

Neuronal Computations Underlying the Firing of Place Cells and Their Role in Navigation

Neil Burgess and John O'Keefe

*Department of Anatomy, University College London,
London, England*

ABSTRACT: Our model of the spatial and temporal aspects of place cell firing and their role in rat navigation is reviewed. The model provides a candidate mechanism, at the level of individual cells, by which place cell information concerning self-localization could be used to guide navigation to previously visited reward sites. The model embodies specific predictions regarding the formation of place fields, the phase coding of place cell firing with respect to the hippocampal theta rhythm, and the formation of neuronal population vectors downstream from the place cells that code for the directions of goals during navigation. Recent experiments regarding the spatial distribution of place cell firing have confirmed our initial modeling hypothesis, that place fields are formed from Gaussian tuning curve inputs coding for the distances from environmental features, and enabled us to further specify the functional form of these inputs. Other recent experiments regarding the temporal distribution of place cell firing in two-dimensional environments have confirmed our predictions based on the temporal aspects of place cell firing on linear tracks. Directions for further experiments and refinements to the model are outlined for the future. © 1997 Wiley-Liss, Inc.

KEY WORDS: hippocampus, rat, spatial, model, network

INTRODUCTION

The Goal of Modeling

This work reviews our recent attempts to construct a model of the role of the rat hippocampus in spatial navigation, encompassing the relevant behavioral, anatomical and electrophysiological data (see Burgess et al. 1993, 1994, 1996; Burgess and O'Keefe, 1995). Our aim is to provide a quantitative framework in which to examine the neuronal mechanisms underlying both 1) the spatial and temporal properties of cellular activity in the hippocampus of the freely moving rat, and 2) the rat's ability to perform the type of sophisticated spatial navigation that is impaired by hippocampal lesions. To achieve this aim we must simultaneously consider the data available at the cell and system level of description. These data are outlined in the following two subsections.

Our model concentrates on the role of the hippocampus in place representation and place navigation. It has been argued (e.g., Eichenbaum et al., 1994; Squire, 1992) that a more general function (such as declarative or relational memory) needs to be ascribed to the hippocampus in order to account for all of the data. We believe that an adequate spatial model will show some of these seemingly nonspatial phenomena as emergent second-

ary properties. For example, as we shall see, the model described here predicts that a small proportion of cells in the hippocampus could respond to only a single cue in a single sensory modality (e.g., olfaction). On the other hand, even if the hippocampus were involved in the encoding of more general relationships, we would argue that understanding its role regarding spatial relationships would be the optimal route to follow since 1) the most cellular and behavioral data comes from the spatial domain, and 2) constructing a working spatial system strongly constrains a model, which could then be used to discriminate between more general relational functions, as it would be compatible with some, but not others.

Behavioral correlates of neuronal activity in and around the hippocampus

Recordings from single units in and around the hippocampus of freely moving rats have revealed striking behavioral correlates of cell firing. Place cells in regions CA3 and CA1 of the hippocampus (see Fig. 1) fire whenever the rat is in a restricted portion of its environment, and the location of firing (the "place field") can be controlled with respect to cues around the environment (see O'Keefe, 1976). Thus, rotating salient cues around an environment can cause the place fields to rotate (O'Keefe and Conway, 1978; O'Keefe and Speakman, 1987; Muller and Kubie, 1987), and increasing the size of the environment causes some place fields to expand in size (Muller and Kubie, 1987).

When rats move about in open-field environments (i.e., large open arenas) the firing rate of place cells is independent of the direction in which the rat is facing; however, it is modulated by direction when the rat is constrained to run back and forth along narrow-armed mazes (McNaughton et al., 1983; O'Keefe and Recce, 1993; Muller et al., 1994). By contrast to place cells, "head direction cells" have been identified in the dorsal presubiculum of the brain which fire only when the rat's head is pointing in a particular direction (Taube et al., 1990), regardless of the rat's location within its environment. Directionally tuned cells have also been recorded in retrosplenial cortex (Chen et al., 1994) as well as in

Accepted for publication July 15, 1996.

Address correspondence and reprint requests to Neil Burgess, Department of Anatomy, University College London, London WC1E 6BT, England.

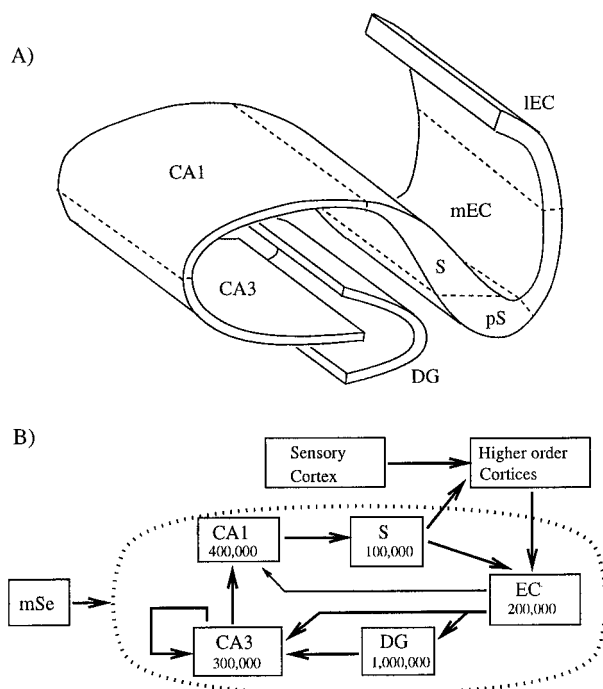


FIGURE 1. The hippocampus is formed from sheets of cells. **A:** A schematic section cut perpendicular to the longitudinal axis of the hippocampus. EC, entorhinal cortex (mEC, medial EC; IEC, lateral EC); S, subiculum; pS, pre- and para-subiculum; DG, dentate gyrus. **B:** The major projections and approximate numbers of the major cell type in each subfield (for the rat). mSe, medial septum. In the DG-CA3 projection a single mossy fiber projects from each granule cell, making very large synapses onto only 14 or so pyramidal cells. All the other projections have large divergence and convergence (many thousands to one), and involve the type of synapse in which "Hebbian" LTP has been observed. A variety of interneurons provide feedforward and feedback inhibition. The mSe plays a role in producing the θ rhythm of the hippocampal EEG. Cells in CA3 and CA1 also project out to the lateral septum via the fornix. The human hippocampus contains about one order of magnitude more cells. Adapted from Burgess et al. (1995).

the lateral dorsal thalamus (Mizumori and Williams, 1993) and anterior thalamic nuclei (Taube, 1995). The directional tuning appears to be governed by visual cues, but vestibular inputs to the system may also play an important role, particularly during initial experience of an environment (McNaughton et al., 1995).

Areas adjacent to the hippocampus have also revealed cells with spatially correlated firing. Quirk et al. (1992) reported finding cells in the entorhinal cortex with spatial firing fields that were broader and more diffuse than hippocampal place fields. This finding contributed to our choice of the inputs representation for our model (see below). More recently Sharp and Green (1994) reported that most subicular cells also showed a robust locational signal although their firing pattern differed from place cells in that it tended to cover much of the environment. As we shall see below, the existence of large firing fields was predicted by our initial model; however, the data regarding the firing patterns of subicular cells are not clear cut (see also Barnes et al., 1990; and Muller et al., 1991a).

Place fields possess spatial structure beyond that required simply to mark a location. The firing rate map of a place cell, showing the cell's firing rate as a function of the rat's location (typically averaged over a trial of several minutes), most often resembles an approximately symmetrical Gaussian with a single peak at a particular location in the environment. Thus, the firing rates of hippocampal place cells could be used as radial basis functions in spatial computations downstream of CA1 (see Burgess et al., 1994). In fact, the firing pattern of an ensemble of many place cells can be used to predict the rat's position (Wilson and McNaughton, 1993). If the firing rate f_i of a place cell i is considered to be a radial basis function peaked at position \tilde{x}_i , then the firing rate weighted vector average of these positions (the "population vector," see Georgopoulos et al., 1988),

$$\frac{\sum_i f_i \tilde{x}_i}{\sum_i f_i} \quad (1)$$

estimates the rat's position. However, note that near to the boundary this simple estimate becomes incorrectly biased toward the center of the environment (Muller et al., 1987). Place field shapes may also deviate from the simple symmetric Gaussian form, in a manner that appears to depend on the geometry of the walls surrounding an environment (O'Keefe and Burgess, 1996). We will return to these recent results in the discussion.

Place cell firing also shows interesting temporal properties. As the rat moves around its environment, place cells tend to fire in short bursts of between one and four action potentials. The timing of the firing of these bursts has a systematic relationship to the concurrently recorded electroencephalogram (EEG). Whenever the rat is involved in head-displacement movements or locomotion, the EEG exhibits a roughly sinusoidal oscillation of 6–11 Hertz, called the "theta" rhythm. As the rat runs through a place field on a linear track, the phase at which spikes are fired is not constant, but shifts in a systematic way (O'Keefe and Recce, 1993). When the animal enters the field the firing occurs late in the cycle but shifts to progressively earlier phases as it runs through the field (see Fig. 2). A subsidiary aim of the modeling was to explore the possible functional significance of this phase relationship. With this in mind, it is interesting to note that the ease of induction of long-term potentiation appears to be modulated by the phase of the EEG theta cycle at which stimulation occurs (Pavlidis et al., 1988).

Rat navigation

Rats rapidly learn the layout of an environment, both before the introduction of a goal (i.e., "latent learning"), and after it (see Tolman, 1948). Having explored the environment, they are able to return to the location of an unmarked goal after visiting it only a small number of times (see, e.g., Morris, 1981). Furthermore, if an experimenter introduces a shortcut to the goal in a maze with which a rat has become familiar, it is able to recognize and exploit the shorter route (see, e.g., Benhamou and Seguinot, 1995; Tolman, 1948). However, we note that the ability to take short cuts can be confounded by the rat's desire to explore the changes made to its environment, and that the extent of this ability has been questioned by some authors (see, e.g., Sutherland et al.,

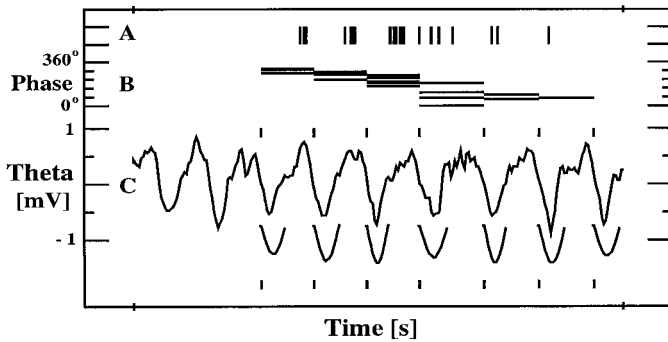


FIGURE 2. Typical pattern of place cell firing relative to the EEG θ rhythm as a rat runs through the firing field on a linear track: 1 second of the EEG θ rhythm is shown in C; vertical ticks immediately above and below mark the positive to negative zero-crossings of the EEG which we use to divide it into cycles. A shows the times of firing of the place cell. B shows the phase within a θ cycle at which each spike fired. Adapted from O'Keefe and Recce (1993).

1987) along with that of latent learning (see, e.g., Whishaw, 1991). In summary, rats are capable of sophisticated navigation, beyond simple skills like following a familiar route or approaching a visible cue (see O'Keefe and Nadel, 1978).

The hippocampus plays a crucial role in the ability of rats to perform spatial tasks of the type outlined above. Solving this type of spatial task is severely impaired after hippocampal lesions in rodents (O'Keefe et al., 1975; Olton et al., 1978; Morris et al., 1982; see Jarrard, 1993, for a review), including finding novel shortcuts (Harley, 1979). More recently, lesion of the hippocampus in primates has been shown to impair the solving of spatial tasks (Parkinson et al., 1988; Murray et al., 1989). Finally, evidence from neuropsychology and functional brain imaging has begun to implicate the human hippocampal formation in tasks involving spatial memory (Smith et al., 1995; Feigenbaum et al., 1996) or real-world navigation (Aguirre et al., 1995; Maguire et al., 1996).

How To Evaluate the Model

The success with which our model addresses the above data should be evaluated in terms of the insights that it generates into how the anatomy and electrophysiology of the hippocampal region produces the behavior observed in spatial tasks of both the place cells on the one hand and the animal on the other. Our major initial concern was to demonstrate that place cell firing could be used for more than just self-localization, i.e., that it could also be used to guide navigation. A second, parallel, objective was to elucidate the neuronal mechanisms underlying the observation of place fields. A third, subsidiary, interest was to explore the possible functional use of the *phase coding* of place cell firing and EEG theta rhythm, as opposed to the *rate coding* used in most conventional connectionist models. In the long term, the model should be judged by the new experimental directions indicated by it, and by whether the process of new experiments and subsequent refinements to the model makes a cumulative contribution to our understanding of the working of the hippocampus.

THE MODEL

In this section we provide a synopsis of the current state of our computational model of the hippocampal region and its role in spatial navigation. Details of simulations can be found elsewhere (most exhaustively in Burgess et al., 1993, 1994).

Introduction, a Simple Model

A simple model of navigation based on place cell firing could work in the following way. When the rat encounters a "goal" (i.e., a location in its environment that is associated with reward), a goal cell downstream from the place cells is strongly excited by the attributes of the goal (e.g., food, water, safety, etc.). At the goal, a one-shot Hebbian increment is induced in the synaptic connections to the goal cell from the place cells that are active at the goal location. Then, as the rat moves away from the goal location, the net activity of place cells with strong connections to the goal cell will be a monotonically decreasing fraction of the total place cell activity. Consequently the activation of each goal cell (as a fraction of net place cell activity) will code for the proximity of a goal location, and thus could, for example, be used as an evaluation function in a gradient-ascent-type search for that goal. That is, the animal could find its way back to a goal location simply by moving around so as to increase the firing rate of the appropriate goal cell. This simple navigational model is illustrated in Figure 3.

Notice that the construction of place cell responses during exploration (see below) enables "latent learning" of environmental structure, independently of the existence of any reward. The "map" represented by the place cells can then be applied to code for a particular goal location after one visit to that location. Both of these features of this simple model reflect the navigational ability of real rats in a way that learning algorithms requiring hundreds of visits to a goal do not.

A Population Vector Model

The simple model outlined above has the disadvantage that the rat would need to hunt around at any given location in order to determine the direction in which to move so as to increase the goal cell's firing rate. Behavioral evidence for such a "hunting" strategy is scarce (but see the work of Brown, 1992, on microchoices). On the contrary, there is evidence that rats know the direction of a goal location even though they may not have previously been able to move directly toward it, i.e., they can take short-cuts (see, e.g., Benhamou and Seguinot, 1995; Tolman, 1948). We propose instead that the output of the hippocampal system is a set of population vectors indicating the direction and proximity of the rat from the various goal locations it has encountered during exploration.

In the simple model, a single goal cell represents the location of the goal by taking a "snapshot" of the place cell activity at the goal. The goal cell subsequently fires at a rate proportional to the net activation of all of the place cells active at the time of the snapshot. The goal cell's firing rate map (the function relating its firing rate

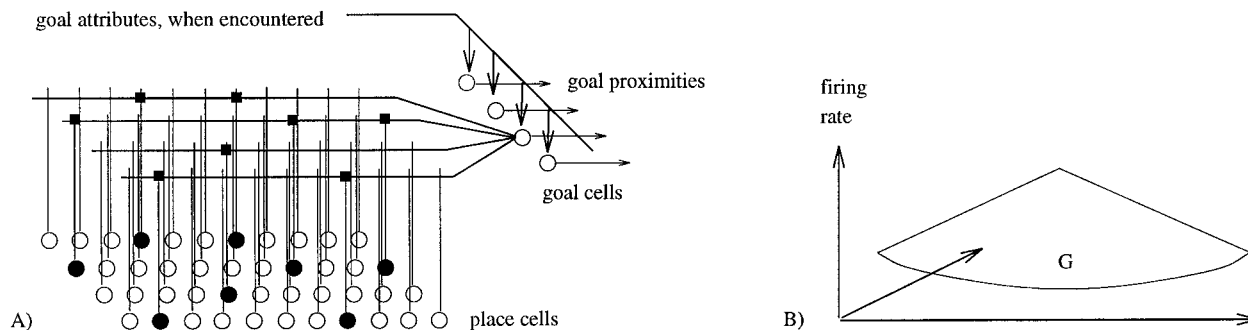


FIGURE 3. Simple model of place cells and navigation. A: A “goal” cell stores a goal’s location by taking a snapshot of place cell activity via Hebbian synaptic modification when the goal cell is excited by the attributes of a particular goal location. Solid circles are active place cells; open circles are inactive place cells and solid squares

mark potentiated synapses. B: The firing rate map of the goal cell during subsequent movements of the rat (roughly an inverted cone) codes for the proximity of the goal (G). Adapted from Burgess and O’Keefe (1995).

to the rat’s position) is thus peaked at the goal location, and decreases in magnitude uniformly with distance from that location (see Fig. 3B). The population vector model replaces each goal cell by a set of goal cells, each cell having a firing rate map that peaks near to, but not exactly at, the goal location. Each goal cell in a set (cell i , say) represents the allocentric direction \tilde{u}_i (e.g., north) in which the location of its peak firing is displaced relative to the goal location. Note that this does not mean that the goal cell’s firing rate is modulated by the rat’s heading direction. The vector sum of directions, weighted by the instantaneous firing rates of the corresponding cells (i.e., $\sum_i f_i \tilde{u}_i / \sum_i f_i$) codes for the direction of the rat from the goal, and the net firing rate ($\sum_i f_i$) codes for the goal’s proximity (see Fig. 4A).

Phase coding and construction of a direction-from-goal population vector

In Burgess et al. (1993) we interpreted the relationship between place cell firing and the phase of theta during runs on a linear track (see O’Keefe and Recce, 1993) as evidence of a more general phase coding in which cells firing “late” in a theta cycle code for places ahead of the rat, and cells firing “early” code for places behind it (see Fig. 4B). This interpretation has since been supported by analyses of place cell firing in two-dimensional (2-D) environments (see Burgess et al., 1994; Skaggs et al., 1996; and Fig. 4C). The use of the extra information carried by the phase of firing makes it easy for a population vector coding for the instantaneous direction of a goal to be formed. We describe our model for this below.

Each goal location is represented by a set of goal cells, all receiving projections from a large number of CA1 place cells. In addition, each goal cell also receives a projection coding for both the rat’s head direction and for the sensory attributes (e.g., food, water, etc.) of a particular goal. This latter connection serves to deliver a “learn now” type reinforcement signal to a goal cell whenever the rat is at the appropriate goal location and facing in the appropriate direction. If this learning signal arrives at a “late” phase of the EEG, or if synaptic plasticity is restricted to this phase, then the goal cell associated with, e.g., the direction north, will form active connections from place cells with receptive fields

centered to the north of the location of the goal. In this way each of the cells in a set of goal cells become associated with a particular allocentric direction. The maximum firing rate of each cell in a set then occurs at a location displaced from the goal in the direction with which that cell is associated (see Fig. 4A). Thus, the population vector of the directions associated with each goal cell, weighted by their firing rates, provides the direction of the rat from the goal. This can be used by navigational systems downstream of the hippocampus during subsequent movement about its environment. Furthermore, the net firing rate of the set of goal cells indicates the goal’s proximity. For example, whenever the rat is to the north of a previously encountered goal location, the goal cell representing north will be firing more than one representing South (i.e., pulling the population vector northward), and the net firing rate of all of the cells representing that goal will increase as the rat moves nearer to the goal (indicating its increasing proximity).

Inputs to the model and the neural basis of place cell firing

Zipser (1985) explored the idea that the response of place cells might result from a matching function between the value of a sensory stimulus (e.g., retinal size) corresponding to a particular cue, and a stored value of that stimulus. This type of model was later implemented as a neural network by Sharp (1991). Inspired by these approaches, the inputs to our model were taken to be sets of “sensory cells,” drawn from a range of sensory modalities, each set responding to a particular cue. The firing rate of the cells in a set has a broad tuning curve response to the distance of the cue from the rat, the tuning curve for each cell being peaked at a different distance. The transformation of the broad tuning curve of the sensory input cells into well-localized place fields is mediated by the entorhinal cells (ECs). Each of these cells receives connections from two sensory cells.

In our model we do not consider the observed phase coding of place cell firing to be generated by a mechanism that is completely internal to the hippocampus. Rather we assume that it results from the temporal properties of the inputs to the hippocampus. One possible scenario is that the direct septal input is modulated

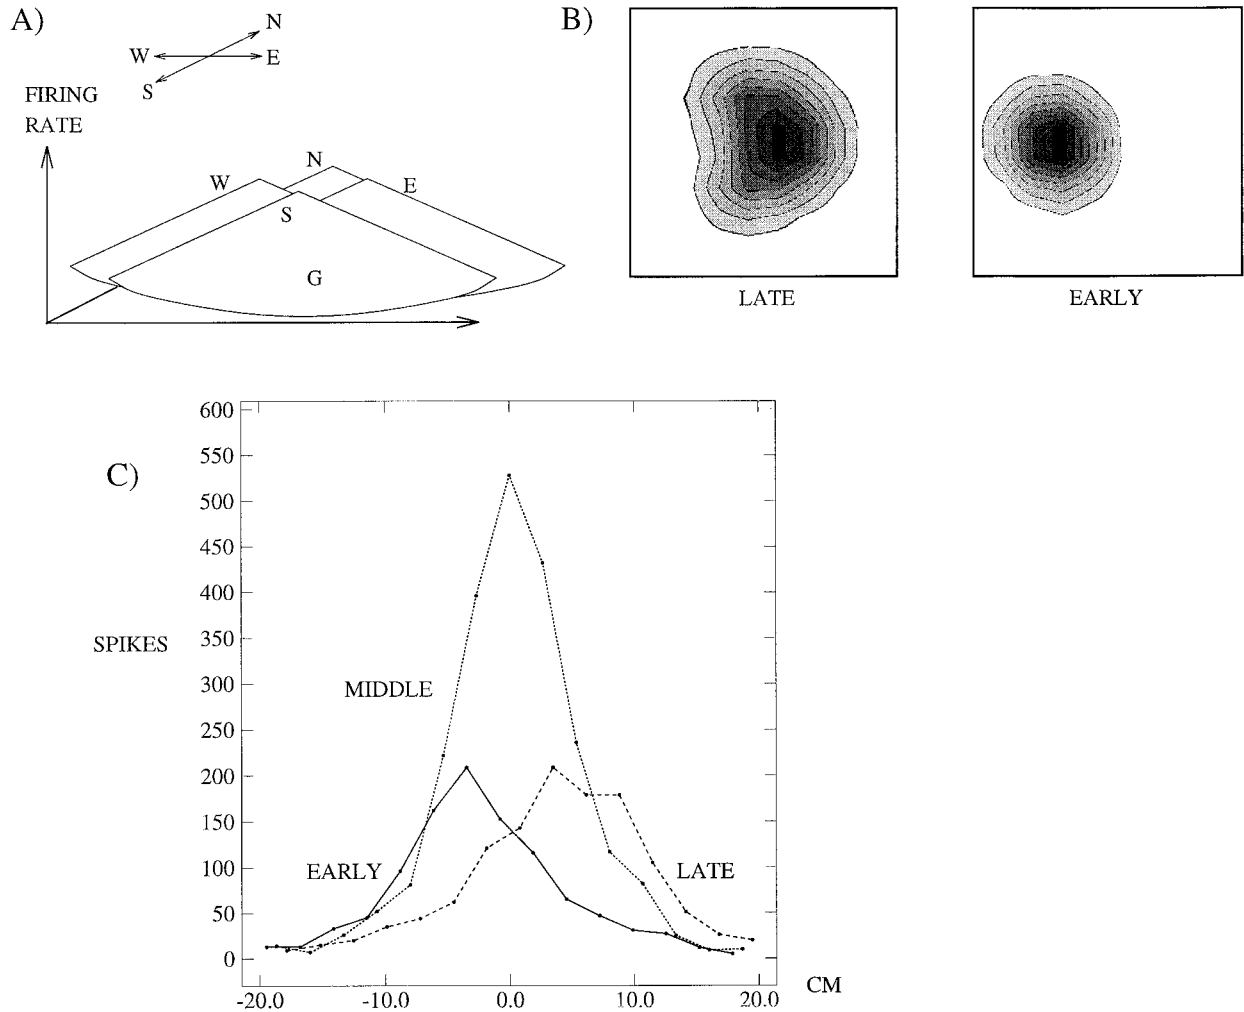


FIGURE 4. Population vector model of navigation. A: The firing rate maps of 4 goal cells whose population vector codes for the goal location (G). Each is associated with the allocentric direction \vec{u}_i in which the location of its peak firing rate is displaced from G; thus the vector sum of the directions \vec{u}_i weighted by the instantaneous firing rates f_i of the goal cells (i.e., $\sum_{i=1}^4 f_i \vec{u}_i / \sum_{i=1}^4 f_i$) codes for the direction of the rat from the goal, and the net firing rate of the goal cells (i.e., $\sum_{i=1}^4 f_i$) codes for the goal's proximity. B: The predicted effect in two-dimensions of the phase coding of place cell firing with respect to the EEG theta rhythm observed on linear tracks (adapted from Burgess et al., 1993). If the rat is in the center of a square environment and moving to the right, box "LATE" shows the net firing rate map of all the place cells active at a late phase, and box "EARLY" shows the net firing rate map of all the place cells active at

an early phase. In other words, place cells firing at a late phase should have place fields centered ahead of the rat, and those firing at an early phase should have place fields centered behind it. C: Data showing the effect of phase coding in two-dimensions (adapted from Burgess et al., 1994; see also Skaggs et al., 1996). The plot shows separate histograms for spikes fired at early, middle, and late phases, binned in terms of the distance ahead (positive numbers, in centimeters) or behind the rat of the location of the cell's peak firing rate. Those spikes fired at a late phase come mostly from cells with place fields ahead of the rat, as predicted. Four thousand nine hundred thirty-two spikes were recorded from 13 place cells in three rats: 312 spikes could not be assigned a phase due to lack of θ rhythm (i.e., the fitted frequency was below 3 Hz or above 20 Hz), 132 spikes fall outside the domain of the histogram—the remaining 4,488 spikes are shown.

at the observed theta frequency f_θ , whereas the sensory input from entorhinal cortex is modulated at a slightly higher frequency f_b . The additive interaction of these inputs might cause the pyramidal cell membrane potential to oscillate according to their interference pattern, with the envelope of the net signal modulating the firing rate and the peaks in the signal determining the times at which bursts of spikes are fired. The bursting frequency of the place cell would then be $(f_a + f_b)/2$, resulting in a systematic phase shift with respect to f_θ . While this scenario remains highly speculative, a need for sensory or motor signals to play a role in the

generation of phase coding is indicated by the observation that the phase of firing codes more accurately for the rat's position on the linear track than for the time elapsed since entering the place field (O'Keefe and Recce, 1993). Evidence that phase coding occurs at least before region CA3 has been provided by Skaggs et al. (1996), who report a small amount of phase precession in the firing of granule cells in the dentate gyrus. Nicolelis and Chapin (1994) provide some evidence for oscillatory modulation of sensory signals (from the whisker pad), and Macrides (1975) indicates a temporal relationship between myostatial sniffing and theta.

As a simple model of phase precession, relying on phase-coded sensory inputs, the approximate egocentric direction of the centroid of the two cues providing the input to an entorhinal cell from the rat (ahead, to the side, or behind) is coded by its phase of firing (late, middle, or early, respectively). The selection of the pairs of sensory cells that are used to drive entorhinal firing is weighted toward those that respond maximally at cue-distances approximately equal to half of the separation of the two cues concerned. This biases entorhinal receptive fields to be peaked near to the centroid of the two cues, and thus produces the desired phase coding, with a tendency for cells firing at a late phase to have receptive fields centered ahead of the rat. This property propagates through the model.

The selection of the pairs of sensory cells that are used to drive entorhinal firing is also arranged so that plenty of entorhinal cells have fields at or beyond the edge of the environment. This enables the goal cell population vector to code accurately for goals near the edge (for details see Burgess et al., 1994). Note that if enough cells code for locations up to one place field diameter beyond the edge of the environment, the central bias problem of a population representation of space (see Muller et al., 1987) is avoided.

The Neuronal and Behavioral Simulation

In order to evaluate the performance of this model, we have constructed a computer simulation. We make the simplifying assumptions that different subsets of place cells are active in different environments, and that the role of the dentate gyrus (DG) and CA3 recurrent collaterals (see Fig. 1) is to generate the different subsets as necessary. We consider only a single environment and simulate an active subset of place cells for that environment, ignoring DG, the CA3 recurrent collaterals, and merging fields CA3 and CA1 into a single set of place cells.

The firing of the “sensory cells” is meant to be analogous to the visual, olfactory, auditory, or interoceptive signals a real rat experiences. These inputs vary as a function of the rat’s position and propagate through three “hidden” layers of neurons, corresponding to entorhinal cortex (EC), place cells (PC, the combined CA1 and CA3), and subicular cells (SC), to an output layer which is a population of eight “goal” cells. The θ rhythm of the EEG is taken to be 10 Hz; each θ cycle is divided into three phases: “early,” “middle,” and “late.” The rat’s direction of motion is updated at the end of each θ cycle. The model has two behavioral states: during the “exploration” state the rat’s next direction is simply a random variable within 30° of its previous direction; during “searching” its direction is determined by the average of its previous direction and that indicated by the goal cells. The model is schematized in Figure 5 (see Burgess et al., 1994, for details).

Learning in the model

The activations of place and subicular cells, and the connections between them, are governed by competitive learning rules (detailed in Burgess et al., 1993, 1994) which transform coarse position-related sensory inputs into robust, place-specific cell firing fields. The competitive learning paradigm used is able to build stable firing fields very quickly, a major empirical constraint (see Hill, 1978; Wilson and McNaughton, 1993) which rules out

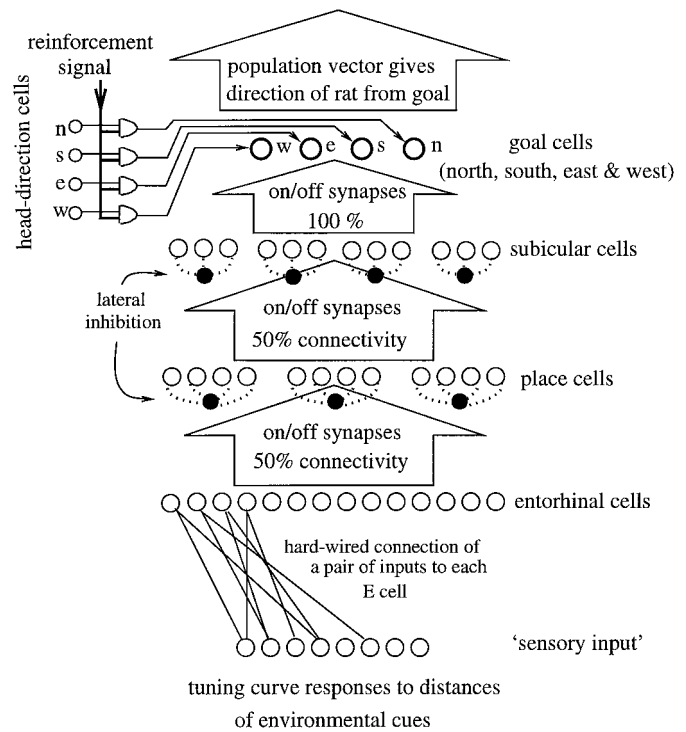


FIGURE 5. The population vector model of the role of the hippocampus in navigation. The position of the rat relative to cues placed around the environment determines the input to the network; the output determines what direction the rat should move in to reach the goal. Adapted from Burgess et al. (1994).

slower learning algorithms such as error back-propagation (Rumelhart et al., 1986). In brief, connections between layers have weights of 0 or 1 and initially only a few connections have weight 1, “seeding” the future locations of place fields. Once the rat enters the environment, connections between simultaneously active units are switched to weight 1, and layers are divided into groups within which units compete to be active (cf. Rumelhart and Zipser, 1986). The unsupervised nature of the scheme eliminates the need for something that already knows where the animal is to act as a “teaching input,” as required by supervised learning algorithms such as error back-propagation.

The firing fields of place cells end up being relatively small because, to fire at all, a place cell must overcome competition from many other cells. If place fields cover only a small fraction of the environment there is a problem in the locality of access to the place information: The firing of place cells far from a goal will contain no information regarding the goal location unless there are some cells that fire at both locations. To get around this, in the model subicular cells each compete with fewer cells than place cells, and can therefore have larger firing fields. These are built up from the superposition of many smaller place fields during exploration. However, we note the recent discovery of large place fields in ventral hippocampus (Jung et al., 1994) which indicate that the model’s layer of subicular cells may not be necessary, and that the goal cells themselves might occur directly in the subiculum (but see also Poucet et al., 1994, who found fields of a similar size to those in dorsal hippocampus).

Goal cells

All goals are taken to be distinguishable by the characteristics of the reward, and a different population of eight goal cells codes for each location. Whenever the rat encounters a goal location (food, water, safety, etc.) the goal cell corresponding to its current head direction receives a reinforcement signal at the “late” phase of θ . The rat looks around in eight directions 45° apart (north, northeast, east, etc.), for one θ cycle each, allowing the appropriate connections to be incremented to each cell in the set. Connections to a goal cell from a subicular cell are switched on whenever the goal cell receives a reinforcement signal and the subicular cell fires at the same time (this connection modification is made even if the subicular cell fires only weakly, so that goal firing fields end up big enough to cover the entire environment).

There can be many goals in one environment: which population of goal cells is used to guide navigation is presumed to depend on the desired sensory attributes of the goal, in the same way that the population of goal cells that receives the reinforcement signal at a goal location depends on the perceived sensory attributes of the goal. How the rat chooses among differing goals is not modeled here. “Obstacles,” in the sense of locations to be avoided, can be coded for in the same way as goals. Navigation is implemented at present simply as directly approaching the goal if there are no obstacles in the way. Obstacle circumnavigation can be accomplished by subtracting the population vector of cells coding for an obstacle that is estimated to be in the way from the population vector for the goal. However, this use of the hippocampal output is designed simply to demonstrate that it could support navigation, not as a serious model of path planning.

One possible limitation of using a goal cell population vector representation of goal direction is that a discernible difference in the firing rates of the goal cells representing opposite directions must be maintained whenever the rat is displaced from the goal along that direction (see Fig. 6). It is interesting to speculate on the maximum range over which this type of coding could be maintained. The limited maximum firing rate of a goal cell (which might cause goal cell firing rate maps to have very shallow peaks in large environments) is not a constraint in itself, as many goal cells might code for each direction with a downstream system looking at the net firing of all of them. A more serious limitation is the acuity required of a downstream cell that makes the comparison of firing rates. If a difference in rate of 1% can be detected (i.e., $\Delta f/f > 0.01$, where f is the firing rate of one goal cell and Δf is the absolute difference between the cells’ firing rates) then goal cells with conical firing rate maps whose peaks are displaced from the goal location by a distance d would work over a range of up to $2df_{max}/\Delta f$ (see Fig. 6A). This implies a limit of about 10 meters since $d \approx 5$ cm (from Fig. 4C) and the maximum discernible value of $f_{max}/\Delta f$ is 100. However, if a goal cell’s firing rate falls exponentially with distance from the location of peak firing, then a fixed fractional difference $\Delta f/f$ in firing rates can be maintained, irrespective of distance from the goal (see Fig. 6B). In this case the goal cell representation works over the range within which the absolute level of goal cell firing can be detected (i.e., $f > f_{min}$). This

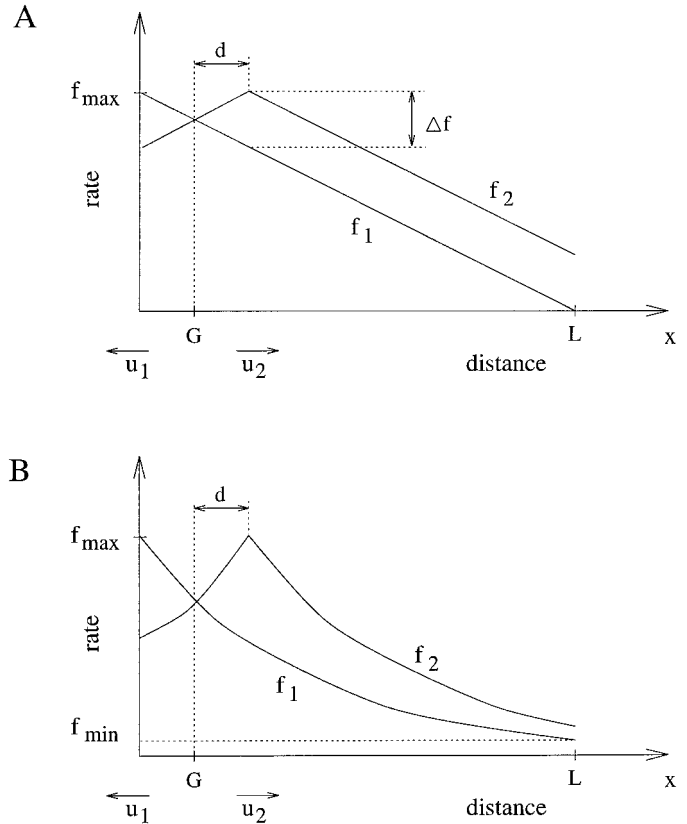


FIGURE 6. The range over which a goal cell population vector can guide navigation. The graphs show the firing rates f_1 and f_2 of two goal cells (associated with opposite directions \vec{u}_1 and \vec{u}_2) coding for goal location G, as a function of the projection of the rat’s position onto \vec{u}_2 . The peak firing rates are f_{max} occurring at a distance d from the goal location, and the difference in the rates is $\Delta f = |f_1 - f_2|$ when the rat is displaced from the goal by more than d along direction \vec{u}_2 . **A:** Conical (linearly decreasing) goal cell firing rate maps. In this case the absolute difference in rates Δf is constant. If the rat moves beyond L then goal cell 1 will stop firing, $(L - 2d)/L = (f_{max} - \Delta f)/f_{max}$, so if A is the acuity of the downstream system that reads the goal cell output (i.e., the minimum discernible value of $\Delta f/f_1$) then $L = 2d/A$ is the range over which the goal cell representation will work ($L \approx 10$ m; see text). **B:** Exponentially decreasing goal cell firing rate maps. In this case the relative difference in rates $\Delta f/f_1$ can be kept constant with distance from the goal, for example $f_1(x) = f_{max} \exp(\ln(1 - A)x/2d)$ and $f_2(x) = f_1(x - 2d)$ maintains $\Delta f/f_1 = A$. If a minimal firing rate of f_{min} is required, then the range over which the goal cell representation will work is given by $L = -2d \ln(f_{max}/f_{min})/\ln(1 - A)$, i.e., for $A \ll 1$ we have $L \approx 2d \ln(f_{max}/f_{min})/A$ giving $L \approx 46$ m, see text.

suggests a range of approximately $2d(f/\Delta f) \ln(f_{max}/f_{min})$ (see Fig. 6B). This implies a limit of about 46 meters if $f_{max} = 100 f_{min}$.

A second concern for population vector representations is how a putative downstream system could de-reference the information it contains. That is, how does the downstream system know which location \vec{x}_i place cell i codes for (see equation 1), or, in our model, which direction \vec{u}_i goal cell i codes for? In the case of the goal cells there are two possibilities: 1) when the connections to a goal cell (representing north, say) from active place cells are switched on (i.e., when the rat faces north at the goal location), connections are

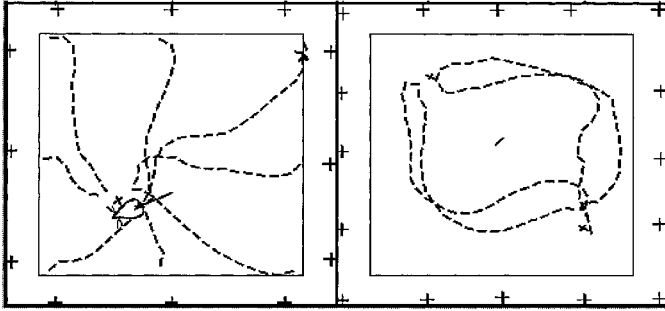


FIGURE 7. Left: simulated trajectories from 8 novel starting positions, to a goal encountered after 30 s of exploration (at 60 cm/s) in a $135 \times 135\text{-cm}^2$ environment; the rat is shown to scale. Right: simulated navigation between two goals with an “obstacle” in between; cues are marked by “+,” goals by “x,” and obstacles by “/.” Adapted from Burgess et al. (1995).

also modified associating that goal cell to the appropriate cells in the downstream system (which might also be active at that time due to an input from the head direction cells); 2) the connection of goal cells and cells in the downstream system that corresponding to the same direction (i.e., with connections from the same group of head-direction cells) might occur during development.

Performance

We have described a model for hippocampal function in which three intermediate layers of cells, representing entorhinal cortex, CA3/1, and subiculum, map the sensory input (from a number of intramaze and extramaze cues) to a population vector output (which determines the direction the rat will move in). Within the restrictions imposed by realistic simulation and the fact that we have left out parts of the hippocampal formation which we posit to have a role in distinguishing between multiple environments, the model is successful: it reproduces quite well the firing properties and interrelationships between entorhinal cells, place cells, and subicular cells (although experimental data on the latter are limited).

Moreover, place cell responses are set up rapidly during exploration, and in the absence of any particular goal. Goal locations encountered only once can reliably be returned to from novel starting positions (see Fig. 7, left). During the brief exploration necessary for navigation, the simulated rat does not cover the entire environment, but is nevertheless able to take shortcuts during subsequent navigation. Different sets of goal cells code independently for different goal locations in an environment so that a single spatial location can be used to guide navigation to different goals. Finally, obstacles can be accommodated in the same way as goals (see above and Fig. 7, right).

DISCUSSION

What Have We Learned?

The model successfully demonstrates a scheme for neural-network-based localization and navigation. More importantly, by

its adherence to known biological data it also affords a basis for systematically examining the role of individual hippocampal cells in terms of their effect on spatial behavior and suggests directed experiments which could contribute to a fuller understanding of the neural mechanisms involved.

We have proposed a model of how the hippocampus of the rat might be involved in spatial behavior; the model addresses observed behavioral data on rat navigation, including the speed with which the layout of an environment appears to be learned, even in the absence of “goals” (latent learning). Perhaps its major contribution is the demonstration that the place cell output of the hippocampus could be used to provide a useful spatial map. This was the first model to simulate both cells and the resulting behavior—answering the criticism that place cells could only tell you where you are, not where to go. Moreover, the use of computational modeling has forced us to begin to quantify our assumptions about the mechanisms underlying the spatial and temporal aspects of place cell firing, and this in turn has led to progress toward revealing these mechanisms experimentally. In the remainder of this section we look at the specific properties and predictions of the mechanisms employed in the model, and at what we have learned from each one.

Formation of place fields

A major assumption of the model was that the input to place cells is functionally equivalent to several broad tuning curves peaked at particular distances from particular cues. We have begun to test this experimentally by recording from place cells in a series of rectangular boxes, varying from each other only in their width or length or both (see O’Keefe and Burgess, 1996, and Fig. 8, left). The place fields appear to change shape and size in a systematic way under these conditions, consistent with each being formed by a thresholded sum of Gaussian tuning curves in the distance of the rat from each box wall. This has allowed us to begin to quantify this aspect of the model.

Specifically, if the four corners of a w cm by h cm box are at $(0, 0)$ ($w, 0$) $(0, h)$ and (w, h) then we can model a place cell’s firing rate map $f(x, y)$, where (x, y) is the position of the rat and f is the cell’s firing rate, as

$$f(x, y) \approx A \left(\frac{\exp [-(h-y) - d_N]^2 / 2\sigma^2(d_N)]}{\sqrt{2\pi\sigma^2(d_N)}} + \frac{\exp [-(y-d_S)^2 / 2\sigma^2(d_S)]}{\sqrt{2\pi\sigma^2(d_S)}} + \frac{\exp [-(w-x) - d_E]^2 / 2\sigma^2(d_E)]}{\sqrt{2\pi\sigma^2(d_E)}} + \frac{\exp [-(x-d_W)^2 / 2\sigma^2(d_W)]}{\sqrt{2\pi\sigma^2(d_W)}} \right) - T \quad (2)$$

where A gives the overall amplitude, T is a threshold, d_N is the value at which the tuning curve in the distance from the north wall peaks in cm, d_S , d_W , and d_E are defined similarly, and the width $\sigma(x)$ is an increasing function of x . We used $\sigma(x) =$

$\sigma_0(61^2 + x^2)/61^2$ so that if the rat perceives the distance x to a wall as, say, the angle from vertical to the top of the wall (61 cm high), then the width of the Gaussian is constant under this measure of distance ($\sigma_0 = 20$ cm). Notice that Gaussians peaked close to a wall will be higher and narrower than those peaked far from one, due to the variation in σ .

The model contains an arbitrary threshold value T , which we assume depends on local interactions with other place cells (e.g., lateral inhibition) and on global variables such as familiarity with the environment, circadian rhythms or satiation. Using a threshold proportional to the peak rate, the model can explain fields that peak at a fixed distance from a box wall, fields that appear to peak at a fixed proportion of the distance between two walls, or fields that become stretched or pulled apart as the separation of two opposite walls is varied. If an absolute threshold is used across all trials, it can explain fields that disappear at, above, or below critical values of the separation of opposing walls (see Fig. 8, right). All of these types of place fields were observed in the experiment (Fig. 8, left).

Notice that the model places no restriction on the number of tuning curve inputs to a place cell and does not address the means by which they are generated (e.g., path integration or external sensory stimulation). The place fields described above were modeled using three or four inputs; however, in the extreme case of only one sensory input, the cell would respond to that single sensory stimulus.

Phase coding

We demonstrated that the phase of firing with respect to the EEG theta rhythm of each individual place cell can shift systematically as the rat runs through the place field, even if the net firing rate of all of the place cells correlates with theta (see Burgess et al., 1993, Fig. 2d). Thus, the bursting frequency of each cell can be greater than the theta frequency (see O'Keefe and Recce, 1993, Fig. 9), even if the theta frequency is reflected in the oscillation in net firing rate. The simulations revealed that the modulation of the number of spikes per burst within the field (greatest at the field center), and the fact that the rat does not run right through all of the fields that it enters, result in the total number of spikes fired peaking at a fixed phase of the EEG. This is in accordance with the experimental observation of Fox et al. (1986). Second, we demonstrated one possible use of phase coding, which is coming to be seen as a viable complement to rate coding in neurons.

Was Computational Modeling Justified, or Would Verbal-Qualitative Reasoning Have Sufficed?

The uses of computational modeling can be divided into (at least) two types: 1) development of a conceptual model after consideration of the experimental data, and use of computational modeling to demonstrate that the conceptual model actually has the properties ascribed to it (which is often not the case!); 2) development of a computational model ad hoc, and examination of it in the hopes of finding unexpected analogues of the experimental data. Our model is of type 1; thus, given unlimited

powers of verbal reasoning it would be possible to predict its performance. However, it is a detailed model, so it would be hard to be sure that a purely conceptual analysis of its likely properties would be correct. Thus, proof that the model actually works as a navigational mechanism is valuable. Conversely we suspect that many of the purely conceptual models proposed for navigation may not be practicable. In this regard, implementation as a mobile robot would be an even better demonstration of real-world navigational ability than a computational simulation. It is also difficult to anticipate the effects of phase coding with any confidence by using conceptual argument alone. For example, it initially seems counterintuitive that each cell can have a higher bursting frequency than the frequency of the net activity of the population of cells. Finally, a mathematical description of place fields provides a framework for analysis of their shape and location which is not possible with verbal descriptions, or even information theory. We note that field shape remains problematically dependent on the bin size used in data collection, and that information theoretic approaches (of which we are aware) cannot address field shape since bins are treated independently of their spatial relationships.

Relation to Other Models

In this work we have restricted our consideration to neuronal models of spatial processing in the rat hippocampus. Although the hippocampus may have a wider function, we consider that models at the neuronal level, based on the least controversial experimental data, are the most likely to provide insight into hippocampal function with respect to higher levels of description and wider function (see The Goal of Modeling above). In this field, our model was the first simulation to incorporate both the firing of individual cells and actual navigation by the rat. A more extensive review of spatial models of the hippocampus can be found in Burgess et al. (1995).

O'Keefe's (1991) "centroid" model also has an external basis for phase coding (angle to a cue) and was the precursor to our models of place fields occurring at distances from several cues in particular allocentric directions (we extended the cue centroid to arbitrary combinations of cue distances). In terms of implementation, many very different connectionist models can produce cell behavior resembling place fields (see, e.g., Treves et al., 1992; Shapiro and Hetherington, 1993). We decided to build upon the type of mechanism used by Sharp (1991) for generating place fields because of its biological plausibility.

Hetherington and Shapiro (1993) propose a model in which the CA3 recurrent collaterals can be used to store the sequences of place cell activity corresponding to a rat moving along a route, even in the absence of any input from outside of region CA3. Levy (1996) uses a more realistic learning algorithm to propose a model of CA3 that can also recall sequences of activation (and in some cases cut out those sets of activation corresponding to a loop in the trajectory). However, even in an experiment that rules out external sensory cues, any mechanism totally internal to CA3 still ignores the interoceptive/path integration systems probably in use. Muller et al. (1991b) and Blum and Abbott (1995) also propose models

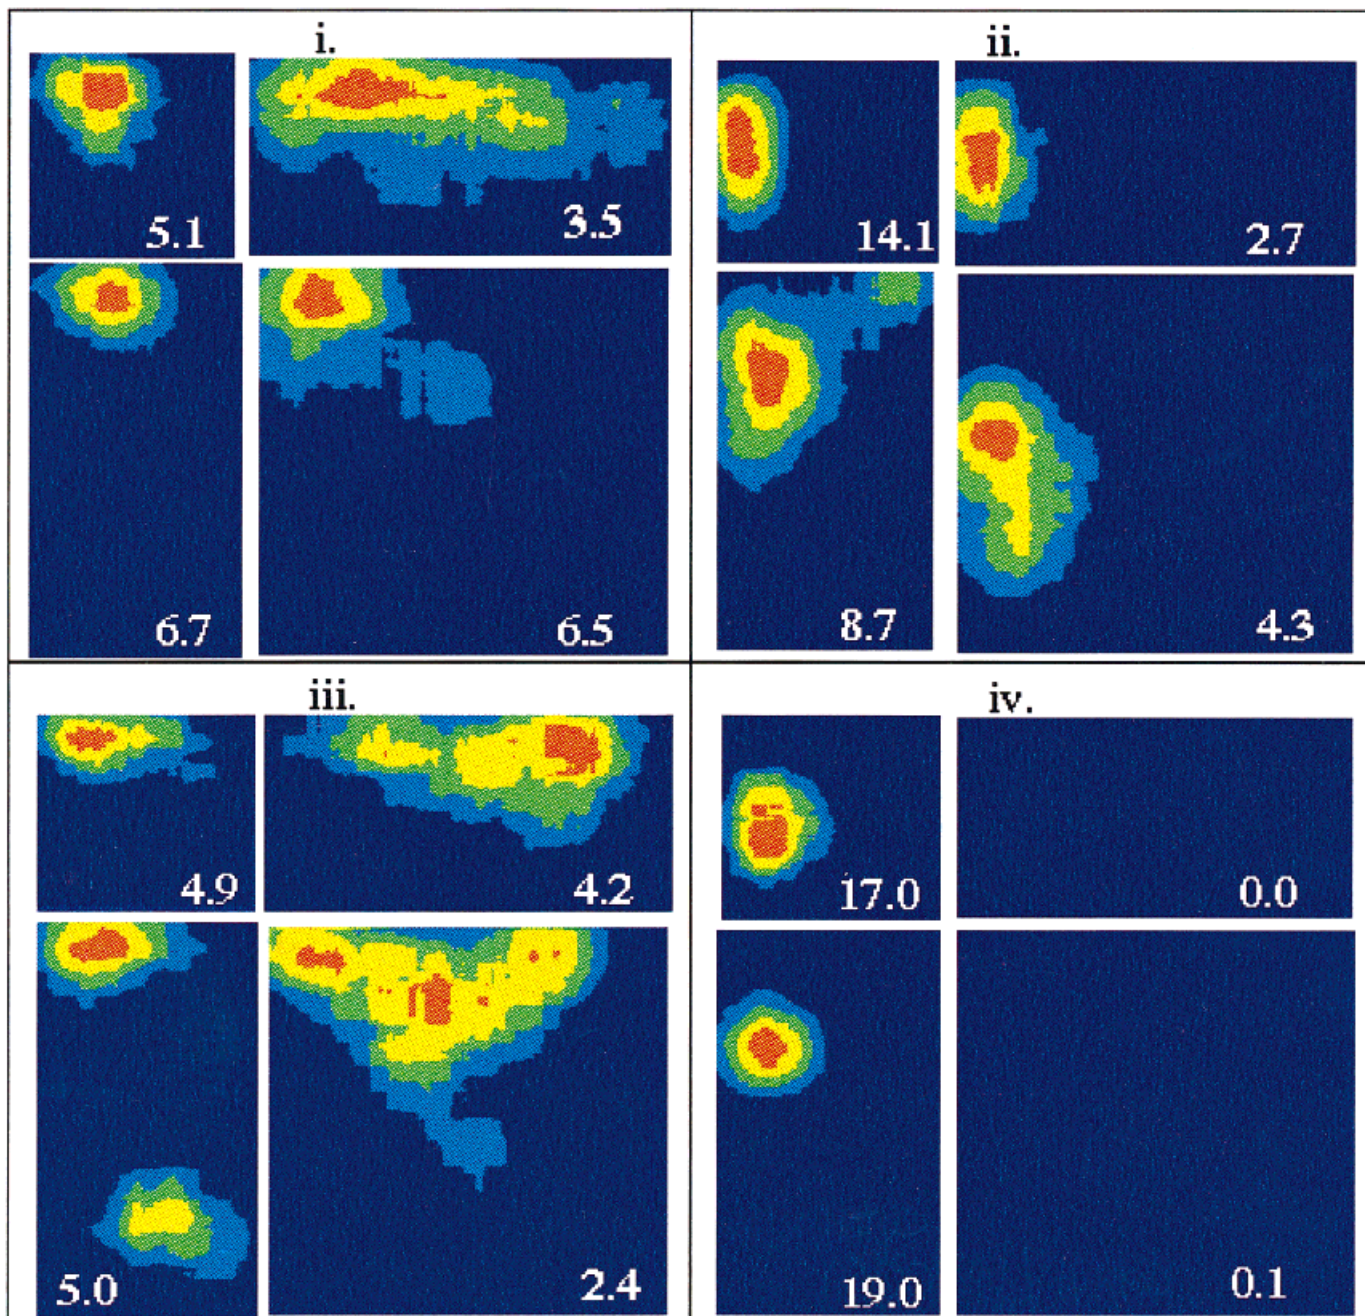


FIGURE 8. A: Place fields and environmental geometry. Firing rate maps are shown for three CA1 place cells (i, iii, and iv) and one CA3 place cell (ii), each recorded in a small square box, a large square box, and rectangular boxes oriented along or across the testing room. Boxes were constructed from four identical grey planks (61 cm high by 122 cm long) fastened at the top by clamps. The planks and the paper floor covering were frequently changed to ensure that local olfactory and tactile cues made no consistent contribution to cell firing. The small square measured 61 × 61 cm; the large square

122 × 122 cm; and the two rectangles 61 × 122 cm. A 128 × 128 grid was placed on the viewing area of the camera, and square bins $\frac{1}{24}$ of the area of a box were placed around each grid point. The number of spikes fired by each cell and the time spent by the rat in each bin were recorded. Firing rate maps show the firing rate at each grid point (number of spikes divided by dwell time in the associated bin) as a color plot of five equally spaced intervals (from deep blue up to red), with linear interpolation between grid points. The peak rate is given in white with each plot; fields with peak rate < 1.0 Hz are not shown.

which put the CA3 recurrent collaterals to ingenious use. In all of these models it is not yet clear how the spatial information in the CA3 recurrent collaterals could plausibly get out of the hippocampus to affect behavior.

McNaughton et al. (1995) suggest that each place cell fires when the rat is at a given distance and allocentric direction from a particular cue. This fits in with previous models (Zipser, 1985; Sharp, 1991; Burgess et al., 1994) but specifically stresses the role

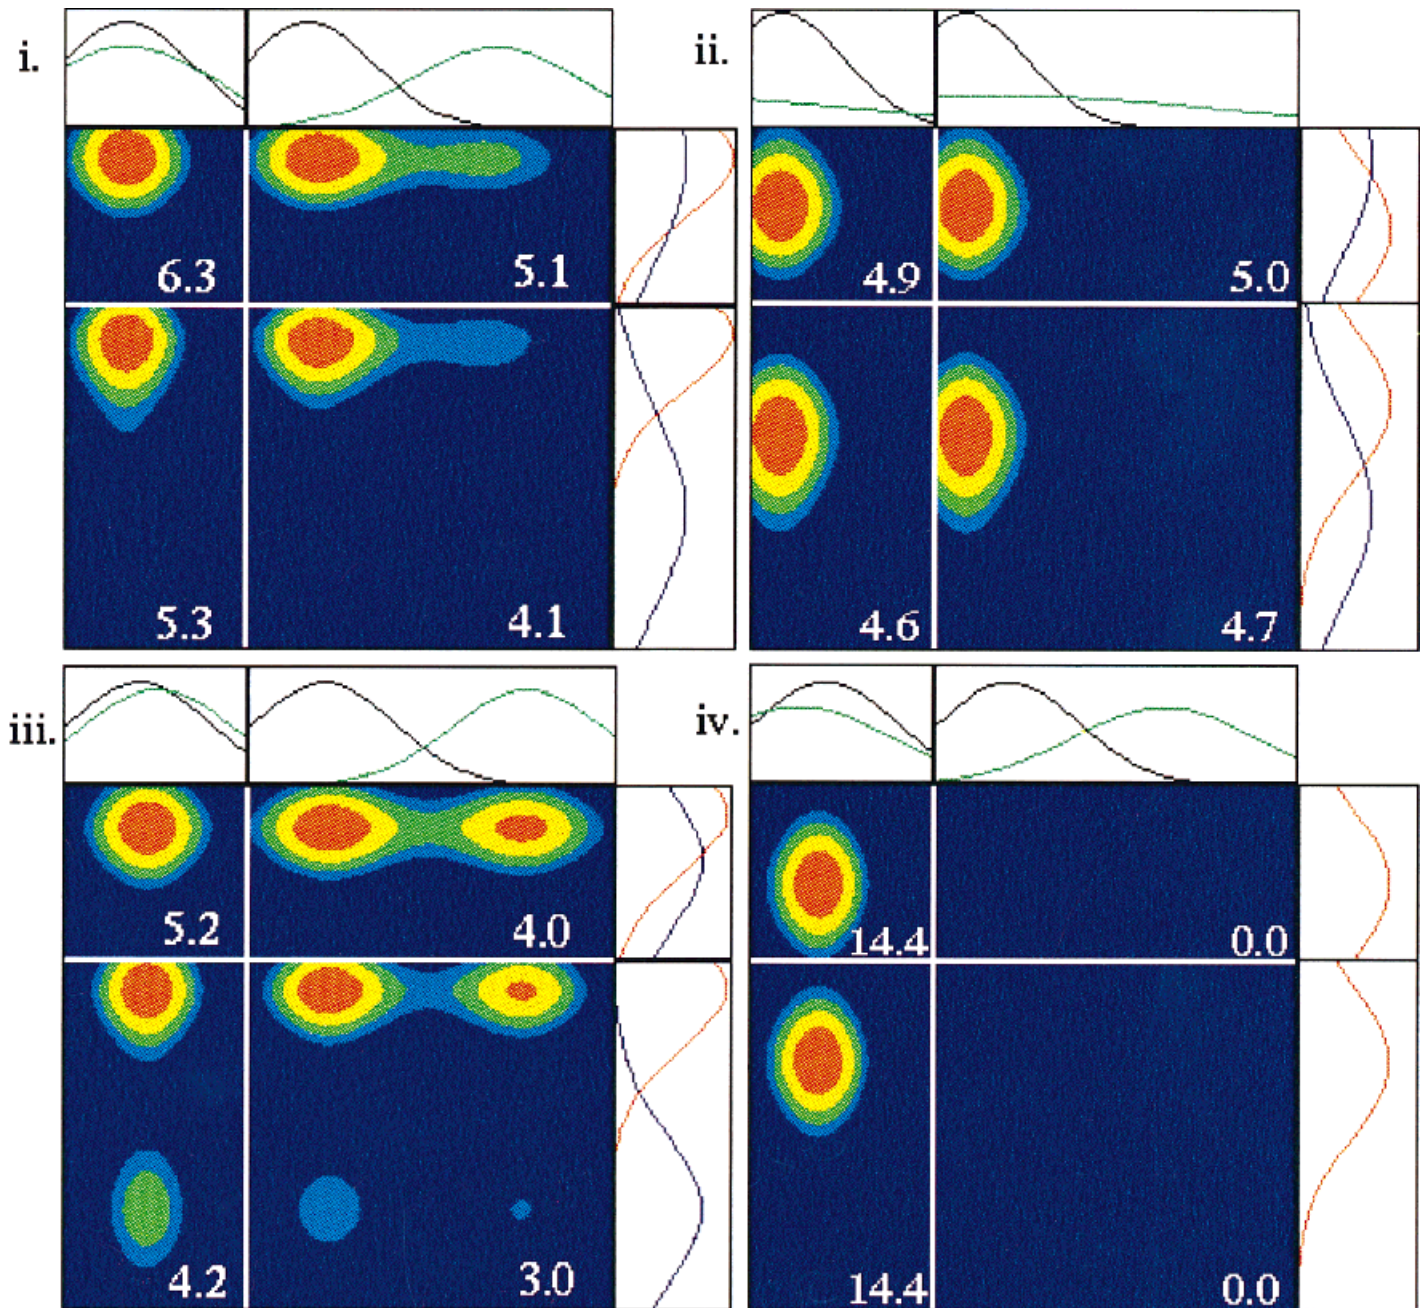


FIGURE 8. B: Simulated firing rate maps according to equation 2. The borders show the shape of the tuning curves peaked at fixed distances from the box walls. i, ii, and iii have threshold T equal to 80% of the peak rate; iv has $T = 75$. Other parameter values were as follows: (i): $A = 500$, $d_N = 10$, $d_S = 50$, $d_E = 40$, $d_W = 20$; (ii): $A = 500$, $d_N = 35$, $d_S = 50$, $d_E = 100$, $d_W = 10$; (iii) $A = 400$, $d_N = 10$, $d_S = 35$, $d_E = 30$, $d_W = 25$; (iv) $A = 2,000$, $d_N = 35$, $d_S = 1,000$, $d_E = 45$, $d_W = 25$. Adapted from O'Keefe and Burgess (1996).

of the vestibular system in the underlying mechanism. McNaughton et al. (1996) return to the idea that place fields are preconfigured to form a distinct and fixed spatial reference frame that can be rapidly attached to a particular environment and whose metric is defined by the animal's movements (i.e., path integration; see also O'Keefe and Nadel, 1978). We agree that

ideothetic information provides an important input (e.g., contributing to the measurement of the distance to a recently visited cue), and that place fields appear in a novel environment too rapidly for the system to be completely dependent on learning. We found that movement of the walls of the environment changed the distances between place fields, but not that of the fields to the

walls (see O'Keefe and Burgess, 1996). While not ruling out the possibility that there is a topological relationship between place fields that survives environmental manipulations, these experiments demonstrate that the absolute distances between them are not preserved. These results can be interpreted in terms of our model, in which the "seeding" of place field locations by the connections initially set to 1 implies a preconfigured relationship between place fields and environmental stimuli rather than among the place fields themselves. By contrast, the McNaughton et al. (1996) model would be required to posit separate reference frames (and origins for path integration) fixed to each wall.

Redish and Touretzky (1996) propose a symbolic implementation of systems-level theories of rodent navigation which provides a useful test-bed, without corresponding directly to neuronal mechanisms. Recce and Harris (1996) also propose a symbolic implementation of a model of hippocampal function. In their model (sonar) local views of the environment are stored in "cortex," with each view associated to a particular "place cell." Place cells then fire according to the match between the current local view and the stored view to which each is associated. Active place cells serve to enable "pattern completion" of the current local view by reactivating the relevant stored views (i.e., providing an autoassociative memory function for cortex). When a goal is encountered, its location is automatically marked on all of the stored local views (even if it can only be sensed at the goal location). Thus pattern completion of a local view makes the goal appear, enabling navigation toward it. While not providing a realistic neuronal model of navigation, the model can address the role of the hippocampus in exploration and in more general memory-related functions.

With similar motivation to ourselves, Brown and Sharp (1995; see also Sharp et al., 1996) also propose a connectionist model of the possible use of place cell firing in navigation. The main difference with our work is in their output representation in which egocentric directions (body-turns) are associated to places over many runs to a goal. This produces testable differences in behavior: their model would not show latent learning, performance would not be affected by whether the rat can look around from the goal location or not (cf. Arolfo et al., 1994), and navigation to the goal would be strongly affected if stereotyped routes were used during learning (e.g., a clockwise spiral into the goal). These predictions also apply to Blum and Abbott (1995), Muller et al. (1991b), and Hetherington and Shapiro (1993).

Predictions and New Directions for Research

The model generated a number of testable predictions:

1. That there are pyramidal cells with large place fields in subiculum, for which there is preliminary evidence (Sharp and Green, 1994). Large fields have also been found in ventral hippocampus (Jung et al., 1994).
2. That there exist goal cells in or downstream from the subiculum with firing fields covering the entire environment, but peaked near to the goal. It is not known at present if subicular cells in real rats serve this purpose (but see Barnes et al., 1990; Muller et al., 1991a,b). We note recent reports of hippocampal cells that

fire near to a goal location (which itself moves around within the environment; see Gothard et al., 1996), which may bear some relation to the predicted goal cells.

3. That the phase relationship between entorhinal cell firing and the ongoing hippocampal EEG θ rhythm is governed by the relative positions of the cues to which they respond.

4. That if the rat is prevented from looking around at a goal location, it will be impaired in trying to return to the goal (preliminary evidence supporting this prediction is provided by Arolfo et al., 1994). Conversely, if the rat can look around at the goal location, it is not important which direction(s) it approached the goal from during learning.

5. That the phase coding of place cell firing in open fields would resemble that in Figure 4B. Evidence supporting this is given in Burgess et al. (1994) and Skaggs et al. (1996); see Fig. 4C.

6. That place fields are formed from the conjunction of several tuning curve inputs, each peaked at a fixed distance from an environmental cue. This model has been confirmed and extended: Specifically the environmental boundary itself appears to be used to determine field locations. By contrast, objects at or beyond the boundary appear to be used to provide the directional orientation for the system (which may also be maintained by ideothetic inputs), e.g., enabling identification of boundary walls by their allocentric direction (see *What Have We Learned?* above). Note that the locations of interesting objects within the environment are stored by the "goal cells" but do not affect the model's place fields.

The model also raises many questions that remain unanswered:

1. What provides the distance measure? We have alluded to sensory inputs to place cells and have started to determine their functional form. However, it is not clear exactly what constitutes these inputs, except that the distance from the walls of the environment seems to be a particularly powerful factor. Experiments in which place fields were maintained after spatial cues were removed or after the lights were turned off may indicate a role for an internal metric, e.g., provided by footfalls or vestibular inputs.
2. How are inputs selected by a place cell (this is arbitrary in the present model, with bias for those combinations of inputs providing a field near the centroid of the cues)? Is this selection hard wired (predetermined) or plastic, and why does it usually result in a unitary place field?
3. How could phase coding be generated? Note that the correlation with location rather than elapsed time implies that information concerning the rat's velocity enters the system somehow, e.g., by affecting the interoceptive or exteroceptive input signals.
4. Is there phase coding in the entorhinal cortex, and is the interaction between sensory 10-Hz oscillations and motor-related hippocampal theta a possible source for phase coding? Are there differentiable sensory and motor aspects of theta?
5. Can the model be applied directly to humans (note recent evidence of human hippocampal theta rhythm; Tesche, 1995)? Or should it be applied to other inputs as a model of storing trajectories through memory for events, i.e., storing episodes?

Acknowledgments

We would like to thank Kate Jeffery and Jim Donnett for carefully reading this manuscript. This research was supported by an MRC program grant. N.B. is a Royal Society University Research Fellow.

REFERENCES

- Aguirre GK, Detre JA, D'Esposito M, Alsup D, Rosenquist AC (1995) Maze navigation activates the hippocampal formation during learning and recall: a fMRI study. *Soc Neurosci Abstr* 21(116.2):272.
- Arolfo MP, Nerad L, Schenk F, Bures J (1994) Absence of snapshot memory of the target view interferes with place learning by rats in the water maze. *Behav Neurosci* 108:308–316.
- Barnes CA, McNaughton BL, Mizumori SJY, Leonard BW, Lin L-H (1990) Comparison of spatial and temporal characteristics of neuronal activity in sequential stages of hippocampal processing. *Prog Brain Res* 83:287–300.
- Benhamou S, Seguinot V (1995) How to find ones way in the labyrinth of path integration models. *J Theor Biol* 174:463–466.
- Blum KI, Abbott LF (1995) Functional significance of long term potentiation in recurrent networks. *Int J Neural Syst [Suppl]* 6:25–31.
- Brown MF (1992) Does a cognitive map guide choice in the radial-arm maze? *J Exp Psychol* 18:56–66.
- Brown MA, Sharp PE (1995) Simulation of spatial learning in the Morris water maze by a neural network model of the hippocampal formation and nucleus accumbens. *Hippocampus* 5:171–188.
- Burgess N, O'Keefe J (1995) Modelling spatial navigation by the rat hippocampus. *Int J Neural Syst [Suppl]* 6:87–94.
- Burgess N, O'Keefe J, Recce M (1993) Using hippocampal 'place cells' for navigation, exploiting phase coding. In: *Advances in neural information processing systems 5* (Hanson SJ, Giles CL, Cowan JD, eds), pp 929–936. San Mateo, CA: Morgan Kaufmann.
- Burgess N, Recce M, O'Keefe J (1994) A model of hippocampal function. *Neural Networks* 7:1065–1081.
- Burgess N, Recce M, O'Keefe J (1995) Hippocampus—spatial models. In: *The handbook of brain theory and neural networks* (Arbib MA, ed), pp 468–472. Cambridge, MA: MIT Press.
- Burgess N, Donnett JG, O'Keefe J (1996) Modelling navigation by the rat hippocampus. *Forma* 11:83–101.
- Chen LL, Lin LH, Green EJ, Barnes CA, McNaughton BL (1994) Head-direction cells in the rat posterior cortex. 1. Anatomical distribution and behavioral modulation. *Exp Brain Res* 101:8–23.
- Eichenbaum H, Otto T, Cohen NJ (1994) 2 functional components of the hippocampal memory system. *Behav Brain Sci* 17:449–472.
- Feigenbaum JD, Polkey CE, Morris RG (1996) Deficits in spatial working-memory after unilateral temporal lobectomy in man. *Neuropsychologia* 34:163–176.
- Fox SE, Wolfson S, Ranck JB (1986) Hippocampal theta-rhythm and the firing of neurons in walking and urethane anesthetized rats. *Exp Brain Res* 62:495–508.
- Georgopoulos AP, Kettner RE, Schwartz AB (1988) Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. *J Neurosci* 8:2928–2937.
- Gothard KM, Skaggs WE, McNaughton BL (1996) Binding of hippocampal CA1 neural activity to multiple reference frame in a landmark-based navigation task. *J Neurosci* 16:823–835.
- Harley CW (1979) Arm choices in a sunburst maze: effects of hippocampal resection in the rat. *Physiol Behav* 23:283–290.
- Hetherington PA, Shapiro ML (1993) A simple network model simulates hippocampal place fields: II. Computing goal-directed trajectories and memory fields. *Behav Neurosci* 107:434–443.
- Hill AJ (1978) First occurrence of hippocampal spatial firing in a new environment. *Expl Neurol* 62:282–297.
- Jarrard LE (1993) On the role of the hippocampus in learning and memory in the rat. *Behav Neural Biol* 60:9–26.
- Jung MW, Wiener SI, McNaughton BL (1994) Comparison of spatial firing characteristics of units in dorsal and ventral hippocampus of the rat. *J Neurosci* 14:7347–7356.
- Levy WB (1996) A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal problems. *Hippocampus* 6:579–590.
- Macrides F (1975) Temporal relationship between hippocampal slow waves and exploratory sniffing in hamsters. *Behav Biol* 14:295–308.
- Maguire EA, Burke T, Phillips J, Staunton H (1996) Topographical disorientation following unilateral temporal lobe lesions in humans. *Neuropsychologia* 34:993–1001.
- McNaughton BL, Barnes CA, O'Keefe J (1983) The contribution of position, direction and velocity to single unit activity in the hippocampus of freely-moving rats. *Exp Brain Res* 52:41–49.
- McNaughton BL, Knierim JJ, Wilson MA (1995) Vector encoding and the vestibular foundations of spatial cognition: a neurophysiological and computational hypothesis. In: *The cognitive neurosciences* (Gazzaniga M, ed), pp 585–595. Boston: MIT Press.
- McNaughton BL, Barnes CA, Gerrard JL, Gothard KM, Jung MW, Knierim JJ, Kudrimoti H, Qin Y, Skaggs WE, Suster M, Weaver KL (1996) Deciphering the hippocampal polyglot—the hippocampus as a path integration system. *J Exp Biol* 199:173–185.
- Mizumori SJY, Williams JD (1993) Directionally sensitive mnemonic properties of neurons in the lateral dorsal nucleus of the thalamus of rats. *J Neurosci* 13:4015–4028.
- Morris RGM (1981) Spatial localisation does not depend on the presence of local cues. *Learning Motivation* 12:239–260.
- Morris RGM, Garrard P, Rawlins JNP, O'Keefe J (1982) Place navigation impaired in rats with hippocampal lesions. *Nature* 297:681–683.
- Muller RU, Kubie JL (1987) The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J Neurosci* 7:1951–1968.
- Muller RU, Kubie JL, Ranck JB (1987) Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *J Neurosci* 7:1935–1950.
- Muller RU, Kubie JL, Bostock EM, Taube JS, Quirk GJ (1991a) Spatial firing correlates of neurons in the hippocampal formation of freely moving rats. In: *Brain and space* (Paillard J, ed), pp 296–333. UK: Oxford University Press.
- Muller RU, Kubie JL, Saypoff R (1991b) The hippocampus as a cognitive graph. *Hippocampus* 1:243–246.
- Muller RU, Bostock E, Taube JS, Kubie JL (1994) On the directional firing properties of hippocampal place cells. *J Neurosci* 14:7235–7251.
- Murray EA, Davidson M, Gaffan D, Olton DS, Suomi S (1989) Effects of fornix transection and cingulate cortical ablation on spatial memory in rhesus monkeys. *Exp Brain Res* 74:173–186.
- Nicolelis MA, Chapin JK (1994) Spatiotemporal structure of somatosensory responses of many-neuron ensembles in the rat ventral posterior medial nucleus of the thalamus. *J Neurosci* 14:3511–3532.
- O'Keefe J (1976) Place units in the hippocampus of the freely moving rat. *Exp Neurol* 51:78–109.
- O'Keefe J (1991) The hippocampal cognitive map and navigational strategies. In: *Brain and space* (Paillard J, ed), pp 273–295. Oxford: Oxford University Press.
- O'Keefe J, Burgess N (1996) Geometric determinants of the place fields of hippocampal neurons. *Nature* 381:425–428.
- O'Keefe J, Conway DH (1978) Hippocampal place units in the freely moving rat: why they fire where they fire. *Exp Brain Res* 31:573–590.

- O'Keefe J, Nadel L (1978) *The hippocampus as a cognitive map*. Oxford: Oxford University Press.
- O'Keefe J, Recce M (1993) Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3:317–330.
- O'Keefe J, Speakman A (1987) Single unit activity in the rat hippocampus during a spatial memory task. *Exp Brain Res* 68:1–27.
- O'Keefe J, Nadel L, Keightley S, Kill D (1975) Fornix lesions selectively abolish place learning in the rat. *Exp Neurol* 48:152–166.
- Olton DS, Walker JA, Gage H (1978) Hippocampal connections and spatial discrimination. *Brain Res* 139:295–308.
- Parkinson JK, Murray EA, Mishkin M (1988) A selective mnemonic role for the hippocampus in monkeys: memory for the location of objects. *J Neurosci* 8:4159–4167.
- Pavrides C, Greenstein YJ, Grudman M, Winson J (1988) Long-term potentiation in the dentate gyrus is induced preferentially on the positive phase of θ -rhythm. *Brain Res* 439:383–387.
- Poucet B, Thinus-Blanc C, Muller RU (1994) Place cells in the ventral hippocampus of rats. *Neuroreport* 5:2045–2048.
- Quirk GJ, Muller RU, Kubie JL, Ranck JB (1992) The positional firing properties of medial entorhinal neurons: description and comparison with hippocampal place cells. *J Neurosci* 12:1945–1963.
- Recce M, Harris KD (1996) Memory for places: a navigational model in support of Marr's theory of hippocampal function. *Hippocampus* 6:735–748.
- Redish AD, Touretzky DS (1996) Navigating with landmarks: computing goal locations from place codes. In: *Symbolic visual learning* (Ikeuchi K, Veloso M, eds), Oxford: Oxford University Press, in press.
- Rumelhart DE, Zipser D (1986) Feature discovery by competitive learning. In: *Parallel distributed processing 1* (Rumelhart DE, McClelland JL, eds), pp 151–193. Cambridge, MA: MIT Press.
- Rumelhart DE, Hinton GE, Williams RJ (1986) Learning internal representations by error propagation. In: *Parallel distributed processing 1* (Rumelhart DE, McClelland JL, eds), pp 318–362. Cambridge, MA: MIT Press.
- Shapiro ML, Hetherington PA (1993) A simple network model simulates hippocampal place fields: I. parametric analyses and physiological predictions. *Behav Neurosci* 107:34–50.
- Sharp PE (1991) Computer simulation of hippocampal place cells. *Psychobiology* 19:103–115.
- Sharp PE, Green C (1994) Spatial correlates of firing patterns of single cells in the subiculum of the freely moving rat. *J Neurosci* 14:2339–2356.
- Sharp PE, Blair HT, Brown MA (1996) Neural network modeling of the hippocampal formation spatial signals, and their possible role in navigation: a modular approach. *Hippocampus* 6:720–734.
- Skaggs WE, McNaughton BL, Wilson MA, Barnes CA (1996) Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* 6:149–172.
- Smith ML, Leonard G, Crane J, Milner B (1995) The effects of frontal-lobe or temporal-lobe lesions—on susceptibility to interference in spatial memory. *Neuropsychologia* 33:275–285.
- Squire LR (1992) Declarative and nondeclarative memory: multiple brain systems supporting learning and memory. *J Cog Neurosci* 4:232–243.
- Sutherland RJ, Chew GL, Baker JC, Linggard RC (1987) Some limitations on the use of distal cues in place navigation by rats. *Psychobiology* 15:48–57.
- Taube JS (1995) Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. *J Neurosci* 15:70–86.
- Taube JS, Muller RU, Ranck JB (1990) Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *J Neurosci* 10:420–435.
- Tesche CD (1995) Noninvasive detection of theta activity in human hippocampal formation with a whole-head magnetoencephalographic array. *Soc Neurosci Abstr* 21(117.11):276.
- Tolman EC (1948) Cognitive maps in rats and men. *Psychol Rev* 55:189–208.
- Treves A, Miglino O, Parisi D (1992) Rats, nets, maps, and the emergence of place cells. *Psychobiology* 20:1–8.
- Whishaw IQ (1991) Latent learning in a swimming pool place task by rats—evidence for the use of associative and not cognitive mapping processes. *Q J Exp Psychol B* 43:83–103.
- Wilson MA, McNaughton BL (1993) Dynamics of the hippocampal ensemble code for space. *Science* 261:1055–1058.
- Zipser D (1985) A computational model of hippocampal place fields. *Behav Neurosci* 99:1006–1018.