

## DECIPHERING THE HIPPOCAMPAL POLYGLOT: THE HIPPOCAMPUS AS A PATH INTEGRATION SYSTEM

B. L. McNAUGHTON, C. A. BARNES, J. L. GERRARD, K. GOTHARD, M. W. JUNG, J. J. KNIERIM, H. KUDRIMOTI, Y. QIN, W. E. SKAGGS, M. SUSTER AND K. L. WEAVER

*Arizona Research Laboratories, Division of Neural Systems, Memory, and Aging, University of Arizona, Tucson, AZ 85724, USA*

### Summary

**Hippocampal ‘place’ cells and the head-direction cells of the dorsal presubiculum and related neocortical and thalamic areas appear to be part of a preconfigured network that generates an abstract internal representation of two-dimensional space whose metric is self-motion. It appears that viewpoint-specific visual information (e.g. landmarks) becomes secondarily bound to this structure by associative learning. These associations between landmarks and the preconfigured path integrator serve to set the origin for path integration and to correct for cumulative**

**error. In the absence of familiar landmarks, or in darkness without a prior spatial reference, the system appears to adopt an initial reference for path integration independently of external cues. A hypothesis of how the path integration system may operate at the neuronal level is proposed.**

Key words: navigation, dead reckoning, place cells, head-direction cells, cognitive maps.

### Introduction

‘...Each place cell receives two different inputs, one conveying information about a large number of environmental stimuli or events, and the other from a navigational system which calculates where an animal is in an environment independently of the stimuli impinging on it at that moment. The input from the navigational system gates the environmental input, allowing only those stimuli occurring when the animal is in a particular place to excite a particular cell.

One possible basis for the navigational system relies on the fact that information about changes in position and direction in space could be calculated from the animal’s movements. When the animal had located itself in an environment (using environmental stimuli) the hippocampus could calculate subsequent positions in that environment on the basis of how far, and in what direction the animal had moved in the interim....In addition to information about distance traversed, a navigational system would need to know about changes in direction of movement either relative to some environmental landmark or within the animal’s own egocentric space...’ (O’Keefe, 1976).

Spatial relationships provide the context for most adaptive behaviors, as well as the framework in which episodic memories are encoded. One milestone in the evolution of our understanding of the central mechanisms of spatial information processing was an analysis by O’Keefe and Nadel

(1978) of neurophysiological and neuropsychological literature on the hippocampal formation. Their analysis drew attention to the crucial role played by this system in the development of high-level internal representations of allocentric spatial relationships and in learning to solve problems that require memory for these relationships. A key element in their theory that the hippocampus is the neural substrate of a ‘cognitive map’ was the earlier discovery (O’Keefe and Dostrovsky, 1971) that hippocampal neurons are selectively active in particular locations in an environment (‘place fields’). Subsequent, intensive study of the properties of such ‘place cells’ has attempted to characterize what, if any, perceptual invariants are encoded by their firing and exactly how they might contribute to navigation. The resulting base of experimental data is both rich and complex; yet to date there is no satisfactory, all-encompassing theory that explains the full range of observed phenomena. In spite of O’Keefe’s initial suggestion that the hippocampal navigation system might be fundamentally organized in terms of the integration of self motion, most studies since that time have been conducted with the underlying theoretical assumption that the primary source of spatial information leading to place-specific firing is the input from visual landmarks and other stimuli in the environment.

The idea that animals possess a means of keeping track of their position by integrating ideothetic (self-motion) information dates back at least to Darwin (see Barlow, 1964). It is clear that rodents, from which the bulk of our knowledge

of the properties of place cells has been derived, possess a highly developed 'path integration' capability (Mittelstaedt and Mittelstaedt, 1980; Thinus-Blanc *et al.* 1987; Etienne, 1987; Alyan and Jander, 1994). A growing body of evidence suggests that the hippocampal formation is a fundamental component of this path integration system and that it contains an abstract representation of space that is at least partly preconfigured, based on a self-motion metric, and is initially independent of landmark information.

An essential subcomponent of this hypothetical, horizontal, inertial guidance system is a population of neurons that appears to signal head orientation, irrespective of location, primarily on the basis of angular motion signals. These 'head-direction' cells were predicted by O'Keefe (1976) and have been found in a widespread, but interconnected, set of structures, including the posterior parietal cortex, retrosplenial cortex, dorsal presubiculum and anterior thalamus (Ranck, 1984; Taube *et al.* 1987; Taube, 1995; Mizumori and Williams, 1993; Chen *et al.* 1994*a,b*). They are typically rather narrowly tuned for the horizontal component of head orientation with respect to the earth's gravitational field. Their relative directional tuning appears to be independent of pitch or roll attitude and can be dynamically altered by a complex interaction between visual and angular motion signals that will be discussed presently. In every case reported to date, a manipulation that alters the reference direction for one of these neurons results in a corresponding alteration in the reference direction for the whole system.

In this paper, we review the evidence that leads to the hypothesis that the hippocampal formation and head-direction system are components of a network that has evolved to enable the internal representation of abstract spatial relationships, without the necessity of explicit reference to specific objects. The basis for this capability may be a set of quasi-independent spatial reference frame representations that are preconfigured within the synaptic matrix of the hippocampus and related structures, and which permit position and direction to be updated solely on the basis of ideothetic information. Under this hypothesis, in a disoriented rat released into an unfamiliar environment, an arbitrary preconfigured reference frame would be selected within the hippocampus from those not currently in use. As the rat acquired landmark information, these data would be bound by associative synaptic modification to location representations within the current frame. These learned connections would serve to initialize the rat's location within the appropriate frame on subsequent visits and to correct for cumulative error in the path integrator during perambulation within the frame. Apart from these roles, however, these learned connections would not be essential for maintaining an internal description of location and orientation. Under certain circumstances, it is possible that the same physical space would be represented by different frames, depending on other contextual variables, and this possibility will be shown to provide a plausible explanation for several otherwise perplexing observations concerning hippocampal place cells.

### Sufficiency of ideothetic information

The primary evidence for the assertion that place fields reflect a path integration mechanism comes from the observation that, although they can be initialized by information about location with respect to visual landmarks in a familiar environment, place fields are largely preserved if, in the animal's presence, those landmarks are subsequently removed (O'Keefe and Speakman, 1987; Muller and Kubie, 1987; Muller *et al.* 1987) or the room lights are extinguished (McNaughton *et al.* 1989; Quirk *et al.* 1990; Markus *et al.* 1994). In O'Keefe and Speakman's study, the rats were required to select one of four arms of a plus-shaped maze on the basis of its relationship to a set of landmarks that was placed in one of four orientations at the beginning of the trial, but removed before the rats were allowed to choose. The rats usually chose correctly; however, when errors did occur, the place fields shifted to be consistent with the rat's position with respect to the chosen arm rather than the arm that was correct in terms of its relationship to the most recently presented orientation of the controlled landmarks. Muller and Kubie (1987) and Muller *et al.* (1987) studied the properties of place cells in a gray, cylindrical environment with a single white cue card to establish polarity. When the rats were subsequently placed in the cylinder with the card in a different orientation, place fields shifted accordingly; however, if the cue card was removed altogether, place fields were intact but often located at a random orientation relative to the external environment. Both McNaughton *et al.* (1989) and Quirk *et al.* (1990) found that place fields of CA1 cells were well conserved in darkness, provided that the animal was first allowed some exploration of the apparatus under illuminated conditions. More recently, Markus *et al.* (1994) studied the effects of alternating light and dark trials on a forced-choice version of the radial arm maze. Most, but not all, place cells had similar fields under the two conditions, but these fields were less reliable and conveyed less information about location in darkness than under illuminated conditions. This is consistent with the expected behavior of a path integrator deprived of correction for drift. The observation that a significant number of cells changed their location preferences suggests that, when the rats experienced the light and dark conditions many times, they may have developed partly different representations or reference frames for the two conditions. Interestingly, although the forced-choice version of the eight-arm maze is very easy for the animals, some rats very occasionally, on dark trials, made the error of re-entering the arm from which they had just departed. Overall, there was a significant correlation across animals between the frequency of such errors and the reliability of their place fields.

### Preconfiguration of place fields

An allocentric map of the environment, based on stored relationships among landmarks, would be expected to require considerable time for its acquisition. Moreover, as was argued by McNaughton *et al.* (1991), storing object-centered relationships among  $n$  landmarks requires storage of the order

of  $n^2$  items. In contrast, if direction and distance information are available, then a vector-based storage system can, in principle, be constructed, whose storage requirements increase only linearly with the number of landmarks. These arguments suggest the possibility that the hippocampus may be at least partly preconfigured to represent sets of two-dimensional relationships using place fields linked by self-motion information. The linkage would be such that, given the activity of one small set of place cells and information about the direction and speed of motion, a specific different set of cells would be reliably activated. We will discuss how this linkage might be implemented presently. This description is formally equivalent to saying that place fields within a frame represent distances and bearings (i.e. vectors) from some arbitrary point in the frame. Repeated pairing between these ideothetically updated place representations and landmark information would subsequently allow the landmarks both to correct for drift in the path integrator and to set the origin for path integration. According to this view, however, the fields themselves would be based on a pre-existing synaptic matrix.

Hill (1978) was the first to address the question of whether experience was necessary for place cell firing. He concluded that that little, if any, experience was required. Ten of twelve cells studied showed firing as robust on the first visit to their place fields as on subsequent visits. The other two cases seemed to develop fields over about 10 min. More recently, Wilson and McNaughton (1993), using parallel recording methods that enabled simultaneous study of more than 100 cells, provided evidence that place fields are less robust during the initial 10 min of exploration of a novel 60 cm<sup>2</sup> arena than after this period of exploration. They suggested that place fields might be at least partially a result of rapid learning. A slightly different interpretation of these data, however, is that the fundamental, ideothetic relationships among the place fields were preconfigured. Under this interpretation, the initially lower precision of the place fields in the novel arena was due to the fact that they were initially driven entirely by path integration mechanisms and hence subject to greater drift error. Over about 10 min, however, there would have been gradual associative binding of local-view information to place cells, which would have enabled a more-or-less continual correction for the cumulative error that is an unavoidable component in all path integration systems. Hill's experiments were performed in a more geometrically restricted apparatus where such cumulative error might have been attenuated because of the limitation of possible trajectories. Moreover, the position tracking method available to Hill was not very precise, and it is likely that such subtle effects would have been missed.

In the Wilson and McNaughton (1993) study, the apparatus consisted of two square arenas separated by a partition into familiar and novel halves. In a follow-up to the published experiment, we repeated the partition removal procedure over several days until the rat was well familiarized with both halves of the apparatus. We then conducted a trial in which the second box, together with its associated visual cues, was removed altogether. When the partition was removed, the rat was free

to explore the open tabletop in the region where the second box had stood. Consistent with the cue removal experiments of O'Keefe and Speakman, there were no pronounced changes in the distribution of place fields even though, in this case, there was a rich array of distal cues. A parsimonious interpretation of this observation is that the animal's ideothetically based sense of location relative to the stable, familiar arena exerted a more powerful control over the place fields than the visual stimuli.

Strong support for the ideothetic basis of place-specific firing comes from a study by Quirk *et al.* (1990), who introduced rats into a cylindrical environment in total darkness. Although the animals had experienced the same environment previously under illumination, in many cases introduction in darkness resulted in the establishment of new place fields that persisted in subsequent light. These new fields ultimately came under the control of visual input in that they rotated when the single polarizing stimulus (a white cue card) was rotated. Thus, place fields can arise independently of the visual landmarks which eventually develop control over their expression.

Indirect support for a preconfiguration of place fields can also be derived from observations of the statistical interactions among populations of simultaneously recorded place cells while the animal is either asleep or in quiet wakefulness prior to exposure to a novel spatial situation. Contrary to an initial report based on a small number of experiments (Wilson and McNaughton, 1993), the extent to which two cells will have overlapping place fields in a novel environment can be predicted, to a small degree at least, by the strength of their activity correlation (i.e. the correlation of their firing rates over sequential 100 ms epochs) during sleep prior to the experience (Kudrimoti *et al.* 1995). In agreement with the initial report, however, only those cells that actually end up having overlapping fields in a given environment exhibit enhanced correlations in subsequent sleep. It seems that the activation of specific groups of cells in a given environment results in some priming effect on these cells, possibly, but not necessarily, as a result of synaptic modification during exploration. This priming results in their increased likelihood of co-activation for periods of 30–60 min after the experience.

#### Metastable binding to visual landmarks

Several studies suggest that the binding of place fields to visual landmarks is acquired only slowly and is intrinsically unstable. Visual stimuli that clearly control place fields at one time lack the power to control the fields at other times. In some instances, the change in place field configuration can be accounted for by a simple change of directional reference (i.e. a rotation), whereas in others, place fields can exhibit radically different distributions in the same visual environment. Whether or not this effect corresponds to a translation within the frame of reference or to the selection of a new one is, so far, unclear.

Sharp *et al.* (1990) studied hippocampal place fields in a cylindrical environment similar to the one used by Quirk *et al.* (1990). The question addressed was whether, in a visual

environment with mirror symmetry, place fields would also have twofold symmetry. This result occurred in only three of 18 cases. Typically, fields retained their original shape and orientation relative to one of the cue cards in the cylinder, usually the one to which the rat was nearest at the beginning of the trial. A complementary experiment was performed by Bostock *et al.* (1991), who studied the effects of replacing a white cue card in the cylinder with a black one. The place cells initially behaved in a manner typically observed in darkness in this apparatus. The fields were unchanged in shape or location relative to the radial axis of the cylinder, but they often adopted different, apparently random, orientations relative to the card. After several trials in which the white and black cards were alternated, the place field distributions in the two conditions became completely uncorrelated at any relative orientation.

McNaughton *et al.* (1994) addressed the question of the relative control of visual *versus* vestibular cues over the firing of both place cells and lateral dorsal thalamic head-direction cells. Recordings were conducted in two, high-walled apparatus (a cylinder and a cube with conspicuous polarizing visual landmarks). One apparatus was highly familiar to the animal at the time of the experiment. The other was completely novel, except for the few minutes of initial exploration, during which baseline recordings were made. In the familiar apparatus, a 180° rotation of the entire apparatus (including the floor), at a rate well above vestibular threshold, typically, but not always, resulted in an equal rotation of both place fields and directional preference of head-direction cells, the vestibular input being ignored; however, when the rat's exploratory experience in the environment was limited to a few minutes, the opposite was more typical. This effect is consistent with the postulated role of familiar landmarks as cues for correcting path integration errors. Once the necessary associative binding of landmark information to the path integration network has occurred, the visual information would be expected to override the vestibular signals. More recently, J. J. Knierim, H. Kudrimoti and B. L. McNaughton (unpublished observations) have found that a strong factor in whether visual or vestibular inputs predominate is the magnitude of the mismatch that is induced. For small rotations (45°), the place and direction systems tend strongly to align with the visual cues, whereas for large rotations, there is an increasing tendency for them to remain aligned with the inertial frame and to disregard the landmark information. The realignment with the visual frame, when it does occur, typically does so only after a delay.

In a study of the spatial properties of electrophysiologically identified granule cells of the dentate gyrus, Jung and McNaughton (1993) observed several cases in which place fields of several simultaneously recorded cells exhibited a curious behavior for which, at the time, we had no plausible explanation (Fig. 1). Well-defined place fields that were stable over several trials (sometimes over days) were recorded on the radial maze. Occasionally, however, the complete set of fields on a given trial appeared at the same radial coordinates, but in a different orientation (i.e. different arms of the maze) relative

to the prominent visual landmarks. All cells recorded together exhibited the same relative rotation. On subsequent trials, the orientation typically reverted to the original one.

In a follow-up study, Knierim *et al.* (1995a,b) tested the hypothesis that this rotational instability was due to inconsistent associations between the internal directional reference system, presumably mediated by head-direction cells, and the landmarks. One group of rats was trained in the cylinder apparatus of Muller *et al.* (1987) with an explicit attempt to disorient them during the transfer from the colony room to the recording room. This disorientation was achieved by transporting the rats in a sound-attenuating box and performing a series of very slow rotations of the box *en route*. The other group was transported in the open, with no intentional disorientation. Both groups of animals received the same amount of experience in the cylinder and, during this period, the cue card was maintained in a fixed orientation. During subsequent recording sessions, the cue card was reoriented to one of four directions (including the original one) prior to introducing the rat. In addition, the disorientation procedure was carried out on all animals during recording sessions. The rats that were originally disoriented during training exhibited frequent reorientations of their directional reference relative to the cue card, as defined by the conjoint rotations of place and head-direction cell firing preferences. Typically, one orientation would predominate across sessions, with other orientations occurring occasionally and apparently at random on some trials. Occasionally, azimuthal drift of the place and/or head-direction tuning could be observed within a single trial (Fig. 2). Cells of rats that had been trained without disorientation exhibited strong binding of their orientation preference to the cue card during the first several days of recording under disorientation conditions. With repeated disorientation, however, the behavior of these cells tended to destabilize. In general, once the directional orientation of cells of a given rat became unstable with respect to the cue card, subsequent retraining without disorientation did not result in restabilization. The general conclusion from these data is that, under some conditions, the 'stimulus' that confers polarity on an environment is derived from the state of the angular component of the path integrator at the time the animal enters the environment. Repeated entries with a constant internal orientation reference normally lead to stable, unique associations between visual stimuli and cells of the path integration system, which allow the visual input to realign the system if mismatches subsequently occur. Repeated mismatches, however, lead to multiple conflicting associations and a lack of ability of the array of visual stimuli to define a stable orientation framework.

A complementary conclusion comes from recent experiments (Gothard *et al.* 1995) in which rats were trained to obtain a food reward on a linear track approximately 1.5 m long. One end of the track contained a small (approximately 1 ml) cup for food delivery. At the other end was a movable cardboard box that also contained a reward cup. The animal's task was to leave the box, traverse the track to the fixed reward

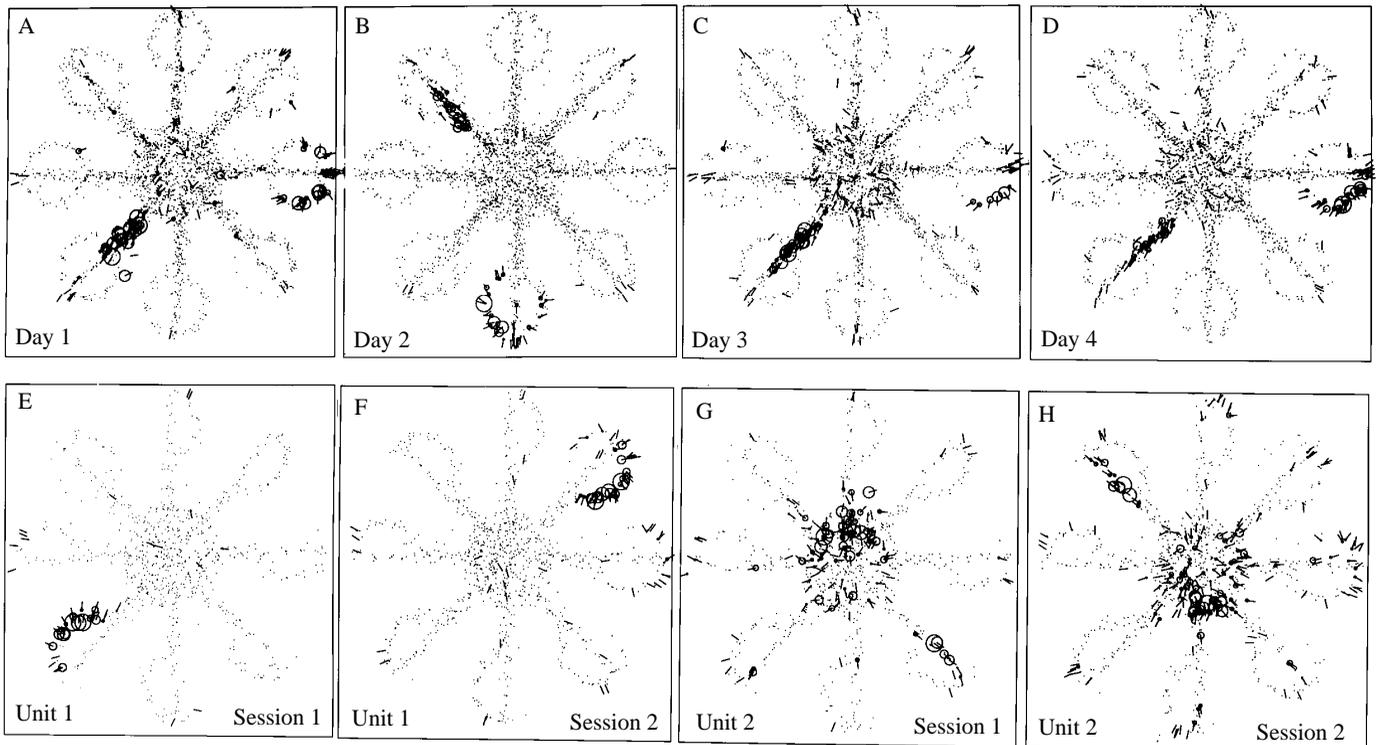


Fig. 1. Metastability of the azimuthal component of the hippocampal place representation. The data are from recordings of place fields of granule cells of the fascia dentata while the animal traversed a radial eight-arm maze for food reinforcement (at arm-ends). Each figure represents superimposed data from 8–10 complete traverses of the maze. The maze was located in a moderately illuminated room with large, high-contrast visual stimuli located around the maze. Dots represent locations visited by the rat. Circles are drawn whose diameter is proportional to the locally computed firing rate on a given trial. Tick marks on circles indicate head orientation. (A–D) A single unit with two preferred firing fields. This cell was studied for 4 days. On day 2, the two fields appeared at positions  $90^\circ$  clockwise to their positions on day 1. They reverted to their original locations on the next 2 days. (E–H) Two separate units recorded simultaneously in two sessions. Both units exhibit  $180^\circ$  rotations of their fields between sessions one and two (from Jung and McNaughton, 1993).

cup and then return to the box which, in the interim, had been shifted at random along the track to one of five different equally spaced locations. In general, two classes of cell were found, those that responded in a fixed relationship to the static reference frame (i.e. the frame in which the track and any other uncontrolled spatial cues remained stationary) and those that responded in a fixed relationship to the last-experienced location of the box. Many box-related cells had place fields on the track that shifted according to the position of the box from which the rat had just departed. In other words, these cells were virtually uninfluenced by either the remote or local spatial cues of the static reference frame. Some cells did exhibit an apparent interaction between distance from the box and absolute location on the tracks, such that the place field could be described as reflecting the relative distance between box and goal. An alternative explanation of this phenomenon will be suggested in the last section. It is well known that, in tasks involving repeated, stereotyped trajectories between fixed locations, hippocampal place fields are highly directionally sensitive, typically having non-significant firing when the rat faces in the direction opposite to the preferred direction in the place field (McNaughton *et al.* 1983). In the present experiments, the box-bound cells fired only on the outward

journey, and the cells bound to the static reference frame fired only on the return journey to the box. The most parsimonious interpretation of these results is that, in this situation, the two sources of food become associated with two separate reference frames and that a reference frame-shift occurs about the time that the rat reaches each source. In this case, the place cells would encode the distance the rat has moved away from a goal. On the return journey, the same cells would not fire, because a different reference frame would be active. This plausibly explains the directional specificity of hippocampal neurons on linear mazes in general.

A completely different sort of metastable binding occurs, under some conditions, between place fields and visual stimuli. As already discussed, rats introduced into an environment in darkness sometimes exhibit a new distribution of place fields that is completely different from those observed during previous sessions in the same environment under illumination; yet, when the lights are turned on, the new fields often persist (Quirk *et al.* 1990). Similarly, following replacement of the white cue card by a black one, Bostock *et al.* (1991) observed that the fields were the same, except for a random rotation, on the first trial, but usually completely different on subsequent trials. Thus, as concluded by Quirk *et al.* (1990), the locations

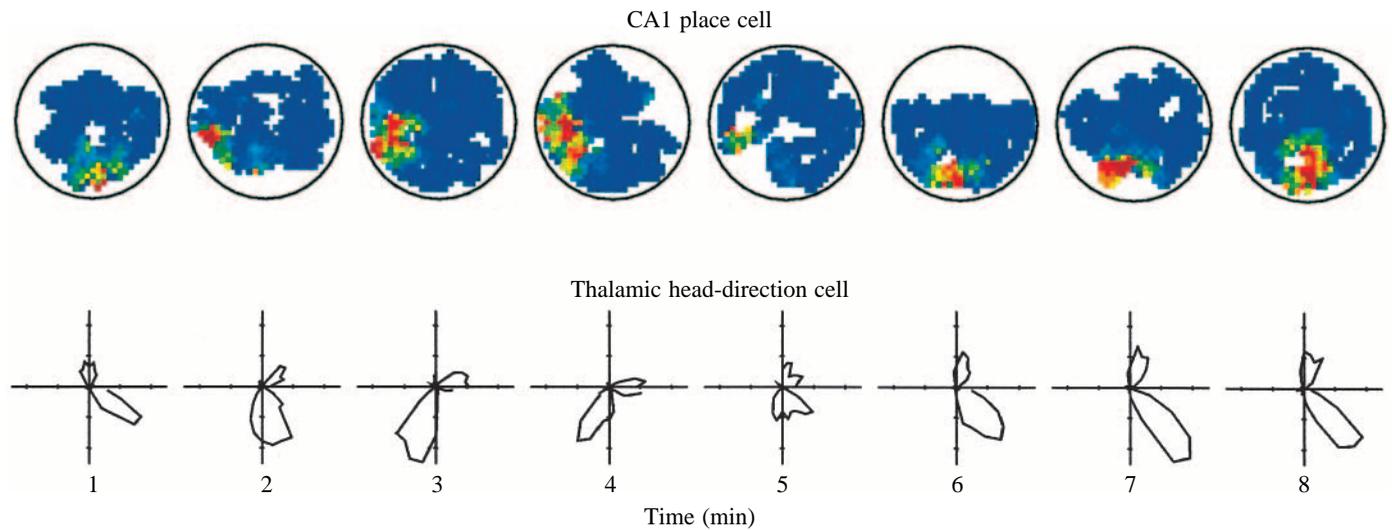


Fig. 2. Illustration of drift of the azimuthal component of place and head-direction cells during foraging behavior on the floor of a high-walled, cylindrical chamber, 76 cm in diameter. A head-direction cell in the anterior thalamus and a hippocampal place cell (CA1) were recorded simultaneously. The data from the first 8 min of the recording session are divided into 1 min periods. Both cells exhibit a drift of approximately  $90^\circ$  in the azimuthal coordinate of their peak firing rate over the first several minutes, followed by a return to the original value. For the place cell, sampled positions at which no firing was recorded are illustrated by dark blue. Peak firing is illustrated by red. The directional tuning of the head-direction cell is plotted in polar coordinates. The head-direction cell data are slightly contaminated by spikes from a different head-direction cell, giving rise to the small lobe in the polar plot. Notice that both lobes of the polar plot rotate in synchrony (from Knierim *et al.* 1995a).

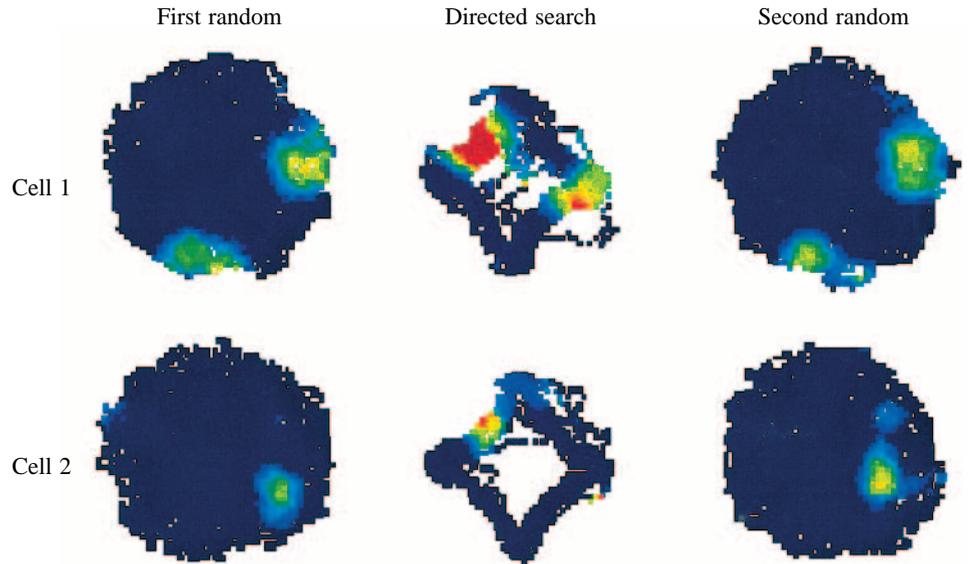
at which place cells fire appear to depend more on the animal's previous experience than on the visual stimuli *per se*, although the nature of this dependence is, so far, uncertain. Knierim *et al.* (1995a,b) observed a similar effect in their combined recordings of place and head-direction cells. Typically, when the head-direction system adopted an unusual orientation with respect to the cue card, the place fields rotated accordingly but were otherwise unchanged. In some cases, however, one distribution of place fields was observed when the head-direction system was in its normal orientation and a completely different distribution when the direction system adopted an alternative orientation. It was as if the animal believed that there were two different environments, one with the cue card at one bearing relative to its internal reference and another with the cue card at a different bearing.

Three studies have addressed the manner in which the spatial task performed by the rat affects place-specific firing of hippocampal cells. In a study by Markus *et al.* (1995), rats were first trained to forage more or less randomly for chocolate pellets on a large, round, open platform. After recording the place field configuration in this context, the food was placed only at four equally spaced sites around the circumference of the circle. First one site was baited, then the next, and so forth. The rats were required to go sequentially from one site to the next, in the same direction for many cycles. The animals typically learned the new task within 15–20 min. By the time performance was reliable, there were pronounced changes in the firing locations of about one-third of the cells and a significant increase in directional dependency. In rats pretrained to switch back and forth between these tasks, the task-specific configurations appeared immediately after the

task was switched (Fig. 3). Although there are several alternative interpretations of these data, they demonstrate that place fields can be independent of visual stimuli and suggest that changes in spatial attention may be an important variable in the control of place-specific firing. Similar results were reported by Fukuda *et al.* (1992), who used medial forebrain bundle stimulation as reward, and also by Breese *et al.* (1989). In the light of the results of Knierim *et al.* (1995a,b) discussed above, however, it is unclear whether the change in firing location of the cells in the Breese *et al.* (1989) study reflected a true remapping or a rotation of the reference direction, resulting from the change in the location of the most salient cue, the water reward.

In recent experiments in our laboratory, we observed an even more dramatic metastability effect. These experiments involved parallel recording of groups of up to 80 hippocampal cells. The studies were performed in a moderately illuminated room in which the animals ran repeatedly around an elevated track (either triangular or rectangular) for food reinforcement. The animals were adapted to their particular task and environment for more than a week prior to the actual experiment. The purpose of this experiment was to observe the effects of an extensive episode of exploration of novel spaces on the distribution of place fields in the familiar environment. The procedure involved an initial recording session in the familiar environment, followed by a 1 h period in which the rats were permitted 10 min of exploration, each in a different room in the laboratory facility. The animals were then returned to the familiar environment for further recording. In many cases, the distribution of place fields was completely rearranged following the intervening period of exploration. This effect was not a result of recording instability.

Fig. 3. Dependence of hippocampal place representation on behavioral context. Two hippocampal cells were recorded simultaneously. In the first approximately 20 min of the recording session, the rat foraged randomly for morsels of food on an open circular platform. During the next 20 min, food was placed sequentially only at four equally spaced locations around the periphery of the platform. Finally, the food was again scattered randomly over the platform for an additional 20 min. This animal had previously been trained to perform both the directed and random search tasks in sequential order. Notice that there are radical changes in the spatial distribution of firing according to the behavioral context (from Markus *et al.* 1995).



Moreover, the effect was typically all-or-nothing. Either the system adopted a completely new configuration or the original one was preserved (Fig. 4). A possible explanation of these data is that the opportunity to explore regions of the laboratory outside the recording room leads to the activation of a new reference frame. One might suppose that the rat attends to its linear and angular motion while being carried back to the recording room, even if completely inaccurately, and hence enters the recording room while the new frame is still active. This experience may lead to a conflict between reference frames, in which sometimes the input from the path integration dominates, with a resulting 'remapping' of the familiar environment, and sometimes the familiar visual stimuli dominate and the place fields are unchanged.

#### The necessity of a 'motor set'

A path integration system would presumably receive substantial input from the motor system in the form of either motor efference copy or simple gating signals. Participation of the motor system in the path integration process is suggested by the finding that, under conditions of movement restraint, both hippocampal place cells (Foster *et al.* 1989) and thalamic head-direction cells (Knierim *et al.* 1995a,b) become virtually silent. In these experiments, the rats were trained, with food reinforcement, to tolerate restraint for periods of 20–30 min without struggling. The rats were free to move their heads and to engage in myostatial sniffing, but were completely unable to move their limbs. Movement *per se* did not appear to be essential for place cell activity when the animal was free to move. For place cells to fire, it was sufficient that the animals were free to move if they wanted to, even if they did not actually move. The rats learned to refrain voluntarily from moving during restraint, and we assume that, in the absence of intention (or preparedness) to move, the path integration system becomes inactive. The effect is not likely to be a stress response involving hormonal modulation, because place-specific firing

typically resumes almost instantaneously when the restraint is removed. Although further study is needed, the basic result is intuitively consistent with the overall thesis presented here.

#### A theoretical framework

Much theoretical work on the hippocampus has focused on the possibility that it may act as an autoassociative memory, capable of rapidly encoding arbitrary associations, as a result of Hebbian synaptic modification (Marr, 1971; McNaughton and Morris, 1987; Treves and Rolls, 1994; McClelland *et al.* 1995). Of major concern in such work is the number of stable activity states that could be encoded in systems such as hippocampal subfield CA3, which possesses an extensive network of modifiable recurrent collaterals. A 'state', in this context, is typically defined as a distribution of firing rates over a population of neurons. It may have stability as a consequence of cooperative synaptic interactions. State stability is similar to Hebb's (1949) concept of a 'cell assembly' and also to the concept of an 'attractor' in dynamic systems theory (e.g. Hopfield, 1982; Amit, 1989). If the network is forced by the input into a non-stable state, it tends to be attracted to (i.e. to converge on) the most similar of the stable states. In associative memory networks, these states correspond to memories encoded through synaptic modification; however, there is no *a priori* requirement for learning. The stable states could be determined by patterns of connection strengths established through some developmental process. The usual conclusion is that the number of unique stable states is vast if each state involves activity of only a small fraction (about 0.01) of the total population of cells and each cell is involved in many different states. This is known as sparse, distributed coding.

Suppose that, rather than forming through the storage of arbitrary events, the synaptic weight matrix was preconfigured in a manner that implicitly defined a large set of two-dimensional surfaces. We shall call these surfaces 'path integration reference frames' and envision them as consisting

of a random subset of the total population of place cells in CA3 in which the synaptic connections define metrical, two-dimensional neighborhood relationships. The easiest way to visualize the structure of such a reference frame is to imagine that each of its component cells is assigned at random to one or a small number of locations on a disk and that the connection strengths from each of the cells to its neighbors are defined by a two-dimensional, approximately Gaussian, function of distance. It is important, however, to remember that the disk is merely a conceptual device we are using to define the synaptic interactions in the network. It is well known that there is little or no correlation among neighboring hippocampal neurons as to where they might fire in a given environment. Given appropriate global inhibition to control the threshold, each neuron in a group defines the center of an 'attractor', i.e. a stable pattern of activity involving the given neuron firing maximally and its neighbors firing at rates that fall off with distance (Shen and McNaughton, 1994). Each location on the hypothetical disk is thus defined by a unique pattern of activity, and we can consider a disk to be an abstract internal representation of a two-dimensional surface. Moreover, states involving activity at widespread, disconnected regions of a disk are not stable, unless the cells that are active define a neighborhood in some other reference frame, in which case that reference frame is considered to be the active one. Which reference frame is active, and which location within it, is thus defined not by the activity of any particular cell, but by the particular collection of cells that is active (and, of course, by their relative firing rates). From experimental observations (e.g. Kubie and Ranck, 1983), we know that place cells can have place fields in many but not all environments. For a typical experimental apparatus (say  $1 \text{ m}^2$ ), there is a probability of about 0.3 that a given cell will have at least one (and possibly several) place field somewhere in the environment. We also know that there is essentially no correlation between environments in the two-dimensional distribution of place fields. These observations are consistent with the concept of multiple reference frames; however, we currently have no basis for estimating how many distinct reference frames can be encoded or how often a given cell might be used in a given reference frame. The use of a given cell in multiple reference frames, or several times in the same reference frame, would create potential 'worm holes' in the spatial representation system. Given enough noise, this overlap might lead to undesirable jumps among distant location representations in a manner analogous to quantum mechanical tunneling. The probability of such transitions would be a function of the noise level and the degree of similarity of any given location representation to another more remote one. Because neighboring locations have very similar representations, however, it seems that the most probable effect of small amounts of noise in such a system would be continuous, gradual, random drift.

We have so far proposed an intrinsic, two-dimensional manifold within a high-dimensional neuronal representation space in which locations are defined by stable patterns of neural activity and in which there are orderly proximity

relationships among the locations. This 'state' space topology should not be confused with the sort of anatomical topology seen in most sensory systems, which, according to most evidence, does not occur in hippocampus. There is, however, an interesting anatomical order of a different sort (Jung *et al.* 1994). Cells located more temporally in the hippocampus tend to have a lower probability of having a place field in the typical, rather small, apparatus used in such studies than cells in the septal (i.e. dorsal) region. When a ventral cell exhibits a field, it is typically substantially larger than those recorded more dorsally. There are at least two interesting possible functional consequences of this gradient of place field size. One is that there may be some advantage to encoding space at several different scales simultaneously. For example, the activity of cells with the largest of place fields (i.e. the most ventrally located) may define the reference frame itself. The alternative is that the center of each of our hypothetical disks is used to encode a reference point for path integration, and that distance from this reference point may be encoded on a nonlinear scale by the size of the place field and by its location on the septo-temporal axis; the larger the distance, the lower the precision. We presume that there is a limit to the magnitude of the distance from a reference point that can be represented in this system and that exceeding this limit is at least one of the factors involved in selecting a new reference frame. When a rat is first disoriented and then placed in a novel environment, we suppose that the center of an arbitrary new reference frame for path integration is activated (presumably one that has not recently been mapped to a different environment by the binding of feature information).

The synaptic matrix we have proposed so far, while defining metric relationships within two-dimensional spaces, is not sufficient to account for path integration. Nor does it provide a basis for cognitive mapping, which can be thought of as the binding of external sensory data into a spatial reference frame. Two additional elements are still required. One is a means of coupling information about the rat's linear and angular motion to the matrix in a manner that causes the appropriate shifts of the activity focus within the reference frame when the rat moves. The other is a mechanism for associating external stimuli with locations in the internal representation. We have previously developed a model for how the head-direction system can be updated on the basis of either angular velocity signals or learned relationships between head orientations and external stimuli (McNaughton *et al.* 1991, 1994; Skaggs *et al.* 1995; see Fig. 5B). The model for place cells proposed here is a simple extension of the head-direction model, from one to two dimensions. We begin with an explanation of the head-direction model because it is easier to understand.

Head direction can be represented as a location in a closed, one-dimensional space (a circle centered on the rat). We assign each head-direction cell a location on this circle, and arrange the synaptic interactions to have a limited, Gaussian spread in both directions. Global feedback inhibition to the whole network performs the standard threshold setting (e.g. Marr, 1971). If the  $n$  cells are evenly distributed, then we have  $n$  stable

states or 'attractors', each defining a unique direction (i.e. the directional resolution is  $2\pi/n$ ). This system is referred to as the H system in Fig. 5. We next assume that there is an adjunct layer of neurons (the H'H system in Fig. 5) that receives information about the current location from the head-direction ring and information about rotational motion from the vestibular system and other sources of such information (the H' system in Fig. 5). These cells encode the interaction between current location and the sign of rotation and, depending on the sign of rotation, feed this information to cells to either one side or the other of the current focus of activity in the direction circle. Cells with these response properties have been observed in the rat posterior cortex (Chen *et al.* 1994a,b). The result is a shift in the focus of activity on the head-direction circle, and a consequent shift in the internal representation of head orientation. We assume that this preconfigured network of head-direction cells also receives rich, modifiable connections from other sensory systems, particularly the visual system. By a modified Hebbian (i.e. conjunctive) learning rule, the system can learn to associate activity foci on the head-direction ring with visual landmarks and other external events. If the associations are consistent and unique, the landmarks can eventually shift the focus of activity in the case of mismatches arising from error in the vestibular updating mechanism.

The extension of this model to a two-dimensional representation space is, in principle, straightforward. In addition to place cells (the P system in Fig. 5), we require an adjunct set of neurons with input from both the place cells and the direction ring, whose cells are tuned (possibly rather broadly) for both direction and location in a given frame (i.e. place cells that are directionally specific in their place fields). Hippocampal cells sometimes behave in this way, yet will not serve our purpose because, in random foraging situations, they have essentially no directional tuning. As discussed above, we believe that the apparent directionality of place cells in some situations results from reference frame shifts. Recently, however, Sharp and Green (1993) reported that some cells in the subiculum and dorsal presubiculum have broad, but significant, directional tuning in situations where directionality is absent from hippocampal cells. These cells would fit our model well, because the subiculum receives place information from CA1 (the output layer of the hippocampus proper) and is likely to receive head-direction information from the anterior thalamus. Subicular cells are also rather strongly modulated by self-motion, which would be necessary in order for the location focus to remain stationary when the rat sits still. This system, whose cells are both place- and direction-selective and whose activity is globally modulated by self-motion, is referred to as the PHxM system in Fig. 5. The final requirement is that the projections of the PHxM system back onto each location reference frame in the P system be asymmetrical, with a bias corresponding to the direction of the current head orientation. This asymmetry would result in a shift in the focus of activity in the reference frame in the correct direction during linear motion.

The foregoing scheme describes a mechanism for updating the internal representation of location according to the animal's

motion. As in the case of the directional system, we assume that the location system receives rich modifiable connections from other sensory systems, particularly the visual system, that enable both correction for cumulative error and the setting of the origin for path integration following shifts in the spatial reference frame (for example, when the animal is removed from its home cage and placed in the recording apparatus). The model easily explains the apparent preconfiguration of place fields, as well as their preservation in darkness and after cue removal. It also explains, in principle, why, depending on the rat's recent experience, the same visual world can lead to radically different place field distributions and why it is sometimes possible to induce radical changes in the visual input without changing place fields.

One of the long-standing problems in the place cell literature that is particularly easy to account for with this conceptual framework is the directional dependence of place fields during stereotyped trajectories among fixed sites and the virtually complete lack thereof during random foraging. As outlined above, if each behaviorally significant location in the environment becomes a reference center, then the reference center, and hence the distribution of place cells, will differ according to the direction in which the animal is moving. This dichotomy will be reinforced by the binding of visual and local sensory information to the corresponding frameworks so that, with a little experience, the reference frame will shift even if the rat executes an occasional direction reversal in the middle of its trajectory. In contrast, during random foraging, in which there are no places of particular significance, the environment is represented within a single reference frame, and the visual associations derived from the different views from a given location increasingly reinforce this. This explanation is also compatible with the observed rapid switching of place fields when sequential and random foraging tasks are alternated in a fixed environment.

In the moving box experiment, some of the cells bound to the box framework exhibit an interaction between distance from the box and absolute spatial location. This interaction could be expected in a situation in which the same view information would have been bound to several different points on the trajectory. If we assume that the view information for a particular feature has a Gaussian distribution, then one can imagine that the association with the current reference framework would approximate the sum of several overlapping distributions, one for each start location. The visual input would thus tend to 'correct' the path integration system towards the center of this input weight distribution. This effect would lead to the observed interaction between distance from the box and actual location in the laboratory frame.

The recent history of thought concerning the function of the hippocampal formation has been dominated by two apparently orthogonal themes: the representation of space (O'Keefe and Nadel, 1978) and the initial encoding and subsequent consolidation (Marr, 1971; Squire, 1992; McClelland *et al.* 1995) of a class of memory, of which spatial memory is often considered to be just a subclass. As indicated in the

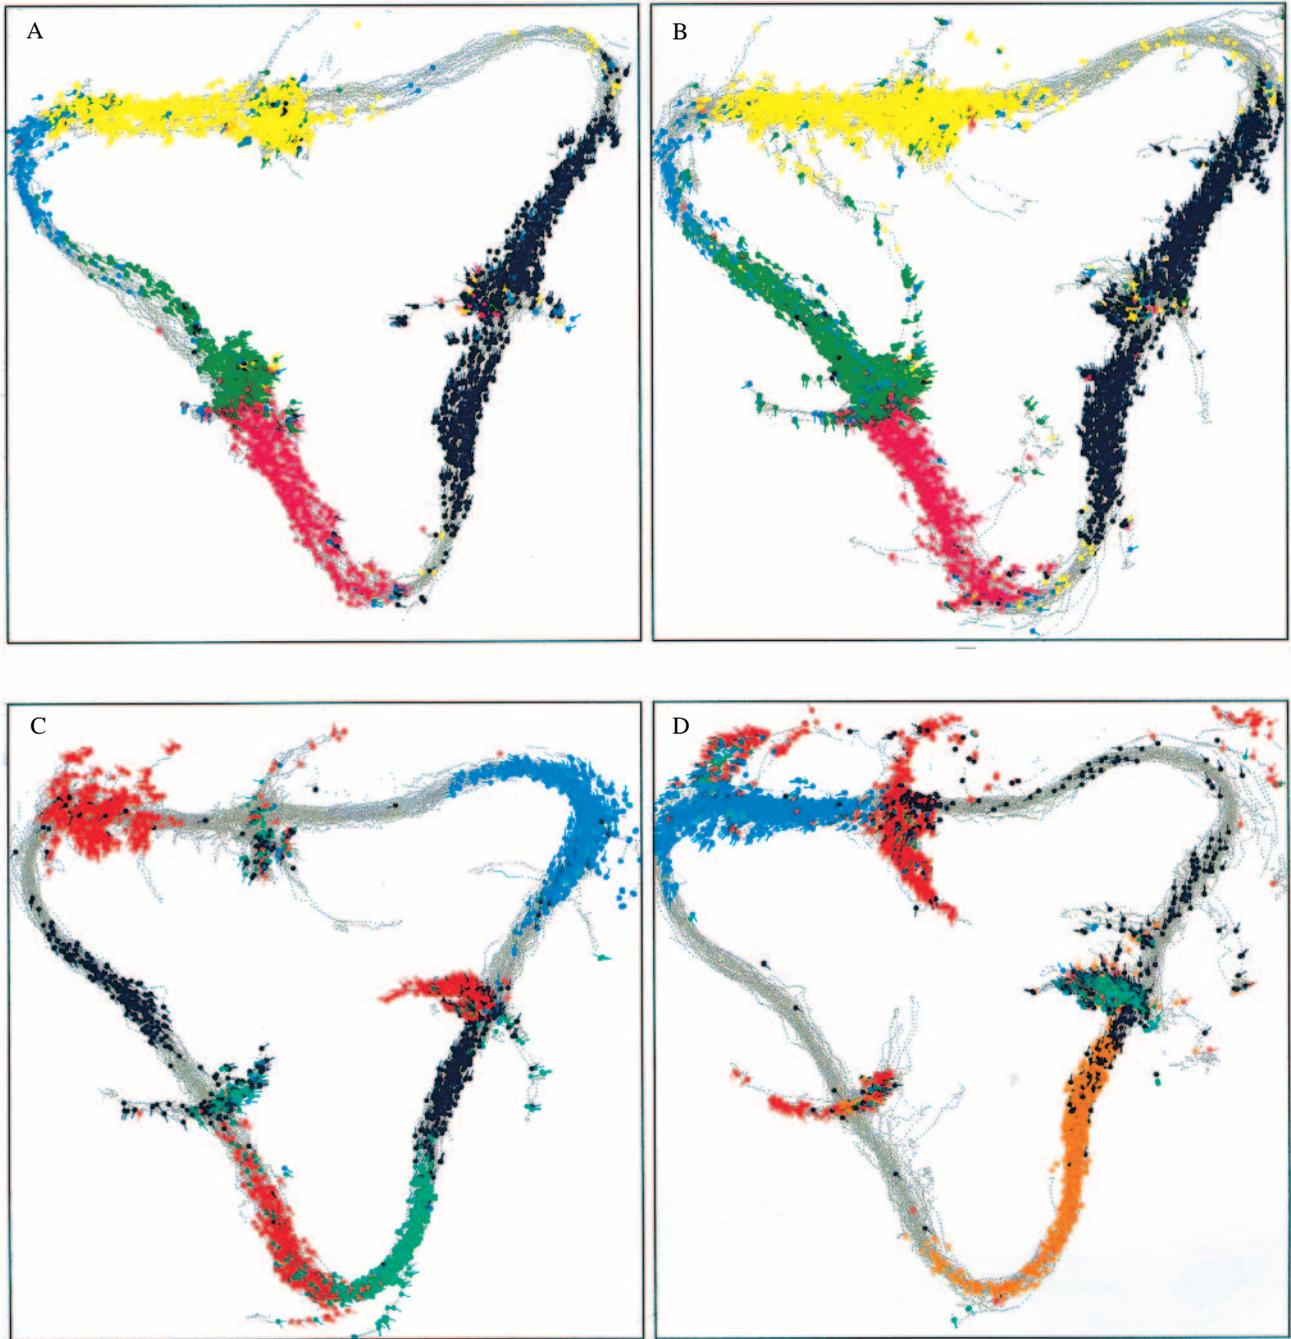
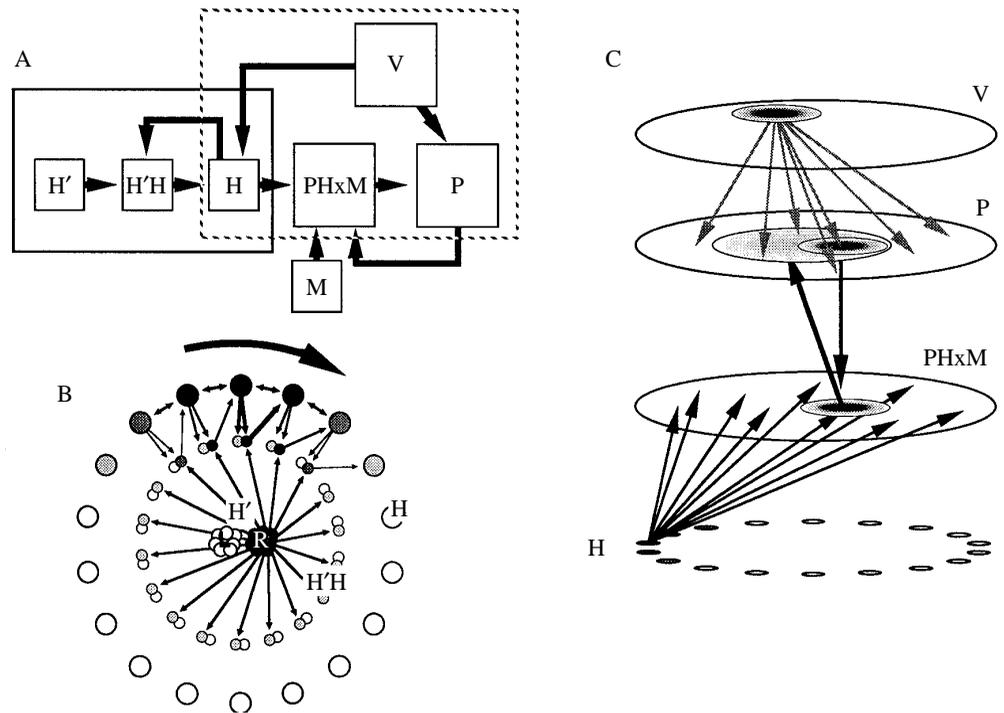


Fig. 4. All-or-nothing reconfiguration of the hippocampal spatial representation of the familiar recording environment, following a 1 h period of extensive experience in six different rooms, which the animal had never previously visited. The recording environment consisted of a moderately illuminated room, with a raised triangular track located in the middle. The rat's task was to move continuously in one direction around the track to receive food reinforcement at the mid-point of each side. A control session is illustrated in A and B. For this session, the rat was returned to its home environment for the 1 h period between the two recordings. An experimental session is illustrated in C and D. Place fields are illustrated by representing the discharges of individual units as colored dots, with the color corresponding to the particular unit. These recordings were conducted on different days, and the units shown in A and B are not the same as those shown with the corresponding colors in C and D. The gray lines indicate locations visited by the rat. The five different units shown in A and B maintained their place fields between recordings. In contrast, the firing distribution during the experimental session was radically altered. Of the four place cells illustrated in C, three exhibited completely different firing patterns following the experience (D) and one ceased firing altogether on the maze (although it was subsequently active while the animal sat quietly outside the maze). Additional cells, which were not active on the triangle in the initial session, developed fields on the platform following the novel spatial experience. The examples shown constitute a small proportion of the total number of cells recorded simultaneously in the respective experiments and were chosen for illustration purposes.

Fig. 5. Hypothetical organization of the path integration system at the neuronal level. (A) A general outline of the presumed excitatory connections among the various components of the path integration system. The head-direction system is indicated by H. These cells signal the azimuthal component of the animal's head orientation with respect to the gravitational horizontal plane. Their firing is independent of either pitch or roll and also independent of the location of the rat. Information concerning the direction and magnitude of rotation is provided by the H' system, which is presumed to be primarily, but not exclusively, vestibular. The H and H' systems converge on the H'H system, which contains cells that are selective for both azimuth and the direction of rotation in the horizontal plane. The hippocampal place cells, particularly those of CA3, are indicated by P.



Both the H and P systems converge on a separate population of cells, the PHxM system, whose activity depends on both head direction and location and is nonspecifically modulated by locomotion (M). Finally, additional sensory information (V), largely, but not exclusively, visual in nature, projects to both the H and P systems *via* modifiable connections that enable the learning of associations between external sensory stimuli and the preconfigured, abstract representations of head direction and location. The portions of the system enclosed in the solid and broken rectangles are illustrated in more detail in B and C respectively. (B) The angular velocity integration system is composed of three groups of cells. The head-direction cells can be thought of as constituting a ring-like structure in which adjacent points in the ring are strongly interconnected and more distant points are weakly or not connected. The H'H layer comprises a second ring-like structure whose inputs from the H layer are in point-to-point register. In the illustration, it is presumed that the animal is rotating to the right, and the corresponding R cells in the H' layer are active. The projections from the H'H layer to the H layer are asymmetrical. Those cells receiving input from right-turn-selective cells in the H' layer project to the right of the cells in the H ring from which they receive input. A right turn thus causes the focus of activity in the H ring to rotate to the right. The modifiable connections from the external sensory systems (V) to the H layer are not shown, but it is presumed that there is rich input, from cells conveying information about specific landmarks, distributed uniformly throughout the H layer. (C) The representation of location in the P layer is updated by information about the rat's heading and whether it is moving forward. This is accomplished in the following manner: the current focus of activity in the P layer results in a corresponding focus of input in the PHxM layer. Within this focus, cells receive differential input from the H layer and project asymmetrically back to the P layer, leading to an asymmetry of the excitation in the P layer in the direction of motion. This causes the focus to shift in the direction of motion of the animal. The intensity of output from the PHxM layer is presumed to be proportional to the forward velocity, which thus controls the rate of translation of the focus of activity in the P layer. Information about specific landmarks or external events is projected, *via* widespread, modifiable connections (stippled arrows), throughout the P layer. Where these connections encounter a focus of activity, they are strengthened. Subsequently, the similar input from the V layer can reactivate the same focus. This serves both as a means for establishing the animal's initial location and to correct for accumulated error in the path integration mechanism.

Introduction, however, space provides the contextual background within which most episodic memory is encoded (Nadel *et al.* 1985), and the proposed view of the organization of the hippocampus is entirely consistent with both themes. The preconfiguration of hippocampal state-space into a two-dimensionally organized system of stable attractors would provide a robust mechanism for the spontaneous, off-line reactivation of recent experience that is thought to be necessary for memory consolidation and which has been shown to occur in the hippocampus during sleep (Pavlidis and Winson, 1989; Wilson and McNaughton, 1994). It has been shown (Shen and McNaughton, 1994), for example, that such a preconfigured system of spatial reference frames will spontaneously

reactivate representations of locations within a given frame in the presence of random noise input and that representations associated with recently experienced locations can be preferentially reactivated either by increasing the excitability of recently active neurons or by superimposing a transient increase in the connection weight among coactive neurons with overlapping place fields. The present model proposes that landmark (and event) information becomes secondarily bound to these preconfigured representations by Hebbian learning. The same sort of plasticity between the hippocampal return projections and the neocortical zones from which the external input was derived would permit the spontaneous reinstatement of the original events in the neocortex. This organization would

enable the eventual development of an efficient representation of episodic and categorical knowledge within the neocortex using slower but more powerful learning mechanisms (McClelland *et al.* 1995). One advantage of such a system is that it would provide an automatic mechanism of selecting orthogonal representations for events that are similar to each other, but occur in different locations. Such orthogonalization is widely believed to be essential for the maximization of event storage capacity of any autoassociative memory system (Marr, 1971; McNaughton and Nadel, 1989).

The theory outlined above has much in common with the ideas of several other authors, such as O'Keefe and Nadel (1978) or Worden (1992), who have attempted to provide an explanation of functional role of the hippocampus *per se* or who, like Touretzky and his colleagues (Wan *et al.* 1994), have sought a more comprehensive account of the neural mechanisms of rodent navigation in general. Our hypothesis differs from most other suggestions in that it proposes, within the hippocampus, an explicit synaptic matrix for path integration and orientation in space, which would preserve and possibly enhance the general associative memory capabilities of the system. It also appears to explain many of the perplexing phenomena in the place cell literature that were reviewed above and that have hitherto been unexplained or explained only with difficulty. The hypothesis should be regarded as preliminary as there exist and will, no doubt, continue to be discovered data that will necessitate its revision. Moreover, the hypothesis raises as many questions as it answers. For example, what governs the selection of a new reference frame in a familiar environment when task variables or reinforcement sites change? What initiates the transition between one reference frame and another? What are the limits on the extent of a reference frame and do representations of extended environments constitute a single frame or a mosaic of frames? Finally, how might different reference frames be integrated to enable efficient, long-range navigation?

He had bought a large map representing the sea,  
Without the least vestige of land:

And the crew were much pleased when they found it to be  
A map they could all understand.

'What's the good of Mercator's North Poles and Equators,  
Tropics, Zones, and Meridian Lines?'

So the Bellman would cry: and the crew would reply  
'They are merely conventional signs!'

Other maps are such shapes, with their islands and capes!  
But we've got our brave Bellman to thank'

(So the crew would protest) 'that he's bought us the best –  
A perfect and absolute blank!'

(From Lewis Carroll's *The Hunting of the Snark*, as cited in Skaggs, 1995.)

Supported by NS20331, the Office of Naval Research, MH01227 and a McDonnell-Pew Center Grant for Cognitive Neuroscience. We also acknowledge the contribution of M. A. Wilson, R. D'Monte, K. Moore and L. Frank to the collection

of some of these data, and K. Stengel, M. Williams and J. Wang for hardware, software and/or analysis contributions.

## References

- ALYAN, S. AND JANDER, R. (1994). Short-range homing in the house mouse, *Mus musculus*: Stages in the learning of directions. *Anim. Behav.* **48**, 285–298.
- AMIT, D. J. (1989). *Modeling Brain Function: The World of Attractor Networks*. New York: Cambridge University Press.
- BARLOW, J. S. (1964). Inertial navigation as a basis for animal navigation. *J. theor. Biol.* **6**, 76–117.
- BOSTOCK, E., MULLER, R. U. AND KUBIE, J. L. (1991). Experience-dependent modifications of hippocampal place cell firing. *Hippocampus* **1**, 193–206.
- BREESE, C. R., HAMPSON, R. E. AND DEADWYLER, S. A. (1989). Hippocampal place cells: stereotypy and plasticity. *J. Neurosci.* **9**, 1097–1111.
- CHEN, L. L., LIN, L.-H., BARNES, C. A. AND MCNAUGHTON, B. L. (1994a). Head-direction cells in the rat posterior cortex. II. Contributions of visual and ideothetic information to the directional firing. *Exp. Brain Res.* **101**, 24–34.
- CHEN, L. L., LIN, L.-H., GREEN, E. J., BARNES, C. A. AND MCNAUGHTON, B. L. (1994b). Head-direction cells in the rat posterior cortex. I. Anatomical distribution and behavioral modulation. *Exp. Brain Res.* **101**, 8–23.
- ETIENNE, A. S. (1987). The control of short-distance homing in the golden hamster. In *Cognitive Processes in Spatial Orientation in Animal and Man* (ed. P. Ellen and C. Thinus-Blanc), pp. 223–251. Dordrecht, The Netherlands: Martinus Nijhoff.
- FOSTER, T. C., CASTRO, C. A. AND MCNAUGHTON, B. L. (1989). Spatial selectivity of rat hippocampal neurons: dependence on preparedness for movement. *Science* **244**, 1580–1582.
- FUKUDA, M., KOBAYASHI, T., ONO, T. AND TAMURA, R. (1992). Effects of learning place significance on rat hippocampal place cells. *Soc. Neurosci. Abstr.* **18**, 1426.
- GOTHARD, K. M., SKAGGS, W. E. AND MCNAUGHTON, B. L. (1995). Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. *J. Neurosci.* (in press).
- HEBB, D. O. (1949). *The Organization of Behavior*. New York: Wiley.
- HILL, A. J. (1978). First occurrence of hippocampal spatial firing in a new environment. *Exp. Neurol.* **62**, 282–297.
- HOPFIELD, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proc. natn. Acad. Sci. U.S.A.* **79**, 2554–2558.
- JUNG, M. W. AND MCNAUGHTON, B. L. (1993). Spatial selectivity of unit activity in the hippocampal granular layer. *Hippocampus* **3**, 165–182.
- JUNG, M. W., WIENER, S. I. AND MCNAUGHTON, B. L. (1994). Comparison of spatial firing characteristics of units in dorsal and ventral hippocampus of the rat. *J. Neurosci.* **14**, 7347–7356.
- KNIERIM, J. J., KUDRIMOTI, H. S. AND MCNAUGHTON, B. L. (1995a). Hippocampal place fields, the internal compass and the learning of landmark stability. *J. Neurosci.* **15**, 1648–1659.
- KNIERIM, J. J., KUDRIMOTI, H. S., SKAGGS, W. E. AND MCNAUGHTON, B. L. (1995b). The interaction between vestibular cues and visual landmark learning in spatial navigation. *Ann. N.Y. Acad. Sci.* (in press).
- KUBIE, J. L. AND RANCK, J. B. J. (1983). Sensory-behavioral correlates in individual hippocampus neurons in three situations: space and

- context. In *Neurobiology of the Hippocampus* (ed. W. Seifert), pp. 443–447. New York: Academic Press.
- KUDRIMOTI, H. S., MCNAUGHTON, B. L., BARNES, C. A. AND SKAGGS, W. E. (1995). Recent experience strengthens preexisting correlations between hippocampal neurons during sleep. *Soc. Neurosci. Abstr.* **21** (in press).
- MARKUS, E. J., BARNES, C. A., MCNAUGHTON, B. L., GLADDEN, V. L. AND SKAGGS, W. E. (1994). Spatial information content and reliability of hippocampal CA1 neurons: Effects of visual input. *Hippocampus* **4**, 410–421.
- MARKUS, E. J., QIN, Y., LEONARD, B. J., SKAGGS, W. E., MCNAUGHTON, B. L. AND BARNES, C. A. (1995). Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *J. Neurosci.* (in press).
- MARR, D. (1971). Simple memory: a theory for archicortex. *Phil. Trans. R. Soc. Lond. B* **262**, 23–81.
- MCCLELLAND, J. L., MCNAUGHTON, B. L. AND O'REILLY, R. C. (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. *Psych. Rev.* **102**, 419–457.
- MCNAUGHTON, B. L., BARNES, C. A. AND O'KEEFE, J. (1983). The contributions of position, direction and velocity to single unit activity in the hippocampus of freely-moving rats. *Exp. Brain Res.* **52**, 41–49.
- MCNAUGHTON, B. L., CHEN, L. L. AND MARKUS, E. J. (1991). 'Dead Reckoning,' landmark learning and the sense of direction: A neurophysiological and computational hypothesis. *J. cog. Neurosci.* **3**, 190–202.
- MCNAUGHTON, B. L., KNIERIM, J. J. AND WILSON, M. A. (1994). Vector encoding and the vestibular foundations of spatial cognition. In *The Cognitive Neurosciences* (ed. M. Gazzaniga), pp. 585–595. Boston: MIT Press.
- MCNAUGHTON, B. L., LEONARD, B. AND CHEN, L. (1989). Cortical–hippocampal interactions and cognitive mapping: A hypothesis based on reintegration of the parietal and inferotemporal pathways for visual processing. *Psychobiol.* **17**, 230–235.
- MCNAUGHTON, B. L. AND MORRIS, R. G. M. (1987). Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends Neurosci.* **10**, 408–415.
- MCNAUGHTON, B. L. AND NADEL, L. (1989). Hebb–Marr networks and the neurobiological representation of action in space. In *Neuroscience and Connectionist Theory* (ed. M. A. Gluck and D. E. Rumelhart), pp. 1–63. Hillsdale: Lawrence Erlbaum Associates.
- MITTELSTAEDT, M. L. AND MITTELSTAEDT, H. (1980). Homing by path integration in a mammal. *Naturwissenschaften* **67**, 566–567.
- MIZUMORI, S. J. Y. AND WILLIAMS, J. D. (1993). Directionally selective mnemonic properties of neurons in the lateral dorsal nucleus of the thalamus of rats. *J. Neurosci.* **13**, 4015–4028.
- MULLER, R. U. AND KUBIE, J. L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J. Neurosci.* **77**, 1951–1968.
- MULLER, R. U., KUBIE, J. L. AND RANCK, J. B., JR (1987). Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *J. Neurosci.* **77**, 1935–1950.
- NADEL, L., WILNER, J. AND KURZ, E. M. (1985). Cognitive maps and environmental context. In *Context and Learning* (ed. P. Balsam and A. Tomie), pp. 385–406. Hillsdale, NJ: Lawrence Erlbaum Associates.
- O'KEEFE, J. (1976). Place units in the hippocampus of the freely moving rat. *Exp. Neurol.* **51**, 8–109.
- O'KEEFE, J. AND DOSTROVSKY, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* **34**, 171–175.
- O'KEEFE, J. AND NADEL, L. (1978). *The Hippocampus as a Cognitive Map*. Oxford: Clarendon Press.
- O'KEEFE, J. AND SPEAKMAN, A. (1987). Single unit activity in the rat hippocampus during a spatial memory task. *Exp. Brain Res.* **68**, 1–27.
- PAVLIDES, C. AND WINSON, J. (1989). Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. *J. Neurosci.* **9**, 2907–2918.
- QUIRK, G. J., MULLER, R. U. AND KUBIE, J. L. (1990). The firing of hippocampal place cells in the dark depends on the rat's recent experience. *J. Neurosci.* **10**, 2008–2017.
- RANCK, J. B., JR (1984). Head direction cells in the deep cell layer of dorsal presubiculum in freely moving rats. *Soc. Neurosci. Abstr.* **10**, 599.
- SHARP, P. E. AND GREEN, C. (1993). Spatial correlates of firing patterns of single cells in the subiculum of the freely-moving rat. *J. Neurosci.* **14**, 2339–2356.
- SHARP, P. E., KUBIE, J. L. AND MULLER, R. U. (1990). Firing properties of hippocampal neurons in a visually symmetrical environment: contributions of multiple sensory cues and mnemonic properties. *J. Neurosci.* **10**, 3093–3105.
- SHEN, B. AND MCNAUGHTON, B. L. (1994). Simulation of the spontaneous reactivation of experience-specific hippocampal cell assemblies during sleep. *Soc. Neurosci. Abstr.* **20**, 1206.
- SKAGGS, W. E. (1995). Relations between the theta rhythm and activity patterns of hippocampal neurons. PhD dissertation, Program in Neuroscience, University of Arizona, Tucson, USA.
- SKAGGS, W. E., KNIERIM, J. J., KUDRIMOTI, H. S. AND MCNAUGHTON, B. L. (1995). A model of the neural basis of the rat's sense of direction. In *Advances in Neural Information Processing Systems*, vol. 7 (ed. G. Tesauro, D. Touretzky and T. Leen), pp. 173–180. Cambridge: MIT Press.
- SQUIRE, L. R. (1992). Memory and the hippocampus: a synthesis from findings with rats, monkeys and humans. *Psychol. Rev.* **99**, 195–231.
- TAUBE, J. S. (1995). Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. *J. Neurosci.* **15**, 70–86.
- TAUBE, J. S., MULLER, R. U. AND RANCK, J. B. J. (1987). A quantitative analysis of head-direction cells in the postsubiculum. *J. Neurosci.* **13**, 1332.
- THINUS-BLANC, C., BOUZOUBA, L., CHAIX, K., CHAPUIS, N., DURUP, M. AND POUCEY, B. (1987). A study of spatial parameters encoded during exploration in hamsters. *J. exp. Psychol.: Anim. Behav. Proc.* **13**, 418–427.
- TREVES, A. AND ROLLS, E. T. (1994). A computational analysis of the role of the hippocampus in memory. *Hippocampus* **4**, 373–391.
- WAN, H. S., TOURETZKY, D. S. AND REDISH, A. D. (1994). Towards a computational theory of rat navigation. In *Proceedings of the 1993 Connectionist Models Summer School* (ed. M. Mozer, P. Smolensky, D. S. Touretzky, J. Elman and A. Wigend), pp. 11–19. Hillsdale: Lawrence Erlbaum Associates.
- WILSON, M. A. AND MCNAUGHTON, B. L. (1993). Dynamics of the hippocampal ensemble code for space. *Science* **261**, 1055–1058.
- WILSON, M. A. AND MCNAUGHTON, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science* **265**, 676–679.
- WORDEN, R. (1992). Navigation by fragment fitting: a theory of hippocampal function. *Hippocampus* **2**, 165–188.