Psychological Review
2007, Vol. 114, No. 2, 340–375

0033-295X/07/$12.00 DOI: 10.1037/0033-295X.114.2.340

Remembering the Past and Imagining the Future: A Neural Model of Spatial Memory and Imagery

Patrick Byrne and Suzanna Becker
McMaster University

Neil Burgess
University College, London

The authors model the neural mechanisms underlying spatial cognition, integrating neuronal systems and behavioral data, and address the relationships between long-term memory, short-term memory, and imagery, and between egocentric and allocentric and visual and ideothetic representations. Long-term spatial memory is modeled as attractor dynamics within medial–temporal allocentric representations, and short-term memory is modeled as egocentric parietal representations driven by perception, retrieval, and imagery and modulated by directed attention. Both encoding and retrieval/imagery require translation between egocentric and allocentric representations, which are mediated by posterior parietal and retrosplenial areas and the use of head direction representations in Papez's circuit. Thus, the hippocampus effectively indexes information by real or imagined location, whereas Papez's circuit translates to imagery or from perception according to the direction of view. Modulation of this translation by motor efference allows spatial updating of representations, whereas prefrontal simulated motor efference allows mental exploration. The alternating temporal–parietal flows of information are organized by the theta rhythm. Simulations demonstrate the retrieval and updating of familiar spatial scenes, hemispatial neglect in memory, and the effects on hippocampal place cell firing of lesioned head direction representations and of conflicting visual and ideothetic inputs.

Keywords: navigation, path integration, representational neglect, hippocampus, computational model

One of the most intriguing challenges in cognitive neuroscience is to understand how a higher cognitive function such as memory arises from the action of neurons and synapses in our brains. Such an understanding would serve to bridge between the neurophysiological and behavioral levels of description via systems neuroscience, allowing for the reinforcement of convergent information and the resolution of questions at one level of description by inferences drawn from another. Moreover, a theory that bridges the cellular and behavioral levels can lead to the development of experimental predictions from one level to another and improved ability to relate behavioral symptoms to their underlying pathologies. In terms of developing such an understanding of memory, *spatial memory* provides a good starting point due to the ability to use similar paradigms in humans and other animals.

We are often faced with the challenging task of deciding how to act in the absence of complete sensory information, for example, when navigating toward an unseen goal. To solve such tasks, we must rely on internal representations of object locations within their environment. Here we attempt to develop a model of the uses of these internal representations in spatial memory, incorporating data from single-unit recording systems, neuroscience and behavioral studies, and describing how each relates to the other. Central questions in the cognitive neuroscience of spatial memory concern the frames of reference used for representations of location, for example whether they are egocentric (relative to parts of the body) or allocentric (relative to the external environment), the durations over which different representations are maintained, the uses they are put to, and how they interact with each other. However, there is currently no clear consensus, with various investigators stressing one or the other type of representation (e.g., cf. Poucet, 1993; Wang & Spelke, 2002). To address these questions, we propose a general organizational structure for spatial memory (see also Burgess, 2006; Mou & McNamara, 2002) encompassing encoding and retrieval of spatial scenes as well as some aspects of spatial navigation, imagery, and planning. We then implement the key components of this structure in a neurophysiologically plausible simulation, to provide a quantitative model relating behavior to the actions of networks of neurons. We provide example simulations of four key test situations, showing that the model can account for aspects of representational neglect, as well as spatial updating and mental exploration in familiar environments, and can place cell firing patterns seen in rats with lesions to the head direction system and in normal rats navigating through environments that unexpectedly change shape (Gothard, Skaggs, & McNaughton, 1996). First, we briefly review some of the data at each of these levels of description that motivate the design of the model.

Patrick Byrne and Suzanna Becker, Department of Psychology, Neuroscience and Behaviour, McMaster University, Hamilton, Ontario, Canada; Neil Burgess, Institute for Cognitive Neuroscience and Department of Anatomy, University College London, London, England.

We thank John O'Keefe, Tom Hartley, and Lynn Nadel for useful discussions, and Allen Cheung for pilot simulations. Neil Burgess is supported by the Medical Research Council and Wellcome Trust, United Kingdom, and Suzanna Becker is supported by Natural Sciences and Engineering Research Council, Canada. Code for the model presented herein, along with detailed comments, can be retrieved from http://psycserv.mcmaster.ca/beckerlab/ByrneBeckerBurgessModel/

Correspondence concerning this article should be addressed to Suzanna Becker, Department of Psychology, Neuroscience and Behaviour, McMaster University, Hamilton, Ontario L8S 4K1, Canada. E-mail: becker@mcmaster.ca

Neuronal Representations

Data from electrophysiological recordings in behaving animals provide perhaps the most direct evidence of the nature of the representations at work in spatial cognition. We start with the apparently allocentric representations associated with the mammalian medial temporal lobe. View-invariant hippocampal "place cells" fire selectively for an animal's location in space (e.g., O'Keefe, 1976), but show little dependence on the animal's orientation during random, open field foraging. We refer to this representation as *allocentric*, representing location relative to the environment, even though the location represented is that of the animal itself. In a linear track, place cells tend to be direction specific, however, when the track environment is enriched with place-unique cues, the place cells are much less directionally selective (Battaglia, Sutherland, & McNaughton, 2004). O'Keefe and Nadel (1978) argued that this collection of place-selective neurons forms the basis of a cognitive map and provides the rat's internal allocentric representation of location within the environment. Evidence for the existence of place cells has also been found in the hippocampus in nonhuman primates (Matsumura et al., 1999; Ono, Nakamura, Nishijo, & Eifuku, 1993) and in humans (Ekstrom et al., 2003). The representation of the complementary spatial information—orientation independent of location—has also been found; "head direction cells" (see, e.g., Taube, 1998) are found along an anatomical circuit largely homologous to Papez's circuit (Papez, 1937) leading from the mammillary bodies to the presubiculum via the anterior thalamus. A representation related to place cells has also been found in the parahippocampal and the hippocampal region of both nonhuman (Rolls & O'Mara, 1995) and human primates' (Ekstrom et al., 2003) "view cells," which fire when an animal is looking at a given location from a range of vantage points.

The location of a place cell's response depends on large, extended local landmarks rather than on discrete objects, whereas the orientation of the overall place and head direction representations depend on landmarks at or beyond the reachable environment (see Barry et al., 2006; Burgess & O'Keefe, 1996; Cressant, Muller, & Poucet, 1997). Thus, the location and shape of the firing fields of hippocampal place cells can be explained if it is assumed that their firing is driven by the activity of a population of *boundary vector cells* (BVCs; Hartley, Burgess, Lever, Cacucci, & O'Keefe, 2000; O'Keefe & Burgess, 1996). These neurons, hypothesized to exist within parahippocampal cortex, show maximal firing when an animal is at a given distance and allocentric direction from an environmental landmark or boundary. The direct or indirect reciprocal connectivity of the hippocampal formation and parahippocampal regions with each other and with the perirhinal cortex (for a review, see Burgess et al., 1999), an area that is known to be important for object recognition (Davachi & Goldman-Rakic, 2001; Murray & Bussey, 1999; Norman & Eacott, 2004), probably allows for the positions and identities of landmarks visible at a particular location to be bound to that location.

In parallel to the above allocentric representations, egocentric representations, which are ubiquitous throughout the sensory, motor, and parietal cortices, are clearly directly involved in all aspects of spatial cognition. Sensory representations will be egocentric, reflecting the reference frame of the receptor concerned (e.g., retinotopic in the case of visual input), whereas motor output will reflect the reference frame appropriate for the part of the body to be moved (see, e.g., Georgopoulos, 1988). Coordinating these representations, the posterior parietal cortices are heavily involved in sensorimotor mappings. The posterior parietal cortex is known to contain neurons that respond to stimuli in multiple reference frames, especially areas near or within the intraparietal sulcus. In particular, Galletti, Battaglini, and Fattori (1995) have found neurons in the anterior bank of the parietal–occipital sulcus (V6A) in the ventromedial parietal cortex that represent the positions of visual stimuli in a craniotopic reference frame. Also, area 7a contains neurons that exhibit egocentrically tuned responses that are modulated by variables such as eye position and body orientation (Andersen, Essick, & Siegel, 1985; Snyder, Grieve, Brotchie, & Andersen, 1998). Such coding can allow transformation of locations between reference frames (Pouget & Sejnowski, 1997; Zipser & Andersen, 1988). Furthermore, head direction selective neurons that exhibit responses tuned to various different reference frames have been found in the posterior cortices of the rat (Chen, Lin, Barnes, & McNaughton, 1994). Such properties might allow for the establishment of the angular relationship between different representational frames.

A number of single-unit recording studies have shown that areas of the primate posterior parietal cortex, again in and around the intraparietal sulcus, contain neurons that exhibit firing patterns modulated by various combinations of head position, velocity, acceleration, and visual stimuli (Andersen, Shenoy, Snyder, Bradley, & Crowell, 1999; Bremmer, Klam, Duhamel, Hamed, & Graf, 2002; Klam & Graf, 2003). The nature of these interactions appears to be complex, but Bremmer et al. (2002) suggested that this idiothetic modulation of parietal neuron firing might be related to object tracking during self-motion. This argument is indirectly supported by Duhamel, Colby, and Goldberg (1992), who have shown that eye movements that bring the location of a previously flashed stimulus into the receptive field of a parietal neuron elicit a response from that neuron, even though the stimulus is no longer present (see also Colby, 1999). Area 7a is the part of the parietal cortex most strongly connected with the medial temporal lobe, including efferent projections into the parahippocampus, presubiculum, and CA1 (Ding, Van Hoesen, & Rockland, 2000; Rockland & Van Hoesen, 1999; Suzuki & Amaral, 1994) and afferent connections from entorhinal cortex and CA1 (Clower, West, Lynch, & Strick, 2001). In addition, single-unit recordings from monkey dorsolateral prefrontal and posterior parietal cortices suggest that spatial working memory is, indeed, egocentric in nature (Chafee & Goldman-Rakic, 1998; Funahashi, Bruce, & Goldman-Rakic, 1989).

Finally, some hints of the temporal dynamics of neural processing during navigation come from the observation that the theta rhythm (i.e., 4–12 Hz) of the electroencephalogram invariably accompanies voluntary displacement motion of the rat (O'Keefe & Nadel, 1978). In addition, the phase of firing of place cells correlates strongly with the rat's location within the firing field (O'Keefe & Reece, 1993) and independently of firing rate or running speed (Huxter, Burgess, & O'Keefe, 2003). Recent results indicate a possible role for theta in human navigation (Caplan et al., 2003; Kahana, Sekuler, Caplan, Kirschen, & Madsen, 1999), and several experiments indicate a role for theta phase (e.g., Pavlides, Greenstein, Grudman, & Winson, 1988) in modulating hippocampal synaptic plasticity and theta power (Sederberg et al.,

2003) or theta coherence between hippocampus and nearby neocortical areas (Fell et al., 2003) in modulating encoding into memory.

Lesions, Neuropsychology, and Functional Neuroimaging

The medial temporal lobes, and hippocampus in particular, have long been known to be crucial for long-term memory (Eichenbaum & Cohen, 1988; Scoville & Milner, 1957; Squire, 1986), together with other elements of Papez's circuit (Aggleton & Brown, 1999). Within the spatial domain, neuropsychological studies have left little doubt that the medial temporal lobe, particularly in the right hemisphere, is critical for remembering the locations of several objects within a visual scene over a significant delay (Crane & Milner, 2005; Piggott & Milner, 1993; Smith & Milner, 1989). Within a broader memory deficit, hippocampal damage seems to specifically impair performance in tasks likely to require allocentric representations of location or representations that can be flexibly accessed from novel points of view rather than being directly solved by use of egocentric representations. For example, where locations must be remembered from a different point of view to presentation, performance is impaired relative to location memory from the same view even over short timescales (Abrahams, Pickering, Polkey, & Morris, 1997; Hartley et al., 2007; Holdstock et al., 2000; King et al., 2002). More generally, accurate spatial navigation to an unmarked goal location is impaired by hippocampal damage in rats (e.g., Jarrard, 1993; Morris, Garrard, Rawlins, & O'Keefe, 1982) and in humans (Bohbot et al., 1998; Maguire, Burke, Phillips, & Staunton, 1996; Spiers et al., 2001). Human neuroimaging studies also show involvement of the hippocampus in accurate navigation (Hartley, Maguire, Spiers, & Burgess 2003; Iaria et al., 2003; Maguire et al., 1998). Additionally, neuroimaging of the perceptual processing of spatial scenes, including plain walled environments, implicates the parahippocampal cortex (Epstein & Kanwisher, 1998), a region associated with landmark recognition (Aguirre & D'Esposito, 1999) and navigation (Bohbot et al., 1998). See Burgess, Maguire, and O'Keefe (2002) for a review.

Human neuropsychology has long recognized the parietal lobes as playing a major role in spatial cognition. Parietal damage leads to deficits in sensorimotor coordination such as optic ataxia, deficits in spatial manipulation such as mental rotation, and deficits in spatial working memory (see, e.g., Burgess et al., 1999; Haarmeier, Thier, Repnow, & Petersen, 1997; Karnath, Dick, & Konczak, 1997). Visual processing in the temporal and parietal lobes has been generally characterized respectively in terms of the dorsal and ventral "what and where" (Ungerleider & Mishkin, 1982) or "what and how" (Goodale & Milner, 1992) processing streams. The parietal region in the dorsal stream is concerned with representing the locations of stimuli in the various egocentric reference frames appropriate to sensory perception and motor action and translating between these frames to enable sensorimotor coordination. In contrast, the occipital and temporal visual regions in the ventral stream are concerned with visual perceptual processes related to object recognition, see neuronal representations above.

Unilateral damage to the parietal cortex (most often on the right) and surrounding areas commonly results in the syndrome of hemispatial neglect: a reduced awareness of stimuli and sensations on the contralateral side of space (*perceptual neglect*). Of particular interest here is the phenomenon of *representational neglect*—a lack of awareness of the contralateral side of internal representations derived from memory. In the classic demonstration (Bisiach & Luzzatti, 1978), patients were asked to imagine the Piazza del Duomo in Milan (with which they were very familiar) and to describe the scene from two opposite points of view. Buildings to the left of the given point of view (e.g., facing the Cathedral) were neglected, but those same buildings were described when given the opposite point of view (e.g., facing away from the Cathedral), indicating intact long-term memory of the entire Piazza, despite neglect of the left of each imagined scene. Perceptual and representational neglect depend, at least in part, on different neural systems and can be dissociated, even within the same patient (Beschin, Basso, & Della Sala, 2000). It is interesting that representational, but not perceptual, neglect is associated with impaired navigation to an unmarked location (Guariglia, Piccardi, Iaria, Nico, & Pizzamiglio, 2005). Consistent with these findings of parietal involvement in imagery, neuroimaging experiments have shown heightened activity within the precuneus (i.e., medial parietal cortex) during mental imagery (e.g., Fletcher, Shallice, Frith, Frackowiak, & Dolan, 1996) and visuospatial working memory (e.g., Wallentin, Roepstorff, Glover, & Burgess, 2006). Transcranial magnetic stimulation and fMRI studies also indicate that areas surrounding the right intraparietal sulcus, including areas 7a and 40, are essential in the generation and manipulation of egocentric mental imagery (Formisano et al., 2002; Knauff, Kassubek, Mulack, & Greenlee, 2000; Sack et al., 2002).

Behavioral and single-unit studies indicate that memory for locations in general, and the place cell representation of location in particular, is automatically updated by self-motion, a process more generally known as *path integration* or *spatial updating* (see below). This process may reflect an interaction between the parietal and hippocampal systems, as the parietal cortex appears to be centrally involved (Alyan & McNaughton, 1999; Commins, Gemmel, Anderson, Gigg, & O'Mara, 1999; Save, Guazzelli, & Poucet, 2001; Save & Moghaddam, 1996). For example, Save, Paz-Villagran, Alexinsky, and Poucet (2005) have shown that lesions to the associative parietal cortex of rats result in altered place cell firing, suggesting that egocentric sensory information must travel through the parietal cortex in order to elicit appropriate place cell firing. This is consistent with a number of experiments that demonstrate that mental exploration/navigation depends on the posterior parietal and extrahippocampal medial temporal regions in primates and on homologous regions in the rodent brain (Ghaem et al., 1997; Pinto-Hamuy, Montero, & Torrealba, 2004). The interaction between the parietal and medial temporal areas likely involves the retrosplenial cortex, lesions of which selectively disrupt path integration (Cooper, Manka, & Mizumori, 2001), and the parietal–occipital sulcus, which has been associated with topographical disorientation (Ino et al., 2002) and cells coding for locations in space (Galletti et al., 1995).

Prefrontal regions, as well as parietal ones, are implicated in spatial working memory, with parietal areas predominantly associated with storage and prefrontal areas with the application of control processes, such as active maintenance or planning (Shallice, 1988; R. Levy & Goldman-Rakic 2000; Oliveri et al., 2001), the use posterior spatial representations. Thus, fMRI studies have shown activation in both of these areas when subjects were re-

quired to remember the locations of various objects for short periods of time (Galati et al., 2000; Sala, Rämä, & Courtney, 2003). Manipulations of working memory may also involve making or planning eye movements in order to direct attention to spatial locations in imagery. In support of this notion, voluntary eye movements disrupt spatial working memory (Postle, Idzikowski, Della Sala, Logie, & Baddeley, 2006), whereas left-hemispatial neglect patients show abnormal eye movements that deviate about 30° rightward during visual search (Behrmann, Watt, Black, & Barton, 1997) as well as while at rest (Fruhmann-Berger & Karnath, 2006). Moreover, adapting prisms that shift the neglected visual field toward the good side of space, which would compensate for a rightward bias in gaze direction, ameliorate both perceptual and representational neglect (Rode, Rosetti, & Boisson, 2001). Studies involving mental navigation and route planning consistently find elevated activation in frontal regions, especially on the left side (Ghaem et al., 1997; Ino et al., 2002; Maguire et al., 1998). For example, Maguire et al. (1998) found additional activation in the left prefrontal cortex associated with the planning of detours when subjects were navigating in a familiar virtual town in which the most obvious route had suddenly been blocked. This suggests that left prefrontal areas contribute to route planning, perhaps guiding egocentric mental imagery within the temporal–parietal systems activated by the basic navigation condition.

Cognitive Psychology

Given the electrophysiological and lesion evidence for parallel egocentric and allocentric representations of location, we next consider converging evidence from cognitive psychology in which one, the other, or both may contribute to behavior. Simons and Wang (1998; Wang & Simons, 1999) performed an elegant series of experiments in which subjects were required to remember an array of objects presented on a circular table. During the delay period preceding the memory test, the table would either remain stationary or rotate through a fixed angle. At the same time, the subject would either remain stationary or walk through the same angle around the table. Thus, the test stimuli could be aligned with the studied view, with a rotated view consistent with the subject's motion, with both, or with neither. Subjects' performance on a memory task (detecting which object had moved) provided evidence for the use of both (a) a visual-snapshot representation of the presented array, and (b) an egocentric representation that is updated to accommodate self-motion by showing an advantage whenever the test array was aligned with either representation. The latter *spatial updating* ability (Rieser, 1989) can be thought of as a generalization of path integration, allowing an organism to keep track of several locations, including its origin of motion during real or imagined navigation in the absence of visual cues. The results suggest that both of types of representation exist in the brain. It is interesting to note that evidence suggests that allocentric representations of object locations (i.e., relative to visual landmarks external to the array) are also used in this type of experiment, as shown by a subsequent study incorporating a rotatable landmark (Burgess, Spiers, & Paleologou, 2004). Parallel influences of egocentric and allocentric representations are also indicated by human search patterns within deformable virtual reality environments (Hartley, Trinkler, & Burgess, 2004). In these experiments, the locus of search can be predicted by a model based on the firing of hippocampal place cells, indicating allocentric processing of location. However, subjects also tended to adopt the same orientation at retrieval as at encoding, indicating egocentric processing of orientation.

Further evidence for the use of both egocentric and allocentric representations of space can be found in reaction time data from a number of experiments involving the recognition/recall of previously presented object configurations from novel viewpoints. Diwadkar and McNamara (1997) had subjects learn the locations of objects on a desktop from a number of viewpoints before taking part in a recognition test. When presented with a novel view of the same or a different object configuration, subjects' reaction time was found to vary linearly with the angular distance between the observed view and the closest trained view. Related results were found when blindfolded subjects had to point to where a given object would be from a specific imagined viewpoint: Accuracy and/or reaction time reflected the distance and angle between the studied viewpoint and the imagined viewpoint (Easton & Sholl, 1995; Rieser, 1989; Shelton & McNamara, 2001). These results are consistent with spatial updating of an egocentric representation. However, the additional use of allocentric representations in these tasks is indicated by improved performance for viewpoints aligned with the walls of the room or the sequence of learning (Mou & McNamara, 2002), with external landmarks (McNamara, Rump, & Werner, 2003), and with the absence of a relationship to distance or angle for objects configured into a regularly structured array (Easton & Sholl, 1995; Rieser, 1989). In possibly related findings, Wang and Spelke (2000) suggested that the high variance of the error in pointing to different objects after blindfolded disorientation indicates independent egocentric representations for the location of each object. In the same experiment, the lower variance in errors when pointing to features of the testing room indicated a single coherent (allocentric) representation for the layout of the room. Similarly, judgments of relative direction between objects from an imagined location at a third object do not increase in variance with disorientation, indicating use of a more coherent representation in this task than that used for egocentric pointing (Waller & Hodgson, 2006). See Burgess (2006) for further discussion.

Theoretical Analyses

It has been proposed (e.g., Milner, Paulignan, Dijkerman, Michel, & Jeannerod, 1999) that the relative contribution of egocentric and allocentric representations to spatial memory depends on the timescale of the task concerned. Short-term retention of perceptual information for the purpose of immediate action will be best served by egocentric representations appropriate to the corresponding sensory and motor systems. By contrast, long-term memory for locations will be best served by allocentric representations (i.e., relative to stable landmarks) because the location and configuration of the body at retrieval typically will be unrelated to that at encoding (see Burgess, Becker, King, & O'Keefe, 2001, for further discussion). This observation is consistent with the evidence for the role of the parietal and prefrontal areas in supporting egocentric representations and short-term memory and the role of medial temporal lobe areas in supporting allocentric representations and long-term memory, reviewed above.

For intermediate timescales (e.g., tens of seconds), it may be possible to relate the configuration of the body at retrieval to that at encoding via the egocentric process of path integration or spatial updating referred to above. Pierrot-Deseilligny, Müri, Rivaud-Pechous, Gaymard, and Ploner (2002) reviewed evidence suggesting that spatial memory may have at least three important timescales. For the first approximately 20 s, they claim that a frontal–parietal spatial working memory system is the dominant mechanism, followed for approximately 5 min by a medium-term, parahippocampally dependent memory system, and finally by a hippocampally dependent long-term memory system that operates only after delays of several minutes. Spatial scale might also be a factor in determining which representations are used. For example, in mammals, path integration becomes unreliable over long or convoluted paths (see, e.g., Etienne, Maurer, & Seguinot, 1996), whereas egocentric parietal and premotor representations may be preferentially recruited for representations of locations in "peripersonal" space that can be directly acted upon (e.g., Duhamel, Colby, & Goldberg, 1998; Goodale & Milner, 1992; Graziano & Gross, 1993; Ladavas, di Pellegrino, Farne, & Zeloni, 1998).

Along the above lines, Mou, McNamara, Valiquette, and Rump (2004) proposed a transient egocentric representation of object locations for immediate action and an allocentric representation of the environment, including the subject's own location, for actions supported by information from long-term memory. On the basis of the experiments probing memory for object location as a function of differences between the studied, imagined, and actual views, they argued that two types of spatial updating occur: spatial updating of egocentric representations of object locations, and spatial updating of the subject's own location in the environmental representation. A related proposal suggested transient egocentric representations of single objects in parallel with a more coherent enduring representation (Waller & Hodgson, 2006). (For a discussion of the neural mechanisms supporting the integration of self-motion and sensory information, see Guazzelli, Bota, & Arbib, 2001; Redish, 1999.)

In summary, evidence from psychology and neuroscience indicates that spatial cognition involves multiple parallel frames of reference, with short-term/small-scale tasks more likely to recruit egocentric representations and long-term/large-scale tasks more likely to recruit additional allocentric representations. However, this proposed division of labor involving different reference frames is neither absolute nor uncontroversial. Thus, Wang and Brockmole (2003) have also argued that even long-term spatial memory is purely egocentric. They found the current view to influence the ability of students to point to an occluded but very familiar landmark on the campus. Conversely, even short-term memory can be shown to depend on the hippocampus when the viewpoint is changed between study and test (King et al., 2002, 2004; Hartley et al., 2007) and on allocentric representations when landmarks are parametrically manipulated (Burgess et al., 2004); see Burgess (2006) for further discussion.

The Model: Overview

From the forgoing discussion, it appears that mammalian spatial memory can make use of both egocentric and allocentric representations in parallel, depending on the nature of the task. We now propose a model of spatial cognition that accounts for the interaction between long- and short-term memory processes in encoding, retrieval, imagery, and planning. The model addresses data at multiple levels of analysis, from single-unit recordings to large-scale brain systems to behavior, and the relative roles played by egocentric and allocentric representations and by visual and idiothetic inputs. We first provide a brief overview of the functional architecture of our model, with further details of its implementation given in the next section and fully elaborated in the Appendix.

In our model, long-term spatial memory formation involves the generation of allocentric representations in the hippocampus and surrounding medial temporal lobe structures (perirhinal and parahippocampal cortices). The hippocampal place cell representation is driven by convergent inputs from the dorsal and ventral visual pathways. The ventral stream input consists of object features in the perirhinal cortex, whereas the dorsal stream input consists of BVCs in the parahippocampal cortex. These medial temporal lobe areas are all mutually interconnected to permit pattern completion. When cued with a partial representation of a place, such as a specific landmark, the model thereby automatically retrieves the full representation of that place, comprising the location of the observer as well as the surrounding landmarks and their visual appearance.

Both short-term spatial memory and imagery are modeled as egocentric representations of locations in the precuneus, which can be driven by perception or by reconstruction from long-term memory (see below). The neural activations within this medial parietal representation can be modulated by directed attention, to capture the fact that one can attend sequentially to the spatial locations of items in imagery just as in perception, presumably via planned eye movements (see Postle et al., 2006). Both encoding and retrieval require translation between the egocentric precuneus and allocentric parahippocampal representations of landmarks. This occurs via a coordinate transformation mediated by the posterior parietal and retrosplenial cortices, reflecting the current head direction.

Retrieval from long-term memory, cued by knowledge of position and orientation relative to one or more landmarks, corresponds to pattern completion of the parahippocampal representation of the allocentric locations of landmarks around the subjects via its connections with the hippocampal and perirhinal representations. Thus, the medial temporal lobe acts as an attractor network within which a representation of the visual features, distances, and allocentric directions of landmarks can be retrieved, which is consistent with perception from a single location (represented in the hippocampus). This representation is translated into the egocentric precuneus representation, within which directed attention can boost the activation of egocentrically defined locations of interest. Finally, the additional activation can feed back to the parahippocampal representation, again via posterior parietal translation, and thence to the perirhinal representation so as to activate the visual features of the attended landmark.

Motor efference drives the spatial updating of the egocentric representation of the locations of landmarks. Specifically, modulation of the posterior parietal egocentric–allocentric transformation by motor efference causes allocentric locations to be mapped to the egocentric locations pertaining after the current segment of movement. The reactivation of the BVCs by this shifted egocentric representation then updates the medial temporal representation to be consistent with the parietal representation. The *bottom up*

(parietal to temporal) and *top down* (temporal to parietal) flows of information are temporally organized into different phases of the theta rhythm. Additionally, the generation of mock motor efference in the prefrontal cortex allows mental exploration in imagery via mock spatial updating.

A central component of our model is circuitry that transforms between different representations of the space surrounding an animal. This proposed egocentric–allocentric transformation suggests a solution to two puzzles regarding the functional anatomy of memory and navigation. The first is the observation that Papez's circuit (including the mammillary bodies, anterior thalamus, retrosplenial cortex and fornix, as well as the hippocampus) is both crucial for episodic recollection, which is impaired by lesions anywhere along it (see, e.g., Aggleton & Brown, 1999), and provides the neural basis for head direction cells (Taube, 1998). A second, related puzzle is the ubiquitous involvement of retrosplenial cortex and the anterior parietal–occipital sulcus in both navigation (reviewed in Maguire, 2001) and memory (see, e.g., Burgess, Maguire, Spiers, & O'Keefe, 2001). We propose (see also Burgess, Becker, et al., 2001; Burgess, Maguire, et al., 2001) that the segment of Papez's circuit from the mammillary bodies to the hippocampal formation via the anterior thalamus carries the head direction information needed to transform the allocentric directional tuning of the BVC representation into an egocentric (head-centered) representation suitable for mental imagery and that the retrosplenial cortex/parietal–occipital sulcus may mediate or buffer the stages of transformation between egocentric and allocentric representation (see also Ino et al., 2002). A related proposal is that the retrosplenial cortex serves to integrate mnemonic and path-integrative information (Cooper & Mizumori, 2001), which maps onto our own proposal given the assumption of allocentric long-term memory and egocentric spatial updating.

The Model: Architecture and Dynamics

In this section, we discuss the architecture of our model and then describe the model dynamics and how spatial updating, mental exploration, and learning are simulated. A simplified version of our model with preliminary simulation results was described by Becker and Burgess (2001). By lesioning the parietal region of the model, the authors were able to simulate aspects of hemispatial neglect. The model presented here builds on this earlier work by deriving, in a more principled manner, the neural circuits for allocentric representation and allocentric–egocentric transformations, and augments this work with parietal neural circuitry to support spatial updating and mental navigation. The architecture of our model rests upon three key assumptions:

1. The parietal window hypothesis: An egocentric window provides exclusive access into long-term spatial memory in the service of mental imagery, planning, and navigation.

2. Allocentric coding in the medial temporal lobe: Allocentric BVC representations are constructed in the parahippocampal region and project to hippocampal place cells where long-term spatial memories are stored.

3. Transformation circuit: Access by the parietal window into allocentrically stored spatial representations is mediated by a transformation circuit; the same circuit also operates in the inverse direction, such that the products of recall are mapped from allocentric into egocentric representations of space.

The Parietal Window Hypothesis

We hypothesized that a population of neurons maintains a head-centered, egocentric map of space that can be driven either by bottom-up sensory input or by top-down inputs from long-term memory. This map represents the locations of all landmarks/objects that are visible from an animal's current location in space or from a location that the animal recalls from previous experience. This neuronal population, assumed to exist within the posterior parietal cortex and very likely within the precuneus, will henceforth be referred to as the *parietal window*. We claim that the contents of the parietal window are generated on the basis of some combination of information from the senses (e.g., dorsal visual stream) and from allocentric long-term spatial memory, with the exact combination depending on the demands of the current task. Manipulation of spatial information for the purposes of planning or navigation, including spatial updating, occurs within the parietal window.

The network model also includes circuitry that can manipulate the contents of the parietal window so as to allow for spatial updating or mental exploration. In the case of spatial updating, this circuitry is activated by idiothetic information (proprioceptive cues signaling the observer's change in direction and location), whereas in the case of mental exploration, it is activated by some mentally generated equivalent (e.g., imagined rotation and translation during path planning). The former ability allows the model to maintain an internal representation of its surroundings even with degraded or absent sensory input, whereas the latter provides a means of recalling the locations of occluded landmarks and generating navigational strategies for reaching them.

Allocentric Representations in the Medial Temporal Lobe

In contrast to the parietal window's egocentric frame of reference, we postulate that an allocentric frame of reference is used in the medial temporal lobe. The model's egocentric reference frame has its origin bound to the observer's location, with its *y*-axis fixed along the observer's heading direction. The model's allocentric reference frame has its origin bound to the observer's location (in this sense, like place cell firing, it is not fully allocentric), but its orientation is fixed relative to the external environment. Therefore, both reference frames are similar in that they remain fixed with respect to the observer so long as the observer undergoes translational motion only. However, when the observer's head rotates within the environment, the egocentric frame rotates with it, but the allocentric frame remains stationary with respect to the environment. An example of an object in the allocentric frame and its corresponding location in the egocentric frame is shown in Figure 1.

Consider the situation depicted in Figure 2 in which an observer surrounded by six walls is located at the position marked "X," with a heading direction indicated by the arrow. If the walls of this "two-room" environment are discretized uniformly into a set of "landmark segments" (to simplify later calculations), then the egocentric frame positions of the segments viewable from "X" can be inferred readily. These positions are depicted by open circles in the top panel of Figure 3. Representation of this egocentric information by the parietal window neurons is accomplished by first forming a one-to-one correspondence between the set of neurons

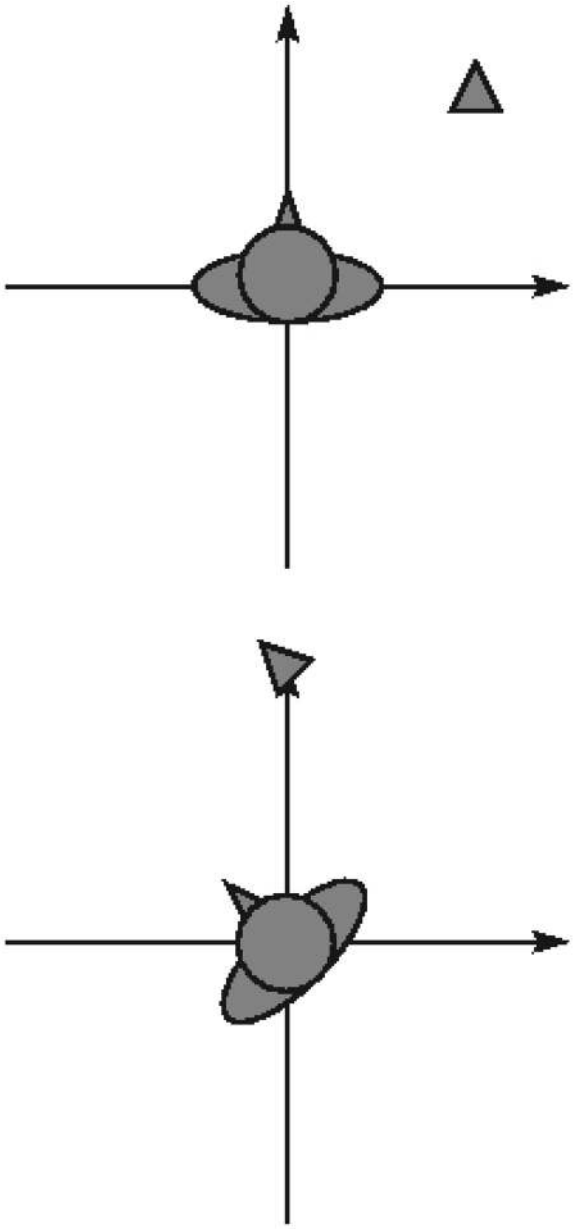

Figure 1. Top: Egocentric reference frame in which the observer is always at the origin, facing along the positive *y*-axis. A triangular landmark sits in front and to the right of the observer in this frame. Bottom: The same situation as above, but depicted in the allocentrically aligned reference frame. In this frame, the observer is always at the origin, but the direction of the *y*-axis is fixed to the external environment instead of the observer's heading direction. With the heading direction depicted (approximately 45° away from the positive *y*-axis in the counterclockwise direction), the triangular landmark lies directly on the positive *y*-axis and is rotated 45° in the counterclockwise direction.

and a polar grid covering the egocentric reference frame. This grid is depicted by the closed circles in the top panel of Figure 3. Each neuron in the grid is tuned to respond most strongly to an object or landmark at a particular direction and distance relative to the organism's head, which is at the origin of the grid. The neuron's response falls off exponentially for objects located further away from the neuron's preferred distance and direction (see the Appendix for details). When multiple segments are present within a neuron's receptive field, they contribute additively to its firing rate, up to a maximum firing rate of 1. The parietal window representation of the information depicted in the top panel of Figure 3 is shown in the bottom panel of the same figure, where the firing rate of each neuron is plotted at the location of its corresponding grid point.

We assume that the observer in Figure 2 aligns its allocentric frame such that the *y*-axis is perpendicular to the wall labeled *1* and the *x*-axis is parallel to the same wall. The locations of the landmark segments in this frame, which will not depend on the observer's heading direction, are depicted in the top panel of Figure 4. By forming a one-to-one correspondence between a set of neurons and a polar grid centered at the origin of the allocentric reference frame, it becomes possible to represent the configuration of landmark segments by the firing rates of this neural population. In analogy with the egocentric parietal window neurons, each allocentric neuron in the grid is tuned to respond most strongly to an object or landmark at a particular distance from the organism's head, which is fixed to the origin of the grid and *allocentric* direction (relative to the fixed environment). Again, the neuron's response falls off exponentially for objects located farther away from the neuron's preferred distance and direction. Note that these allocentrically tuned neurons are essentially the same as the BVCs described in the introduction and are referred to as such from this point on. The BVC representation of the information depicted in the top panel of Figure 4 is shown in the bottom panel of the same figure, where the firing rate of each neuron is plotted at the location of its corresponding grid point. Although we assume that these BVCs exist within the parahippocampal cortex, we note that cells with BVC-like responses have been found in the subiculum (Barry et al., 2006; Sharp, 1999), an alternative location to the parahippocampal cortex but one that is less consistent with neu-

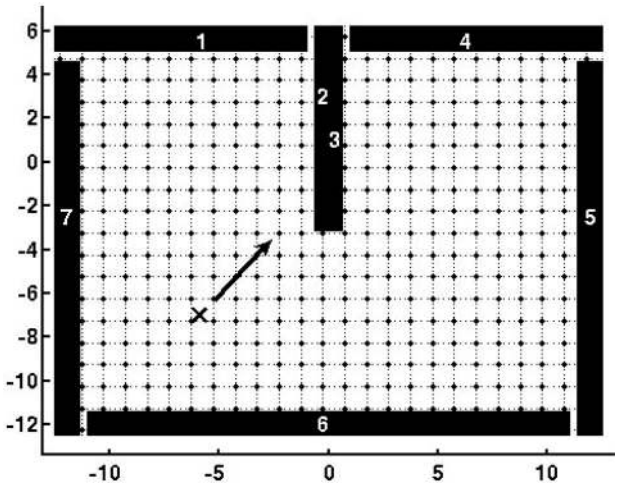

Figure 2. Map of the "two-room" environment used in the second set of simulations. Solid rectangles represent environmental boundaries/landmarks. Each grid point corresponds to a maximal firing location for one hippocampal place cell. The "X" represents the model's current location and the arrow its heading direction.

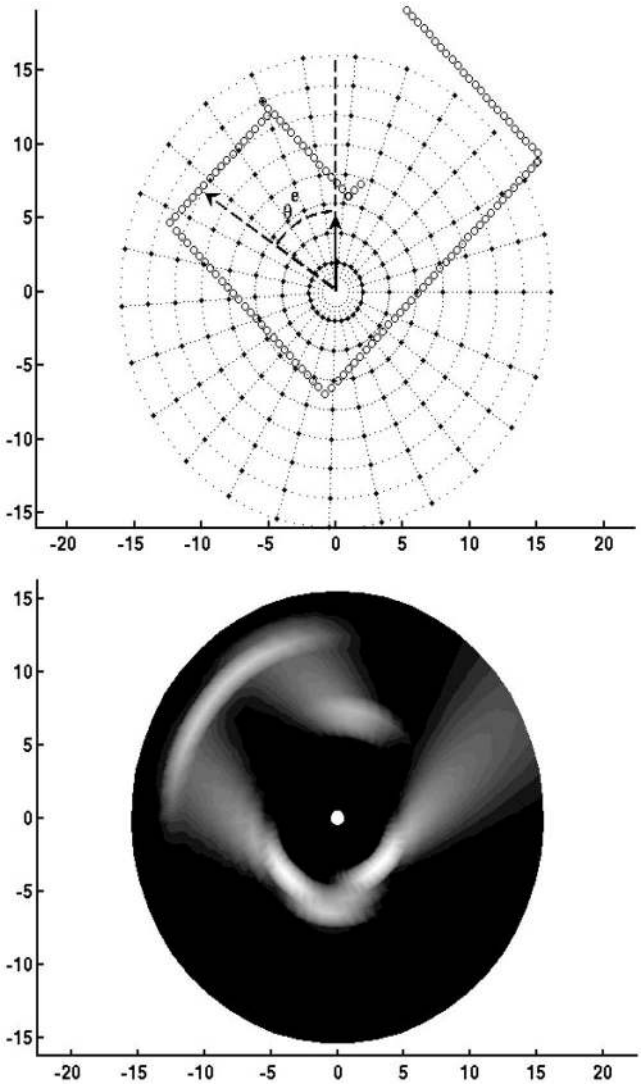

Figure 3. Top: Egocentric reference frame. Each grid point corresponds to the preferred boundary/landmark location of a parietal window neuron, which fires maximally when a landmark segment is located at that grid point's coordinates. The landmark segments for the discretized "two-room" environment, as viewed from the model's current location, are also shown. The landmark segment at egocentric direction, θ^e, is indicated by the dashed arrow. Finally, the model's heading direction, which is always the same in egocentric space, is indicated by the solid arrow. Bottom: Activation of parietal window neurons corresponding to the landmark segment configuration. The firing rate of each neuron is plotted at that neuron's corresponding grid point, with lighter shades indicating higher firing rate.

roimaging results in humans showing parahippocampal processing of spatial scenes including plain walled environments (Epstein & Kanwisher, 1998).

To form long-term memories for specific spatial locations, spatial input from BVCs and visual input from the perirhinal layer are combined into a place cell representation. Although, in reality the hippocampal formation consists of multiple spatially selective regions (dentate gyrus, CA3, CA1), for simplicity, our model hippocampus contains a single layer of recurrently connected place cells. Their place preferences are arranged uniformly over a Cartesian grid that covers the relevant allocentric space for a given environment (see Figure 2). In particular, a one-to-one correspondence is formed between each of the model place cells and the set of grid points so that a given place cell fires maximally when the model is located at that cell's corresponding grid point. These model hippocampal neurons are reciprocally connected to the layer of BVCs and to a layer of perirhinal identity neurons, thus allow-

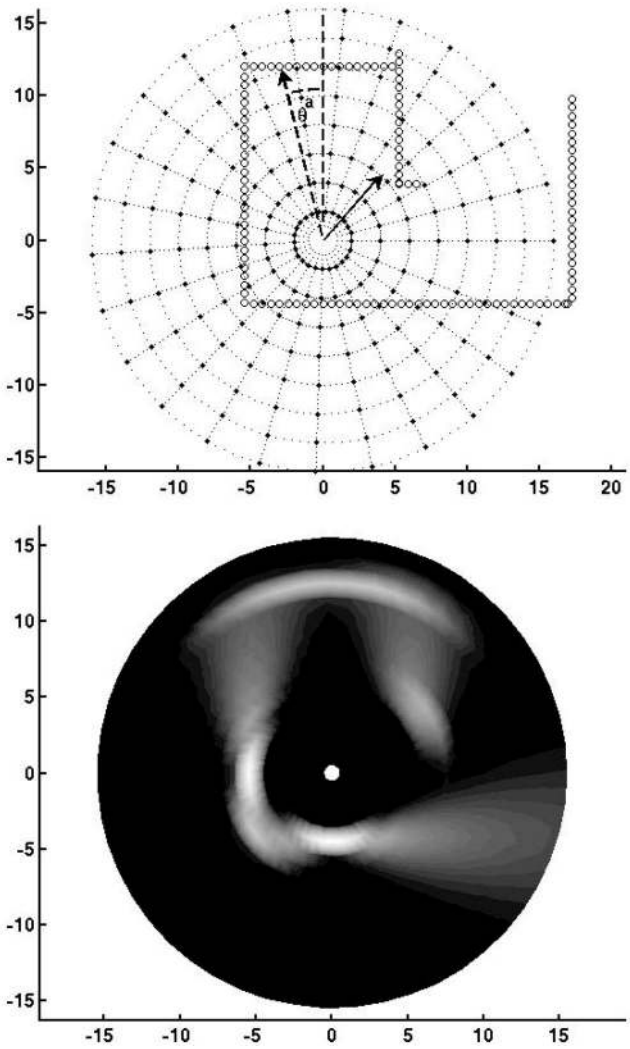

Figure 4. Top: Allocentric reference frame. Each grid point corresponds to the preferred boundary/landmark location of a BVC, which fires maximally when a landmark segment is located at that grid point's coordinates. The landmark segments for the discretized "two-room" environment, as viewed from the model's current location, are also shown. The dashed vector points to the landmark segment at egocentric direction. In this map, it is located at the same distance from the model, but its direction, θ^a, is equal to θ^e plus the model's current heading direction. Finally, the model's heading direction within the allocentric reference frame is indicated by the solid arrow. Bottom: Activation of BVCs corresponding to the landmark segment configuration. The firing rate of each neuron is plotted at that neuron's corresponding grid point, with lighter color indicating higher firing rate.

ing environmental geometry and landmark identities to be bound simultaneously to a given "place." In addition, the layer of BVCs is reciprocally connected to the layer of perirhinal neurons, thereby allowing the association of landmark identities with allocentric locations (see Figure 5 for a schematic of the full model). The full reciprocal connectivity between the three medial temporal lobe components of the model allows for the recall of a landmark's identity when attention is directed toward the parietal window representation of that landmark's location. This process of recall is described in the next section.

Within our gross simplification of hippocampal circuitry, the model's single layer of place cells is most consistent with area CA3, an area that is heavily recurrently connected, and that exhibits place-selective firing. In our model, this recurrent connectivity allows for recall/pattern completion, as it is often argued to do in CA3 (Brun et al., 2002; Nakazawa et al., 2002). Another gross simplification in our model is the strictly spatial function of the hippocampus. Although the hippocampus is known to be important in spatial memory, its more general contribution to episodic memory is well established (for a review, see Burgess, Maguire, & O'Keefe 2002).

Transformation Circuit

The assumption in our model of egocentric access to allocentrically stored spatial information has an important implication: There must be circuitry that transforms between these representations. In order to be able to recall the locations and identities of environmental boundaries relative to one's own location and orientation, long-term allocentric internal representations of space must be transformed into egocentric representations. Conversely, in order for sensory input to cue such recall, or for it to enter long-term allocentric storage in the first place, the inverse transformation from egocentric to allocentric representation must be performed. That is, a visual stimulus at a retinocentrically encoded location must be transformed into an allocentrically encoded location in order to match against or store within spatial long-term memory. We assume that sensory information is first transformed into the head-centered egocentric parietal window reference frame and then to the allocentric BVC representation. The transformation from the parietal window representation to the BVC representation, and its inverse, can be accomplished very simply if absolute heading direction is known. Consider, for example, that you are facing west (90° in allocentric angular coordinates, where north is

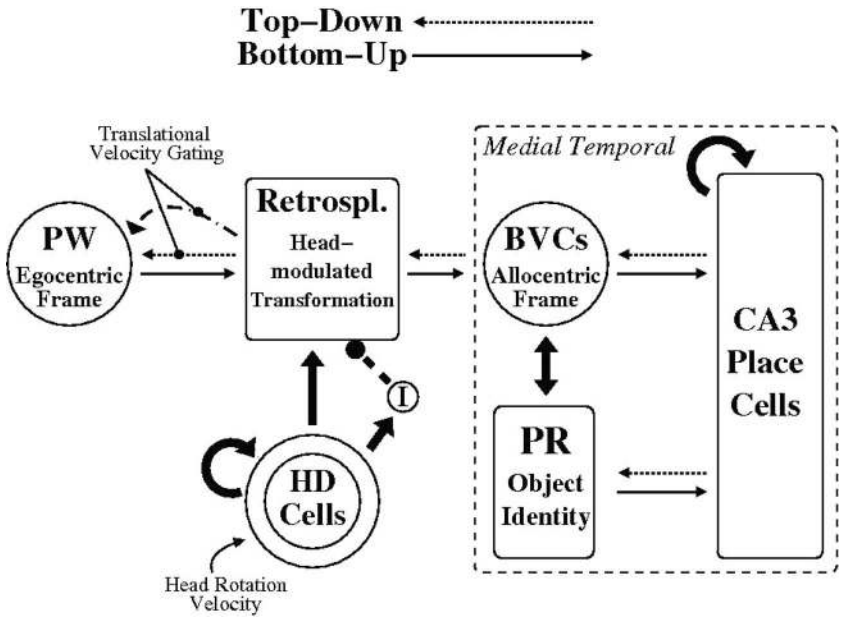

Figure 5. Schematic of the model. Each box or oval represents a set of neurons in a different brain region. Thin, solid arrows represent full bottom-up interconnectivity between the neurons in the connected regions, whereas the dashed arrows represent full top-down interconnectivity. Thick, solid arrows represent full connectivity, which is unaffected by the bottom-up/top-down cycling. The thick dashed line from the inhibitory interneuron population (I) represents inhibition that is unaffected by the bottom-up/top-down phases. A given perirhinal (PR) neuron fires maximally when the model attends to a landmark segment with a particular identity. Hippocampal neurons are associated with a Cartesian grid covering allocentric space such that a given neuron fires maximally when the model is localized at its corresponding grid point. Boundary vector cells (BVCs) or parietal window (PW) neurons are associated with a polar grid covering allocentric/egocentric space. A given BVC/PW neuron fires maximally when a landmark segment is a certain distance and allocentric/egocentric direction away from the model. A given head direction (HD) neuron fires maximally for a given head direction. The transformation layer neurons are responsible for transforming allocentric BVC representations of space into egocentric PW representations. A second set of top-down weights (curved, dashed arrow) from the transformation layer to PW are gated by egocentric velocity signals to allow for spatial updating/mental exploration. Retrospl. = retrosplenial transformation layer.

0°), and there is an object to your left (90° in egocentric angular coordinates, where straight ahead is 0°); the object's allocentric direction can be calculated simply by adding the heading direction to the object's egocentric direction to obtain 180°—similarly, if the object is known to be located to the south (an allocentric angle of 180°) then its egocentric direction can be calculated by subtracting the heading direction from the object's allocentric direction. Thus, in our model the egocentric–allocentric transformations are mediated by input from head direction cells that provide the necessary modulation of firing rates by head direction (Snyder et al., 1998), and the same neural circuitry can then perform the transformation in either direction. The computation is a bit more complicated than a simple subtraction or addition of angles because angular directions are encoded across populations of narrowly direction-tuned neurons; nonetheless, it can be accomplished in a single layer of neurons whose activities are nonlinearly modulated by head direction (cf. Pouget & Sejnowski, 1997). See Figure 6 for a schematic of the full transformation circuit.

When an animal first enters a new environment, we assume that salient perceptual features reliably orient the head direction system. We model the head direction system as a set of neurons configured in a ring via lateral connections to behave as a one-dimensional continuous attractor, as in previous models (e.g., Skaggs, Knierim, Kudrimoti, & McNaughton, 1995; Stringer,

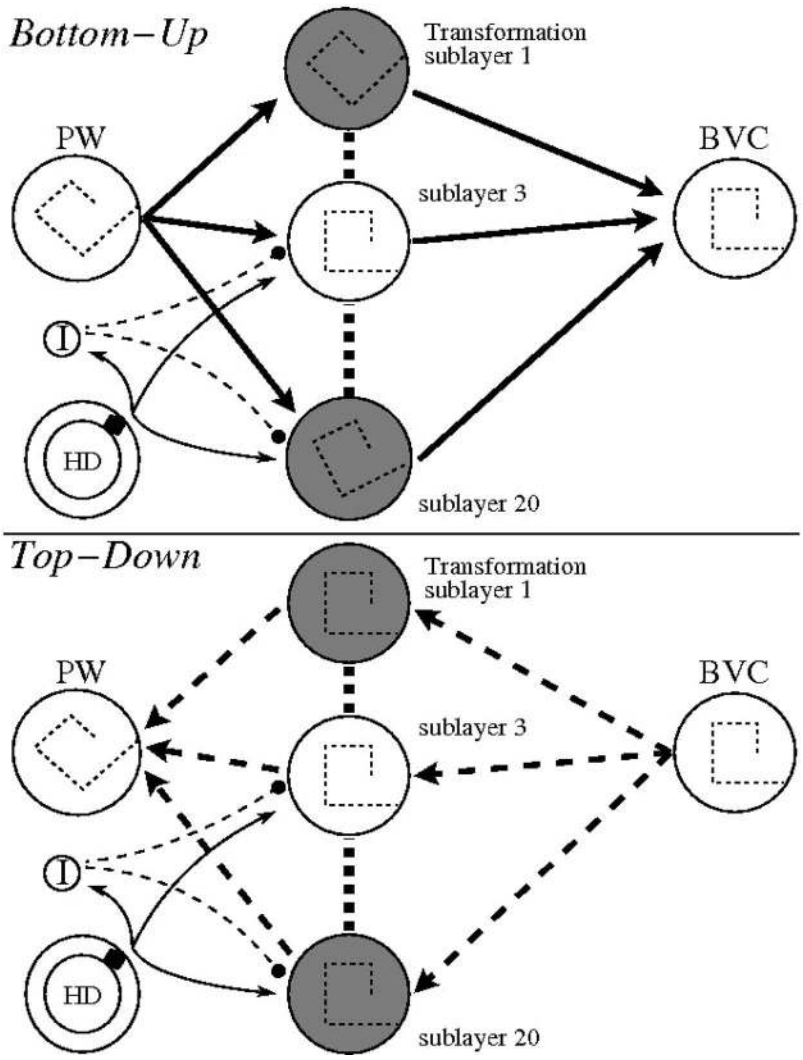

Figure 6. Top: Transformation circuit in bottom-up mode. A representation of the egocentric positions of all viewable landmark segments is shown in the parietal window (PW). Rotated representations are projected onto the various transformation sublayers, which are inhibited by current head direction (HD) activity via a population of inhibitory interneurons (I). One transformation sublayer receives direct excitation from the HD system, thus allowing its representation to project forward to the boundary vector cells (BVCs). Bottom: Transformation circuit in top-down mode. The allocentric BVC representation of the environment is projected identically onto each of the transformation sublayers. Each of these identical representations would be rotated through different angles by the transformation to PW weights, but excitation and inhibition from the head direction system allows only the correct sublayer to maintain sufficient activity to drive PW neurons.

Trappenberg, Rolls, & de Araujo, 2002; Zhang, 1996). The continuous attractor property implies that the network will stabilize on a single bump of activity corresponding to a single head direction, and this bump can move continuously through 360° to reflect self-motion or perceptual inputs. Moreover, the reliability of the input mapping implies that if the animal returns to the same environment in the future, the head direction system will be oriented in exactly the same fashion and will exhibit the same firing pattern as it did on the first exposure to the environment.

The egocentric-to-allocentric transformation is accomplished by a circuit that combines head direction information with egocentric spatial input from the parietal window. The transformation circuit, assumed to be in the retrosplenial cortex/intraparietal sulcus, is comprised of a set of *N* identical neural subpopulations, each tuned to a specific head direction. Each subpopulation encodes a rotated egocentric map consistent with the direction of its preferred heading. Thus, connections between the parietal window and any one of the transformation subpopulations are weighted such that a rotated version of the egocentric spatial information contained in the parietal window is projected onto that transformation sublayer. In our model, there are 20 such sublayers corresponding to evenly spaced allocentric directions. Each transformation sublayer then projects an identical copy of its activation pattern onto the layer of BVCs. By setting connections from the layer of head direction cells to the transformation neurons such that only the sublayer corresponding to the current head direction is active, the transformation from egocentric to allocentric coordinates is accomplished. See Figures 5 and 6. In this way, when the animal's head rotates within the environment, head direction cell activity and parietal window activity vary in time, but so long as the animal undergoes no translation, activity projected to BVC neurons remains constant. The gating function of the head direction cells is accomplished via a combination of direct excitation from the head direction cells to the appropriate transformation sublayer and indirect uniform inhibition of all transformation layers by a population of inhibitory interneurons driven by head direction cell activity. This circuitry allows a localized bump of activity in the head direction layer to select the set of transformation units corresponding to that head direction.

The egocentric–allocentric transformation results in a single viewpoint-independent representation of each location in an environment. The allocentric representation consists of a distributed pattern of activation across the boundary vector cell layer. To encode this pattern as a distinct place memory, and to permit subsequent cued recall, this pattern can be learned by an autoassociative memory system. A retrieval cue, such as incomplete egocentric sensory or mentally generated spatial information, can then feed forward through the transformation circuit and reactivate the correct allocentric representation of the model's real or imagined surroundings. Conversely, the place memory can generate a viewpoint-specific mental image if we assume that the connections in the transformation circuit operate with equal weights in both directions. The recalled allocentric representation can thereby be converted back into egocentric mental imagery of the environment via the same neural circuitry.

Model Dynamics

Neurons in our model are rate coded (i.e., their activations represent average neural firing rates rather than individual spikes) and exhibit a continuous dynamic governed by "leaky-integrator" equations. The complete mathematical details of the model, along with these dynamical equations, can be found in the Appendix. Here we present a more intuitive description of the model's overall behavior. For now, the issue of biologically realistic learning is ignored and it is assumed that the model has already learned about the spatial environments it encounters. The actual ad hoc training procedure used to set the model weights for this work will also be described briefly in a subsequent section, with full details presented in the Appendix. In a subsequent section, we also discuss general principles that might underlie the learning of egocentric–allocentric transformations in biological systems.

At the highest level of dynamics, our model operates in alternating bottom-up and top-down stages, each lasting for 15 arbitrary time units. This periodic alternation in dynamics is based on modeling work by Hasselmo, Bodelón, and Wyble (2002), who argued that the hippocampal theta rhythm regulates the communication of this structure with interconnected brain regions. In particular they argued that during troughs in the rhythm, the hippocampus primarily receives input from surrounding structures but that during peaks, it primarily transmits information to these structures. We implement this alternating dynamics in our model both because of the evidence supporting its existence and because it allows the model to account for more experimental data than it otherwise could. In particular, without these distinct phases the model would have to engage in both bottom-up and top-down processing at the same time. We have found that a functional version of such a model exhibits states that strongly resist change in response to external inputs.

During the top-down phase, activity from the hippocampal layer feeds back to perirhinal cortex and also to the parietal window via the BVC and transformation layers. In addition, during this phase, the parietal window receives input from the senses, which we assume can be down regulated if the model is performing mental exploration or recall of a familiar environment without actually changing its vantage point (see Figures 5 and 6). During the bottom-up phase, the activity of the window is "frozen" to the last pattern present during the top-down phase. This activity pattern, which is the model's current representation of the geometry of egocentric space, is hypothetically maintained by a frontal–parietal short-term memory system (which we do not model here), consistent with evidence presented earlier. The frozen information from the parietal window feeds forward during the bottom-up phase to the hippocampal layer along with information from perirhinal cortex, thus influencing the current hippocampal attractor state. In principle, rigid freezing of the parietal window representation during the bottom-up phase is not necessary, but such an approach eliminates the need for additional neural circuitry in the model.

An animal would need to recall the details of an environment stored in long-term memory for two main reasons. First, there could be transient environmental conditions that impede sensory input and thus leave the animal with little direct access to spatial information. Second, the animal might need to remember what would be around it at an imagined location for the purposes of planning. For the former case, we assume that the model has enough sensory information to orient the head direction system. Although we only deal with visual information here, the model could be extended easily to include other cues such as vestibular input for this purpose as well. Once the head direction system is

oriented, the available but incomplete sensory input to the parietal window and perirhinal cortex can flow to the hippocampus in a bottom-up phase and activate an attractor state for the complete corresponding allocentric representation. During the next top-down phase, this attractor state reconstructs the environmental geometric information in the parietal window. Once the model has reconstructed the geometry of the environment, it must be able to identify the boundaries/landmarks that surround it. This is assumed to occur via directed attention to a spatial location. We simulate this in our model as extra activation (calculated from Equation A17) being directed to the area of interest in the parietal window. The boundary within the focus of attention in the parietal window will generate a corresponding focus of activation on its allocentric location within the BVC layer. The associative pathways within the medial temporal lobe can then retrieve the object's identity in the perirhinal cortex.

As a concrete example of spatial attention, if the model is instructed (perhaps by some prefrontal brain region controlling planned eye movements, not modeled here) to identify a boundary to its egocentric left, then extra activation is directed to the parietal window neurons that represent space to the egocentric left. This activation then flows through the transformation circuit, to the BVC layer, and finally to the perirhinal layer. The extra activation from the parietal window increases the firing rate of all perirhinal neurons corresponding to boundary identities that the model could encounter to its left when it has the current heading direction. The correct boundary identity, consistent with the subject's current location, can then be disambiguated by allowing the top-down connections in the model to operate at a low level (5% of the normal top-down value) even during a bottom-up phase. In this way, the place cell activity can provide the requisite disambiguation. For consistency, we also allow bottom-up connections to operate at the same reduced level during top-down phases.

In cases in which an animal needs to recall the details of its surroundings from a particular imagined point of view, we assume that the suggestion of (in the case of humans) or the memory of a highly salient environmental feature located at some point in the animal's egocentric space might be enough to orient the head direction system. The correct perirhinal units could also be activated by this process, and activity corresponding to the location of the feature could be sent to the parietal window. During the next bottom-up phase, the processes of pattern completion and directed attention would then follow as described above.

Spatial Updating and Mental Exploration

The recall processes described in the previous section are useful only if an animal requires stationary "snapshots" of an environment. However, a moving animal, often faced with partially or fully occluded sensory information, requires an accurate, real-time representation of its surroundings. Similarly, if an animal wishes to plan a route through a familiar environment, the ability to perform mental exploration of the surrounding space would be useful.

A key part of our overall theory is that parietally generated egocentric mental imagery can be manipulated via real or mentally generated idiothetic information in order to accomplish spatial updating or mental exploration in familiar environments. A detailed neural mechanism for accomplishing such tasks in the case of pure short-term or working memory has been described elsewhere (Byrne & Becker, 2004). Here we are concerned primarily with the updating process applied to medial temporal lobe dependent long-term memory. For this case, we assume that rotational and forward-translational egomotion signals act upon the egocentric parietal window representation of space via different mechanisms. In the case of rotation, the egomotion signal causes head direction cell activity to advance sequentially through the head direction map, thus rotating the image that is projected into the parietal window from the BVCs. This velocity-modulated updating of head direction is similar to the model described by Stringer et al. (2002). The potential for such one-dimensional continuous attractor networks to account for multiple aspects of the head direction cell assembly has been investigated in detail by Conklin and Eliasmith (2005); Goodridge and Touretzky (2000); Hahnloser (2003); Redish, Elga, and Touretzky (1996); among others. However, a detailed summary of such work is beyond the scope of this article. For the case of forward translation, the egomotion signal gates the top-down connections from the parietal transformation layer to the parietal window such that the "normal" top-down weights connecting these regions are down regulated, whereas a second, alternate set of top-down weights are up regulated. With no forward velocity signal, the normal top-down connections perform reconstruction of a head-centered egocentric representation of the model's current spatial surroundings in the parietal window by using information originating from place cell activity. Once up regulated by the velocity signal, the alternate set of top-down connections performs an almost identical function, except that the representation of space reconstructed in the parietal window is of the model's current surroundings but shifted backwards slightly in the model's egocentric space. When the next bottom-up phase begins, the shifted spatial information, represented as parietal window activity, flows through the transformation and BVC layers to activate place cells that correspond to the location slightly ahead of the model's current location. This process repeats itself during the next top-down/bottom-up cycle until the velocity signal dissipates, resulting in a continuous relocation of the model's internal representation of its location in space. Further details of this updating procedure can be found in the Appendix.

Learning in the Model

The purpose of our model is to reproduce experimental data and to generate novel predictions of spatial behavior in adult animals, rather than to account for learning in a biologically realistic manner. Hence, we use a simplistic Hebbian learning procedure that associates together prespecified activation patterns in each layer of the model, in order to train all of the model connection strengths except for those involved with spatial updating/mental exploration. The latter connection strengths are calculated as described in the Appendix. Briefly, learning for the remainder of the weights involves positioning the model at numerous random locations and heading directions within an environment while, at each of these locations, sequentially directing attention to each landmark segment viewable from the current location. For each attending event at each location, appropriate activation patterns are imposed upon the model layers and connection strengths between neurons are updated via a simple correlational rule. Once training is complete, weights are normalized. A detailed description of the training procedures is provided in the Appendix.

It should be noted that the transformation circuitry in our model is only trained once, but the medial temporal component is retrained on each unique environment in the simulations reported here. Training on multiple environments with the relatively small-scale models used here can result in a degradation of information when it travels through the transformation circuitry and activation of an incorrect hippocampal attractor state. This problem could be addressed by including a greater number of model neurons in the transformation layer. Additionally, a larger scale version of the medial temporal lobe portion of the model should, in principle, be capable of storing multiple environments in distinct subsets of place cells (a possible role for the dentate gyrus and CA3 recurrent connections; McNaughton & Morris 1987; Samsonovich & McNaughton, 1997). There is no reason to expect that the simultaneous storage of attractor states corresponding to multiple environments would affect any of the results we obtain from the model in this article.

Simulation 1: Recall of Landmarks and Geometry in Hemispatial Neglect

Method

In order to simulate representational neglect (see introductory section and Bisiach & Luzzatti, 1978), we first tested the ability of the intact model to recall environmental geometry and landmark identity. This was accomplished by first training the medial temporal component of the model on the simplified cathedral square depicted in the upper left panel of Figure 7. During training, the allocentric reference frame was taken to be aligned with this depiction of the environment so that its *y*-axis would be perpendicular to the inward facing walls of Buildings 1 and 3 but parallel to the inward facing walls of Buildings 2 and 4. In reality, it is likely that the orientation of the allocentric reference frame within the environment would be set by the head direction system alignment when the animal first experiences the environment. Once training was complete, the model was cued to imagine itself facing the cathedral in the trained environment by injecting appropriate activation into the head direction, parietal window, and perirhinal identity layers. Cuing activation for the parietal window was calculated by applying Equation A5 to a discretized linear boundary, representing the front of the cathedral, located directly in front of the model in the egocentric reference frame. Similarly, cuing activation for the perirhinal neurons was calculated from Equation A3, with the cathedral (Building Identity 1) being the attended landmark. Finally, it was assumed that the cathedral is sufficiently salient that cuing its location relative to the subject is enough to orient the head direction system. Thus, activation for the head direction layer was calculated from Equation A6, with the heading direction (ϕ) set to zero, indicating perfect alignment between egocentric and allocentric reference frames. The cuing activations were applied to the model for two full bottom-up/top-down cycles, after which they were down regulated, and the retrieved attractor states in the head direction system and the hippocampal place cell layer maintained the model's parietal window representation of the imagined geometry of the environment.

In order to "ask" the model to identify the boundaries that would be visible from the current viewpoint (see Figure 7), we simulated the focus of attention along four different directions: left, right, forward, and backward. In each direction, the corresponding activation calculated from Equation A17 was injected directly into the parietal window. During a subsequent bottom-up phase, this activation flowed forward through the transformation and parahippocampal layers to activate the correct perirhinal identity neuron. For example, in the case of rightward attention, the correct response would be perirhinal activity corresponding to Building 2 (see the Appendix for details).

Next, the model was cued to imagine itself in the square facing away from the cathedral. This was accomplished by focusing attention on a boundary directly behind the model in the parietal window, while simultaneously activating the perirhinal neurons representing the visual features of the cathedral and the allocentric head direction 180° away from the current egocentric frame.

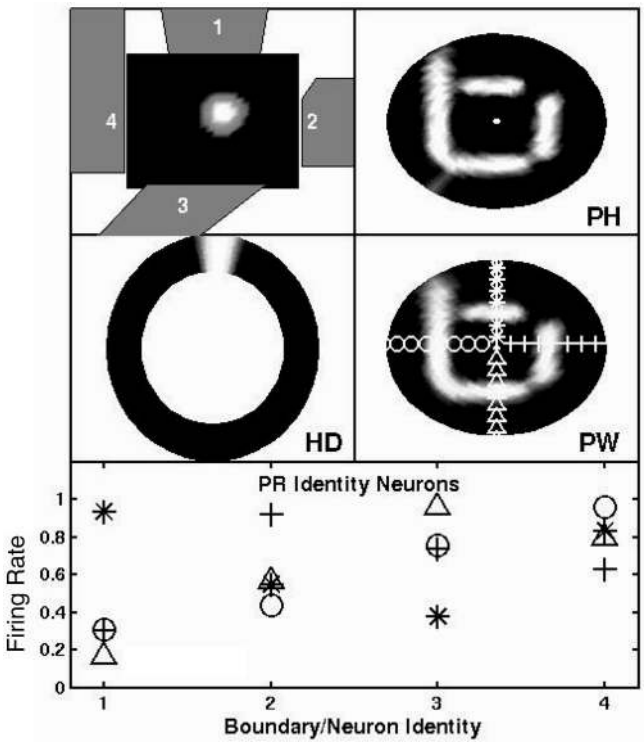

Figure 7. Top four panels: Activation in the various model layers averaged over a full cycle after it was cued to face the cathedral (Building 1). Upper left: Environmental boundaries are represented by gray walls superimposed upon the hippocampal place cell representation. Here, the firing rates of all hippocampal place cells are presented, with each shown at its corresponding grid point within the environment. Bottom left: The head direction (HD) activity peak indicates that the model was facing "forward" relative to the stored allocentric map. Therefore, parietal window (PW) activity (bottom right), which is the model's representation of its surrounding egocentric space, was highly similar to parahippocampal (PH) boundary vector cell activity (upper right), which corresponds to the model's allocentric representation of space. The various symbols superimposed upon the egocentric PW representation indicate the attention directions. Bottom: Activation in perirhinal (PR) identity neurons at the end of the first bottom-up phase after attention is directed in the PW. For example, when attention is directed to the egocentric right ("+"), PR neuron 2, which corresponds to Boundary/Building 2, is the most active identity neuron.

Once it was confirmed that the model could identify surrounding landmarks from different viewpoints, hemispatial neglect was simulated by performing a random knock out of 50% of the parietal window neurons representing the left side of egocentric space and then repeating exactly the same procedures as just described for testing the intact model.

Results and Discussion

The ability of the intact model to recall environmental geometry and landmark identity, when cued that it was facing the cathedral, is shown in Figure 7. The top four panels show the activity in the various network layers averaged over one full cycle after the removal of the cuing activity. Although the spatial resolution of the model's representation of the environment is coarse, the geometry represented in the parietal window is roughly correct. The bottom panel of Figure 7 shows the activity of perirhinal neurons at the end of a bottom-up phase. Perirhinal activity is plotted with open circles for leftward attention, asterisks for forward attention, crosses for rightward attention, and triangles for backward attention, indicating that the model can identify all landmarks correctly. Performance of the intact model when cued that it was facing away from the cathedral is shown in Figure 8. The resultant activities of the various network layers averaged over a full cycle after down regulation of cuing inputs are shown in the top four panels. Once again the model formed the correct egocentric representation of spatial information in the parietal window and directed attention resulted in the correct identification of the surrounding boundaries. For example, when attention was directed to the egocentric right, the identity of Building 4 was activated in the perirhinal layer. Building 4 would be to the right of the model if it were facing away from the cathedral.

Results of the simulations with the lesioned model, simulating hemispatial neglect, are shown in Figures 9 and 10 and corresponding to Figures 7 and 8, respectively. From these results, it is clear that the model could identify landmarks to its right, but not to its left, regardless of its imagined heading direction. These simulation results are consistent with a central tenet of our model, namely, that allocentric representations of space are formed in long-term memory and are transformed into egocentric views as needed, in the service of memory recall and imagery. Moreover, our model provides a mechanistic explanation for patterns of deficits observed in perceptual and representational neglect patients, a previously perplexing phenomenon in neuropsychology. Both the long-term memory representation and the transformation mechanism are intact, whereas the egocentric representation projected from long-term memory, and/or the transformation mechanism itself, is faulty. This could arise in patients either from a lesion to the pathway from the transformation circuit to the parietal window (resulting in pure representational neglect) or from a lesion to the parietal window itself (resulting in both perceptual and representational neglect). Pure perceptual neglect in the absence of representational neglect could arise from a lesion along the sensory or motor pathways projecting into and out of posterior parietal cortex. Testing of these predictions based on currently available data is difficult because of the extensive lesions suffered by most patients suffering from unilateral neglect. For the case of perceptual neglect, recent studies indicate that a disconnect between parietal cortex and prefrontal areas (Doricchi & Tomaiuolo, 2003; Thiebaut de Schotten et al., 2005) or between parietal cortex and medial temporal regions (Bird et al., 2006) is critical to a realization of the phenomenon. However, we are unaware of any data that so clearly indicate which regions of the brain must be damaged in order to induce pure representational neglect, the focus of the current set of simulations.

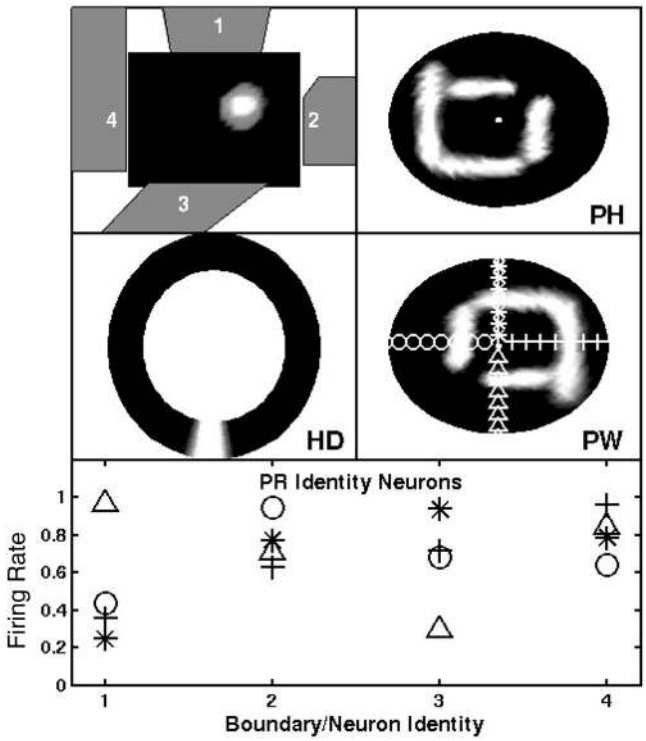

Figure 8. Top four panels: Activation in the various model layers averaged over one full cycle after it was cued to face away from the Cathedral. The head direction (HD) activity peak indicates that the model was facing "backwards" relative to the stored allocentric map. Therefore, parietal window (PW) activity is rotated 180° relative to boundary vector cell activity. The various symbols superimposed upon the egocentric PW representation indicate the directions in which attention was directed. Bottom: Activation in perirhinal (PR) neurons at the end of the first bottom-up phase after attention is directed in the PW. PH = parahippocampal.

Simulation 2: Spatial Updating During Physical and Mental Navigation

One of the key functions of the model is its ability to perform spatial updating of its internal representations of location, given a motion signal. Spatial updating is critical for navigation in the absence of perceptual input (path integration), for mental imagery involving viewpoint changes, and for path planning. Spatial updating should allow relatively normal navigation and place cell firing over short durations in the absence of perceptual input, and it should account for data on spatial updating such as that of Wang and Brockmole (2003), described in the introductory paragraphs. In our model, path integration occurs outside of the hippocampus through updating the parietal egocentric representation. Further, the same machinery accounts for the process of mental navigation

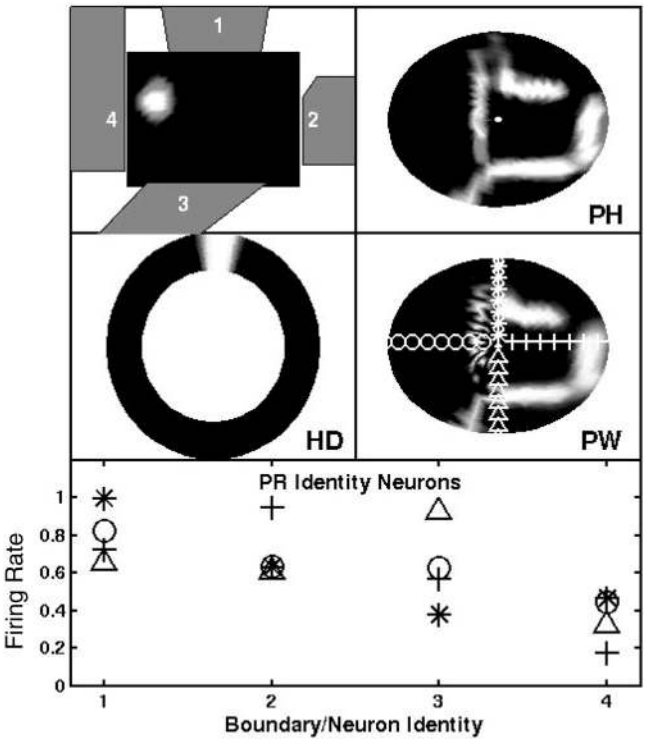

Figure 9. Top four panels: Activation in the various model layers averaged over one full cycle after the lesioned model was cued to face the cathedral. Bottom: Activation in perirhinal (PR) neurons at the end of the first bottom-up phase after attention is directed in the parietal window (PW). PH = parahippocampal; HD = head direction.

by generating an imagined motor signal in place of the efference–proprioceptive–vestibular signal generated by actual motion. This should allow the model to address performance and reaction time data in tasks in which the subject is asked to respond from a different imagined viewpoint and/or location (e.g., Diwadkar & McNamara, 1997; Easton & Sholl, 1995; Rieser, 1989; Shelton & McNamara, 2001) or asked to simulate some aspects of spatial planning.

Method

In order to simulate spatial updating or mental navigation, the medial temporal component of the model was trained on the "two-room" environment shown in the upper left panel of Figure 11, with the allocentric reference frame taken to be aligned with the vertical axis of the environment as depicted. The training procedure and architecture for this component of the model were identical to those used in the previous set of simulations, except that in addition, within the parietal window, the velocity-gated translational weights given by Equation A9, and the rotational head direction weights, trained as described in the Appendix, now come into play.

After training was complete, the model was first cued to a location near to and directly facing Wall 1. Such cuing would be equivalent to asking the model to imagine itself facing Wall 1 in the two-room environment. This was accomplished as in the previous simulations by injecting appropriate activations into perirhinal, head direction, and parietal window neurons for two full cycles. Attention was then focused along four different directions, leftward, rightward, forward, and backward, to demonstrate that the model could identify the surrounding landmarks from memory.

Next, we simulated spatial updating after several steps of imagined egomotion. The same situation could arise during real navigation if an animal spontaneously loses sensory information about its real surroundings (e.g., navigating in the dark). In either case, attractor states in the head direction system and in the hippocampal formation of our model are able to maintain an internal representation of the real/imagined surroundings. Mental exploration or spatial updating based on this self-sustaining internal representation was simulated in the model by a series of eight egomotion steps. This egomotion, if assumed to be generated by real idiothetic information, would correspond to spatial updating, or if generated by a mental equivalent, would correspond to mental exploration. In the first step, to simulate making a 180° turn, a counterclockwise rotational velocity signal lasting for 150 time units gated the rotational head direction weights until the model's egocentric representation of space rotated by a full 180°. In the second step, to simulate forward egomotion, a translational velocity signal lasting 135 time units gated the transformation to parietal window translational weights, causing the model's egocentric representation of the locations of boundaries to translate backwards. Similarly, a further six egomotion steps were performed to complete the simulation.

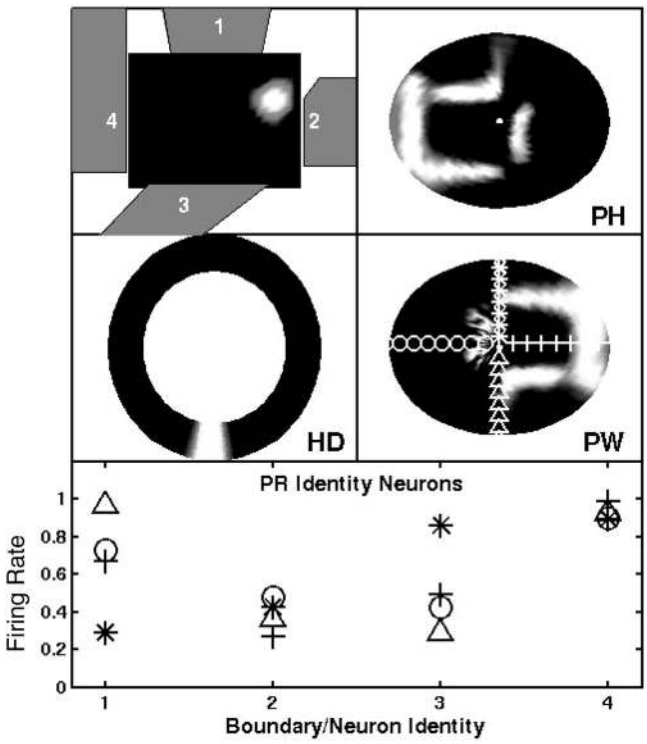

Figure 10. Top four panels: Activation in the various model layers averaged over one full cycle after the lesioned model was cued to face away from the cathedral. Bottom: Activation in perirhinal (PR) neurons at the end of the first bottom-up phase after attention is directed in the parietal window (PW). PH = parahippocampal; HD = head direction.

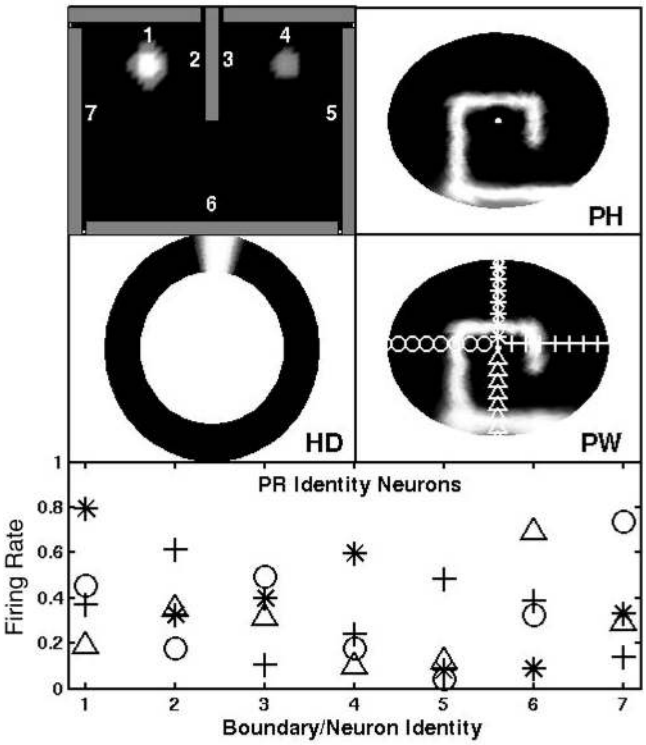

Figure 11. Top four panels: Activation in the various model layers averaged over one full cycle after it was cued to localize itself in the "two-room" environment facing Wall 1. Environmental boundaries are represented by gray walls superimposed upon the hippocampal representation. The various symbols superimposed on the parietal window (PW) representation indicate the sequential attention directions. Bottom: Activation in perirhinal (PR) neurons for the various attention conditions at the end of the first bottom-up phase after attention is directed in the PW. PH = parahippocampal; HD = head direction.

As a control, we compared spatial updating in imagined versus sensory-driven navigation. Although the model's ability to perform spatial updating/mental exploration on internally maintained representations of space is of primary interest, it must also function in a consistent way during real navigation through a familiar environment with intact sensory information. Thus, we simulated the same situation as above but in the presence of accurate sensory cues during the eight steps of egomotion. In this case, sensory information corresponding to visible boundaries calculated from Equation A5 was simultaneously injected into the parietal window during egomotion.

Results and Discussion

The ability of the model to retrieve the appropriate context in the two-room environment, when asked to imagine itself facing Wall 1, is shown in Figure 11. Network activity averaged over a full cycle after down regulation of the cuing inputs can be seen in the top four panels of Figure 11. The results of the four directed attention events are shown in the bottom panel of Figure 11, indicating that the model could also identify the surrounding landmarks.

The performance of the model after several steps of imagined egomotion is shown in Figures 12 and 13. Figure 12 shows activation in the various network layers averaged over one full cycle following the first two egomotion steps. The remaining six steps brought the model's internal representation of space to that shown in Figure 13, where it was nearby and facing Wall 2. Three directed attention events show that the model could correctly identify surrounding boundaries from this new viewpoint (see bottom panel of Figure 13).

In the case of sensory-driven navigation, the analogous results to Figures 11, 12, and 13 are shown in Figures 14, 15, and 16, respectively. Results of the sensory-driven simulations after eight steps of egomotion are nearly indistinguishable from the corresponding results with imagined egomotion.

The fact that an egocentric translational velocity signal causes spatial updating/mental navigation to occur at a constant velocity is discussed in more detail with respect to Simulation 4 and in the General Discussion. Here we simply note that it is consistent with the reasonably accurate (if scaled) correspondence between mental navigation times and actual navigation times (see, e.g., Ghaem et al., 1997; Kosslyn, 1980).

Simulation 3: Place Cell Firing With Head Direction Cell Lesions

In Simulations 1 and 2, we compared our model against behavioral data. The purpose of Simulations 3 and 4 was to evaluate the adequacy of our model in explaining and predicting data at the level of single-unit recordings. For this third set of simulations, the static model, that is in the absence of egomotion, is evaluated with respect to place cell firing after lesions to the head direction system. In Simulation 4, the model is evaluated under conditions of cue conflict between direct sensory and path-integrative inputs.

Calton et al. (2003) have shown that rats with lesions to the anterodorsal thalamic nuclei or to the postsubiculum, two locations where head direction cells have been found, show altered place cell

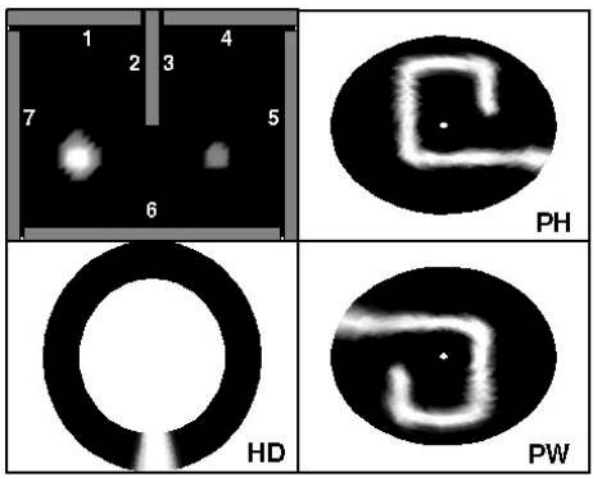

Figure 12. Activation in the various model layers averaged over one full cycle after the application of the rotational velocity signal for 150 time units followed by a forward translational velocity signal for 135 time units. PH = parahippocampal; HD = head direction; PW = parietal window.

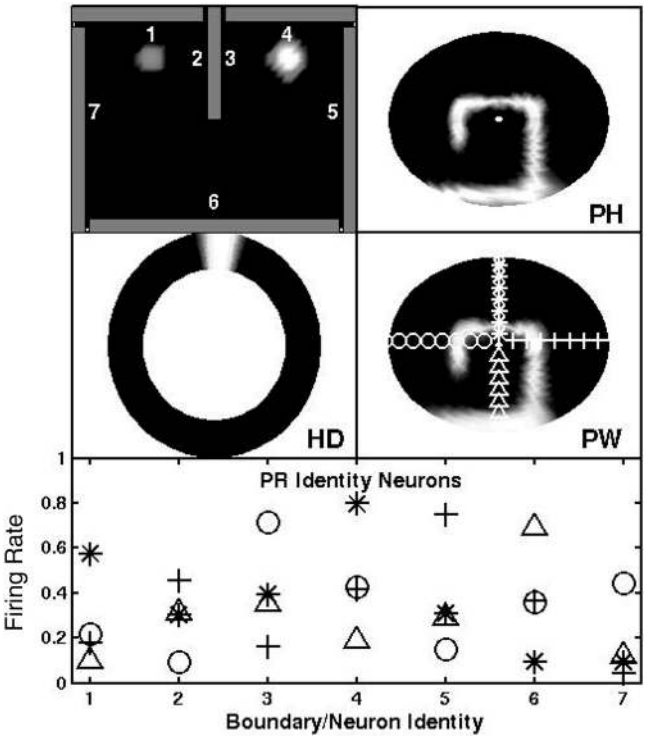

Figure 13. Top four panels: Activation in the various model layers averaged over one full cycle at the end of the eight step sequence of egomotion. Bottom: Activation in perirhinal (PR) neurons for the various attention conditions at the end of the first bottom-up phase after attention is directed in the parietal window (PW). PH = parahippocampal; HD = head direction.

firing characteristics when compared with intact animals. Although variations in place cell firing properties between the two lesioned groups were seen, there were a number of characteristics in common to both groups. Specifically, place cells in both groups showed roughly normal in-field firing but elevated out-of-field firing. Additionally, this out-of-field firing showed dependence on heading direction.

In order to understand how our model could address the results of Calton et al. (2003), it is useful to return briefly to the description of how incoming sensory information activates the correct place cell attractor states. Recall, we have assumed that incoming information about environmental geometry first reaches the egocentric parietal window representation before being transformed via the transformation layer into an allocentric BVC representation. The BVC pattern, in conjunction with perirhinal activity, then activates the appropriate hippocampal attractor state. This transformation relies upon a gating mechanism driven by the head direction system that will be clearly disrupted if head direction cells are destroyed. Thus, under normal circumstances, a given pattern of activity in the head direction system allows only one transformation sublayer to project activity onto the BVC layer. However, if the former is damaged, its gating function will be compromised, reducing the activity received by the BVC layer from the correct transformation sublayer and increasing the activity from other sublayers. Depending on the extent of the lesion to

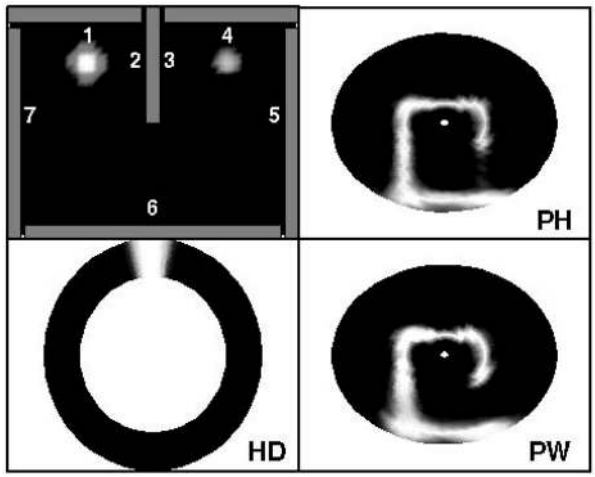

Figure 14. Results for the simulation in which sensory information about the environment is being continuously input to the parietal window (PW) representation throughout the duration of the simulation. PH = parahippocampal; HD = head direction.

the head direction system, the garbled BVC representation could still overlap significantly with the representation required to activate the appropriate attractor state given the model's current sensory information, or it could be that the overlap is very small. In intermediate cases, the correct hippocampal place cells might receive enough activation to fire, but other neurons might be driven past their firing thresholds as well.

Method

A realistic simulation of the effects of lesions to the head direction cells in our model is not possible because of the use of a single inhibitory interneuron that causes each head direction cell to inhibit all transformation sublayers equally. A more realistic cir-

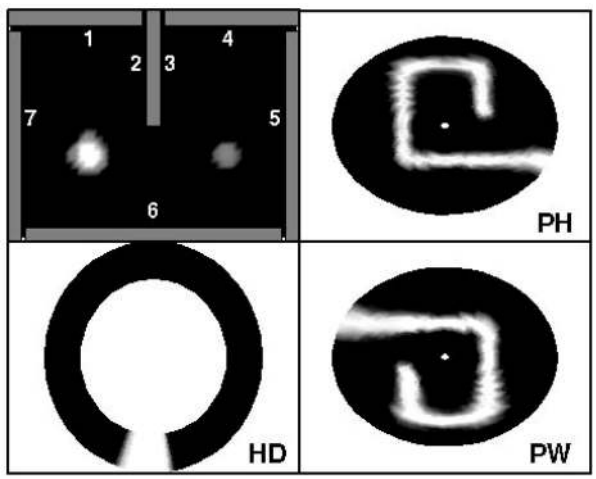

Figure 15. Results for the simulation in which sensory information about the environment is being continuously input to the parietal window (PW) representation throughout the duration of the simulation. PH = parahippocampal; HD = head direction.

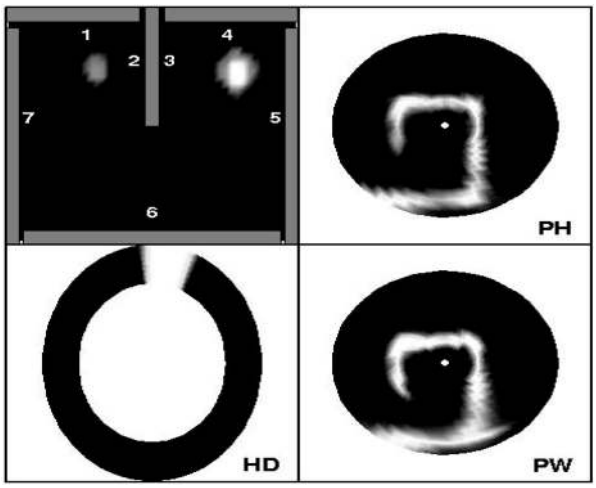

Figure 16. Results for the simulation in which sensory information about the environment is being continuously input to the parietal window (PW) representation throughout the duration of the simulation. PH = parahippocampal; HD = head direction.

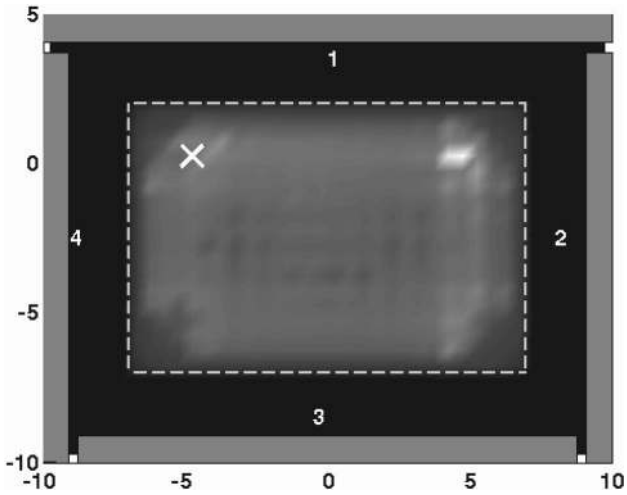

Figure 17. Activity of a single place cell recorded from the model with a simulated head direction cell lesion. Recordings were made when the model was localized at numerous points within the dashed rectangle. In this simulation the model's head direction was consistent with perfect alignment between parietal window and boundary vector cell representations of space. Note also that the recorded cell would fire maximally at the "X" for all head directions in the nonlesioned model.

cuit would use a population of inhibitory interneurons that were connected randomly within the constraint that they would achieve the same gating function (in combination with excitatory head direction connections to the transformation layer). We did not use such a population because, given the unnatural training methods used, it would have behaved like a single unit anyway. With a more natural configuration, partial lesions to the head direction system would result in reduced excitation to the selected transformation sublayer and decreased inhibition to random regions of the overall transformation layer. To simulate the equivalent effect in our model, for each lesioned head direction, the excitatory head direction input to the corresponding transformation sublayer was reduced, whereas the inhibitory input to a random selection of other transformation sublayers was decreased (see the Appendix for details).

Because the lesioning procedure does not involve the medial temporal structures, the latter region was trained once on the "box" environment shown in Figure 17. The model was then localized at numerous positions within the environment by injecting appropriate egocentric sensory information from all of the environmental boundaries into the parietal window neurons. At each location, the sensory input was maintained for one top-down/bottom-up cycle, and the activity of a selected place cell was recorded and averaged over the bottom-up cycle. This procedure was performed for two simulated head directions, one of which corresponded to perfect alignment between egocentric and allocentric representations and the other of which corresponded to perfect antialignment between the two representations.

Results and Discussion

The average firing rates for a model place cell recorded when the lesioned model was localized at numerous locations within a rectangular subregion of the "box" environment are depicted in Figures 17 and 18. In Figure 17, these rates correspond to the aligned heading direction, whereas in Figure 18 the results correspond to the antialigned simulation condition. Clearly, the firing field of the model neuron varied with simulated head direction, and moreover, its peak-firing location for either head direction did not correspond to the location where the cell would have attained its maximal firing rate in the nonlesioned model (marked with an "X" in both figures). In addition, for the aligned condition, the cell exhibited a firing maximum in one location but with an additional area of elevated firing near "X." These data are qualitatively similar to the data shown in Figure 4B of Calton et al. (2003).

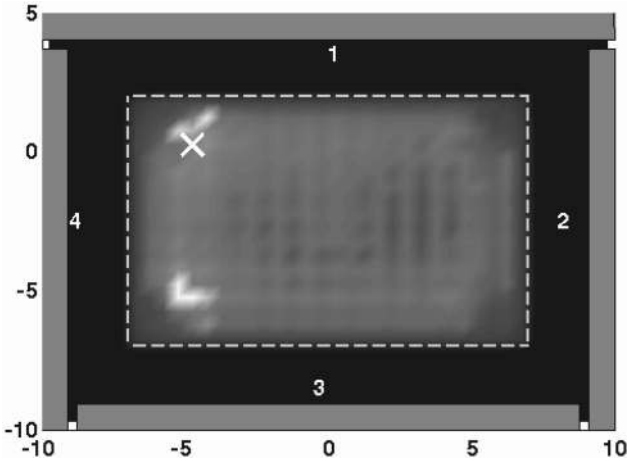

Figure 18. Activity of a single place cell from the model with a simulated head direction consistent with perfect antialignment between parietal window and boundary vector cell representations of space. Note also that the recorded cell would fire maximally at the "X" for all head directions in the nonlesioned model.

Our model makes two unique predictions regarding the outcome of experiments similar to those of Calton et al. (2003). First, a place cell that has a prelesion preference for a location about which there is a high degree of rotational symmetry (e.g., the center of a cylinder) should maintain its place preference postlesion. Conversely, place cells that show prelesion preferences for locations of low rotational symmetry should tend to show shifts in their preferred locations after a lesion. An example of this latter effect is seen clearly in the simulation presented in Figures 17 and 18. Second, the relative firing rates for place cells when measured at locations of high rotational symmetry should demonstrate little dependence on head direction after a lesion. For example, if Cell A demonstrates a high postlesion firing rate at the center of a cylinder for a given head direction, and if Cell B demonstrates a low firing rate at that location and head direction, then for all other head directions Cells A and B should show similar relative firing rates at that location. Conversely, the relative firing rates for place cells when measured at locations with lower levels of rotational symmetry should exhibit higher levels of head direction dependence after a lesion.

In order to understand these predictions, one only needs to note that each transformation sublayer contains a representation of the same egocentric space but rotated about the origin. Therefore, if the egocentric parietal window representation shows a reasonable degree of rotational symmetry at a given location, then allowing extra regions of the overall transformation layer to project to the BVCs will not have a large effect on the resultant geometric information represented there, regardless of head direction. Hence, a place cell that fires maximally/minimally at such a location before a head direction system lesion would still receive high/low levels of stimulation there after a lesion; moreover, because of the rotational symmetry, it will do so for all head directions.

Simulation 4: Place Cell Firing With Conflicting Visual and Path-Integrative Inputs

The basis of the medial temporal component of our model was derived from a simple feed-forward model of place cell firing (Hartley et al., 2000; O'Keefe & Burgess, 1996) driven by input from BVCs. This earlier model included a number of simplifications, one of which was that BVCs and therefore place cell firing rates were independent of firing history. However, memory in general, and path integration in particular, make important contributions to place cell firing, in addition to immediate sensory perception such as vision, olfaction, et cetera. For example, place cells can continue to fire normally in the dark (O'Keefe, 1976); path integration, distant visual cues, and multimodal local cues can be pitted against each other to control the orientation of place cell firing (Jeffery, Donnett, Burgess, & O'Keefe, 1997; Jeffery & O'Keefe, 1999); and congenitally blind rats show normal place fields once they have explored the polarizing environmental cues (Save, Cressant, Thinus-Blanc, & Poucet, 1998).

Here we have coupled the medial temporal model to a parietal system capable of spatial updating. An obvious test of this extended model is to determine whether it can capture the joint effects of path integration and sensory perception on place cell firing, thereby extending the simple feed-forward place cell model. Another line of evidence for the differential contributions of path integration and sensory perception to place cell firing comes from Gothard et al. (1996), who examined the activity of hippocampal place cells in rats running along a linear track. By varying the track length during recording sessions, they were able to pit sensory and locomotor cues against each other. In our final set of simulations, we sought to compare the performance of the model to Gothard et al.'s data.

Gothard et al. (1996) trained rats to run back and forth along a narrow, elevated track with food cups at either end. One food cup was fixed directly to one end of the track, and the other was fixed to the floor of a sliding box that could be in any one of five locations (Box 1–Box 5), thereby changing the overall track length (see the left panel of Figure 19). Rats were habituated to the apparatus in the maximum length, or Box 1 state, for 3 to 5 days prior to recording. During a recording session, an animal was placed in the box at one of the five positions and allowed to run to the fixed food cup (outbound journey). The box was then moved to a new position before the rat turned around to make the return journey (inbound journey). Most cells fired prederentially in one direction of running, consistent with previous experiments on linear tracks (McNaughton, Barnes, & O'Keefe, 1983; O'Keefe & Reece, 1993). The firing profile for each cell was calculated separately for all types of journey (e.g., Box 1–out, Box 2–out, Box 1–in, Box 2–in) and was compared with the corresponding Box 1 profile. Specifically, the amount by which the peak firing location for a given cell was shifted from its preferred location in the Box 1 condition was plotted against the corresponding shift of the box relative to its Box 1 position (see Figure 19). This measure is sensitive to whether the place field shifts with the movable box or remains at a fixed location relative to stationary cues, but note that deformations in firing field shape occurred, such as bimodal fields as well as simple shifts. By fitting a regression line to the data for a given cell across box positions, a displacement slope, normalized to range between 0 and 1, was calculated. A slope of 0 corresponds to firing peaked at the same location relative to the fixed food cup in all conditions, whereas a slope of 1 corresponds to peak firing at the same location relative to the movable box, regardless of its position. Thus the movable box controls the location of firing fields with a large displacement slope, whereas the fixed food cup and other room cues control the location of fields with small displacement slopes.

Gothard et al.'s (1996) displacement slope results for inbound and outbound selective neurons are shown in Figure 20 along with some sample firing fields. Neurons that fired near the box or the cup in the original configuration continued to fire near the box or cup in the other configurations. Similarly, cells that fired in between the two cups did so in all configurations, except on the shortest journeys when they did not fire at all. However, for most of the distance traveled on a given journey, place cell firing appeared to be predominantly controlled by the landmark which the animal was moving away from. For outbound journeys, firing peaked near to the box in the Box 1 configuration have displacement slopes around 1, and this value gradually decreases to zero for neurons with peak firing positions farther away from the box. However, the slope value remains above 0.5 for peak firing locations much more than for those halfway down the track from the box. This additional influence of the cue from which the rat is running is also clearly evident for the inbound journeys in which most neurons, excepting those with peak firing very close to the

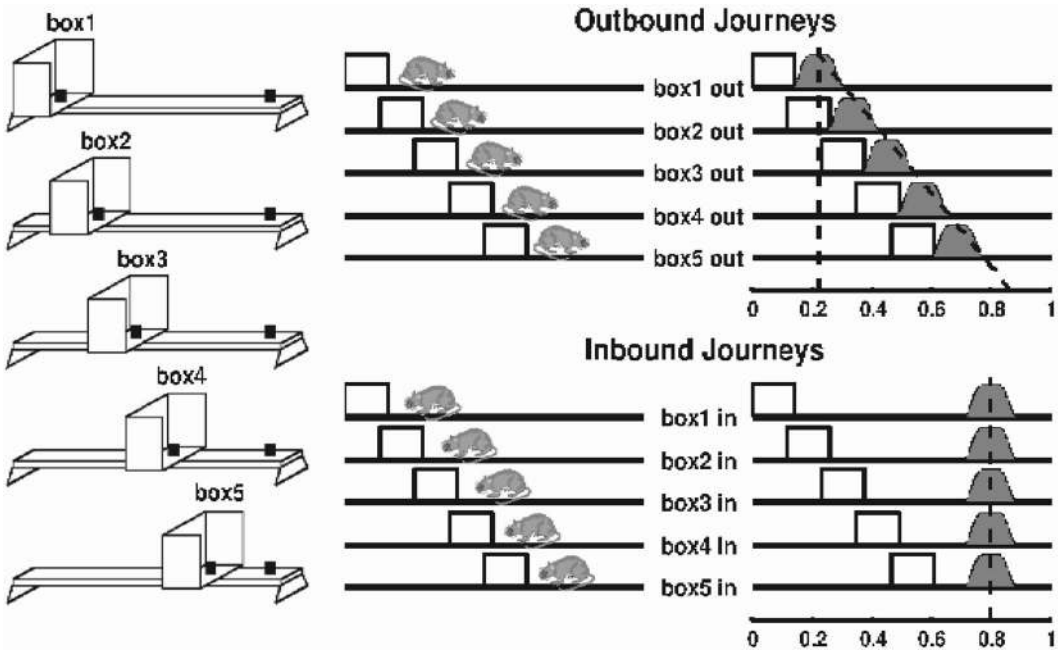

Figure 19. Left: Linear track apparatus used by Gothard et al. (1996). Top middle: Rat on outward journey from box to fixed cup for the five different box positions. Top right: Hypothetical average firing patterns for a place cell in each of the five outward conditions plotted against relative position along the track in the box1 condition (0 is the position of the box in the box1 condition, whereas 1 is the position of the fixed cup). The dashed diagonal line is the regression line used to calculate displacement slope, which is 1.0 for this cell because it fires near the box in all conditions. The vertical dashed line shows the location of peak firing on the box1-out trials. Bottom middle: Rat on inward journey from fixed cup to the box for the five different box positions. Bottom right: Hypothetical average firing patterns for a place cell in each of the five inward conditions plotted against relative position along the track in the box1 condition. This cell fires near the fixed cup in all conditions, giving a displacement slope of 0.0. From "Dynamics of Mismatch Correction in the Hippocampal Ensemble Code for Space: Interaction Between Path Integration and Environmental Cues," by K. M. Gothard, W. E. Skaggs, and B. L. McNaughton, 1996, *Journal of Neuroscience, 16,* p. 8028. Copyright 1996 by the Society for Neuroscience. Adapted with permission.

box, are controlled by the cup, showing displacement slopes close to zero.

The BVC model of place cell firing (Hartley et al., 2000; O'Keefe & Burgess, 1996) predicts much of Gothard et al.'s pattern of data, for example, that the location of maximal firing will tend to remain a fixed distance from the nearer of the two boundaries and how the fields stretch, develop subpeaks, reduce in firing rate, and disappear when the component BVCs fail to coincide in one or other new configuration. However, the increased influence of the boundary behind the rat compared with the one in front is not captured by this model (also noted in O'Keefe & Burgess, 1996). These results appear to require an interaction between BVCs responsive to the inconsistent visual cues and path-integrative locomotor information (see also Redish et al., 2000), consistent with the idea that both path-integrative and perceptual inputs are required to determine the hippocampal representation of location (O'Keefe & Nadel, 1978). Here we investigate the behavior of the model, which now includes both BVCs and motion-related spatial updating, in the Gothard et al. paradigm.

We model initial place cell firing when the animal is placed at either end of the apparatus, as consistent with the place cell firing for that location within the full-length track. This assumption is reasonable given that the majority of local cues available at either location are consistent with this representation. These cues consist of the three box walls for the box and all the other room cues at the fixed food cup. Upon leaving the start position for a given trial, input from both locomotion-related updating and from visual cues combine to update the animal's internal representation of its position. Within the full-length track (Box 1) condition of Gothard et al.'s (1996) experiment, neuronal activity follows a "normal" continuous trajectory through the set of states representing all intermediate locations within the full-length track and terminating with the state corresponding to the destination end of the track. At each stage, the perceptual input from both ends of the track is consistent with the internally updated input from the previous step. In the remaining conditions (Box 2–Box 5) the visible landmark ahead is closer to the rat than would be consistent with the motion-updated representation; this causes previously unimodal place fields to reduce in peak activity and to deform, showing a compromise between firing at a fixed distance from both ends of the track. At the start of an outbound journey, the cues behind the rat and the ideothetically updated internal representation predominantly control place cell firing, but as the rat proceeds along the track there is an increasing influence of the nearer than expected

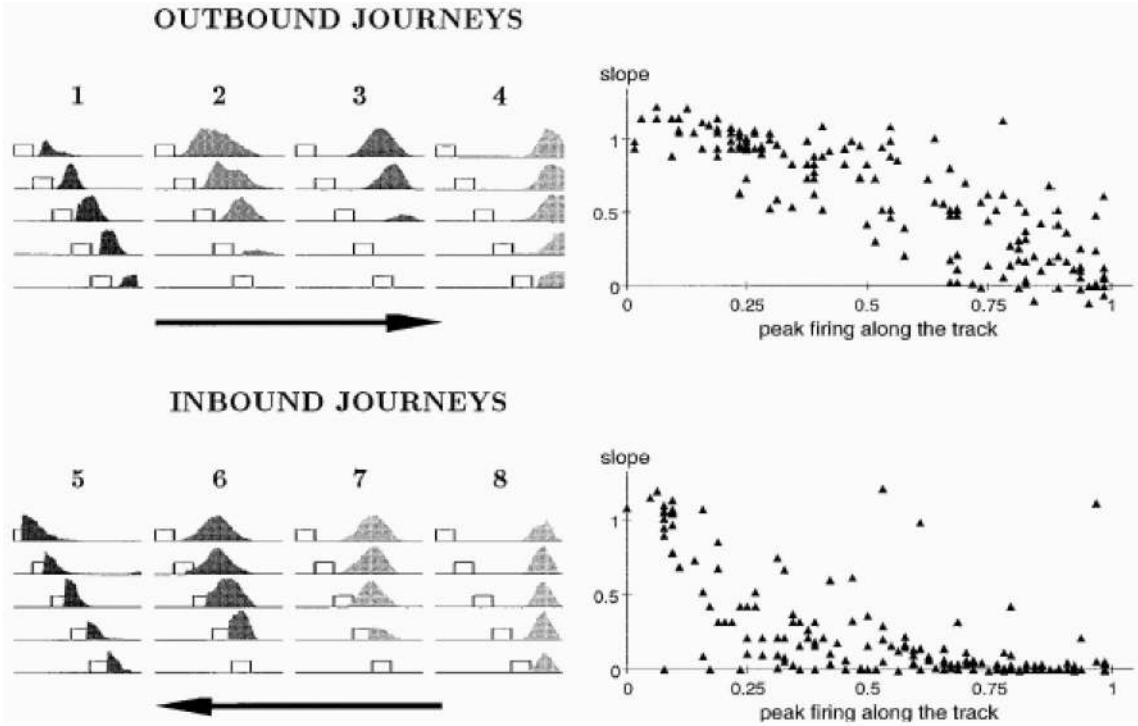

Figure 20. Upper left: Averaged firing profiles of four outward selective neurons in each condition. Rectangles represent the movable box. Upper right: Displacement slopes for multiple outward selective cells plotted against their peak firing positions in the box1 condition. Positions are relative to full track length, with 0 representing the box position in the box1 condition and 1 representing the position of the fixed food cup. Lower panels: Equivalent results for inward selective cells. From "Dynamics of Mismatch Correction in the Hippocampal Ensemble Code for Space: Interaction Between Path Integration and Environmental Cues," by K. M. Gothard, W. E. Skaggs, and B. L. McNaughton, 1996, *Journal of Neuroscience, 16,* p. 8031. Copyright 1996 by the Society for Neuroscience. Reprinted with permission.

destination end. At some point past the midpoint of the track, there will be a transition in the cues—from the cues behind the rat to the cues in front of the rat—controlling place cell firing. For the shortest track conditions, some place cells with fields near to the "transition point" may not fire at all, having roughly equal inputs from both ends on the full-length track, which entirely fail to overlap on the short track. In this case, the inferred location of the rat will jump from one reference frame to the other rather than making a smooth transition.

Before describing our simulations of Gothard et al.'s (1996) experiment in detail, we note one further piece of data. The preceding explanation predicts that if sensory information about the nearer than expected destination end of the track is degraded, then the internally updated representation of landmark positions should take precedence in the control of place cell firing for an even longer portion of the journey. Consistent with this, when rats performed Gothard et al.'s linear track task in darkness, it was found that the cue from which the rat was running maintained control over place cell firing for a greater portion of the journey than it did in the light (Gothard et al., 2001).

Method

To simulate the key aspects of the linear track environment of Gothard et al. (1996), we trained our model on a symmetric environment consisting of two "boxes" that open toward each other, as in the lower left/middle panels of Figure 21. Because of the absence of surrounding room cues, either box can be considered the movable box. In this way, we were able to perform one set of simulations representing both outbound and inbound journeys. Medial temporal and parietal connections were set in the same manner as for the previous simulations. Before performing actual simulations of the Gothard et al. data, the forward translational velocity of the place cell representation under application of an egocentric velocity signal had to be calibrated. This was accomplished by applying the velocity signal after cuing the model to localize itself near Box 1, facing Box 2 (see Figure 21), until place cell firing indicated localization near Box 2. The model's representation of its own location within the environment was calculated at any given instant by averaging the coordinates associated with maximally active place cells. By fitting a regression line to the roughly linear position-time data (see the rightmost panel of Figure 21), a velocity of 0.044 space units per time unit was found. Such a simulation would correspond to the model mentally exploring this familiar environment or performing spatial updating during actual locomotion in the absence of visual cues.

In the next step of the simulation, the model was cued to a location two units away from Box 1 along the direction toward Box 2, facing Box 2. To simulate a shortened track, sensory input

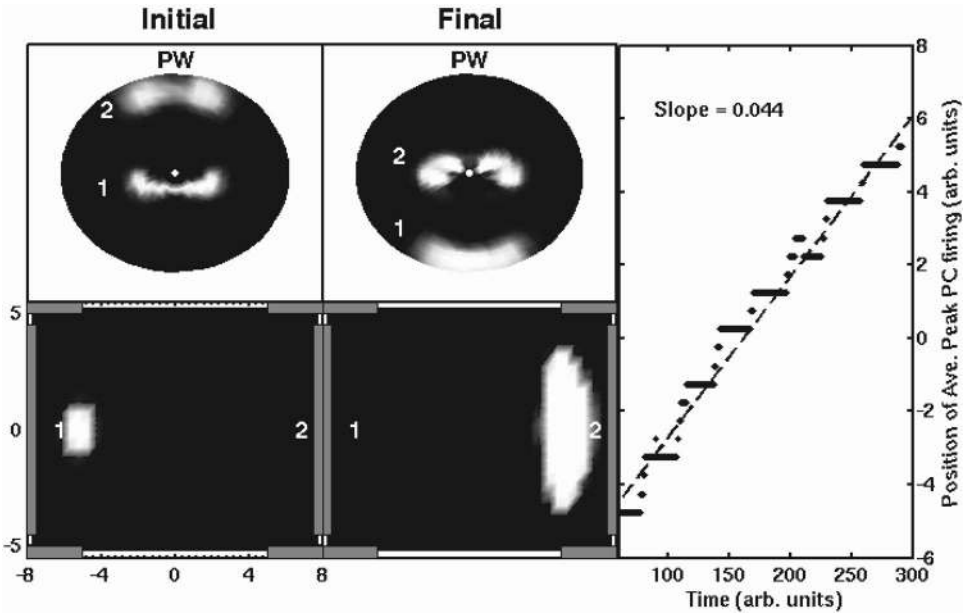

Figure 21. Left, top/bottom: Activation in parietal window (PW)/hippocampal neurons near the beginning of a top-down phase after the model was cued to localize itself 2 units away from Box 1 facing Box 2. Environmental boundaries are represented by gray walls superimposed on the hippocampal representation. Middle, top/bottom: Activation in PW/hippocampal layer near the beginning of a bottom-up phase after application of forward velocity signal. Right: The model's representation of its location within the environment as a function of time. Arb. = arbitrary; Ave. = average; PC = place cell.

corresponding to Box 2 was applied directly to the parietal window layer at either 0, 2, 4, 6, or 7 units closer to the egocentric origin than what would be consistent with the model's learned representation for that location (see the top and bottom panels at the left of Figure 22 for an example). For our initial set of simulations, sensory information corresponding to Box 1 was not applied because we assumed that this landmark did not have the salience of the target landmark and a rat's field of view is only approximately 300°. Locomotion was simulated by turning on the forward velocity signal (corresponding to a velocity of 0.044 space units/time unit) and moving the sensory input corresponding to Box 2 toward the origin of the parietal window coordinate system at the same speed. When this sensory input came within one unit of the origin, its movement was stopped, the velocity signal was turned off, and the model was allowed to relax for 50 time steps before sensory input was down regulated.

During locomotion, the rat's head tends to bob up and down, so that it might receive visual information from Box 1. With this in mind, we performed a second set of simulations identical to those just described but with input representing Box 1 also being applied to the parietal window component of the model. For these simulations, the additional input representing Box 1 was initially configured so as to represent this landmark at 2 units behind the animal. During simulated locomotion, this "sensory" input was moved through the parietal window coordinate system at the same speed and in the same direction as the input representing Box 2.

Finally, we performed simulations identical to those above but with weakened overall connection strengths for the connections terminating on the BVC layer (see Table 1 for parameter values). The motivation for this was that a smaller proportion of space was filled with landmark segments in the linear-track environment than in the previous two environments. This was found to result in a very low-resolution representation of space due to reduced lateral inhibition in the BVC, transformation, and parietal window layers. However, results for both sets of simulations (with and without weakened parameters) are qualitatively similar, except for one difference as discussed below. Furthermore, a more realistic simulation in which the BVC and parietal window layers covered a more extensive region of space would have allowed for the inclusion of distal landmarks (room walls, etc.). Such inclusion would have generated increased lateral inhibition and a sharper representation of space without the need for altering any connection strengths.

Results and Discussion

Results for the 6-unit-closer trial with no Box 1 sensory information are shown in Figure 22. Of particular interest is the fact that the maximum velocity of the place cell activity was 0.058 space units/time unit or about 32% faster than when no inconsistent sensory input was present (see the rightmost panel of Figure 22). Therefore, as with the data reported by Gothard et al. (1996), place cell activity was initially under the control of the nearest landmark, but during locomotion it "caught up" to what it should have been had it been primarily under the influence of the target landmark (Box 2).

In addition to recording the trajectory of place cell activity, the activity of 11 cells, representing equally spaced locations within the environment, were recorded. If the simulation trials are considered as outward journeys, then we can plot the firing profiles in

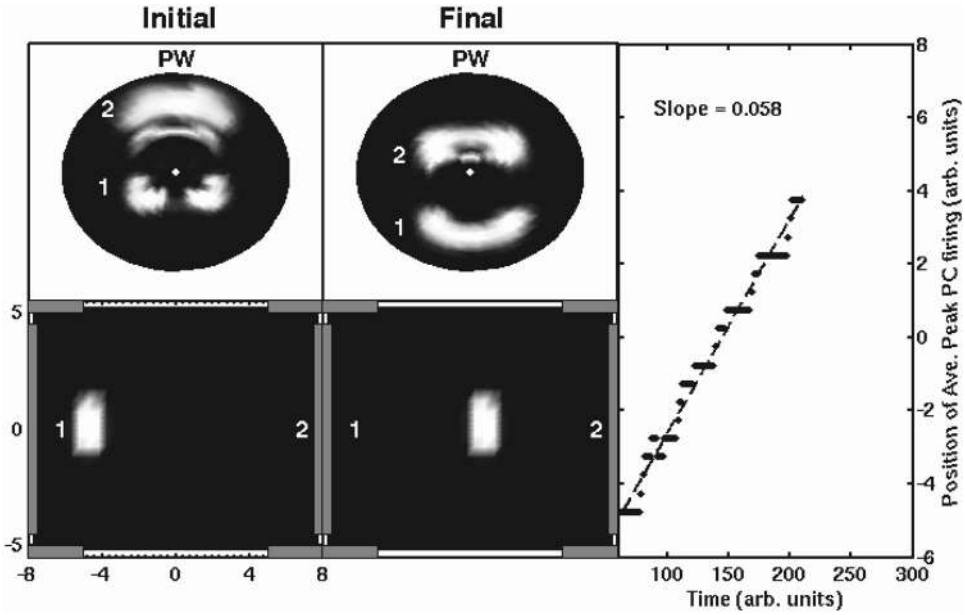

Figure 22. Left, top/bottom: Activation in parietal window (PW)/hippocampal neurons near the beginning of a top-down phase after the model was cued to localize itself 2 units away from Box 1 facing Box 2. Additional activation has been applied directly to PW neurons representing Box 2 at a position 6 units closer to the origin than expected. Environmental boundaries are represented by gray walls superimposed on the hippocampal representation. Middle, top/bottom: Activation in PW/hippocampal layer near the beginning of a top-down phase after the model comes within 1 unit of Box 2. At this point, the velocity signal is switched off, and the sensory input ceases to move. Right: The model's representation of its location within the environment as a function of time. Arb. = arbitrary; Ave. = average; PC = place cell.

Table 1
Model Parameters

Parameter	Value
ν^{α}	5 (50 for the inhibitory interneuron)
$\phi_{inh}{}^{H}$	2.1
$\phi_{inh}{}^{PR}$	9
$\phi_{inf}{}^{BVC}$	0.2
$\phi_{inh}{}^{HD}$	6
$\phi_{inh}{}^{TR}$	0.1
$\phi_{inh}{}^{PW}$	0.1
ϕ^{H}	21
$\phi^{H,BVC}$	140
$\phi^{H,PR}$	25
$\phi^{BVC,H}$	900[a]
$\phi^{BVC,PR}$	1
$\phi^{PR,H}$	6,000
$\phi^{PR,BVC}$	75
$\phi^{TR,BVC}$	54
$\phi^{TR,PW}$	63
$\phi^{BVC,TR}$	900[b]
$\phi^{PW,TR}$	880
ϕ^{HD}	15
$\phi^{TR,HD}$	85
$\phi^{TR,I}$	90
$\phi^{I,HD}$	10
$\phi^{\omega\times HD}$	2
$\phi^{\nu\times TR}$	$\phi^{PW,TR}$

[a] Decreased to 150 for weakened boundary vector cell (BVC) input simulation on linear track. [b] Decreased to 540 for weakened BVC input simulation on linear track.

a way similar to that used by Gothard et al. (1996) to calculate displacement slopes. In Figure 23, the firing profiles for 4 of the 11 recorded place cells in the condition with no Box 1 sensory information are shown along with displacement slopes for all 11 in both conditions. The same information is plotted in Figure 24 for the weak BVC input simulations. For the weak BVC input condition, place cell activity of the navigating model in the shortest track-length trial hopped from one representation of location within the longest environment to another, resulting in a complete lack of firing from one of the four selected cells. Given the symmetry of our environment, displacement slope data can be determined for inward journeys by transforming the data for outward journeys as follows:

$$DS(x) \rightarrow 1 - DS(1 - x), \quad (1)$$

where $DS(x)$ is the displacement slope for a neuron with peak firing position, x, in the Box 1 condition, and x is normalized to range between 0 (at the movable box) and 1 (at the fixed food cup). The transformed curves are shown in the lower right panel of Figures 23 and 24. Notice that both sets of simulation-generated displacement slopes show patterns consistent with Gothard et al.'s results. In particular, the landmark that the animal is moving away from maintains considerable control over place cell firing until the target landmark is nearly reached. For the normal BVC input conditions, this effect is similar regardless of whether we assume the animal has access to sensory information from both Box 1 and Box 2. For the weak BVC input simulations, we obtain a stronger effect if we assume the model has sensory input from both boxes.

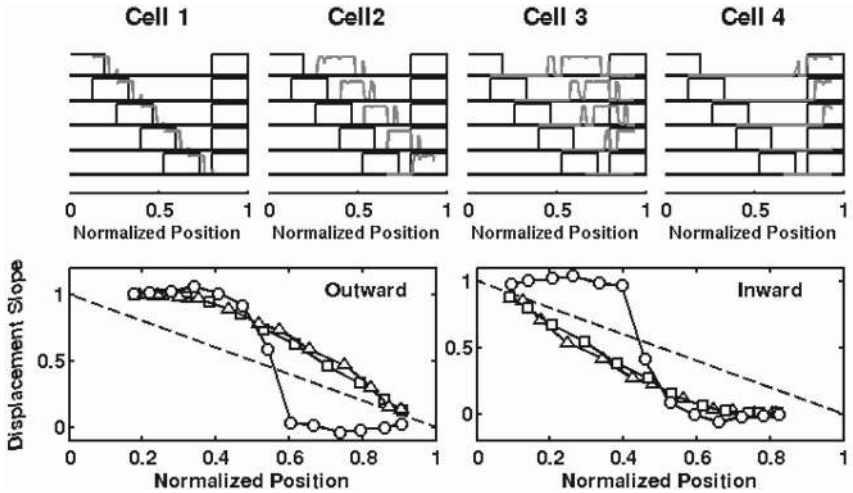

Figure 23. Top: Activity from 4 of 11 selected model place cells (maximal firing coordinates for the selected cells: $x_i \in \{ -4.75 + (i-1)\}_{i=1}^{11}$ and $y_i = 0.25$ for all i in five simulated conditions (Box 1–Box 5, without Box 1 sensory input) plotted against relative position in the longest track-length condition (Box 1 condition). Rectangles represent Box 1 and Box 2. Bottom, left/right: Displacement slopes calculated from the 11 sampled model place cells during outward/inward journeys. Open squares represent results from full-model simulations with only Box 2, and triangles represent results from full-model simulations with Box 1 and Box 2 sensory input. Circles represent results from the simple boundary vector cell explanation. The dashed line is what would be expected if landmarks exerted control over place cell firing in direct proportion to their proximity to the animal.

In summary, our model performs in a manner consistent with the Gothard et al. (1996) data. In a subsequent experiment, the influence of the cue from which the rat is running was seen to last for a constant time, rather than for a constant distance, through the run (Redish et al., 2000). This indicates either a time-limited usefulness for path integration (see, e.g., Etienne, Maurer, & Seguinot, 1996), or (as argued for in Redish et al., 2000) some temporal inertia in place cell firing that is possibly due to attractor dynamics (which can be seen under other experimental circumstances; e.g., Wills, Lever, Cacucci, Burgess, & O'Keefe, 2005). Simulations comparing time and distance in this way were not performed (we used constant velocity) and remain for future work.

Finally, we compared our full model with a model lacking path integration. By considering only the part of the full model con-

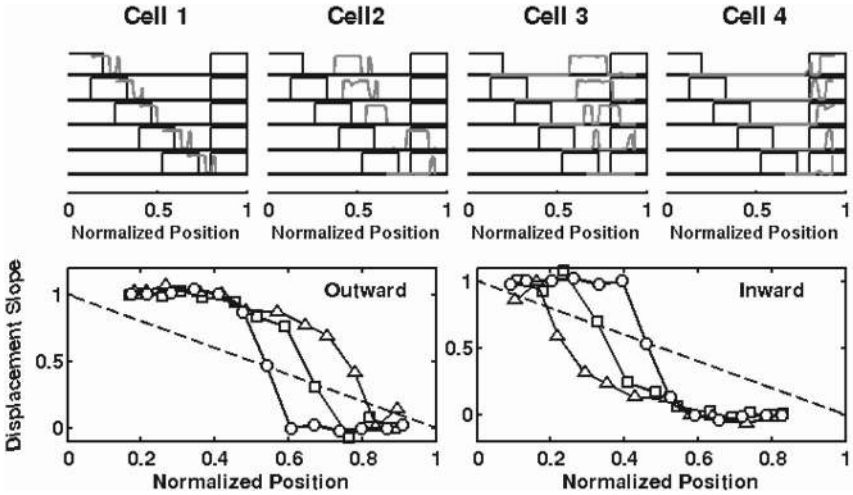

Figure 24. Results for the simulations with weakened boundary vector cell input parameters. Note the hopping behavior of place cell activity in the shortest track-length condition. Open squares represent results from full-model simulations with only Box 2, and triangles represent results from full-model simulations with Box 1 and Box 2 sensory input. Circles represent results from the simple boundary vector cell explanation. The dashed line is what would be expected if landmarks exerted control over place cell firing in direct proportion to their proximity to the animal.

sisting of the BVCs, the place cells, and the feed-forward connections from the BVC to the place cell layer, we were able to verify that the simple BVC explanation of Gothard et al.'s (1996) results does not produce the noted asymmetry. Specifically, we simulated navigation along each track length by providing direct input to the BVC neurons corresponding to the Box 1 and Box 2 landmarks and then translated this input through the BVC coordinate system at 0.044 space units/time unit. In this way, BVCs, and hence place cells, were driven directly by sensory input, and the model's current representation of space was not affected by previous representations of space or idiothetic information. Displacement slope curves for these simulations were calculated as above and plotted in the lower two panels of Figure 23 and 24. Notice that these curves are approximately symmetric about the midpoint of the full-length track. Thus the simple BVC model, in which distances to boundaries in allocentric directions are the only concern, is insufficient to produce the dependence on running direction noted in Gothard et al. (1996), O'Keefe and Burgess (1996), or Redish et al. (2000).

In the current model, perceptual inputs and motion-related updating combine to influence the animal's internal representation of location, and the operation of this mechanism seems to be consistent with the relevant existing data from place cell recording. The functional architecture of the current model was largely informed by thinking about imagery and planning in human spatial memory; however the simulations reported here indicate that it is also able to explain data at the single-unit level of description.

General Discussion

We have outlined a model of the neural mechanisms underlying spatial cognition, focusing on long-term and short-term spatial memory and imagery, egocentric and allocentric representations, visual and ideothetic information, and the interactions between them. We proposed specific mechanisms by which long-term spatial memory results from attractor dynamics within a set of medial temporal allocentric representations, whereas short-term memory results from egocentric parietal representations driven by perception, retrieval, and imagery, and can be investigated by directed attention. However, perhaps our main novel contribution is to propose specific mechanisms by which these systems interact. Thus we propose that encoding and retrieval require translation between the egocentric and allocentric representations, which occurs via a coordinate transformation in the posterior parietal and retrosplenial cortices and reflects the current head direction. In our model, the hippocampus effectively indexes information by real or imagined location, allowing reconstruction of the set of visual textures and distances and allocentric directions of landmarks consistent with being at a single location (see also King et al., 2004). In turn, Papez's circuit translates this representation into an egocentric representation suitable for imagery according to the direction of view (and also translates from egocentric perception during encoding of the allocentric representation). For partially related models, see Becker and Burgess (2001); Burgess, Becker, et al. (2001); Recce and Harris (1996); and Redish (1999). We further propose that modulation of the allocentric-to-egocentric translation by motor efference allows "spatial updating" of egocentric parietal representations, which in turn can feedback to cause updating of the medial temporal representations. Finally, the generation of mock motor efference (e.g., representing planned eye movements) in the prefrontal cortex allows mental exploration in imagery, making a potential contribution to spatial planning. The temporal coordination of the alternating interaction of the temporal and parietal regions was assumed to be provided by the theta rhythm.

For concreteness, and to demonstrate the actual ability of the theory to bridge between single-neuron and systems neuroscience and behavioral data, we implemented it as a fully specified neural network simulation for the case of long-term, hippocampally dependent, spatial memory and its interaction with short-term working memory and imagery. Our simulations provide straightforward explanations for a number of experimental results. The first provides a neural implementation of the idea that representational neglect results from a damaged egocentric window into an intact long-term spatial memory system (see also Baddeley & Leiberman, 1980). From the model architecture, we are able to suggest that unilateral lesions to the precuneus, retrosplenial cortex, parietal area 7a, areas connecting 7a or the retrosplenial cortex with the parahippocampal gyrus, or combinations of these areas have the potential to generate representational neglect. However, currently available patient data makes this prediction difficult to test. The second simulation provides a neural implementation of self-motion-related spatial updating of object locations in memory and of imagined navigation and route planning. The third shows that our interpretation of the role of head direction in memory is consistent with the effects of lesions to the head direction system on single-unit responses in the hippocampus. With this interpretation, we are also able to make two simple predictions about the outcomes of similar experiments, thus allowing the translation component of our model to be tested directly. The final simulation shows that our proposed mechanism for integrating sensory information and self-motion also provides an explanation for single-unit responses in situations of conflicting sensory and ideothetic information (Gothard et al. 1996). In the following, we discuss the implications, predictions, and limitations of the model with respect to the wider literature on the neural bases of spatial cognition and memory more generally.

Temporal–Parietal Interactions, Planning, and Imagery

Our specific model of the temporal–parietal interaction has some straightforward implications for functional anatomy. Thus, it explains why Papez's (mammillar–anterior–thalamic–medial–temporal) circuit is required for episodic recollection into rich visuospatial imagery (Aggleton & Brown, 1999) and also provides the head direction signal in rats (Taube, 1998). It also suggests a functional role for the retrosplenial cortex and the intraparietal sulcus, which are well positioned to integrate or buffer the translation between egocentric and allocentric representations (Burgess, Becker, et al., 2001) or, correspondingly, between path-integrative and mnemonic information (Cooper & Mizumori, 2001). Cooper and Mizumori (2001) and Maguire (2001) have provided evidence that lesions to the retrosplenial cortex, an area interconnected with the parietal and medial temporal regions (Kobayashi & Amaral, 2003; Wyss & Groen, 1992), do indeed impair the navigation of rats and humans under such circumstances. In humans, the intimate link between spatial imagery and navigation is made clear by the correlation of impairments in these two faculties following

unilateral damage (Guariglia et al., 2005). Finally, our model proposes a role for the theta rhythm in coordinating the flow of information between medial temporal and parietal components of the model. Thus, top-down activation from medial temporal to parietal areas occurs at one phase of theta, whereas bottom-up activation from parietal to medial temporal areas occurs at the opposite phase of theta. A related proposal relates hippocampal encoding and retrieval to opposing phases of theta (e.g., Hasselmo et al., 2002), corresponding to our bottom-up and top-down phases, respectively. In our model, spatial updating occurs over repeated top-down and bottom-up cycles, as each (top-down) translation from allocentric to egocentric representations maps to locations adjusted for the subject's velocity and then passes (bottom-up) back to update the allocentric representation.

In order to plan routes through complex environments, the brain must make use of long-term memories of the layout of those environments. Route planning also requires the ability to perform mental navigation: to imagine both moving in a given direction and the consequences of that action. Thus the task in our second set of simulations, involving mentally generating a velocity signal or "mock motor efference," could be viewed as mental exploration of a familiar environment. This exploration would be useful for path planning and many other tasks. For example, this may be how people accomplish the task of Wang and Brockmole (2003). Recall that in this task subjects were led along a path through a familiar environment and asked to point to occluded landmarks at various predetermined times. It was found that when subjects could not accurately point to a given landmark, they often could do so if allowed to walk to some point further along the path from which the landmark was still occluded. Within the framework of our model, subjects may have been mentally navigating from their current location to a location from which the occluded landmark was visible. By integrating the direction of the mentally generated velocity signal, a pointing direction could be generated. However, if the mental path was too long or complex, then the calculation would be swamped by cumulative error. In physically moving further along the path, subjects may have been simplifying the task by reducing the amount of mental navigation required.

Within the framework of route planning, a final prediction of the model presented here is that damage to connections between parietal and medial temporal cortices would impair the ability of an organism to navigate to occluded landmarks in familiar environments. This is because, without access to long-term spatial memory, the parietally supported egocentric window would only have access to short-term memory and direct sensory information, rendering the organism unable to mentally explore the familiar environment beyond regions very recently encountered. Equally, we might expect to see increased theta coherence between temporal and parietal regions as a function of this type of actual, or mental, navigation.

Differences Between Spatial Updating and Path Integration in the Temporal and Parietal Cortices

Path integration can be defined as the ability of an organism to keep track of its current location, on the basis of idiothetic information alone, relative to its starting point as it moves around, whereas *spatial updating* refers to the ability to also keep track of other locations, again on the basis of idiothetic information alone, within the environment (see, e.g., Etienne et al., 1998; Loomis et al., 1993; Mittelstaedt & Mittelstaedt, 2001; Morrongiello, Timney, Humphrey, Anderson, & Skory, 1995). However, either process could operate by individually updating the required egocentric location(s) relative to the organism or by updating an allocentric representation of the organism's location relative to the environment. Both types of updating are probably available in parallel, with the former suitable for small numbers of locations and short movements and the latter for updating multiple locations and longer movements, when perceptual support from the environment is unavailable. Thus spatial updating over short timescales and small movements (e.g., less than 135° rotation) in unfamiliar environments appears to operate on transient egocentric parietal representations, showing independent accumulations of errors in the locations of different objects (Waller & Hodgson, 2006; Wang & Spelke, 2000). In contrast, spatial updating over longer durations or movements or in very familiar environments appears to operate on a coarser but enduring allocentric representation (Mou, McNamara, Rump, & Xiao, 2006; Waller & Hodgson, 2006). See Burgess (2006) for further discussion.

Corresponding to these two types of spatial updating, separate models have been proposed for the mechanisms within each (temporal or parietal) region. Byrne and Becker (2004) proposed a purely parietal mechanism for motion-related updating of the egocentric locations in the parietal window, which would be consistent with single-unit recording and effects of lesions within this region (see the present introductory section). On the other hand, strictly medial temporal mechanisms have been proposed for updating the location of the subject relative to the environment (see, e.g., Howard, Fotedar, Datey, & Hasselmo, 2005; O'Keefe & Nadel, 1978; Redish, Rosenzweig, Bohanick, McNaughton, & Barnes, 2000; Samsonovich & McNaughton, 1997). These latter models are supported by the recently discovered "grid cells" in the entorhinal cortex (Hafting et al., 2005), which appear well suited to this task, with the hippocampus potentially required when path integration has to be tied to environmental locations (O'Keefe & Burgess, 2005; McNaughton et al., 2006). See Whishaw and Brooks (1999) and Save, Guazzelli, and Poucet (2001) for related discussion of the hippocampal contribution to path integration.

Our model primarily concerns the interaction of parietal and medial temporal representations and assumes a single spatial updating mechanism derived as an extension of this interaction. Our second set of simulations provides a detailed mechanism by which the parietal cortex might make use of stored spatial representations in the medial temporal lobe to provide egocentric representations of an arbitrary number of locations within a familiar environment and to update these locations following real or imagined self motion. Other tasks (such as pointing to a recently seen object or imagery for objects or actions as opposed to environmental layout) will be purely parietal and are not addressed by our model. Even within tasks that depend on both regions, such as those simulated, our model will not capture the finer distinctions between spatial updating driven more strongly by one region than the other. Similarly, we do not distinguish the processing of discrete objects, likely more strongly represented in parietal areas, from the processing of extended boundaries, likely key to driving the hippocampal representation. The BVC representation used provides the appropriate dependence of hippocampal representations on

environmental geometry but probably does not correspond so well to some aspects of egocentric parietal representations.

The provenance of the model. We have presented a working model of spatial cognition without really addressing how the brain might have "learned" such a solution. Although a number of models of hippocampal learning have been presented (see, e.g., Becker, 2005), principles underlying the learning of egocentric–allocentric transformations have not been firmly established. In recent work, we have attempted to elucidate more biologically realistic principles upon which such learning could be based (Byrne & Becker, 2006). Specifically, we have proposed two relatively simple learning principles that, when applied to a transformation circuit similar to the one presented here, reliably result in the generation of allocentric representations of space. The first principle is that of minimum reconstruction error. That is, for a given heading direction, the representation produced at the medial temporal lobe level should, through top-down connections, be able to reproduce the corresponding egocentric input. The second principle is the maximization of temporal inertia in medial temporal representations. This is motivated by empirical evidence that both hippocampal pyramidal cells (Redish, McNaughton, & Barnes, 2000) and, under certain circumstances, superficial (Klink & Alonso, 1997) and deep layer (Egorov, Hamam, Fransén, Hasselmo, & Alonso, 2002) entorhinal cells exhibit a resistance to rapid changes in firing rate. We speculate that spatial representations that vary as little as possible in time should maximize accuracy and precision in storage, as well as allowing more rapid spatial updating or mental exploration, because the medial temporal representations would have to vary less rapidly to keep up with the retrieval demands. We have tested the utility of these learning principles in two very different models, one trained by direct minimization of a cost function by using steepest descent learning and one consisting of a coupled network of restricted Boltzmann machines trained sequentially by contrastive Hebbian learning (Hinton, 2002; Hinton et al., 2006). Both models were able to learn allocentric representations of space at the medial temporal lobe output layer and to generate good reconstructions of the egocentric input layer.

Implications beyond spatial memory. Although we have concentrated on the role of the hippocampus in spatial memory, this structure is also known to be important in the maintenance of more general episodic memories (for recent reviews, see, e.g., Becker, 2005; Burgess et al., 2002; Eichenbaum, 2001; for models see Howard et al., 2005; Marr, 1971; McClelland, McNaughton, & O'Reilly, 1995; McNaughton & Morris, 1987; Treves & Rolls, 1992). In our model, hippocampal place cells bind the outputs of various BVCs and visual feature units together to form an allocentric map of an environment. The attractor dynamics of the medial temporal system then performs retrieval by allowing only those conjunctions of visual feature, distance, and allocentric direction that are consistent with being in a single location (represented in the hippocampus). This information is then rotated, with the aid of Papez's circuit, to form an egocentric parietal image for conscious inspection that corresponds to a specific direction of view. Our model is highly consistent with the pattern of fMRI activation in retrieving the spatial context of an event (Burgess, Maguire, et al., 2001; King, Hartley, Spiers, Maguire, & Burgess, 2005). Having defined this functional anatomy in the context of spatial memory, we suspect similar processing occurs much more generally during any detailed mental imagery for environmental layouts derived from long-term knowledge. This would be consistent with reports of deficits in detailed imagery for novel or future events in amnesic patients (Hassabis, Kumaran, Vann, & Maguire, 2006; Klein, Loftus, & Khilstrom, 2002; but see also Bayley, Gold, Hopkins, & Squire, 2005) and similar patterns of activation for thinking about past and future events (Addis, Wong, & Schacter, 2006; Okuda et al., 2003). This function might relate to characterizations of episodic or autobiographical memory in terms of retrieval of rich contextual information or feelings of "reexperiencing," as distinct from the imagery for simple objects and actions which is preserved in amnesia (e.g., Rosenbaum, McKinnon, Levine, & Moscovitch, 2004).

For simplicity, our simulations concerned a single familiar environment. However, retrieval from the best matching of several familiar environments could be mediated, as described by our model, by distinct subsets of place cells (McNaughton & Morris, 1987; Samsonovich & McNaughton, 1997), providing a distinct attractor representation of each environment (Wills et al., 2005). In this way, the hippocampus might be described as providing the spatial context appropriate to recollection (O'Keefe & Nadel, 1978), explaining its role, for example, in context-dependent fear conditioning but not in fear conditioning itself (Kim & Fanselow, 1992; Phillips & LeDoux, 1992). An interesting prediction here is that two situations can be identified as having different "contexts" requiring hippocampal disambiguation, that is if they elicit "remapped" (Muller, 1996) patterns of place cell firing as occurs rapidly with dramatic multimodal changes (Wills et al., 2005) or more slowly with unimodal changes (Lever et al., 2002).

Of course, hippocampal neurons are probably not limited to the spatial functions we have focused on here. For example, rat CA1 and CA3 pyramidal neurons can also respond to various nonspatial cues (see, e.g., Huxter, Burgess, & O'Keefe, 2003; Young, Fox, & Eichenbaum, 1994). This ability to connect nonspatial and spatial information may allow the association of location within an environment to various other elements of experience, that is providing a spatial–temporal context to support context-dependent episodic memory more generally (see, e.g., chaps. 14 and 15 in O'Keefe & Nadel, 1978). We also note that the ability to perform spatial updating of the imagined viewpoint may both aid the process of search during episodic retrieval and the binding of places into remembered trajectories, or sequences, in memory for more extended dynamic episodes (see also Howard et al., 2005; Jensen & Lisman, 1996; W. Levy, 1996; Wallenstein, Eichenbaum, & Hasselmo, 1998). Howard et al.'s (2005) temporal context model (TCM) of memory for lists of items provides an example of how such associations across time might occur. The TCM works by associating items to a slowly varying context representation containing history-dependent information relating to the items themselves. Howard et al. noted that this model is broadly compatible with a spatial function for the medial temporal lobe in providing a mechanism for path integration by representing the recent history of movements. In our model, the medial temporal lobe could be thought of as providing the spatial context of events by representing the actual surrounding spatial scene. Generation of more general representations of context, such as temporal contexts, would be one way in which our model might be extended to include the involvement of the medial temporal lobe in memories for trajectories through space or in nonspatial memory.

Finally, although we have concentrated on spatial memory, the question of how long-term memory and short-term or working memory interact is equally pertinent to nonspatial memory. For example, although much has been learned about both long-term and working memory for verbal stimuli, the interaction of these two systems is a topic of much current interest (e.g., Baddeley, 2000; Burgess & Hitch, 2005). By staying within the spatial domain, where there is much data at the single-unit level, we have provided a detailed model of one form of the interaction between long-term medial temporal and short-term parietal systems. However, our proposals for the functional roles and interactions of the regions in question should generalize to the generation of dynamic visuospatial imagery from stored verbal knowledge. Given the slight lateralization of visuospatial processing to the right hemisphere (e.g., Piggott & Milner, 1993; Smith & Milner, 1989; reviewed in Burgess et al., 2002), we would hope that some of the mechanisms considered here might generalize to the interaction of left medial temporal lobe long-term memory systems for narrative memory (e.g., Frisk & Milner, 1990) and parietal short-term memory systems for verbal working memory.

References

Abrahams, S., Pickering, A., Polkey, C. E., & Morris, R. G. (1997). Spatial memory deficits in patients with unilateral damage to the right hippocampal formation. *Neuropsychologia, 35,* 11–24.

Addis, D. R., Wong, A. T., & Schacter, D. L. (2007). Remembering the past and imagining the future: Common and distinct neural substrates during event construction and elaboration. *Neuropsychologia, 45,* 1363–1377.

Aggleton, J., & Brown, M. (1999). Episodic memory, amnesia and the hippocampal-anterior thalamic axis. *Behavioral and Brain Science, 22,* 425–444.

Aguirre, G. K., & D'Esposito, M. (1999). Topographical disorientation: A synthesis & taxonomy. *Brain, 122,* 1613–1628.

Alyan, S., & McNaughton, B. (1999). Hippocampectomized rats are capable of homing by path integration. *Behavioral Neuroscience, 113,* 19–31.

Andersen, R., Essick, G., & Siegel, R. (1985, October 25). The encoding of spatial location by posterior parietal neurons. *Science, 230,* 456–458.

Andersen, R., Shenoy, K., Snyder, L., Bradley, D., & Crowell, J. (1999). The contributions of vestibular signals to the representations of space in the posterior parietal cortex. *Annals of the New York Academy of Sciences, 871,* 282–292.

Baddeley, A. (2000). The episodic buffer: A new component of working memory. *Trends in Cognitive Sciences, 4,* 417–423.

Baddeley, A., & Leiberman, K. (1980). Spatial working memory. In S. Nickerson, *Attention and performance VIII* (pp. 521–539). Hillsdale, NJ: Erlbaum.

Barry, C., Lever, C., Hayman, R., Hartley, T., Burton, S., O'Keefe, J., et al. (2006). The boundary vector cell model of place cell firing and spatial memory. *Reviews in the Neurosciences, 17,* 71–97.

Battaglia, F. P., Sutherland, G. R., & McNaughton, B. L. (2004). Local sensory cues and place cell directionality: Additional evidence of prospective coding in the hippocampus. *Journal of Neuroscience, 24,* 4541–4550.

Bayley, P. J., Gold, J. J., Hopkins, R. O., & Squire, L. R. (2005). The neuroanatomy of remote memory. *Neuron, 46,* 799–810.

Becker, S. (2005). A computational principle for hippocampal learning and neurogenesis. *Hippocampus, 15,* 722–738.

Becker, S., & Burgess, N. (2001). A model of spatial recall, mental imagery and neglect. In T. Leen, T. Ditterich, & V. Tresp (Eds.), *Advances in neural information processing systems* (Vol. 13, pp. 96–102). Cambridge, MA: MIT Press.

Behrmann, M., Watt, S., Black, S. E., & Barton, J. J. S. (1997). Impaired visual search in patients with unilateral neglect: An oculographic analysis. *Neuropsychologia, 35,* 1445–1458.

Beschin, N., Basso, A., & Della Sala, S. (2000). Perceiving left and imagining right: Dissociation in neglect. *Cortex, 36,* 401–414.

Bird, C. M., Malhotra, P., Parton, A., Coulthard, E., Rushworth, M. F., & Husain, M. (2006). Visual neglect after right posterior cerebral artery infarction. *Journal of Neurology, Neurosurgery & Psychiatry, 77,* 1008–1012.

Bisiach, E., & Luzzatti, C. (1978). Unilateral neglect of representational space. *Cortex, 14,* 129–133.

Bohbot, V., Kalina, M., Stepankova, K., Spackova, N., Petrides, M., & Nadel, L. (1998). Spatial memory deficits in patients with lesions to the right hippocampus and to the right parahippocampal cortex. *Neuropsychologica, 36,* 1217–1238.

Bremmer, F., Klam, F., Duhamel, J.-R., Hamed, S., & Graf, W. (2002). Visual-vestibular interactive responses in the macaque ventral intraparietal area (VIP). *European Journal of Neuroscience, 16,* 1569–1586.

Brun, V., Otnaess, M., Molden, S., Steffenach, H.-A., Witter, M., Moser, M.-B., & Moser, E. (2002, June 21). Place cells and place recognition maintained by direct entorhinal-hippocampal circuitry. *Science, 296,* 2243–2246.

Burgess, N (2006). Spatial memory: How egocentric and allocentric combine. *Trends in Cognitive Science, 10,* 551–557.

Burgess, N., Becker, S., King, J., & O'Keefe, J. (2001). Memory for events and their spatial context: Models and experiments. *Philosophical Transaction of the Royal Society of London B: Biological Sciences, 356,* 1493–1503.

Burgess, N., & Hitch, G. (2005). Computational models of working memory: Putting long-term memory into context. *Trends in Cognitive Sciences, 9,* 535–541.

Burgess, N., Jeffery, K., & O'Keefe, J. (Eds.). (1999). *The hippocampal and parietal foundations of spatial cognition.* Oxford, England: Oxford University Press.

Burgess, N., Maguire, E. A., & O'Keefe, J. (2002). The human hippocampus and spatial and episodic memory. *Neuron, 35,* 625–641.

Burgess, N., Maguire, E. A., Spiers, H., & O'Keefe, J. (2001). A temporoparietal and prefrontal network for retrieving the spatial context of lifelike events. *Neuroimage, 14,* 439–453.

Burgess, N., & O'Keefe, J. (1996). Neuronal computations underlying the firing of place cells and their role in navigation. *Hippocampus, 6,* 749–762.

Burgess, N., Spiers, H., & Paleologou, E. (2004). Orientational manoeuvres in the dark: Dissociating allocentric and egocentric influences on spatial memory. *Cognition, 94,* 149–166.

Byrne, P., & Becker, S. (2004). Modelling mental navigation in scenes with multiple objects. *Neural Computation, 16,* 1851–1872.

Byrne, P., & Becker, S. (2006). *A principle for learning egocentric–allocentric transformation*. Manuscript submitted for publication.

Calton, J. L., Stackman, R. W., Goodridge, J. P., Archey, W. B., Dudchenko, P. A., & Taube, J. S. (2003). Hippocampal place cell instability after lesions of the head direction cell network. *Journal of Neuroscience, 23,* 9719–9731.

Caplan, J. B., Madsen, J. R., Schulze-Bonhage, A., Aschenbrenner-Scheibe, R., Newman, E. L., & Kahana, M. J. (2003). Human theta oscillations related to sensorimotor integration and spatial learning. *Journal of Neuroscience, 23,* 4726–4736.

Chafee, M., & Goldman-Rakic, P. (1998). Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memory task. *Journal of Neurophysiology, 79,* 2919–2940.

Chen, L. L., Lin, L. H., Barnes, C. A., & McNaughton, B. L. (1994). Head direction cells in rat posterior cortex II: Contributions of visual and

idiothetic information to the directional firing. *Experimental Brain Research, 101,* 24–34.

Clower, D., West, R., Lynch, J., & Strick, P. (2001). The inferior parietal lobule is the target of output from the superior colliculus, hippocampus, and cerebellum. *Journal of Neuroscience, 21,* 6283–6291.

Colby, C. (1999). Parietal cortex constructs action-oriented spatial representations In N. Burgess, K. J. Jeffery, & J. O'Keefe (Eds.), *The hippocampal and parietal foundations of spatial cognition* (pp. 104–126). Oxford, England: Oxford University Press.

Commins, S., Gemmel, C., Anderson, M., Gigg, J., & O'Mara, S. (1999). Disorientation combined with parietal cortex lesions causes path-integration deficits in the water maze. *Behavioral Brain Research, 104,* 197–200.

Conklin, J., & Eliasmith, C. (2005). An attractor network model of path integration in the rat. *Journal of Computational Neuroscience, 18,* 183–203.

Cooper, B., Manka, T., & Mizumori, S. (2001). Finding your way in the dark: The retrosplenial cortex contributes to spatial memory and navigation without visual cues. *Behavioral Neuroscience, 115,* 1012–1028.

Cooper, B., & Mizumori, S. (2001). Temporary inactivation of the retrosplenial cortex causes a transient reorganization of spatial coding in the hippocampus. *Journal of Neuroscience, 21,* 3986–4001.

Crane, J., & Milner, B. (2005). What went where? Impaired object-location learning in patients with right hippocampal lesions. *Hippocampus, 15,* 216–231.

Cressant, A., Muller, R., & Poucet, B. (1997). Failure of centrally placed objects to control the firing fields of hippocampal place cells. *Journal of Neuroscience, 17,* 2531–2542.

Davachi, L., & Goldman-Rakic, P. (2001). Primate rhinal cortex participates in both visual recognition and working memory tasks: Functional mapping with 2-dg. *Journal of Neurophysiology, 85,* 2590–2601.

Ding, S., Van Hoesen, G., & Rockland, K. (2000). Inferior parietal lobule projections to the presubiculum and neighboring ventromedial temporal cortical areas. *Journal of Comparative Neurology, 425,* 510–530.

Diwadkar, V., & McNamara, T. (1997). Viewpoint dependence in scene recognition. *Psychological Science, 8,* 302–307.

Doricchi, F., & Tomaiuolo, F. (2003). The anatomy of neglect without hemianopia: A key role for parietal-frontal disconnection. *Neuroreport, 14,* 2239–2243.

Duhamel, J., Colby, C., & Goldberg, M. (1992, January 3). The updating of the representation of visual space in parietal cortex by intended eye movements. *Science, 255,* 90–92.

Duhamel, J., Colby, C., & Goldberg, M. (1998). Ventral intraparietal area of the macaque: Congruent visual and somatic response properties. *Journal of Neurophysiology, 79,* 126–136.

Easton, R., & Sholl, M. (1995). Object-array structure, frames of reference, and retrieval of spatial knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 483–500.

Egorov, A. V., Hamam, B. N., Fransén, E., Hasselmo, M. E., & Alonso, A. A. (2002, November 14). Graded persistent activity in entorhinal cortex neurons. *Nature, 420,* 173–178.

Eichenbaum, H. (2001). The hippocampus and declarative memory: Cognitive mechanisms and neural codes. *Behavioral Brain Research, 127,* 199–207.

Eichenbaum, H., & Cohen, N. J. (1988). Representation in the hippocampus: What do hippocampal neurons code? *Trends in Neurosciences, 11,* 244–248.

Ekstrom, A., Kahana, M., Caplan, J., Fields, T., Isham, E., Newman, E., & Fried, I. (2003, September 11). Cellular networks underlying human spatial navigation. *Nature, 425,* 184–187.

Epstein, R., & Kanwisher, N. (1998, April 9). A cortical representation of the local visual environment. *Nature, 392,* 598–601.

Etienne, A., Maurer, R., Berlie, J., Reverdin, B., Rowe, T., Georgakopoulos, J., & Seguinot, V. (1998, November 12). Navigation through vector addition. *Nature, 396,* 161–164.

Etienne, A., Maurer, R., & Seguinot, V. (1996). Path integration in mammals and its interaction with visual landmarks. *Journal of Experimental Biology, 199,* 201–209.

Fell, J., Klaver, P., Elfadil, H., Schaller, C., Elger, C. E., & Fernandez, G. (2003). Rhinal–hippocampal theta coherence during declarative memory formation: Interaction with gamma synchronization? *European Journal of Neuroscience, 17,* 1082–1088.

Fletcher, P., Shallice, T., Frith, C., Frackowiak, R., & Dolan, R. (1996). Brain activity during memory retrieval: The influence of imagery and semantic cueing. *Brain, 119,* 1587–1596.

Formisano, E., Linden, D., Salle, F. D., Trojano, L., Esposito, F., Sack, A., et al. (2002). Tracking the mind's image in the brain I: Time-resolved fMRI during visuospatial mental imagery. *Neuron, 35,* 185–194.

Frisk, V., & Milner, B. (1990). The role of the left hippocampal region in the acquisition and retention of story content. *Neuropsychologica, 28,* 349–359.

Fruhmann-Berger, M., & Karnath, H. O. (2005), Spontaneous eye and head position in patients with spatial neglect. *Journal of Neurology, 252,* 1194–1200.

Funahashi, S., Bruce, C., & Goldman-Rakic, P. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *Journal of Neurophysiology, 61,* 331–348.

Galati, G., Lobel, E., Vallar, G., Berthoz, A., Pizzamiglio, L., & LeBihan, D. (2000). The neural basis of egocentric and allocentric coding of space in humans: A functional magnetic resonance study. *Experimental Brain Research, 133,* 156–164.

Galletti, C., Battaglini, P. P., & Fattori, P. (1995). Eye position influence on the parieto-occipital area PO (V6) of the macaque monkey. *European Journal of Neuroscience, 7,* 2486–2501.

Georgopoulos, A. (1988). Neural integration of movement: Role of motor cortex in reaching. *FASEB Journal, 2,* 2849–2857.

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. *Neuro Report, 8,* 739–744.

Goodale, M., & Milner, A. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences, 15,* 20–25.

Goodridge, J., & Touretzky, D. (2000). Modeling attractor deformation in the rodent head direction system. *Journal of Neurophysiology, 83,* 3402–3410.

Gothard, K., Hoffman, K., Battaglia, F., & McNaughton, B. (2001). Dentate gyrus and ca1 ensemble activity during spatial reference frame shifts in the presence and absence of visual input. *Journal of Neuroscience, 21,* 7284–7292.

Gothard, K., Skaggs, W. E., & McNaughton, B. (1996). Dynamics of mismatch correction in the hippocampal ensemble code for space: Interaction between path integration and environmental cues. *Journal of Neuroscience, 16,* 8027–8040.

Graziano, M., & Gross, C. (1993). A bimodal map of space: Somatosensory receptive fields in the macaque putamen with corresponding visual receptive fields. *Experimental Brain Research, 97,* 96–109.

Guariglia, C., Piccardi, L., Iaria, G., Nico, D., & Pizzamiglio, L. (2005). Representational neglect and navigation in real space. *Neuropsychologica, 43,* 1138–1143.

Guazzelli, A., Bota, M., & Arbib, M. (2001). Competitive Hebbian learning and the hippocampal place cell system: Modeling the interaction of visual and path integration cues. *Hippocampus, 11,* 216–239.

Haarmeier, T., Their, P., Repnow, M., & Petersen, D. (1997, October 23). False perception of motion in a patient who cannot compensate for eye movements. *Nature, 389,* 849–852.

Hafting, T., Fyhn, M., Molden, S., Moser, M. B., & Moser, E. I. (2005,

August 11). Microstructure of a spatial map in the entorhinal cortex. *Nature, 436,* 801–806.

Hahnloser, R. H. (2003). Emergence of neural integration in the head direction system by visual supervision. *Neuroscience, 120,* 877–891.

Hartley, T., Bird, C. M., Chan, D., Cipolotti, L., Husain, M., Vargha-Khadem, F. & Burgess, N. (2007). The hippocampus is required for short-term topographical memory in humans. *Hippocampus, 17,* 34–48.

Hartley, T., Burgess, N., Lever, C., Cacucci, F., & O'Keefe, J. (2000). Modeling place fields in terms of the cortical inputs to the hippocampus. *Hippocampus, 10,* 369–379.

Hartley, T., Maguire, E. A., Spiers, H. J., & Burgess, N. (2003). The well-worn route and the path less traveled: Distinct neural bases of route following and wayfinding in humans. *Neuron, 37,* 877–888.

Hartley, T., Trinkler, I., & Burgess, N. (2004). Geometric determinants of human spatial memory. *Cognition, 94,* 39–75.

Hassabis, D., Kumaran, D., Vann, S. D., Maguire, E. A. (2006). Patients with hippocampal amnesia can't imagine new experiences. *Proceedings of the National Academy of Sciences, USA, 104,* 1726–1731.

Hasselmo, M. E., Bodelón, C., & Wyble, B. P. (2002). A proposed function for hippocampal theta rhythm: Separate phases of encoding and retrieval enhance reversal of prior learning. *Neural Computation, 14,* 793–817.

Hinton, G. E. (2002). Training products of experts by minimizing contrasting divergence. *Neural Computation, 14,* 1771–1800.

Hinton, G. E., Osindero, S., & Teh, Y. (2006). A fast learning algorithm for deep belief nets. *Neural Computation, 18,* 1527–1554.

Holdstock, J. S., Mayes, A. R., Cezayirli, E., Isaac, C. L., Aggleton, J. P., & Roberts, N. (2000). A comparison of egocentric and allocentric spatial memory in a patient with selective hippocampal damage. *Neuropsychologica, 38,* 410–425.

Howard, M., Fotedar, M., Datey, A., & Hasselmo, M. (2005). The temporal context model in spatial navigation and relational learning: Toward a common explanation of medial temporal lobe function across domains. *Psychological Review, 112,* 75–116.

Huxter, J., Burgess, N., & O'Keefe, J. (2003, October 23). Independent rate and temporal coding in hippocampal pyramidal cells. *Nature, 425,* 828–832.

Iaria, G., Petrides, M., Dagher, A., Pike, B., & Bohbot, V. D. (2003). Cognitive strategies dependent on the hippocampus and caudate nucleus in human navigation: Variability and change with practice. *Journal of Neuroscience, 23,* 5945–5952.

Ino, T., Inoue, Y., Kage, M., Hirose, S., Kimura, T., & Fukuyama, H. (2002). Mental navigation in humans is processed in the anterior bank of the parieto-occipital sulcus. *Neuroscience Letters, 322,* 182–186.

Jarrard, L. (1993). On the role of the hippocampus in learning and memory in the rat. *Behavioral and Neural Biology, 60,* 9–26.

Jeffery, K., Donnett, J., Burgess, N., & O'Keefe, J. (1997). Directional control of hippocampal place fields. *Experimental Brain Research, 117,* 131–142.

Jeffery, K., & O'Keefe, J. (1999). Learned interaction of visual and idiothetic cues in the control of place field orientation. *Experimental Brain Research, 127,* 151–161.

Jensen, O., & Lisman, J. (1996). Hippocampal CA3 region predicts memory sequences: Accounting for the phase precession of place cells. *Learning & Memory, 3,* 279–287.

Kahana, M. J., Sekuler, R., Caplan, J. B., Kirschen, M., & Madsen, J. R. (1999, June 24). Human theta oscillations exhibit task dependence during virtual maze navigation. *Nature, 399,* 781–784.

Karnath, H. O., Dick, H., & Konczak, J. (1997). Kinematics of goal-directed arm movements in neglect: Control of hand in space. *Neuropsychologica, 35,* 435–444.

Kim, J. J., & Fanselow, M. S. (1992, May 1). Modality-specific retrograde amnesia of fear. *Science, 256,* 675–677.

King, J. A., Burgess, N., Hartley, T., Vargha-Khadem, F., & O'Keefe, J. (2002). Human hippocampus and viewpoint dependence in spatial memory. *Hippocampus, 12,* 811–820.

King, J. A., Hartley, T., Spiers, H. J., Maguire, E. A., & Burgess, N. (2005). Anterior prefrontal involvement in episodic retrieval reflects contextual interference. *NeuroImage, 28,* 256–267.

King, J. A., Trinkler, I., Hartley, T., Vargha-Khadem, F., & Burgess, N. (2004). The hippocampal role in spatial memory and the familiarity–recollection distinction: A case study. *Neuropsychology, 18,* 405–417.

Klam, F., & Graf, W. (2003). Vestibular response kinematics in posterior parietal cortex neurons of macaque monkeys. *European Journal of Neuroscience, 18,* 995–1010.

Klein, S. B., Loftus, J., & Kihlstrom, J. F. (2002). Memory and temporal experience: The effects of episodic memory loss on an amnesic patient's ability to remember the past and imagine the future. *Social Cognition, 20,* 353–379.

Klink, R., & Alonso, A. (1997). Ionic mechanisms of muscarinic depolarization in entorhinal cortex layer II neurons. *Journal of Neurophysiology, 77,* 1829–1843.

Knauff, M., Kassubek, J., Mulack, T., & Greenlee, M. (2000). Cortical activation evoked by visual mental imagery as measured by fMRI. *Neuroreport, 11,* 3957–3962.

Kobayashi, Y., & Amaral, D. (2003). Macaque monkey retrosplenial cortex: II. Cortical afferents. *Journal of Comparative Neurology, 466,* 48–79.

Kosslyn, S. (1980). Mental images. *Recherche, 11,* 156–163.

Ladavas, E., di Pellegrino, G., Farne, A., & Zeloni, G. (1998). Neuropsychological evidence of an integrated visuotactile representation of peripersonal space in humans. *Journal of Cognitive Neuroscience, 10,* 581–589.

Lever, C., Wills, T., Cacucci, F., Burgess, N., & O'Keefe, J. (2002, March 7). Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature, 416,* 90–94.

Levy, R., & Goldman-Rakic, P. (2000). Segregation of working memory functions within the dorsolateral prefrontal cortex. *Experimental Brain Research, 133,* 23–32.

Levy, W. (1996). A sequence predicting ca3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus, 6,* 579–590.

Loomis, J., Klatzky, R., Golledge, R., Cicinelli, J., Pellegrino, J., & Fry, P. (1993). Nonvisual navigation by blind and sighted: Assessment of path integration ability. *Journal of Experimental Psychology: General, 122,* 73–91.

Maguire, E. A. (2001). The retrosplenial contribution to human navigation: A review of lesion and neuroimaging findings. *Scandinavian Journal of Psychology, 42,* 225–238.

Maguire, E. A., Burgess, N., Donnett, J., Frackowiak, R. S. J., Frith, C. D., & O'Keefe, J. (1998, May 8). Knowing where and getting there: A human navigation network. *Science, 280,* 921–924.

Maguire, E. A., Burke, T., Phillips, J., & Staunton, H. (1996). Topographical disorientation following unilateral temporal lobe lesions in humans. *Neuropsychologia, 34,* 993–1001.

Marr, D. (1971). Simple memory: A theory for archicortex. *Philosophical Transactions of the Royal Society of London B: Biological Sciences, 262,* 23–81.

Matsumura, N., Nishijo, H., Tamura, R., Eifuku, S., Endo, S., & Ono, T. (1999). Spatial- and task-dependent neuronal responses during real and virtual translocation in the monkey hippocampal formation. *Journal of Neuroscience, 19,* 2381–2393.

McClelland, J., McNaughton, B., & O'Reilly, R. (1995). Why there are 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.

McNamara, T. P., Rump, B., & Werner, S. (2003). Egocentric and geo-

centric frames of reference in memory of large-scale space. *Psychonomic Bulletin & Review, 10,* 589–595.

McNaughton, B. L., Barnes, C., & O'Keefe, J. (1983). The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely moving rats. *Experimental Brain Research, 52,* 41–49.

McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., & Moser, M. B. (2006). Path integration and the neural basis of the 'cognitive map'. *Nature Reviews Neuroscience, 7,* 663–678.

McNaughton, B. L., & Morris, R. G. M. (1987). Hippocampal synaptic enhancement and information-storage within a distributed memory system. *Trends in Neurosciences, 10,* 408–415.

Milner, A., Paulignan, Y., Dijkerman, H., Michel, F., & Jeannerod, M. (1999). A paradoxical improvement of misreaching in optic ataxia: New evidence for two separate neural systems for visual localization. *Proceedings of the Royal Society of London B: Biological Sciences, 26,* 2225–2229.

Mittelstaedt, M.-L., & Mittelstaedt, H. (2001). Idiothetic navigation in humans: Estimation of path length. *Experimental Brain Research, 139,* 318–332.

Morris, R., Garrard, P., Rawlins, J., & O'Keefe, J. (1982, June 24). Place navigation impaired in rats with hippocampal lesions. *Nature, 297,* 681–683.

Morrongiello, B., Timney, B., Humphrey, K., Anderson, S., & Skory, C. (1995). Spatial knowledge in blind and sighted children. *Journal of Experimental Child Psychology, 59,* 211–233.

Mou, W., & McNamara, T. P. (2002). Intrinsic frames of reference in spatial memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28,* 162–170.

Mou, W., McNamara, T. P., Rump, B. & Xiao, C. (2006). Roles of egocentric and allocentric spatial representations in locomotion and reorientation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32,* 1274–1290.

Mou, W., McNamara, T. P., Valiquette, C. M., & Rump, B. (2004). Allocentric and egocentric updating of spatial memories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 30,* 142–157.

Muller, R. U. (1996). A quarter of a century of place cells. *Neuron, 17,* 979–990.

Murray, E., & Bussey, T. (1999). Perceptual-mnemonic functions of the perirhinal cortex. *Trends in Cognitive Sciences, 3,* 142–151.

Nakazawa, K., Quirk, M., Chitwood, R., Watanabe, M., Yeckel, M., Sun, L, et al. (2002, July 12). Requirement for hippocampal CA3 NMDA receptors in associative memory recall. *Science, 297,* 211–218.

Norman, G., & Eacott, M. (2004). Impaired object recognition with increasing levels of feature ambiguity in rats with perirhinal cortex lesions. *Behavioral Brain Research, 148,* 79–91.

O'Keefe, J. (1976). Place units in the hippocampus of the freely moving rat. *Experimental Neurology, 51,* 78–109.

O'Keefe, J., & Burgess, N. (1996, May 30). Geometric determinants of the place fields of hippocampal neurons. *Nature, 381,* 425–428.

O'Keefe, J., & Burgess, N. (2005). Dual phase and rate coding in hippocampal place cells: Theoretical significance and relationship to entorhinal grid cells. *Hippocampus, 15,* 853–866.

O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map.* Oxford, England: Oxford University Press.

O'Keefe, J., & Reece, M. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus, 3,* 317–330.

Okuda, J., Fujii, T., Ohtake, H., Tsukiura, T., Tanji, K., Suzuki, K., et al. (2003). Thinking of the future and past: The roles of the frontal pole and the medial temporal lobes. *Neuroimage, 19,* 1369–1380.

Oliveri, M., Turriziani, P., Carlesimo, G., Koch, G., Tomaiuolo, F., Panella, M., & Caltagirone, C. (2001). Parieto-frontal interactions in visual-object and visual-spatial working memory: Evidence from transcranial magnetic stimulation. *Cerebral Cortex, 11,* 606–618.

Ono, T., Nakamura, K., Nishijo, H., & Eifuku, S. (1993). Monkey hippocampal neurons related to spatial and nonspatial functions. *Journal of Neurophysiology, 70,* 1516–1529.

Papez, J. (1937). A proposed mechanism for emotion. *Archives of Neurology and Pathology, 38,* 725–743.

Pavlides, C., Greenstein, Y. J., Grudman, M., & Winson, J. (1988). Long-term potentiation in the dentate gyrus is induced preferentially on the positive phase of theta-rhythm. *Brain Research, 439,* 383–387.

Phillips, R. G., & LeDoux, J. E. (1992). Differential contribution of amygdala and hippocampus to cued and contextual fear conditioning. *Behavioral Neuroscience, 106,* 274–285.

Pierrot-Deseilligny, C., Müri, R. M., Rivaud-Pechous, S., Gaymard, B., & Ploner, C. (2002). Cortical control of spatial memory in humans: The visuo-oculomotor model. *Annals of Neurology, 52,* 10–19.

Piggott, S., & Milner, B. (1993). Memory for different aspects of complex visual scenes after unilateral temporal- or frontal-lobe resection. *Neuropsychologica, 31,* 1–15.

Pinto-Hamuy, T., Montero, V., & Torrealba, F. (2004). Neurotoxic lesion of anteromedial/posterior parietal cortex disrupts spatial maze memory in blind rats. *Behavioral Brain Research, 153,* 465–470.

Postle, B. R., Idzikowski, C., Della Sala, S., Logie, R. H., & Baddeley, A. D. (2006). The selective disruption of spatial working memory by eye movements. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology, 59*(A), 100–120.

Poucet, B. (1993). Spatial cognitive maps in animals: New hypotheses on their structure and neural mechanisms. *Psychological Review, 100,* 163–182.

Pouget, A., & Sejnowski, T. (1997). Spatial transformations in the parietal cortex using basis functions. *Journal of Cognitive Neuroscience, 9,* 222–237.

Recce, M., & Harris, K. D. (1996). Memory for places: A navigational model in support of Marr's theory of hippocampal function. *Hippocampus, 6,* 735–748.

Redish, A. D. (1999). *Beyond the cognitive map: From place cells to episodic memory*. Cambridge, MA: MIT Press.

Redish, A., Elga, A. N., & Touretzky, D. (1996). A coupled attractor model of the rodent head direction system. *Network: Computation in Neural Systems, 7,* 671–685.

Redish, A., McNaughton, B. L., & Barnes, C. A. (2000). Place cell firing shows an inertia-like process. *Neurocomputing, 32,* 235–241.

Redish, A., Rosenzweig, E., Bohanick, J., McNaughton, B., & Barnes, C. (2000). Dynamics of hippocampal ensemble activity realignment: Time versus space. *Journal of Neuroscience, 20,* 9298–9309.

Rieser, J. (1989). Access to knowledge of spatial structure at novel points of observation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 1157–1165.

Rockland, K., & Van Hoesen, G. (1999). Some temporal and parietal cortical connections converge in ca1 of the primate hippocampus. *Cerebral Cortex, 9,* 232–237.

Rode, G., Rosetti, Y., & Boisson, D. (2001). Adaptation improves representational neglect, *Neuropsychologia, 39,* 1250–1254.

Rolls, E., & O'Mara, S. (1995). View-responsive neurons in the primate hippocampal complex. *Hippocampus, 5,* 409–424.

Rosenbaum, R. S., McKinnon, M. C., Levine, B., & Moscovitch, M. (2004). Visual imagery deficits, impaired strategic retrieval, or memory loss: Disentangling the nature of an amnesic person's autobiographical memory deficit. *Neuropsychologia, 42,* 1619–1635.

Sack, A., Sperling, J., Prvulovic, D., Formisano, E., Goebel, R., Salle, F. D., et al. (2002). Tracking the mind's image in the brain II: Transcranial magnetic stimulation reveals parietal asymmetry in visuospatial imagery. *Neuron, 35,* 195–204.

Sala, J., Rämä, P., & Courtney, S. (2003). Functional topography of a distributed neural system for spatial and nonspatial information maintenance in working memory. *Neuropsychologica, 41,* 341–356.

Samsonovich, A., & McNaughton, B. (1997). Path integration and cognitive mapping in a continuous attractor neural network model. *Journal of Neuroscience, 17,* 5900–5920.

Save, E., Cressant, A., Thinus-Blanc, C., & Poucet, B. (1998). Spatial firing of hippocampal place cells in blind rats. *Journal of Neuroscience, 18,* 1818–1826.

Save, E., Guazzelli, A., & Poucet, B. (2001). Dissociation of the effects of bilateral lesions of the dorsal hippocampus and parietal cortex on path integration in the rat. *Behavioral Neuroscience, 115,* 1212–1223.

Save, E., & Moghaddam, M. (1996). Effects of lesions of the associative parietal cortex in the acquisition and use of spatial memory in egocentric and allocentric navigation tasks in the rat. *Behavioral Neuroscience, 110,* 74–85.

Save, E., Paz-Villagran, V., Alexinsky, T., & Poucet, B. (2005). Functional interaction between the associative parietal cortex and hippocampal place cell firing in the rat. *European Journal of Neuroscience, 21,* 522–530.

Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of Neurology, Neurosurgeory and Psychiatry, 20,* 11–21.

Sederberg, P. B., Kahana, M. K., Howard, M. W., Donner, E. J., & Madsen, J. R. (2003). Theta and gamma oscillations during encoding predict subsequent recall. *Journal of Neuroscience, 23,* 10809–10814.

Shallice, T. (1988). *From neuropsychology to mental structure.* Cambridge, England: Cambridge University Press.

Sharp, P. (1999). Subicular place cells expand or contract their spatial firing pattern to fit the size of the environment in an open field but not in the presence of barriers: Comparison with hippocampal place cells. *Behavioral Neuroscience, 113,* 643–662.

Shelton, A., & McNamara, T. (2001). Systems of spatial reference in human memory. *Cognitive Psychology, 43,* 274–310.

Simons, D., & Wang, R. (1998). Perceiving real-world viewpoint changes. *Psychological Science, 9,* 315–320.

Skaggs, W., Knierim, J., Kudrimoti, H., & McNaughton, B. (1995). A model of the neural basis of the rat's sense of direction. *Advances in Neural Information Processing Systems, 7,* 173–180.

Smith, M., & Milner, B. (1989). Right hippocampal impairment in the recall of spatial location: Encoding deficit or rapid forgetting? *Neuropsychologica, 27,* 71–81.

Snyder, L., Grieve, K., Brotchie, P., & Andersen, R. (1998, August 27). Separate body- and world-referenced representations of visual space in parietal cortex. *Nature, 394,* 887–891.

Spiers, H. J., Burgess, N., Maguire, E. A., Baxendale, S. A., Hartley, T., Thompson, P. J., & O'Keefe, J. (2001). Unilateral temporal lobectomy patients show lateralized topographical and episodic memory deficits in a virtual town. *Brain, 124,* 2476–2489.

Squire, L. R. (1986, June 27). Mechanisms of memory. *Science, 232,* 1612–1619.

Stringer, S., Rolls, E., Trappenberg, T., & de Araujo, I. (2002). Self-organizing continuous attractor networks and path integration: Two-dimensional models of place cells. *Network: Computation in Neural Systems, 13,* 429–446.

Stringer, S., Trappenberg, T., Rolls, E., & de Araujo, I. (2002). Self-organizing continuous attractor networks and path integration: One-dimensional models of head direction cells. *Network: Computation in Neural Systems, 13,* 217–242.

Suzuki, W., & Amaral, D. (1994). Perirhinal and parahippocampal cortices of the macaque monkey: Cortical afferents. *Journal of Comparative Neurology, 350,* 497–533.

Taube, J. (1998). Head direction cells and the neurophysiological basis for a sense of direction. *Progress in Neurobiology, 55,* 225–256.

Thiebaut de Schotten, M., Urbanski, M., Duffau, H., Volle, E., Levy, R., Dubois, B., & Bartolomeo, P. (2005, September 30). Direct evidence for a parietal-frontal pathway subserving spatial awareness in humans. *Science, 309,* 2226–2228.

Treves, A., & Rolls, E. (1992). Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. *Hippocampus, 2,* 189–199.

Ungerleider, L., & Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale, & R. J. W. Mansfield (Eds.), *Analysis of visual behaviour* (pp 549–586). Cambridge, MA: MIT Press.

Wallenstein, G., Eichenbaum, H., & Hasselmo, M. (1998). The hippocampus as an associator of discontiguous events. *Trends in Neurosciences, 21,* 317–323.

Wallentin, M., Roepstorff, A., Glover, R., & Burgess, N. (2006). Parallel memory systems for talking about location and age in precuneus, caudate and Broca's region. *NeuroImage, 32,* 1850–1864

Waller, D., & Hodgson, E. (2006). Transient and enduring spatial representations under disorientation and self-motion. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32,* 867–882.

Wang, R., & Brockmole, J. (2003). Human navigation in nested environments. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29,* 398–404.

Wang, R., & Simons, D. (1999). Active and passive scene recognition across views. *Cognition, 70,* 191–210.

Wang, R., & Spelke, E. (2000). Updating egocentric representations in human navigation. *Cognition, 77,* 215–250.

Wang, R., & Spelke, E. (2002). Human spatial representation: Insights from animals. *Trends in Cognitive Sciences, 6,* 376–382.

Whishaw, I., & Brooks, B. (1999). Calibrating space: Exploration is important for allothetic and idiothetic navigation. *Hippocampus, 9,* 659–667.

Wills, T., Lever, C., Cacucci, F., Burgess, N., & O'Keefe, J. (2005, May 6). Attractor dynamics in the hippocampal representation of the local environment. *Science, 308,* 873–876.

Wyss, J., & Groen, T. V. (1992). Connections between the retrosplenial cortex and the hippocampal formation in the rat: A review. *Hippocampus, 2,* 1–11.

Young, B., Fox, G., & Eichenbaum, H. (1994). Correlates of hippocampal complex-spike cell activity in rats performing a nonspatial radial maze task. *Journal of Neuroscience, 14,* 6553–8563.

Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head direction cell ensemble: A theory. *Journal of Neuroscience, 16,* 2112–2126.

Zipser, D., & Andersen, R. (1988, February 25). A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature, 331,* 679–684.

(*Appendix follows*)

Appendix

Implementation of the Model

Mathematical Details

In presenting the mathematical details of the training procedure for the model, each component (medial temporal, transformation, etc.) is considered separately. Following this, the dynamical equations governing the model's behavior during simulation are presented.

Medial Temporal Component

Before the model was trained on a particular environment, the landmarks/boundaries of that environment were discretized by overlaying them on a Cartesian grid with a linear dimension of approximately 3 grid points/unit length. Any grid point that fell within half a lattice spacing of a boundary was then marked as a *landmark segment*. This set of landmark segments, examples of which have been presented in Figures 3 and 4 in text, constituted the training data for the current environment. Training proceeded with the model being positioned at random locations within the environment, while, at each location, attention was sequentially directed to each landmark segment that was potentially viewable from that location. For each of these attending events at each location, appropriate firing rates were imposed on all neurons in the medial temporal layers, and connection strengths between neurons were incremented via a Hebbian learning rule. The procedure for calculating the firing rates during the training phase are now considered.

For the hippocampal layer, a one-to-one correspondence was established between the model neurons and the points on a Cartesian grid, such that each neuron fired maximally at its preferred location. The grid points were spaced with linear density of 2 grid points/unit length covering the relevant allocentric space for each of the environments simulated (see Figure 2 in text for an example). When the model was located at the location with coordinates (x, y), the firing rate of the i^{th} hippocampal neuron was calculated via

$$R_i^{\text{H}} = \text{e}^{\frac{(x_i-x)^2 + (y_i-y)^2}{0.5^2}}, \tag{A1}$$

where (x_i,y_i) are the coordinates of that neuron's preferred location. Next, for the BVC layer, a one-to-one correspondence between the set of BVCs and a radial grid centered at the model's current location and covering allocentric space (see Figure 4) was formed. For all environments, this grid had a radial resolution of 1 grid point/unit length to a maximum of 16 units and an angular resolution of $51/2\pi$ grid points/rad. The contribution of a landmark segment with allocentric coordinates (r,θ^a) to the firing rate of the i^{th} BVC neuron was calculated via

$$R_i^{\text{BVC}} = \frac{1}{r}\, e^{-\left(\frac{\theta_i^a-\theta^a}{\sigma_\theta}\right)^2} e^{-\left(\frac{r_i-r}{\sigma_r}\right)^2}, \tag{A2}$$

where (r_i,θ_i^a) are the allocentric coordinates of that neuron's corresponding grid point, and σ_θ and σ_r are chosen to have values of $(0.005)^{1/2}$ and $(0.1)^{1/2}$, respectively. The total firing rate of the i^{th} BVC neuron was obtained by summing Equation A2 to a maximum value of 1 over all landmark segments viewable from the current location. The particular values chosen for σ_θ and σ_r allow for reasonable spatial resolution with the model architecture; however, the exact values of these parameters are not critical. In fact, with a sufficiently high number of neurons covering space, the only constraint on these values would be the desired spatial resolution of the model. It should be noted that the above definition of BVCs simplifies that of Hartley et al. (2000) and O'Keefe and Burgess (1996), for which the sharpness of the distance tuning decreased with the preferred distance, r_i, of the cell. However, a similar effect of increased influence for nearby versus distant boundaries is achieved through the increased angle subtended by a nearby boundary, which therefore controls the firing of a larger proportion of the BVC population (see Barry et al., 2006). Finally, boundary/landmark identity neurons were modeled by associating each perirhinal neuron with an environmental landmark identity. Thus, the firing rate of the i^{th} perirhinal neuron is given by

$$R_i^{\text{PR}} = C_{\text{PR}} \times \begin{cases} 1 & \text{if } i \text{ is the index of the attended landmark} \\ 0 & \text{otherwise} \end{cases}, \tag{A3}$$

where C_{PR} is set to 1.

Once firing rates for a given training step (attending event) were imposed upon all medial temporal layers, the model weights were updated via the Hebbian learning rule

$$W_{ij}^{\alpha,\beta}(t+1) = W_{ij}^{\alpha,\beta}(t) + R_i^{\alpha}(t)\, R_j^{\beta}(t), \tag{A4}$$

where α and β are layer labels chosen from {BVC,H,PR}, and $W_{ij}^{\alpha,\beta}(t)$ is the weight connecting the j^{th} neuron in layer β to the i^{th} neuron in layer α at training step t. After the completion of the training session, each neuron's vector of incoming weights from each other layer was normalized to sum to unity. Each hippocampal neuron's vector of incoming weights on recurrent connections was normalized by dividing by its maximum incoming recurrent weight. Note that no learning rate parameter was required in Equation A4 because of the weight normalization after learning.

Parietal Component

The parietal component of the model, including the parietal window, the transformation layer, the head direction system, and the connections within/between these regions and from/to the BVC layer, was trained separately from the medial temporal component because the former needed training only once. For each training step a heading direction, ϕ, was randomly chosen from the set of heading directions, $\{2\pi i/20\}_{i=0}^{19}$, corresponding to the set of transformation sublayers. Next, a linear boundary of random location and orientation in allocentric space was discretized in the same way as landmark boundaries were in the medial temporal training procedure described above. The length of this linear boundary was chosen proportional to the distance between its midpoint and the allocentric origin in order to sample sparsely distributed neurons distant from the origin as frequently as densely distributed neurons

near the origin. BVC firing rates were then calculated for the discretized boundary by using Equation A2 and were identically imposed on the BVC layer and the transformation sublayer corresponding to the randomly chosen rotation angle, ϕ. By rotating the linear boundary through ϕ about the allocentric origin, the egocentric positions of the individual landmark segments for this boundary were then found. As with the BVC layer, firing rates of the parietal window neurons in the presence of the boundary were found by first forming a one-to-one correspondence between the set of parietal window neurons and a radial grid centered at the model's current location and covering egocentric space (see Figure 3 in text). The contribution of a single landmark segment with egocentric coordinates (r,θ^e) to the firing rate of the i^{th} such neuron was calculated via

$$R_i^{\text{PW}} = \frac{C_{\text{PW}}}{r} e^{-\left(\frac{\theta_i^e-\theta^e}{\sigma_\theta}\right)^2} e^{-\left(\frac{r_i-r}{\sigma_r}\right)^2}, \tag{A5}$$

where (r_i,θ_i^e) are the egocentric coordinates of that neuron's corresponding grid point, C_{PW} is set to 1, and σ_θ and σ_r are chosen as in Equation A2. The total firing rate of the i^{th} parietal window neuron was calculated by summing Equation A5 to a maximum value of 1 over all landmark segments viewable from the current location. Finally, the head direction layer is a one-dimensional continuous attractor (e.g., Skaggs et al., 1995; Stringer, Trappenberg, et al., 2002; Zhang, 1996) composed of 100 neurons uniformly covering 360° of angular head direction space, with the firing rate of the i^{th} such neuron calculated via

$$R_i^{\text{HD}} = C_{\text{HD}} e^{\frac{(\phi_i-\phi)^2}{0.1885^2}}, \tag{A6}$$

where ϕ_i is the preferred heading direction of that neuron and where C_{HD} is set to 1.

Once firing rates were imposed on each layer for a given head direction and linear boundary, all connection weights were incremented according to Equation A4. After 400,000 such training iterations, the vector of incoming weights for each parietal neuron from each other layer was normalized to sum to unity. Weights from the transformation layer to the parietal window were clipped so that the smallest 30% were set to zero. This was done so that the weight matrices became sparse, a manipulation that decreased required simulation time considerably. For normalization purposes, all transformation sublayers were taken as part of the same layer. The vector of weights on incoming recurrent connections for each head direction neuron was normalized by dividing by the maximum incident weight value for that neuron. Although all weights in the parietal component of the model were trained on a discrete set of 20 transformation angles, the model was found to interpolate accurately between these values.

Velocity Integration

In order to maintain a localized packet of self-sustaining activity, the head direction system must have a set of recurrent excitatory connections, each originating from a particular head direction cell representing and terminating on another cell that represents a nearby or equal direction. Overall, connections from any given head direction cell must be balanced in such a way that that cell's activity equally excites neurons representing directions to either side of the current direction. The training procedure described in the previous section results in the formation of just such a set of weights. An applied angular velocity signal can move an activity bump around in this network in a continuous fashion by modulating an appropriately formed second set of self-excitatory connections (Zhang, 1996). Any connection in this set also originates from a cell representing a particular direction and terminates on another cell that represents a nearby direction, but these "rotational" connections are asymmetric so that activity in the presynaptic head direction cell preferentially excites cells corresponding to nearby directions that are to one side of the current direction. In principle, the angular velocity of the shift is proportional to the size of the asymmetric component (Zhang, 1996); however, for simplicity, we simulate rotations of fixed velocity, with an angular velocity signal that simply gates the use of a fixed set of "rotational" connections in either sense (clockwise or counterclockwise). We achieved such a weight distribution by moving a bump of activity around the head direction neurons at a constant velocity in order to simulate rotational egomotion. During this simulated rotation, the velocity-gated weights on recurrent connections within the head direction layer were updated by the trace Hebbian learning rule given by

$$W_{ij}^{\omega\times\text{HD}}(t+1) = W_{ij}^{\omega\times\text{HD}}(t) + R_i^{\text{HD}}(t)\bar{R}_j^{\text{HD}}(t), \tag{A7}$$

where $W_{ij}^{\omega\times\text{HD}}(t)$ is the velocity-gated weight from the j^{th} to the i^{th} head direction neuron at training step t, where $\bar{R}_j^{\text{HD}}(t)$ is given by

$$\bar{R}_j^{\text{HD}}(t) = \sum_{k=1}^{100} e^{-k\Delta t} R_j^{\text{HD}}(t-(k-1)\,\Delta t), \tag{A8}$$

and where $\Delta t = 0.05$ time units. After training, the velocity-gated head direction weights were normalized in the same way as the nonvelocity-gated recurrent head direction weights. A similar model of the head direction cell ensemble has been described in detail by Stringer, Trappenberg, et al. (2002).

Translation, which can occur in parallel with rotation in our model, is accomplished by introducing a second set of velocity-gated "translational" weights from the transformation sublayers to the parietal window. The original "static" set of weights is responsible for projecting a rotated image of BVC activity onto the parietal window during top-down phases and becomes inactive during translational motion. Instead, the translational set of weights projects a similar rotated image onto parietal window neurons, but it is displaced by a small amount in egocentric space. This is accomplished by setting the translational weights as

$$W_{ij}^{v\times\text{TR}_n} = \sum_k e^{-\frac{(x_k^e-x_i^e)^2+(y_k^e-y_i^e+1.5)^2}{\sigma(r_i)^2}} W_{kj}^{\text{PW,TR}_n}, \tag{A9}$$

where $(x_i^e,y_i^e) = (r_i\cos\theta_i^e, r_i\sin\theta_i^e)$ are the maximal firing coordinates of the i^{th} parietal window neuron in the egocentric map, and $W_{kj}^{\text{PW,TR}_n}$ is the static weight connecting the j^{th} neuron in the n^{th} transformation sublayer and the k^{th} neuron in the parietal window layer. Although σ in this equation could be set to a constant, we

(*Appendix continues*)

found that with our limited resolution for landmark representation at larger distances, a more practical form was given by

$$\sigma(r_i) = 0.45 \log\left(1 + \frac{5r_i}{16}\right). \quad \text{(A10)}$$

Because feedback connections propagate the displaced parietal window activity resulting from the up-regulated weights of Equation A9 back to the place cell layer during bottom-up phases, BVC and place cell firing shifts to reflect the new parietal window activity. This, in turn, results in a further shifting of the activity projected back onto the parietal window in the next top-down phase. Thus, translation of both the egocentric and allocentric representations of space continues until the velocity signal is removed and the original static weights are up-regulated again. As with the rotational connections, we simulate only a single speed of motion. A more complete model might simulate different speeds of translation by using a number of different sets of connections from the transformation layer to the parietal window, each corresponding to a slightly different displacement, and each gated by separate signals for the corresponding speeds. Alternatively, it might titrate the influence of static and translational weights according to speed of movement. However, because of their intense computational requirements, we have not explored these more detailed models here.

Dynamics

During simulations, all neurons in our model were of the "leaky-integrator" variety and all dynamical equations were integrated by using the simple Euler method with a time step of 0.05 units. For the medial temporal part of the model (perirhinal, BVC, and hippocampal), we have

$$\frac{d\mathbf{A}^{\alpha}}{dt} = -\mathbf{A}^{\alpha} - \phi^{\alpha}_{\text{inh}}\,\hat{\mathbf{1}} \cdot \mathbf{R}^{\alpha} + \delta_{\alpha,\text{H}}\phi^{\text{H}}\,\mathbf{W}^{\text{H,H}} \cdot \mathbf{R}^{\text{H}} + \sum_{\beta \neq \alpha} \chi^{\alpha,\beta}(t)\phi^{\alpha,\beta}\,\mathbf{W}^{\alpha,\beta} \cdot \mathbf{R}^{\beta} + \delta_{\alpha,\text{BVC}}\phi^{\text{BVC,TR}} \sum_{n} \chi^{\text{BVC,TR}_n}(t)\,\mathbf{W}^{\text{BVC,TR}_n} \cdot \mathbf{R}^{\text{TR}_n} + \delta_{\alpha,\text{PR}}\,\mathbf{I}^{\text{PR}}, \quad \text{(A11)}$$

where $\mathbf{A}^{\alpha}$ is the activation vector for layer α; $\mathbf{W}^{\alpha,\beta}$ is the weight matrix connecting layer β to layer α; $\phi^{\alpha,\beta}$ is a scalar, representing the overall strength of the connection from layer β to layer α; δ is the Kronecker delta function (unity for equal arguments, zero otherwise); $\phi^{\alpha}_{\text{inh}}$ represents an inhibitory bath of interneurons to which all neurons in a given layer are reciprocally connected with equal weight; $\hat{\mathbf{1}}$ is a square matrix with all elements equal to one; and $\mathbf{I}^{\text{PR}}$ is an externally applied source of input (see below) representing direct lower level input into the perirhinal layer. Bottom-up/top-down dynamics are governed by the χ functions, of which $\chi^{\text{H},\beta}(t)$ and $\chi^{\text{BVC,TR}_n}(t)$ are 1 during a bottom-up phase and 0.05 during a top-down phase, $\chi^{\alpha,\text{H}}(t)$ is 1 during a top-down phase and 0.05 during a bottom-up phase, and the remaining χs in Equation A11 are always 1. The length of each of the bottom-up/top-down phases is 15 time units. Finally, the firing rate of the i^{th} neuron in layer α is given by a sigmoid function of its activation, as follows

$$R^{\alpha}_i = \frac{1}{1 + \exp\{-0.2(A^{\alpha}_i - \nu^{\alpha})\}}, \quad \text{(A12)}$$

where ν^{α} acts as a threshold. Exact numerical values for all unspecified parameters are presented in Table 1.

The dynamics of the parietal window and head direction layers are given by Equation A13 (see below) and

$$\frac{d\mathbf{A}^{\text{HD}}}{dt} = -\mathbf{A}^{\text{HD}} - \phi^{\text{HD}}_{\text{inh}}\hat{\mathbf{1}} \cdot \mathbf{R}^{\text{HD}} + \phi^{\text{HD}}\,\mathbf{W}^{\text{HD,HD}} \cdot \mathbf{R}^{\text{HD}} + \delta_{\omega,\text{on}}\phi^{\omega \times \text{HD}}\,\mathbf{W}^{\boldsymbol{\omega} \times \text{HD}} \cdot \mathbf{R}^{\text{HD}} + \mathbf{I}^{\text{HD}}, \quad \text{(A14)}$$

respectively, whereas the dynamics of the i^{th} neuron in the n^{th} transformation sublayer are given by

$$\frac{d\mathbf{A}^{\text{TR}_n}}{dt} = -\mathbf{A}^{\text{TR}_n} - \phi^{\text{TR}}_{\text{inh}}\,\hat{\mathbf{1}} \cdot \mathbf{R}^{\text{TR}_n} + \phi^{\text{TR,HD}}\,\mathbf{W}^{\text{TR}_n,\text{HD}} \cdot \mathbf{R}^{\text{HD}} - \phi^{\text{TR,I}}\,\mathbf{1}\mathrm{R}^{\text{I}} + \sum_{\alpha \in \{\text{BVC,PW}\}} \chi^{\text{TR},\alpha}(t)\phi^{\text{TR},\alpha}\,\mathbf{W}^{\text{TR}_n,\alpha} \cdot \mathbf{R}^{\alpha}, \quad \text{(A15)}$$

where $\mathbf{W}^{\nu \times \text{TR}_n}$ and $\mathbf{W}^{\omega \times \text{HD}}$ are the "translational" transformation layer to parietal window weights and the "rotational" recurrent head direction weights respectively, where $\chi^{\text{TR},\alpha}(t)$ is 1 for $\alpha = \text{BVC}$ during a top-down phase or for $\alpha = \text{PW}$ during a bottom-up phase, and 0.05 otherwise, and where $\mathbf{1}$ is a vector of ones. Finally, the dynamics of the inhibitory interneuron are given by

$$\frac{dA^{\text{I}}}{dt} = -A^{\text{I}} + \phi^{\text{I,HD}}\,\mathbf{1} \cdot \mathbf{R}^{\text{HD}}. \quad \text{(A16)}$$

Parameters in the model were chosen so that the fourth term on the right-hand side of Equation A15 was a constant for all head direction cell activity packets maintained in our simulations by either attractor dynamics or injected current. This constant was equal to the maximum value of $\mathbf{W}^{\text{TR}_n,\text{HD}}$. Therefore, the fourth term on the right-hand side of Equation A15 could have been eliminated by simply subtracting a constant from $\mathbf{W}^{\text{TR}_n,\text{HD}}$ so that their maximum value was zero. With such a simplification, the model could be interpreted as having only inhibitory direct con-

$$\frac{d\mathbf{A}^{\text{PW}}}{dt} = \begin{cases} -\mathbf{A}^{\text{PW}} - \phi^{\text{PW}}_{\text{inh}}\,\hat{\mathbf{1}} \cdot \mathbf{R}^{\text{PW}} + \sum_{n}[\delta_{\nu,\text{off}}\phi^{\text{PW,TR}}\,\mathbf{W}^{\text{PW,TR}_n} + \delta_{\nu,\text{on}}\phi^{\nu \times \text{TR}_n}\,\mathbf{W}^{\nu \times \text{TR}_n}] \cdot \mathbf{R}^{\text{TR}_n} + \mathbf{I}^{\text{PW}} & \text{during top-down} \\ 0 & \text{during bottom-up} \end{cases} \quad \text{(A13)}$$

nections from head direction to the transformation layer, without any inhibitory interneurons. Note also that all neurons in the model interact with their connected neighbors in an identical fashion. Apparent differences in the form of the above dynamical equations are superficial and reflect the fact that the various network layers have unique patterns of connectivity with their neighbors.

In addition to calculating neuronal firing rates for training purposes, Equations A3, A5, and A6 were also used to calculate the cuing/sensory or mentally generated inputs $\mathbf{I}_{PR}$, $\mathbf{I}_{PW}$, and $\mathbf{I}_{HD}$. For this purpose, C_{PR}, C_{PW}, C_{HD}, σ_θ and σ_r were set to 60, 60, 40, $(0.01)^{1/2}$, and $(0.1)^{1/2}$, respectively. When the weak BVC terminating weights were used in Simulation 4, C_{PW} was increased to 100 during calculation of sensory input. Again, the exact values of the listed parameters were not critical but were found to generate localization quickly. In fact, a relatively wide range of parameter values would have produced qualitatively similar results.

Finally, after the model has been cued to "imagine" itself in a certain location and orientation, or during mental exploration/spatial updating, attention can be directed in any egocentric direction in order to identify surrounding landmarks. To simulate focused attention in the direction, ψ, an input given by

$$I_i^{PW} = C_{PW} e^{\frac{(\theta_i^e - \psi)^2}{\sigma_A^2}} \quad \text{(A17)}$$

was applied directly to neurons in the parietal window layer, where σ_A was set to $\sqrt{5}$ for all attending events, except during the identification of Building 1 in Simulation 1. In the latter case, an increased value of $\sqrt{45}$ was used for σ_A (this stronger attention signal would have resulted in the correct identification of the remaining buildings as well and would not have affected any of the results presented here). The value C_{PW} was set to 40 for our simulations.

Simulation of Head Direction Cell Lesions

Input from the head direction cell system to transformation neurons was recorded for all head directions by storing the combined value of the third and fourth terms of the right side of Equation A15 in a vector, $\mathbf{I}_{HDrec}(\phi)$. Each element of this vector corresponds to one transformation layer neuron and is a function of the head direction, ϕ. Thus, the third and fourth terms of the right side of Equation A15 could be replaced by $\mathbf{I}_{HDrec}(\phi)$ during simulation. For a given value of ϕ, all values of $\mathbf{I}_{HDrec}(\phi)$ are less than or equal to zero, with only elements corresponding to transformation layer neurons in the "selected" sublayer being close to zero. All other values are strongly negative, reflecting the gating function of the head direction system.

In order to simulate a head direction cell lesion for a "realistic" model in which inhibition for gating is accomplished via a large population of inhibitory interneurons, a two-part modification of $\mathbf{I}_{HDrec}$ was used. First, all values of $\mathbf{I}_{HDrec}$ greater than a cut-off of 33% larger than the minimum value were set to the cut-off (the average minimum value was -96, so the cut-off was -64). This modification was intended to simulate the loss of direct excitation to the "selected" transformation sublayer. Second, random regions of each transformation sublayer were selected (see below) and the $\mathbf{I}_{HDrec}$ elements corresponding to those neurons were increased in value to the level of the cut-off. The exact random transformation layer regions selected for this manipulation varied with head direction. This modification was intended to simulate the loss of inhibition resulting from lowered levels of stimulation to the inhibitory neuron population.

In selecting random regions of the transformation layer for reduced inhibition, a one-to-one correspondence between the neurons in each transformation sublayer and a radial grid was formed (as described earlier in the training section). A circle with randomly located center and a radius of 7.5 units was formed for each sublayer and all neurons corresponding to grid points within the circle were selected for reduced inhibition. These circular regions were randomly reselected for each head direction.

Received June 21, 2005
Revision received October 25, 2006
Accepted November 2, 2006 ■