the stream of thought.

In a discussion of frontal lobe function, Moscovitch (1995) makes an important distinction between associative memory retrieval and strategic memory retrieval. Associative memory retrieval is an almost effortless response to a cue, as when answering the question, “Have you ever visited Paris?” Strategic memory retrieval, on the other hand, occurs when the answer to a question is not immediately available—for example, when answering the question, “Who was the first person you met when you came into the building today?” In order to answer these sorts of questions, appropriate cues must be provided to the associative memory system. Moscovitch hypothesizes that one role of the strategic memory system mediated by the frontal lobes is to organize a successful retrieval strategy by providing these cues to the associative memory system mediated by the hippocampus. As he aptly puts it, the role of the frontal lobe is that of “working with memory.”

This hypothesis has been supported by several neuroimaging studies. For example, Schacter et al. (1996a) observed more blood flow in the frontal lobes during an effortful retrieval task compared to an easy retrieval task, but more blood flow in the hippocampus in the easy retrieval task. Thus, frontal lobe activity corresponded to retrieval *effort*, while hippocampal activity corresponded to retrieval *success*. In another study, Dolan and Fletcher (1997) observed increased blood flow in the left dorsolateral prefrontal region within the frontal lobes when subjects were required to overwrite an old association between two words with a new association, indicating that this region is involved in active manipulation of memory. In a third study, Thompson-Schill et al. (1997) found that activation within the left inferior frontal gyrus increased whenever subjects performed a task that required them to retrieve only some facts about an item while ignoring others. In addition, a number of fMRI studies have identified regions within the frontal lobe that are active when subjects must hold faces (Courtney et al., 1997; Cohen et al., 1997) or locations (Courtney et al., 1998; but see Owen et al., 1998) in working memory. In all these studies, there is a common theme: Frontal activation does not correspond to memory recall per se, but to mechanisms that guide memory recall, such as choosing one memory over another or holding a memory available for imminent use.

Further supporting this hypothesis is the pattern of deficits found in frontal lobe patients. Many of these patients show an inability to use intentional strategies to query memory. They are frequently proficient at recognition tasks, but are severely impaired at tasks that require free recall. In one study, patients had severe difficulty in providing the answer to questions about historical events such as “Who shot John Lennon?” or supplying a name to go with a famous face. On the other hand, they were as good as controls at recognizing the correct answer when it was provided for them. In another study, frontal patients had great difficulty in answering questions such as, “How long is the average man’s necktie?” (see Shimamura, 1995b for discussion) Since the answer to such a question is not available immediately, subjects must presumably use a sophisticated strategy of simultaneously retrieving the image of a necktie from long-term memory while holding the question at hand in working memory. Such intentional navigation of long-term memory seems to be beyond the abilities of these patients.

Frontal patients do not suffer from an inability to associate—if anything, they associate too easily. What they lack is the ability to control their associations. They often suffer from impaired feelings of knowing (FOK) and increased numbers of tip-of-the-tongue (TOT) states (Shimamura, 1995b). This indicates an inability to convert the feeling of being “near” a memory into actual recollection—an act that often involves choosing a specific memory from among a number of related memories. Many of these patients also have problems with source monitoring: While they are able to remember things, they are unable to remember how that information was learned (Shimamura, 1995a). This indicates an impairment in the ability to monitor the relationship between an event and its context. This inability to monitor retrieved memories for episodic consistency may also account for the numerous confabulations produced by Korsakoff’s patients. These generally are not completely false, but contain jumbled pieces of veridical memories.

Schacter and his colleagues (Schacter, 1996b; Schacter et al., 1996c) report a patient (BG) whose deficits nicely illustrate the role of the frontal lobes in mediating the guiding function of the fringe. This patient suffers from right frontal lobe damage. He does not suffer from amnesia and does not spontaneously confabulate. However, when he is tested on his memory for previously learned word lists, he often claims to remember words that were not actually on the list but which are semantically related to words that were on the list. For example, if “honey” was on the list of studied words, he is much more likely than normals to falsely claim that “sweet” was on the list as well. In contrast, BG almost never makes false alarms to words that are from semantic categories not present in this study list. Schacter argues that BG has difficulty distinguishing between memory for the general features of the study session (such as which categories of words were presented) and memory for specific items in the study list. This is similar to the situation we find ourselves in when we are in a TOT state. In such a state, we feel that we are able to recover information about the context of an item without being able to recover the item itself.

Other Brain Regions

Besides the medial temporal lobes and the frontal lobes, other brain regions may play a role in mediating some of James’ fringe experiences. Take for example the

TOT state. I have described this state as an inability to retrieve a desired word even when one can retrieve the correct semantic context. I have also discussed results that indicate that the inferior frontal region plays a critical role in allowing one to choose a single item from among many in semantic memory (e.g., Thompson-Schill, 1997; Gabrieli et al., 1998) as well as neuropsychological results from a patient with frontal damage who appears to be specifically impaired at distinguishing specific items from their context. This suggests that the failure to choose the target in the TOT state results from activity—or lack of activity—in the inferior frontal lobes. However, there are additional aspects of the TOT phenomenology that are unlikely to result from neural activity in this region. For example, one critical difference between patient BG's false alarms and the TOT state is that the TOT state includes a sense of *effort* to reach the correct answer, while BG makes his memory judgments effortlessly and seems to have no awareness that he is missing the mark. A large part of the experience of a TOT state is this awareness of one's intent to retrieve the target. Any neural account of the TOT state must be able to explain this feeling.

This sense of effort may correlate with activity in motor regions such as the cerebellum. Interestingly, left inferior frontal activation is often paired with right cerebellar activation in recall tasks. The distinct roles of these two regions in memory retrieval have been clarified by an experiment by Desmond et al. (1998). In this experiment, subjects were required to complete letter stems with the first word that came to mind. The response to stems in which there were many possible completions (e.g., "sta") was compared to the response to stems in which there were few possible completions (e.g., "psa"). Left frontal regions were more active in the former case (where recovering completions was easy but response choice was necessary) while cerebellar regions were more active in the latter case (where sustained effort was required but few choices had to be made). Based on these results, we can tentatively suggest that TOT state involves at least three different regions of the brain: (1) activation of parts of the medial temporal lobe memory network associated with the target item, (2) a feeling of sustained effort associated with activation within the cerebellum, and (3) the failure of frontal regions—perhaps inferior prefrontal cortex—to choose the target from among items associated with it.

In addition to this feeling of effort, there are other, more kinesthetic or tactile sensations that are a key element of a number of the fringe experiences described by William James. For example, in some cases, the sensation of effort experienced in the TOT state can be so intense that one feels vividly as if one were grasping for something. Mangan reports:

[W]hen we strain to retrieve a word, I suspect we do have a sense of roughly peripheral target-location relative to our body image. In my case, the feeling of grasping for a TOT often seems to be in a region above and behind my right ear and just outside my skull. (Mangan, 1993b, p. 144)

It is also a common experience to sense a sort of tactile or kinesthetic structure while listening to classical music, and there is certainly some residue of a kinesthetic sense inherent in our experience of abstract thought. Einstein once observed that muscular sensations played important role in his thinking about physics:

The words . . . as they are written or spoken, do not seem to play any role in my mechanism of thought. The psychical entities which seem to serve as elements in thought are cer-

tainly signs and more or less clear images, which can be ‘voluntarily’ reproduced and combined . . . this combinatory play seems to be an essential feature in productive thought . . . The above mentioned [nonverbal] elements are, in my case, of visual and some of muscular type. (Albert Einstein, quoted in Baars, 1993, p. 126)

It is apparent that Einstein was trying to convey the idea that his thought processes involved feelings of movement and effort. Similar sensations may also account for our sense of the connecting terms—“and,” “if,” “but,” and so on—that hold together a logical argument (Talmy, 1988).

These sensations are possibly be related to the “motor-images” proposed by Jeannerod (1994), which are representations of actions which have a certain amount of sensory quality. Experiments by Rizzolatti and his colleagues have revealed the existence of cells in premotor area F5 that may be mediating such representations. These cells respond selectively when the monkey makes different grasping hand movements (see Jeannerod et al., 1995, for review). For example, certain cells fire when the monkey grasps a sphere, but not when it grasps a cylinder. Interestingly enough, 20–30% of these F5 neurons respond to visual presentations even when the monkey makes no hand movement. Some of these neurons respond to the appearance of objects that can be grasped using the grip coded by the neuron, while others respond to similar grasping movements made by the experimenter or another monkey. The fact that these neurons fire both to the appearance and to the performance of the appropriate movement in all these situations indicates that they do more than just control specific muscles: They code a general representation of the movement as a whole. Similar activations of motor areas have been observed in humans using PET during both imagined grasping actions and observations of such actions performed by others (Decety et al., 1994).

It is unclear why these motoric representations, whose primary purpose is to support action in the external world, would be involved in abstract thought and memory retrieval. It is possible that—just as spatial navigation mechanisms evolved to also support memory-space navigation—so motor images that originally supported manipulation of real objects subsequently evolved to also support manipulation of abstract objects (“ideas”). One may speculate that whenever such an evolutionary development takes place, some of the phenomenological flavor of the original process is retained.

IV. “BACKGROUND” MENTAL REPRESENTATIONS

Although James’ description of the bipartite nature of the stream of thought captures an important aspect of the phenomenology of consciousness, it is not a complete description. In fact, the “bird’s life” described by James should be viewed an ideal case. In addition to substantive thoughts and the fringe feelings that guide the transitions between them, we also experience imagistic representations that seem to have little to do with the main train of thought. These “background” representations are not easily classified as either nucleus or fringe.

Consider the common experience of humming a tune “in the background” of one’s mind. In my own experience, I usually get a strong sense that the tune was

definitely going on in my head *before* I noticed it. The act of attending to the song and bringing it into the foreground has a very different feel from the act of recalling it whole cloth from memory. In the latter case, one feels that one is encountering a representation that was definitely *not* previously present in consciousness. In the former case, one feels that the song was there but separate from the main stream of thought, like a little piece of consciousness that got spun off from the main body and is working on autopilot. Often, there is an association between the tune and an earlier thought—for example, through the words of the tune. What kind of beast is this “background” representation? It has the sensory quality of the representations that typically make up the nucleus, but is not actually the nucleus of the current substantive thought.⁸

In section II, I postulated that the substantive thoughts were formed by temporary global synchrony between disparate cortical regions—what Hebb referred to as a cell assembly, and Tononi and Edelman call the dynamic core. The involvement of different cortical regions in this cell assembly will certainly vary in a way that corresponds to the material in the focus of consciousness. For example, when one looks at a face, the parts of the brain involved in processing the visual features of faces will play a dominant role, while the auditory system may be little involved. One can imagine that under certain circumstances, regions that are not contributing to the “core” might in fact be operating independently. Certainly, the results discussed in section II indicate that the various cortical regions process information automatically even when their output is not part of the nucleus. If this processing has some conscious correlate—what Zeki and Bartels (1999) refer to as microconsciousness—then one could have representations that are conscious, but “in the background.”

I hypothesize that a dissociation between the current substantive thought and a background representation is more likely to occur when the background representation is of a different modality from that of the nucleus of the current substantive thought. For example, the tune in the background is not likely to occur when one is listening to a symphony. Such an effect would reflect the pattern of connectivity between different regions of the brain. Simply put: Auditory regions are more interconnected with each other than they are with visual regions and vice versa. Scannell,

⁸ A relevant perceptual phenomenon is scene change blindness (Rensink, O'Regan, & Clark, 1997). In this paradigm, subjects view two scenes that are identical except for one detail. The scenes alternate on the screen until observers discover the difference between them. When a blank field is interposed between the presentations of the two scenes (preventing the use of visual transients as a cue) observers take a surprisingly long time to find the change. Once they do notice, the difference becomes obvious and stands out prominently on every alternation. Thus, despite our subjective impression that we can “see” and remember the entire scene, we do not form a representation durable enough to find an often fairly large change in it 100 ms after it disappears. The intriguing question is: What kind of representation do we have of the unattended parts of the scene? Despite radical claims to the contrary (O'Regan, 1992) we *do* seem to have a conscious experience the part of the scene we cannot report, just as we seem to have a conscious experience of the song playing in our head before we attend to it. In the case of the scene, the unattended material is *literally* in the background. The distinction between the attended and unattended parts of the scene can possibly be mapped onto to the distinction between A-consciousness and P-consciousness, introduced by Block (1995): The unattended part of the scene is P-conscious (phenomenally experienced) without being A-conscious (accessible). Similarly, the jingle in the back of the mind might be P-conscious without being A-conscious.

Blakemore, and Young (1995; Scannell et al., 1999) have performed detailed analyses of cortical connectivity in the cat and found that the cortical regions can be divided into four highly interconnected groups: the visual system, the auditory system, the somatosensory-motor system, and the frontal-hippocampal system. Based on the previous discussion, we can postulate that the establishment of focal awareness requires involvement of the frontal-hippocampal system. This system may get dissociated from some of the others: For example, “fleeting” (see below) visual images may occur if the visual system develops an independent firing pattern that does not synchronize with the firing in the frontal-hippocampal system, while “doodling” may occur if the somatosensory-motor system gets similarly disconnected. Consistent with this general view, behavioral evidence suggests that visual and auditory attention can be dissociated from each other (Duncan, Martens, & Ward, 1997).

These “background” mental representations bear a clear resemblance to the “fleeting” mental representations discussed by Baars (1998). Baars proposed that these were mental representations that are present in the global workspace (i.e., nucleus) for only a brief amount of time. During this interval, they are “momentary accessible” by a range of unconscious processors. If these processors engage the representation, then it becomes the new focus of consciousness. Thus, under Baars’ hypothesis, momentary access to the global workspace affords a mechanism by which consciousness can “test the waters” without fully committing to a new representation (and consequent direction of thought). The elusiveness of these fleeting mental representations, under this hypothesis, is a consequence of their short life span. In contrast, I am proposing a different account, in which the elusiveness of these representations comes not from their brevity, but from the fact that they are not part of the main global synchrony that produces the substantive thought. Nevertheless, even under this altered account, background mental representations may play an important functional role in controlling the stream of thought by acting as precursors to the representations in the nucleus.

Note that the primary difference between my account and Baars’ account is that I believe that there can be conscious representations that are not in the global workspace (i.e., not part of the substantive thought), while Baars (at least, under my interpretation of his theory) does not. My arguments for such representations come largely from neurobiology. This is a good example of how neuroscience can be used to explicate our understanding of the phenomenology of consciousness. Background mental representations are a relatively little explored aspect of this phenomenology. I suggest that a detailed investigation of them would potentially greatly expand our understanding of the neural-cognitive model of consciousness described in this article.

V. CONCLUSION

In this article, I have attempted to link William James’ account of the structure of consciousness to modern cognitive neuroscientific results. I have argued that James’ substantive thoughts are formed by synchronous firing between different regions of the brain, while the fringe corresponds to frontal and medial temporal mechanisms that control the contents of these thoughts. I have further argued that background mental representations can be incorporated into this framework and suggest that they

may be a fruitful area for further phenomenological research. Although these conclusions will certainly need to be refined as our knowledge of the brain advances, I believe that we know enough now to draw preliminary connections between phenomenology and neuroscience that can illuminate both fields.

Whenever possible, I have attempted to combine phenomenological descriptions and experimental results into a coherent whole. My claim is that such a mixed investigation actually tells us something important and interesting about the nature of consciousness. Insofar as consciousness has a structure that can be investigated through both phenomenology and experiment, and insofar as both these approaches reveal a similar structure, then consciousness can be investigated in a way that is both scientifically valid and phenomenologically illuminating. In other words, we can simultaneously learn about how our brains work as information processing or neural mechanisms and how those mechanisms translate into specific subjective experiences. Phenomenology and neuroscience are not separate worlds; they should reinforce and inform each other. As William James was so well aware, a mixture of them both is necessary if we are ever to arrive at a fully satisfactory understanding of consciousness.

ACKNOWLEDGMENTS

I am very grateful to David Galin for his encouragement and for several long discussions that helped me to clarify these ideas. This article also benefited from thoughtful reviews provided by Wallace Chafe and Bruce Mangan and incisive editorial comments from Bill Banks. Finally, I thank Donald Dryden, John Duncan, Aaron Lipman, Emily Mechner, Ken Nakayama, and Peter Tse for providing useful and/or encouraging comments on earlier versions of the manuscript.

REFERENCES

- Aguirre, G. K., & D'Esposito, M. (1997). Environmental knowledge is subserved by separable dorsal/ventral areas. *Journal of Neuroscience*, **17**, 2512–2518.
- Aguirre, G. K., Detre, J. A., Alsup, D. C., & D'Esposito, M. (1996). The parahippocampus subserves topographical learning in man. *Cerebral Cortex*, **6**, 823–829.
- Baars, B. J. (1988). *A cognitive theory of consciousness*. Cambridge, UK: Cambridge Univ. Press.
- Baars, B. J. (1993). Putting the focus on the fringe: Three empirical cases. *Consciousness and Cognition*, **2**, 126–136.
- Baylis, G. C., Driver, J., & Rafal, R. D. (1993). Visual extinction and stimulus repetition. *Journal of Cognitive Neuroscience*, **5**, 453–466.
- Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, **18**, 227–247.
- Brown, A. S. (1991). A review of the tip of the tongue experience. *Psychological Bulletin*, **109**, 204–223.
- Brown, R., & McNeill, D. (1966). The “tip of the tongue” phenomenon. *Journal of Verbal Learning and Verbal Behavior*, **5**, 325–337.
- Buckley, M. J., & Gaffan, D. (1998). Perirhinal cortex ablation impairs visual object identification. *Journal of Neuroscience*, **18**, 2268–2275.
- Bunsey, M., & Eichenbaum, H. (1996). Conservation of hippocampal memory function in rats and humans. *Nature*, **379**, 255–257.
- Burton, S., Murphy, D., Qureshi, U., Sutton, P., & O'Keefe, J. (2000). Combined lesions of hippocampus

- and subiculum do not produce deficits in a nonspatial social olfactory memory task. *Journal of Neuroscience*, **20**, 5468–5475.
- Chafe, W. L. (1994). *Discourse, consciousness, and time: The flow and displacement of conscious experience in speaking and writing*. Chicago: Univ. of Chicago Press.
- Chun, M. M. (1997). Types and tokens in visual processing: A double dissociation between the attentional blink and repetition blindness. *Journal of Experimental Psychology: Human Perception and Performance*, **23**, 738–755.
- Chun, M. M., & Potter, M. C. (1995). A two-stage model for multiple target detection in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance*, **21**, 109–127.
- Cohen, J. D., Perlstein, W. M., Braver, T. S., Nystrom, L. E., Noll, D. C., Jonides, J., & Smith, E. E. (1997). Temporal dynamics of brain activation during a working memory task. *Nature*, **386**, 604–607.
- Corkin, S., Amaral, D. G., Gonzalez, R. G., Johnson, K. A., & Hyman, B. T. (1997). H.M.'s medial temporal lobe lesion: Findings from magnetic resonance imaging. *Journal of Neuroscience*, **17**, 3964–3979.
- Courtney, S. M., Petit, L., Maisog, J. M., Ungerleider, L. G., & Haxby, J. V. (1998). An area specialized for spatial working memory in human frontal cortex. *Science*, **279**, 1347–1350.
- Courtney, S. M., Ungerleider, L. G., Kell, K., & Haxby, J. V. (1997). Transient and sustained activity in a distributed neural system for human working memory. *Nature*, **386**, 608–611.
- Crick, F., & Koch, C. (1998). Consciousness and neuroscience. *Cerebral Cortex*, **8**, 97–107.
- Decety, J., Perani, D., Jeannerod, M., Bettinardi, V., Tadary, B., Woods, R., Mazziotta, J. C., & Fazio, F. (1994) Mapping motor representations with positron emission topography. *Nature*, **371**, 600–602.
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, **18**, 193–222.
- Desmond, J. E., Gabrieli, J. D., & Glover, G. H. (1998). Dissociation of frontal and cerebellar activity in a cognitive task: Evidence for a distinction between selection and search. *Neuroimage*, **7**, 368–376.
- Dolan, R. J., & Fletcher, P. C. (1997). Dissociating prefrontal and hippocampal function in episodic memory encoding. *Nature*, **388**, 582–585.
- Duncan, J., Martens, S., & Ward, R. (1997). Restricted attentional capacity within but not between sensory modalities. *Nature*, **387**, 808–810.
- Duncan, J., Ward, R., & Shapiro, K. (1994). Direct measurement of attentional dwell time in human vision. *Nature*, **369**, 313–315.
- Dusek, J. A., & Eichenbaum, H. (1997). The hippocampus and memory for orderly stimulus relations. *Proceedings of the National Academy of Science of the USA*, **94**, 7109–7114.
- Eichenbaum, H., & Bunsey, M. (1995). On the binding of associations in memory: Clues from studies on the role of the hippocampal region in paired-associate learning. *Current Directions in Psychological Science*, **4**, 19–23.
- Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M., & Tanila, H. (1999). The hippocampus, memory, and place cells: Is it spatial memory or a memory space? *Neuron*, **23**, 209–226.
- Engel, A. K., Fries, P., Konig, P., Brecht, M., & Singer, W. (1999). Temporal binding, binocular rivalry, and consciousness. *Consciousness and Cognition*, **8**, 128–151.
- Engel, A. K., Konig, P., & Singer, W. (1991a). Direct physiological evidence for scene segmentation by temporal coding. *Proceedings of the National Academy of Sciences of the USA*, **88**, 9136–9140.
- Engel, A. K., Konig, P., Kreiter, A. K., & Singer, W. (1991b). Interhemispheric synchronization of oscillatory neuronal responses in cat visual cortex. *Science*, **252**, 1177–1179.
- Epstein, R., Harris, A., Stanley, D., & Kanwisher, N. (1999) The parahippocampal place area: Recognition, navigation, or encoding? *Neuron*, **23**, 115–125.

- Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, **392**, 598–601.
- Farah, M. J. (1990). *Visual agnosia: Disorders of object recognition and what they tell us about normal vision*. Cambridge, MA: MIT Press/Bradford Books.
- Gabrieli, J. D. E., Brewer, J. B., Desmond, J. E., & Glover, G. H. (1997). Separate neural bases of two fundamental memory processes in the human medial temporal lobe. *Science*, **276**, 264–266.
- Gabrieli, J. D. E., Poldrack, R. A., & Desmond, J. E. (1998). The role of left prefrontal cortex in language and memory. *Proceedings of the National Academy of Sciences of the USA*, **95**, 906–913.
- Galín, D. (1994). The structure of awareness: Contemporary applications of William James' forgotten concept of "the fringe." *Journal of Mind and Behavior*, **15**, 375–402.
- Ghaem, O., Mellet, E., Crivello, F., Tzourio, N., Mazoyer, B., Berthoz, A., & Denis, M. (1997). Mental navigation along memorized routes activates the hippocampus, precuneus, and insula. *NeuroReport*, **8**, 739–744.
- Graham, K. S., & Hodges, J. R. (1997). Differentiating the roles of the hippocampal complex and the neocortex in long-term memory storage: Evidence from the study of semantic dementia and Alzheimer's disease. *Neuropsychology*, **11**, 77–89.
- Gray C. M., Konig, P., Engel, A. K., & Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, **338**, 334–337.
- Hampson, R. E., Simeral, J. D., & Deadwyler, S. A. (1999). Distribution of spatial and nonspatial information in dorsal hippocampus. *Nature*, **402**, 610–614.
- Hasselmo, M. E. (1999). Neuromodulation: Acetylcholine and memory consolidation. *Trends in Cognitive Sciences*, **3**, 351–359.
- Husain, M., Shapiro, K., Martin, J., & Kennard, C. (1997). Abnormal temporal dynamics of visual attention in spatial neglect patients. *Nature*, **385**, 154–156.
- Jackendoff, R. (1987). *Consciousness and the computational mind*. Cambridge, MA: MIT Press.
- James, W. (1890). *The principles of psychology*, Vol. I. New York: Dover.
- Jeannerod, M. (1994). The representing brain: Neural correlates of motor intention and imagery. *Behavioral and Brain Sciences*, **17**, 187–245.
- Jeannerod, M., Arbib, M. A., Rizzolatti, G., & Sakata, H. (1995). Grasping objects: The cortical mechanisms of visuomotor transformations. *Trends in Neuroscience*, **18**, 312–320.
- Kanwisher, N. (1987). Repetition blindness: Type recognition without token individuation. *Cognition*, **27**, 117–143.
- Kanwisher, N. (1991). Repetition blindness and illusory conjunctions: Errors in binding visual types with visual tokens. *Journal of Experimental Psychology: Human Perception and Performance*, **17**, 404–421.
- Llinas, R. R., & Ribary, U. (1993). Coherent 40-Hz isolation characterizes dream state in humans. *Proceeding of the National Academy of Sciences of the USA*, **90**, 2078–2081.
- Luck, S. J., Vogel, E. K., & Shapiro, K. (1996). Were meanings can be accessed but not reported during the attentional blink. *Nature*, **383**, 616–618.
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why are there complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, **102**, 419–457.
- McIntosh, A. R., Rajah, M. N., & Lobaugh, N. J. (1999). Interactions of prefrontal cortex in relation to awareness in sensory learning. *Science*, **284**, 1531–1533.
- Maguire, E. A., Frackowiak, R. S. J., & Frith, C. D. (1996). Learning to find your way: A role for the human hippocampal region. *Proceedings of the Royal Society of London B* **263**, 1745–1750.
- Maguire, E. A., Frackowiak, R. S. J., & Frith, C. D. (1997). Recalling routes around London: activation of the right hippocampus in taxi drivers. *Journal of Neuroscience*, **17**, 7103–7110.

- Mangan, B. B. (1991). *Meaning and the structure of consciousness: An essay in psycho-aesthetics*. Unpublished Ph.D. thesis, University of California, Berkeley.
- Mangan, B. B. (1993a). Taking phenomenology seriously: The “fringe” and its implications for cognitive research. *Consciousness and Cognition*, **2**, 89–108.
- Mangan, B. B. (1993b). Some philosophical and empirical implications of the fringe. *Consciousness and Cognition*, **2**, 142–154.
- Marcel, A. J. (1983). Conscious and unconscious perception: An approach to the relations between phenomenal experience and perceptual processes. *Cognitive Psychology*, **15**, 238–300.
- Marr, D. (1971). Simple memory: A theory for archicortex. *Proceedings of the Royal Society of London B*, **262**, 23–81.
- Marshall, J. C., & Halligan, P. W. (1988). Blindsight and insight into visuo-spatial neglect. *Nature*, **336**, 766–767.
- Mummery, C. J., Patterson, K. E., Wise, R. J. S., Price, C. J., & Hodges, J. R. (1999). Disrupted temporal lobe connections in semantic dementia. *Brain*, **122**, 61–73.
- Murray, E. A., & Mishkin, M. (1998). Object recognition and location memory in monkeys with excitotoxic lesions of the amygdala and hippocampus. *Journal of Neuroscience*, **18**, 6568–6582.
- Moscovitch, M. (1995). Models of consciousness and memory. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences*. Cambridge, MA: MIT Press.
- O’Keefe, J., & Dostrovsky, J. (1973). The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely-moving rat. *Experimental Brain Research*, **31**, 573–590.
- O’Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford: Oxford Univ. Press.
- O’Regan, J. K. (1992). Solving the “real” mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology*, **46**, 461–488.
- O’Reilly, R. C., & McClelland, J. L. (1994). Hippocampal conjunctive encoding, storage, and retrieval: Avoiding a tradeoff. *Hippocampus*, **4**, 661–682.
- Owen, A. M. (1997). The functional organization of working memory processes within human lateral frontal cortex: The contribution of functional neuroimaging. *European Journal of Neuroscience*, **9**, 1329–1339.
- Owen, A. M., Stern, C. E., Look, R. B., Tracey, I., Rosen, B. R., & Petrides, M. (1998). Functional organization of spatial and nonspatial working memory processing within the human lateral frontal cortex. *Proceedings of the National Academy of Sciences of the USA*, **95**, 7721–7726.
- Potter, M. C. (1975). Meaning in visual search. *Science*, **187**, 965–966.
- Potter, M. C. (1976). Short term conceptual memory for pictures. *Journal of Experimental Psychology: Human Learning and Memory*, **2**, 509–522.
- Potter, M. C. (1993). Very short-term conceptual memory. *Memory & Cognition*, **21**, 156–161.
- Potter, M. C. (1999). Understanding sentences and scenes: The role of conceptual short-term memory. In V. Coltheart (Ed.), *Fleeting memories*. Cambridge, MA: MIT Press.
- Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance*, **18**, 849–860.
- Redish, A. D. (1999). *Beyond the cognitive map: From place cells to episodic memory*. Cambridge, MA: MIT Press.
- Rensink, R. A., O’Regan, J. K., & Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, **8**, 368–373.
- Rodriguez, E., George, N., Lachaux, J.-P., Martinerie, J., Renault, B., & Varela, F. J. (1999). Perception’s shadow: Long-distance synchronization of human brain activity. *Nature*, **397**, 430–433.
- Roelfsema, P. R., Engel, A. K., Konig, P., & Singer, W. (1997). Visuomotor integration is associated with zero time-lag synchronization among cortical areas. *Nature*, **385**, 157–161.
- Rolls, E. T., Robertson, R. G., & Georges-Francois, P. (1997). Spatial view cells in the primate hippocampus. *European Journal of Neuroscience*, **9**, 1789–1794.

- Scannell, J. W., Blakemore, C., & Young, M. P. (1995). Analysis of connectivity in the cat cerebral cortex. *Journal of Neuroscience*, **15**, 1463–1483.
- Scannell, J. W., Burns, G. A., Hilgetag, C. C., O'Neil, M. A., & Young, M. P. (1999). The connectional organization of the cortico-thalamic system of the cat. *Cerebral Cortex*, **9**, 277–99.
- Schacter, D. L., Alpert, N. M., Savage, C. R., Rauch, S. L., & Albert, M. S. (1996a). Conscious recollection and the human hippocampal formation: Evidence from positron emission tomography. *Proceedings of the National Academy of Sciences of the USA*, **93**, 321–325.
- Schacter, D. L. (1996b). Illusory memories: A cognitive neuroscience analysis. *Proceedings of the National Academy of Sciences of the USA*, **93**, 13527–13533.
- Schacter, D. L., Curran, T., Galluccio, L., Milberg, W., & Bates, J. (1996c). False recognition and the right frontal lobe: A case study. *Neuropsychologia*, **34**, 793–808.
- Shapiro, K., Driver, J., Ward, R., & Sorensen, R. E. (1997). Priming from the attentional blink: A failure to extract visual tokens but not visual types. *Psychological Science*, **8**, 95–100.
- Shimamura, A. P. (1995a). Memory and frontal lobe function. In M. S. Gazzaniga, (Ed.), *The cognitive neurosciences*, Cambridge, MA: MIT Press.
- Shimamura, A. P. (1995b). Memory and the prefrontal cortex. In J. Grafman, K. J. Holyoak, & F. Boller (Eds.), *Structure and functions of the human prefrontal cortex*. New York: New York Academy of Sciences.
- Singer, W., Engel, A. K., Kreiter, A. K., Munk, M. H. J., Neuenschwander, S., & Roelfsema, P. R. (1997). Neuronal assemblies: Necessity, signature and detectability. *Trends in Cognitive Sciences*, **1**, 252–261.
- Skaggs, W. E., & McNaughton, B. L. (1996). Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science*, **271**, 1870–1873.
- Squire, L. R. (1992). Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychological Review*, **99**, 195–231.
- Squire, L. R., & Zola-Morgan, S. (1991). The medial temporal lobe memory system. *Science*, **253**, 1380–1386.
- Suzuki, W. A., & Amaral, D. G. (1994). Perirhinal and parahippocampal cortices of the macaque monkey: Cortical referents. *Journal of Comparative Neurology*, **350**, 497–533.
- Suzuki, W. A., Miller, E. K., & Desimone, R. (1997). Object and place memory in the macaque entorhinal cortex. *Journal of Neurophysiology*, **78**, 1062–1081.
- Suzuki, W. A., Zola-Morgan, S., Squire, L. R., & Amaral, D. G. (1993). Lesions of the perirhinal and parahippocampal cortices in the monkey produce long-lasting memory impairment in the visual and tactile modalities. *Journal of Neuroscience*, **13**, 2430–2451.
- Talmy, L. (1998). Force dynamics in language and cognition. *Cognitive Science*, **12**, 49–100.
- Teng, E., & Squire, L. R. (1999). Memory for places learned long ago is intact after hippocampal damage. *Nature*, **400**, 675–677.
- Thompson-Schill, S. L., D'Esposito, M., Aguirre, G. K., & Farah, M. J. (1997). Role of left inferior prefrontal cortex in retrieval of semantic information: A reevaluation. *Proceeding of the National Academy of Sciences of the USA*, **94**, 14792–14797.
- Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature*, **381**, 520–522.
- Tononi, G., & Edelman, G. M. (1998). Consciousness and complexity. *Science*, **282**, 1846–1851.
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, **12**, 97–136.
- Tulving, E. (1983). *Elements of episodic memory*. New York: Oxford Univ. Press.
- Ungerleider, L. G., & Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale, and R. J. W. Mansfield (eds.), *Analysis of visual behavior*, pp. 549–585. Cambridge, MA: MIT Press.
- Vargha-Khadem, F., Gadian, D. G., Watkins, K. E., Connelly, A., Van Paesschen, W., & Mishkin, M.

- (1997). Differential effects of early hippocampal pathology on episodic and semantic memory. *Science*, **277**, 376–380.
- Wilson, F. A. W., Scaldie, S. P. O., & Goldman-Rakic, P. S. (1993). Dissociation of object and spatial processing domains in primate prefrontal cortex. *Science*, **260**, 1955–1958.
- Wilson, M. A., & McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, **265**, 676–679.
- Wood, E. R., Dudchenko, P. A., & Eichenbaum, H. (1999). The global record of memory in hippocampal neural activity. *Nature*, **397**, 613–616.
- Zeki, S., & Bartels, A. (1999). Toward a theory of visual consciousness. *Consciousness and Cognition*, **8**, 225–259.

Received November 3, 1997