



A Model of Hippocampal Function

NEIL BURGESS, MICHAEL RECCE, AND JOHN O'KEEFE

Department of Anatomy, University College London

(Received and accepted 27 January 1994)

Abstract—The firing rate maps of hippocampal place cells recorded in a freely moving rat are viewed as a set of approximate radial basis functions over the (2-D) environment of the rat. It is proposed that these firing fields are constructed during exploration from “sensory inputs” (tuning curve responses to the distance of cues from the rat) and used by cells downstream to construct firing rate maps that approximate any desired surface over the environment. It is shown that, when a rat moves freely in an open field, the phase of firing of a place cell (with respect to the EEG θ rhythm) contains information as to the relative position of its firing field from the rat. A model of hippocampal function is presented in which the firing rate maps of cells downstream of the hippocampus provide a “population vector” encoding the instantaneous direction of the rat from a previously encountered reward site, enabling navigation to it. A neuronal simulation, involving reinforcement only at the goal location, provides good agreement with single cell recording from the hippocampal region, and can navigate to reward sites in open fields using sensory input from environmental cues. The system requires only brief exploration, performs latent learning, and can return to a goal location after encountering it only once.

Keywords—Hippocampus, Navigation, Place cell, Rat, Reinforcement, Population vector, Phase coding, EEG.

1. INTRODUCTION

The hippocampus has been the focus of a vast, and growing, amount of study. It lies deep in the brain, many synapses removed from sensory transducers or motor effectors, and has a relatively simple structure, containing one projection cell type confined to a single layer. Interest stems from the putative role of hippocampal damage in human amnesia, the discovery of long-term potentiation (the most widely studied model of synaptic plasticity; Bliss & Lømo, 1972), and the discovery that cells are spatially coded in rats (O'Keefe & Dostrovsky, 1971; Taube, Muller, & Ranck, 1990) and in monkeys (Ono, Nakamura, Fukuda, & Tamura, 1991). Lesions studies show clear deficits in the performance on spatial tasks of rats (Morris, Garrud, Rawlins, & O'Keefe, 1982; O'Keefe, Nadel, Keightley, & Kill, 1975; Olton, Walker, & Gage, 1978; see O'Keefe

& Nadel, 1978) and monkeys (Parkinson, Murray, & Mishkin, 1988; Murray, Davidson, Gaffan, Olton, & Suomi, 1989). The popularly held view that the hippocampus is a more general store for declarative memories is undergoing revision because much of the deficit in short-term memory tasks appears to be accounted for by lesions of the perirhinal cortex rather than the hippocampus (see Suzuki, Zola-Morgan, Squire, & Amaral, 1993).

The rat hippocampus contains roughly 700,000 pyramidal cells in regions CA3 and CA1. Extracellular recordings of pyramidal cells in freely moving rats show that most are “place cells” (O'Keefe, 1979), that is, they fire only when the rat is in a particular portion of its environment (O'Keefe & Dostrovsky, 1971; O'Keefe & Nadel, 1978). This paper considers the function of cells in the hippocampal formation in terms of their receptive fields and asks how place cell firing fields (“place fields”) are constructed and what general computational function they could serve. More specifically we suggest how the place fields could be used for navigation, bearing in mind that rats typically show rapid learning of suboptimal trajectories, and are capable of “latent learning” (i.e., learning during exploration in the absence of goals or motivation) and of finding shortcuts (see Tolman, 1948). Finally we present a neuronal simulation of rat navigation that takes sensory

Acknowledgements: We would like to acknowledge useful discussions with Andrew Speakman, Kate Jeffery, and Jim Donnett. N.B. was supported by a Joint Councils Initiative on Human Computer Interaction research fellowship, and as a Royal Society University Research Fellow.

Requests for reprints should be sent to Neil Burgess, Dept. of Anatomy, University College, London WC1E 6BT, UK; E-mail: N.Burgess@ucl.ac.uk

information as its input, and in which the spatial correlates of cell firing mimic those found in the various regions of the hippocampus, thus relating behavioural and electrophysiological data.

Related findings that we shall refer to are the existence of "head-direction" cells (whose firing rate has a tuning curve response to the direction of the rat's head; Taube et al., 1990) in regions neighbouring the rat hippocampus, and the observation of "population vectors" in (monkey) motor cortex coding for the direction of reaching (each cell has a "preferred" reaching direction for which it fires maximally: the vector sum of the preferred directions of each cell weighted by their firing rates is the population vector; Georgopoulos, Kettner, & Schwartz, 1988). We pay particular attention to the possible computational uses of the phase of firing of place cells (O'Keefe & Recce, 1993). Some of the questions related to producing a navigational output from place cells were examined in Burgess, O'Keefe, and Recce (1993) and an early version of the model described in Section 4 is outlined in Burgess, O'Keefe, and Recce (1994).

2. A SCHEMATIC MODEL OF HIPPOCAMPAL FUNCTION

2.1. Properties of Place Fields and a Functional Interpretation

1. Region CA1 as a whole provides many place fields, which densely cover the environment. In any one environment there are tens of thousands of active place cells; an upper limit of 12% for the percentage of all place cells that are active in one environment is implied by the data in Thompson and Best (1990).
2. Each place field is restricted to a portion of the environment and tends to have a characteristic single-peaked shape, postulated by O'Keefe and Nadel (1978). This becomes particularly clear with recent advances in extracellular recording, using a "tetrode" of four channels (each channel simultaneously records the neuronal spike, and the amplitude on each channel varies as a function of the distance between the source and the electrode tip; four distance measurements is the minimum geometrical requirement to localise a source in three-dimensional space; see Recce & O'Keefe, 1989). Prior methods have found that over 50% of the cells are active in more than one part of the environment (e.g., Pavlides & Winson, 1989), and that some cells have as many as five active regions (Wiener, Paul, & Eichenbaum, 1989). In contrast, our data show that less than 5% of the cells recorded have multiple place fields, and that the place fields are single peaked, smooth functions (see Recce, Speakman, & O'Keefe, 1991).
3. Cells are controlled by sensory cues from different

modalities and rotation of cues in a restricted environment causes the fields to rotate (O'Keefe & Conway, 1978; Muller, Kubie, & Ranck, 1987). It is believed that the sensory inputs come from the entorhinal cortex, which has access to multimodal sensory information (see e.g., Amaral & Witter, 1989).

4. Stable place fields are built up rapidly on entering a new environment, but not instantaneously (Wilson & McNaughton, 1993).
5. In open field environments the place fields are non-directional, whereas on more structured mazes (e.g., mazes of narrow arms) they are directional (i.e., the cell may fire only when the rat is running in a particular direction, see McNaughton, Barnes, & O'Keefe, 1983; Bostock, Taube, & Muller, 1988).
6. A small percentage of place fields tend to be "edge fields," that is, the firing field follows a section of the edge of the environment (see Muller et al., 1987).

The firing rate maps of four place cells are shown in Figure 1.

A simple model for the formation of place fields is to suppose that the sensory information incident on the hippocampus is coded by cells whose response is a tuning curve function (see e.g., Ballard, 1986) of the distances of cues from the rat (see Figure 2; see also McNaughton, Knierim, & Wilson, 1993; Wan, Touretzky, & Redish, 1993). Thus a cell receiving projections from two such sensory cells might effectively multiply the firing rate on each input (see Bugmann, 1991), producing an approximate radial basis function place field. This is true if there is good angular separation between the two cues to which the sensory cells respond (i.e., the two cues subtend approximately 90° at all points in the place field); place fields would become less radially symmetric for two cues close together, and be more likely to have double fields for two cues on opposite sides of the environment.

In this paper we suggest that place cell firing rate maps could be used as a set of radial basis functions $\{\phi_i(\|\vec{x} - \vec{c}_i\|)\}_{i=1,N}$, where \vec{x} is the position of the rat, \vec{c}_i is the centre of the i th place field and $\phi_i(r) = \exp(-r^2/2\sigma_i^2)$. An arbitrary function $f(\vec{x})$ could then be approximated by

$$\sum_{i=1}^N \lambda_i \phi_i(\|\vec{x} - \vec{c}_i\|),$$

for example, see Powell (1985). Thus an artificial neuron with a linear transfer function and a connection of weight λ_i from each place cell i would have a firing rate map over the environment approximating $f(\vec{x})$.

2.2. Phase Coding

Evidence suggests that the computational mechanisms at work in the hippocampus include a temporal aspect.




FIGURE 1. Firing rate maps of place fields from 4 CA1 cells recorded simultaneously with a tetrode (Recce & O'Keefe, 1989) in a freely moving rat. The firing rate map is constructed by laying a 32×32 grid of square bins on the environment, calculating the number of spikes fired divided by the occupancy time for each bin, smoothing (a 2-D low pass filter of the spatial frequencies in the resulting histogram), and linearly interpolating between each grid point. Each contour represents 10% of the peak firing rate, peak firing rates vary between 4.8 and 22.8 Hz.

The hippocampal EEG exhibits a sinusoidal oscillation of 5–12 Hz (the “ θ rhythm”), that in the rat is associated with particular behaviours, such as walking, running, or swimming, but not grooming, eating, or remaining stationary (Vanderwolf, 1969). Lesions to the medial septum (removing the θ rhythm pacemaker input) destroy both the θ rhythm and the rat's ability to solve spatial tasks (Winson, 1978).

Place cell firing has a systematic phase relationship to the local EEG (O'Keefe & Recce, 1993): when a rat

on a linear track runs through its place field, the cell fires groups of spikes, with each successive group occurring at an earlier phase of the θ cycle (see Figure 3). Another way of looking at the phase shift in place cell firing is to realise that those cells whose place field the rat is entering will fire “late” in a θ cycle, whereas those whose place field the rat has traversed will fire earlier in the cycle. Because place fields cover the entire environment, and place cells fire at all phases of a θ cycle, it is useful to consider place fields as a function




FIGURE 2. Contour map of a place field constructed from sensory input consisting of tuning curve responses to the distance from the rat of two cues.

of the phase at which the associated place cell fired. During the approximately 100 ms of a single θ cycle, the set of active place cells changes progressively from those with fields centred behind the rat to those centred ahead of the rat.

Using data from freely moving rats in open field environments, Figure 4 shows that the mean phase of firing of a place cell within a θ cycle contains the information as to whether the centre of the corresponding place field is ahead of the rat (late phase) or behind it (early phase), as suggested in Burgess et al. (1993). Notice that taking the (linear) mean phase of firing over a θ cycle always classifies messy firing (i.e., spikes spread over an entire cycle, rather than clumped to-

gether in one group) into the "middle" phase. Finally, assigning three phases (early, middle, and late) to the spikes in each θ cycle requires only modest temporal discrimination (≈ 15 ms) of a system downstream of the place cells.

This phase coding may enable the system to separate out the firing of place cells whose place field lies predominantly ahead of the rat from the firing of those whose place field lies predominantly behind the rat (see Figure 4). The available data show a peak in the mean phase of firing of place cells, averaged over many cells, coinciding with the peak of the θ rhythm. This has been interpreted as a preferred phase of firing of individual cells, although in general these data do not refer to the high firing rate occurring in the place field (Fox, Wolfson, & Ranck, 1986), and are not inconsistent with the phase of each individual cell shifting as described above (see Burgess et al., 1993, figure 2d).

3. NAVIGATION

Lesion studies clearly implicate the rat hippocampus in navigation tasks that cannot be solved by following a well used route or approaching a single cue (see Morris et al., 1982). Place cells have long been thought to provide a "map" to enable navigation (O'Keefe & Nadel, 1978). The question remains as to how information in the place cell map could be read out in a useful way. In this section we explore how radial basis function place fields could be used to endow a set of cells downstream of the hippocampus with firing rate maps that support navigation. We show how the output of the hippocampal navigation system could be a set of "goal" cells providing a population vector coding for the instantaneous direction of the rat from interesting locations encountered in an environment. We consider one environment in which there is a single goal location to which the rat must return (e.g., the escape platform in a water maze or food location in an open field). Figure




FIGURE 3. Typical pattern of place cell firing relative to the EEG θ rhythm as a rat runs through the firing field on a linear track: one 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 & Recce, 1993).




FIGURE 4. The position of the place field relative to the rat, as a function of the phase of firing of the place cell. Place cell firing and EEG were recorded from freely moving rats chasing pellets in an open field. The EEG was divided into θ cycles (see Figure 3C; and O'Keefe & Recce, 1993). Each spike from a place cell was assigned a phase equal to the mean phase of all the spikes fired by the cell in the same θ cycle. At the moment that each place cell fired, the distance to the peak of its firing rate map was projected onto the rat's direction of movement (positive meaning ahead of the rat). The x axis represents this projected distance in centimeters with the rat's head at 0. Shown are the histograms for those spikes fired at early, middle, or late phases, the bin width was 2.66 cm. Spikes fired late in a θ cycle tend to have place fields peaked ahead of the rat, whereas those fired early have fields peaked behind the rat. 4,932 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 or above 20 Hz), 132 spikes fall outside the domain of the histogram, the remaining 4,488 spikes are shown.

5 shows the architecture of our model. Henceforth entorhinal cells will be referred to as ECs, place cells as PCs, and subicular cells as SCs.

We postulate a population of goal cells one synapse downstream of the subiculum and one synapse upstream of motor areas that receive a strong "reinforcement" signal (as opposed to "supervised learning") whenever a particular type of reward is encountered (e.g., the nucleus accumbens, layer IV of EC, or lateral septum). Learning takes place in the connections from SCs, at the goal location, as described below. The firing rate of each goal cell is interpreted by the motor system to be the (probabilistic) weight that the rat is a particular direction from the goal (e.g., north, southwest, etc). We refer to the goal cell representing the rat being north of the goal as GC_N , similarly GC_S represents the

rat being south of the goal, etc. The population vector is the vector sum of the direction represented by each goal cell, weighted by its firing rate. For this vector to approximate the direction of the rat from the goal, GC_N must fire more than GC_S whenever the rat is north of the goal, and similarly for the goal cells representing "east" and "west," etc. In terms of the firing rate maps of the set of goal cells, the firing rate map of GC_N must be higher than GC_S s north of the goal and lower south of the goal, and similarly for other directions.

There are many different possible configurations of goal cell firing rate maps consistent with producing a population vector for navigation; the one we choose here (see also Burgess et al., 1993) is for GC_N 's firing rate map to be a cone, positive over the entire environment and peaked to the north of the goal. Similarly




FIGURE 5. Functional 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.

GC_S's firing rate map is a cone of equal height, peaked an equal distance south of the goal (see Figure 6). Any number of goal cells (>2) could be used, given that the directions represented have an even distribution over the set of possible directions. A bonus of this choice, if the cones' peaks are fairly close to the goal, is that the net firing rate of all the goal cells increases with the proximity of the goal, and could be used (by the motor system) to estimate the distance to the goal as well as the rat's direction from it (see Burgess et al., 1993, figure 6c).




FIGURE 6. One possible set of firing rate maps giving rise to a population vector representing the position of the rat from the goal at G.

3.1. Subicular Representation

If we suppose that information about the goal location is stored when the rat is at the goal (triggered by a "reinforcement" signal), then there is a problem of locality in the information available from the set of PCs. There is little or no overlap between the pattern of PC activity at two points on opposite sides of the environment. Thus, using only the PCs, there would be no information available far from the goal with which to determine the direction of travel required to return to the goal. We postulate that a solution is provided by SCs, building up large firing fields from the superposition of many place fields during exploration (the limited data available from extracellular recording in subiculum are not clear (see Barnes, McNaughton, Mizumori, Leonard, & Lin, 1990; Muller, Kubie, Bostock, Taube, & Quirk, 1991)).

3.2. Population Vector From Reinforcement at Goal, Using Phase Coding

Our model for the creation of a population vector for navigation is as follows. Each goal cell receives a strong reinforcement signal whenever the rat encounters the goal (food reward, escape platform, etc.); however we assume further that the reinforcement signal to GC_S is

gated by the firing of a head-direction cell tuned to fire maximally when the head direction is north; similarly reinforcement to GC_S is gated by a south head-direction cell, and so on. Thus a big excitatory input is received by a goal cell representing a particular direction only when the rat is at the goal site and facing in that direction. We suppose that a connection from a subicular cell to a goal cell is "switched on" only if the goal cell is receiving the reinforcement signal and the subicular cell is active (see Section 4.4). Thus the firing rate map of a goal cell will look like a superposition of the firing rate maps of all subicular cells that are active when connection modification occurs.

Now, the phase of firing of place cells and (hence) subicular cells in our model is such that those cells active late in a θ cycle have firing fields peaked *ahead* of the rat (cf. Figure 4). Thus if the reinforcement signal arrives at a late phase of each θ cycle, the firing rate map of, say, GC_N will be peaked to the north of the goal location, the firing rate map of GC_S will be peaked to the south of the goal location, and so on. When a rat encounters an interesting location it looks around in many directions, the connections to goal cells representing many directions are modified, and the population vector of the goal cells is sufficient to enable navigation (provided the rat looked in two or noncolinear directions when at the goal location). This is the functional use made of the phase coding of cell firing in this model.

Notice that different sets of goal cells could be used to code for the instantaneous direction of all interesting places encountered by the rat during exploration. Which set is used to control navigation depends on which place the rat wants to revisit. Thus, in the map view of the hippocampal system (O'Keefe & Nadel, 1978) the firing rates of place cells contain information as to the location of the rat, which drives the firing of goal cells that contain the information as to where interesting places are (i.e., the rat's direction and distance from them), allowing navigation.

4. A NEURONAL SIMULATION OF HIPPOCAMPUS

To simplify the model, our initial working hypothesis is that different subsets of place cells are active in different environments (see Wilson & McNaughton, 1993), and that the functional role of the dentate gyrus (DG) and the CA3 recurrent collaterals (rcs) is to produce and recover the different subsets as necessary. We further assume that the role of CA1 is to produce a cleaner and more robust copy of the place cell firing in CA3. We consider only one environment and simulate an active subset of place cells for that environment, ignoring the DG, the CA3 rcs, and merging fields CA3 and CA1. Within these assumptions, we have a functional model of the hippocampus.

We have simulated the movement of a rat in an environment containing a set of discrete sensory cues. A set of "sensory" cells fires with broad tuning curve responses to the distance of each cue. These stimuli could be visual, olfactory, or auditory. The input from them varies as a function of the rat's position and propagates through three "hidden" layers corresponding to ECs, PCs, and SCs, to the output layer, which is a population of 8 "goal" cells, see Figure 5. The simulated rat moves at a constant speed of 60 cm/s around a 150×150 cm box (including a 15 cm border in which to place extra-maze cues). It would be unusual for a real rat to maintain this speed for more than a few seconds, but allowing the rat to stop and start would make modelling much more complex. The environment is represented on a 500×500 pixel array, so the maximum spatial resolution is 0.3 cm.

The θ rhythm of the EEG is taken to be 10 Hz, and the basic timestep of the model is $1/30$ s, that is, 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 only two behavioural states. During "exploration" the rat's next direction is simply a random variable within 30° of its previous direction. During "searching" its direction is determined by the goal cells, although it has some momentum: its next direction is the average of its previous direction and that indicated by the goal cells. Cue and rat positions are considered to be single points in simulations. The rat automatically "encounters" any goal (or obstacle, see later) within 10 cm of its position and looks around in eight directions. As it moves around the environment it bounces off the edge of the environment or any obstacles that it collides with.

4.1. Sensory Input

Up to 16 cues are placed in and around the rat's environment. For each cue there is a set 15 "sensory" cells, each with a broad tuning curve response peaked at different distances from that cue. The distribution of distances at which each cell responds maximally covers the environment uniformly. Rather than using graded responses, each sensory cell i (where $i = 0-14$) responds by firing a number of spikes $n_i = 0, 1, 2,$ or 3 . The tuning curve responses of the i th cell, responding to a cue x centimeters away, is approximated by:

$$n_i = \begin{cases} 0 & \text{for } |x - iL| \geq 6L, \\ [(7L - |x - iL|)/2L] & \text{otherwise.} \end{cases} \quad (1)$$

(a triangular bump peaked at cue-distance iL , with a base of $12L$), where we have discretised distance in steps of $L = 14$ cm, and $[y]$ is the integer part of y . The receptive field of the i th cell is an annulus around the cue, peaked at a radius of iL centimeters. In most




FIGURE 7. Escape latencies for blocks of trials (see Figure 10): the effect of the number of cues in the environment. The first block of trials after encountering the goal, following 30 s exploration, is shown on a log scale from 1 to 10 s. For the extra-maze condition (□) there was one cue in each corner plus a further 0, 1, 2, or 3 evenly spaced cues on each side of the environment. For the extra + intra-maze condition (△), there was one cue in each corner and one in the centre (5), and a 3 × 3 (9) and a 4 × 4 (16) grid of cues. $\langle C_m \rangle = 1$.

simulations the majority of cues are arranged around the edge of the environment (see Figure 7).

4.2. Entorhinal Cells and Phase of Firing

The transformation of broad tuning curve input to well-localized place fields occurs via the intermediate layer of entorhinal cells. Each EC receives connections from two sensory cells (i and j) and fires $[n_i n_j / 2]$ spikes. We simulate a maximum of 1000 ECs: insufficient to simulate an EC receiving an input from every possible pair of sensory cells. Instead, following the idea of cells responding at the centroid of subsets of cues (O'Keefe, 1991), we only simulate those ECs that receive input from pairs of sensory cells that respond to two distinct cues, and have tuning curve responses centred at the $m + 1$ distances iL from a cue where:

$$\left[\frac{s}{2L} - \frac{m}{2} \right] \leq i \leq \left[\frac{s}{2L} + \frac{m}{2} \right],$$

where s is the separation of the two cues, and $[y]$ is the integer part of y . That is, input from sensory cells with peak responses near to the half-separation $s/2$ of the cue pair are selected (notice that these tuning curves will tend to overlap near to the two cues' centroid). The total number of ECs is controlled by m which is varied heuristically with the number of cues c and their

separation according to $m = 160/[c(1 + 6s/212)]$ to ensure roughly equal numbers of ECs as c varies, and a roughly even distribution of EC receptive field centres over the environment (i.e., avoiding a concentration of fields at the centre of the environment). The most common result of the EC setup is a large spatial firing field (see Figure 8), consistent with extracellular recording (Quirk, Muller, Kubie, & Ranck, 1992).

For the simulation to reflect our proposed mechanism for navigation, the phase of firing of place cells must have the property described in Section 2.2, that is, cells firing at a late phase tend to have firing fields centred ahead of the rat, those firing early have fields centred behind the rat. We must also specify the phase of firing (w.r.t. the θ rhythm) of entorhinal cells in response to their sensory input (as with place cells, there are data showing a peak in the histogram of the phase of firing (Stewart, Quirk, Barry, & Fox, 1992), but this is not necessarily inconsistent with phase coding). The rule we have adopted is to relate the phase of firing of an EC to the average angle α , from the rat's heading direction, of the two cues that provide the sensory input to the EC: if α is ahead of the rat the EC fires at a late phase, if α is behind the EC fires early, and if α is to the side of the rat the EC fires at a middle phase (see Figure 9).

The above rule means that those ECs that fire late in a θ cycle are responding to sensory input from pairs of cues whose centroid is *ahead* of the rat. The peak firing rate of an EC will also occur near to the centroid of the two cues providing its input, where the receptive fields of its two sensory inputs overlap most strongly (see above). Hence the net firing rate map for all ECs active late in a θ cycle is peaked *ahead* of the rat. Thus, we have built in a tendency for the phase of firing of cells throughout the model (as the activation of ECs feeds forward to the other layers in the model) to code for whether the cell's firing field is predominantly ahead of or behind the rat. Notice that we have made a strong prediction as to the phase of firing of ECs to produce the phase shift of place cell firing in our model; however, it is equally possible that the real cause does not involve ECs but stems from an intrinsic property of PCs.

4.3. Learning in Place and Subicular Cells

The dynamics of the place and subicular cells, and of the connections to them, are described below. Both feedforward and feedback inhibition are presumed to occur in the hippocampus (see e.g., Schwartzkroin, Scharfman, & Sloviter, 1990), and to model this we use an adapted form of competitive learning (Rumelhart & Zipser, 1986). Differences from competitive learning are that connection weights are binary, cells are characterised by the number of spikes fired per time step ($\frac{1}{30}$ s), and the input to each cell is divided by an adaptive threshold that increases proportionally to the




FIGURE 8. Typical firing rate maps of cells in the different layers of the model after 60 s exploration, showing qualitative agreement with known extracellular recordings from entorhinal cells and place cells, and predicting firing rate maps for subicular cells and goal cells. The simulated rat moves evenly across the environment; spikes were binned in a 10×10 grid, each contour represents 10% of the peak firing rate. Top row: entorhinal cell, peak rate is 40 Hz (left), place cell peak rate is 30 Hz, (right). Bottom row: subicular cell, peak rate is 40 Hz (left), goal cell representing east (the goal is at the centre), peak rate 101 Hz (right).

number of active inputs to the cell (see Appendix A).

Connection weights take value 0 (off) or 1 (on). The majority of connections are initialised to 0, but there is a small probability that any given connection is on, so that on average each cell initially receives $\langle C_{in} \rangle$ on connections. We assume that $\langle C_{in} \rangle$ is at least 1.0, and the size of $\langle C_{in} \rangle$ reflects the amount of interference from previously learned environments, that is, if $C_{in} > 1$ for a given PC, that PC will initially have more than one peak: transmitting ambiguous spatial information.

This turns out to be an important parameter of the model (see Figures 7 and 10). The net input to a cell in a given timestep is the sum of the spikes fired by all cells from which it receives an on connection (its “active inputs”).

A connection is switched on whenever the pre- and postconnection cells both fire the maximum number of spikes (4) in the same time step. The net input to each cell is divided by $C_{sh}(1 + m)$, where m is the number of times a connection to the cell is switched




FIGURE 9. The phase, within each θ cycle, at which a model entorhinal cell fires a group of spikes is determined by the average angular position α of the two cues (marked by +) providing its sensory input. The phase of firing is late if $|\alpha| < 60^\circ$, middle if $60^\circ < |\alpha| < 120^\circ$, and early if $|\alpha| > 120^\circ$. The number of spikes fired depends on the strengths of the two sensory inputs.

on and C_{sh} is a constant (i.e., there is some shunt inhibition close to the cell that increases whenever connection modification occurs). Thus the input to a cell with $C_{in} = 1$ is proportional to the average number of spikes fired by its active inputs. This avoids the global normalisation of connection weights required by standard competitive learning (Rumelhart & Zipser, 1986).

To model feedback inhibition PCs and SCs are arranged in groups: the cells with the largest inputs in each group are active, the others remain silent. At each time step the four cells with the largest input fire a number of spikes equal to the size of their input, provided that it is not more than a maximum number that depends on its rank (see Appendix A): the cell with the largest input can fire a maximum of four spikes, the one with the second largest can fire a maximum of three spikes, the third largest two spikes and so on. However, we assume that there is more severe inhibition among the PCs than the SCs (to encourage larger firing fields), hence PCs are arranged in five groups of 50, whereas SCs are arranged in 10 groups of 25; and $C_{sh} = 1$ for a PC, whereas $C_{sh} = \frac{1}{3}$ for an SC.

Notice that the learning within each cell is self-limiting: (a) once a connection has been switched on it cannot be further altered, and (b) introduction to an environment is accompanied by frequent connection modification, which increases the shunt term in the cell's input and reduces firing rates, also reducing the frequency of connection modification. The firing fields of PCs and SCs arise as follows. Suppose that the first

time that connections to a typical PC or SC are switched on occurs at \bar{x} . Then the only place where the PC's input is strong enough for more connections to be modified (i.e., where the average number of spikes fired by its active inputs is at least three) is near to \bar{x} . Hence place fields tend to be peaked at a position that is "seeded" by one of the connections that are initially on.

Place fields are relatively small because, to fire at all, the PC must overcome the competition from other PCs. By contrast, the input to an SC can be strong enough for more connections to be modified in a wider area around \bar{x} (as its input is equal to twice the average number of spikes fired by its active inputs). The firing field of an SC is larger because each competes with fewer cells. We have not attempted to model the long-term decrement of synaptic connections because there is not yet a consensus as to the biological data.

4.4. Goal Cells and Navigation

Whenever the rat encounters a goal location (escape platform, a place with food, etc.) the goal cell corresponding to its head direction receives a reinforcement signal at the late phase of θ . At the goal the rat looks around in eight directions (north, northeast, east, etc.) for one θ cycle each. Connections to a goal cell from




FIGURE 10. Escape latencies: average time to the goal from eight starting positions on the edge of an environment. Each point represents the mean over five different goal positions (one block of trials; one goal is in the centre, and one is in the centre of each quadrant), each run was limited to a maximum of 10 s. The effect of performing successive blocks of trials after the goal was encountered (no exploration). The duration of each trial is $5 \times 8 \times y$, where y is the average escape latency shown for that block of trials. There were 16 evenly spaced extra-maze cues. $\langle C_{in} \rangle$ was 1 (Δ) or 5 (\square).

an SC are switched on whenever the goal cell receives a reinforcement signal and the SC fires one or more spikes in the same time step (the fact that SC does not have to fire the maximum number of spikes for connection modification to occur results in goal cell firing fields big enough to cover the entire environment, see Section 3.2, and Figure 8). A goal cell's input is also divided by the number of on connections it receives, and each cell fires a number of spikes equal to 25 times its input (relatively high firing rates reduce rounding effects in calculating the population vector from integer numbers of spikes). For efficient navigation, the contribution of each goal cell to the population vector should also be normalised by its firing rate the last time the rat was at the goal, otherwise the population vector can be distorted by variations in the density of SC receptive fields (however, this is not crucial: without normalisation the mean escape latency after 30 s exploration with 16 extra-maze cues is 1.96 ± 0.57 s as opposed to 1.47 ± 0.15 s with normalisation, see Figure 11).

All goal locations are distinguishable: a different population of eight goal cells codes for each one. Thus there can be many goals in one environment: which population of goal cells is used to guide navigation (or which receives the reinforcement signal) depends on which goal the rat is navigating to (or which goal the rat encountered). "Obstacles," in the sense of locations to be avoided, can be coded for in the same way as




FIGURE 11. Escape latencies (see Figure 10): the effect of time spent exploring the environment, before encountering the goal (latent learning). The first block of trials after encountering the goal is shown. There were 16 evenly spaced extra-maze cues. $\langle C_{in} \rangle = 1$.

goals. Our crude model of navigation is limited to approaching the goal directly, if there are no obstacles, or subtracting the population vector of cells coding for an obstacle that is "in the way" of the population vector for the goal. An obstacle is in the way if the rat thinks it is within 45° of the goal direction (judged by comparing population vectors), and nearer than the goal (judged by comparing the net firing rate of the two sets of eight cells).

5. PERFORMANCE

We have restricted our model to navigation in a single environment, and have ignored a large part of the hippocampal system, which we propose has the function of storing distinguishable representations of different environments (see Section 4). Within this restriction the model has been successful in terms of roughly reproducing realistic firing properties of the cells representing entorhinal cells, place cells, and (as far as we know) subicular cells. It has also achieved realistic navigation (i.e., rapid learning of suboptimal, but successful, trajectories), and does not rely on biologically implausible learning rules involving thousands of iterations to learn the position of one goal location.

5.1. Cell Firing Fields

In the model, place cells have relatively small place fields and receive input from many entorhinal cells, becoming virtually independent of the position of any one cue. In contrast ECs receive input from only two cues and have much larger firing fields (extracellular recording shows ECs to be more "sensory bound" than PCs; Quirk et al., 1992). SCs, having larger fields that depend on the trajectories taken during exploration, are even less dependent on individual cue positions (see Figure 8). PC firing fields are less well localised than real ones (see Figure 1) and are modulated by the rat's direction (which affects the phase at which PCs and SCs fire, and therefore affects which other cells they compete with). We have not yet fully analysed this directionality. Because goals are not used as cues, the environmental representation is not affected by changing the position of the goal.

Functionally speaking, small place fields could be required to reduce the interference between stored place cell representations for different environments, that is, using a sparse coding scheme in which different "active subsets" of cells code for different environments. This picture of hippocampal processing is consistent with that in Barnes et al., 1990.

5.2. Navigation

Of the order of 30 s (i.e., 300 movements of θ cycles) to 1 min of exploration, in the absence of a goal, is

sufficient to enable navigation back to the goal after encountering it once (see Figures 11, 12). Thus the model is capable of latent learning, as are rats (Tolman, 1948). Figures 7, 10, and 11 characterise navigational performance in terms of the average time taken to reach the goal location ("escape latency") from eight starting positions around the edge of the environment. All escape latencies shown are averaged over five different goal positions (one in the centre of the environment and one in the centre of each quadrant). Each set of

runs to five goal positions, each from eight starting positions, is referred to as a block of trials. Each run to goal is limited to 10 s, that is, not finding the goal is scored as 10 s. Lack of reliability in finding the goal is reflected in the large variance about mean escape latencies. The minimum possible average escape latency is 1.226 s, the escape latency from random movement (as in exploration) is 7.38 ± 0.93 s.

Navigation improves with time spent in the environment (whether exploring or searching) (see Figures

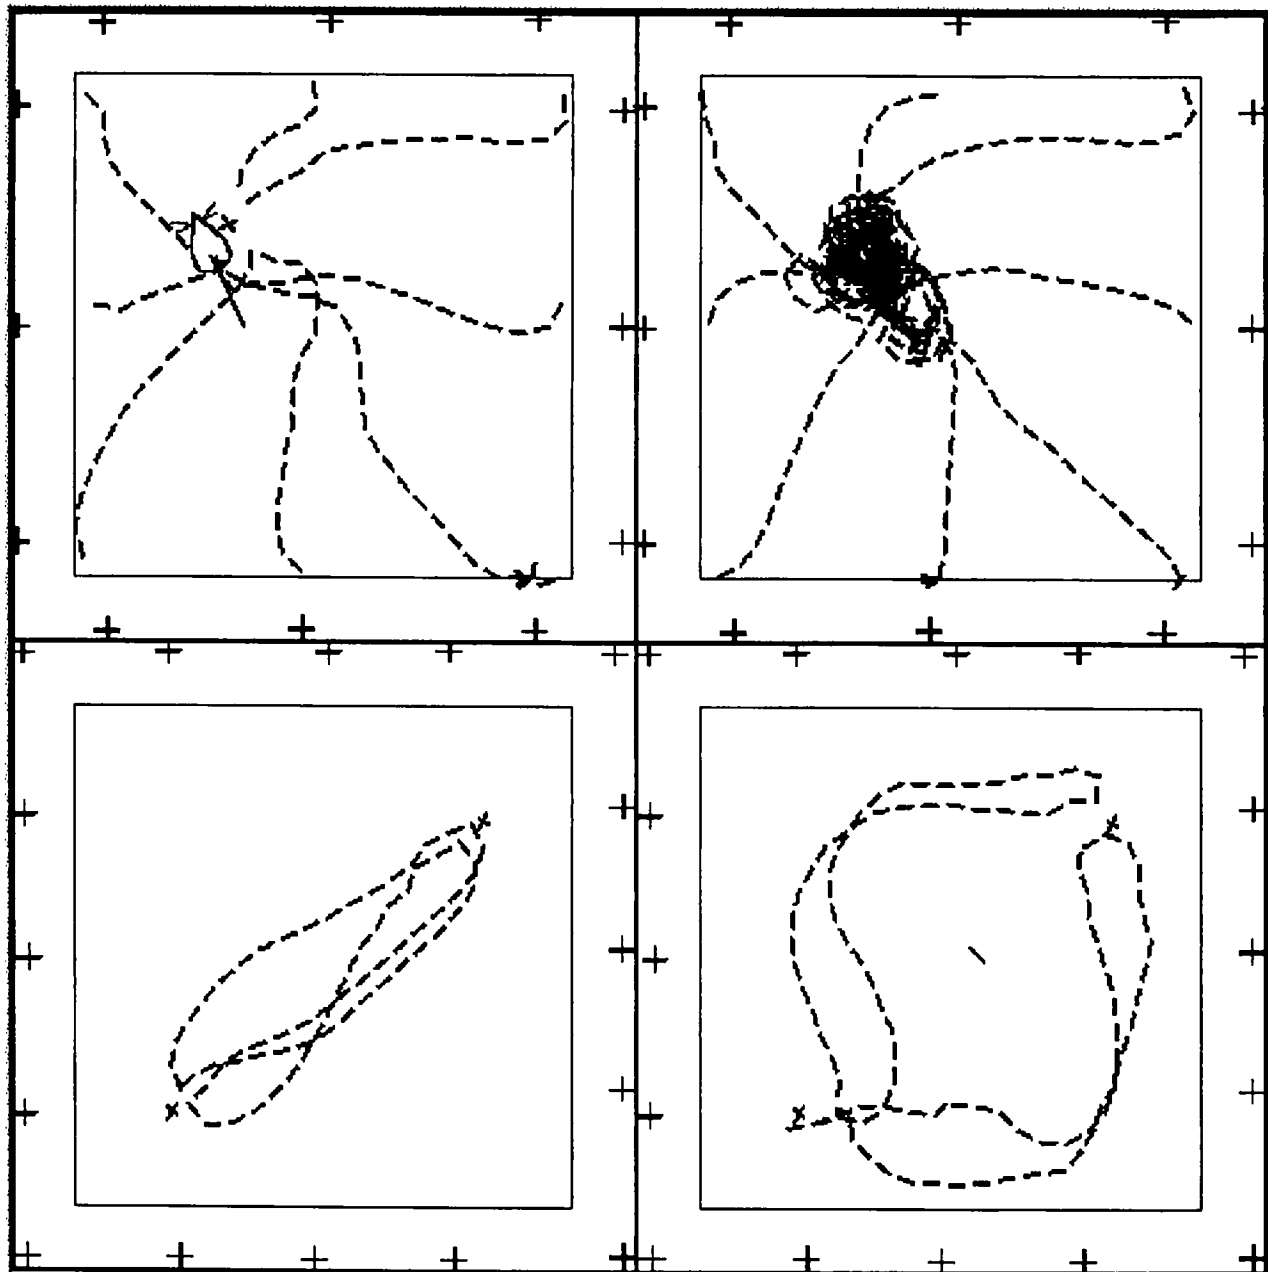


FIGURE 12. Trajectories taken by the simulated rat. Goals (marked by x) were encountered after 30 s exploration in the goalless environment. Each dash represents one movement (0.1 s), cues are marked by +. The average initial number of on connections to a PC or SC, $\langle C_{in} \rangle$, was 1. Top: navigating to the goal from eight starting positions on the first block of trials after encountering the goal, average escape latency is 1.4 s (left), navigation with the goal removed, showing localisation of search (right). Bottom: navigation between two goals (left), and after an obstacle (marked by \) is encountered (right).

10, 11). During the brief exploration necessary for navigation the rat does not cover the entire environment, but is subsequently capable of navigating over stretches of novel territory (notice that rats are capable of taking short cuts; Tolman, 1948).

The simulated rat navigates successfully with cues distributed evenly across the entire environment, or with cues only outside its perimeter (extra-maze cues). Navigation becomes poorer as the total number of cues decreases, and usually fails if there are four or less (see Figure 7). Although the escape latency is very long for four cues, the rat searched in the correct quadrant every time, but too close to the environment's edge. Performance deteriorates when there are uneven numbers of cues on each side of the environment (e.g., average escape latency for cues in each corner plus 0, 1, 1, and 2 extra cues on each side is 3.38 ± 2.79 s; cf. 1.65 ± 0.37 s for eight evenly spread extra-maze cues). A consequence of the sensory input-EC model: the goal location must be inside the convex hull of the set of cues for successful navigation.

Performance also depends on the time that the goal is encountered during exploration: if the goal is encountered before 30 s of exploration the average escape latency on the first block of trials is 1.78 ± 0.23 s compared to 1.47 ± 0.15 s if it is encountered after the exploration ($\langle C_{in} \rangle = 1.0$; 16 extra-maze cues were present). Notice that a hungry rat in a novel environment will often ignore food the first few times it is encountered, in favour of exploration; that is, prior exploration is very important to a rat, although solely being placed on the escape platform of a water-maze does aid navigation (Keith & McVety, 1988).

5.3. Discussion

Searching and exploration make no difference to the dynamics of the neuronal network (there are not two separate learning and recall phases): connection modification continues during trials. Learning at PCs and SCs occurs during movement in the environment, irrespective of the presence of a goal or whether the rat is exploring or searching. Because of this it is hard to separate effects of pure exploration and of running to the goal: a block of trials itself contains a total time of 40 times the average escape latency of moving through the environment (see Figure 10), so that substantial learning occurs within one block of trials.

When a goal location is encountered once, the one-shot modification of connections to goal cells is enough to enable subsequent navigation to the goal, provided that the rat looks around in many directions when at the goal site, and has spent enough time in the environment (see Figures 10, 11). By recruiting different sets of goal cells to code for the direction of different objects encountered in an environment, multiple goals within an environment can be handled (see

Figure 12). Similarly a limited number of obstacles can be avoided during navigation to a goal (see Section 4.4 and Figure 12).

We have not yet considered navigation in a restricted environment such as a radial arm maze or a linear track. The simulations of Sharp (1991) indicate that competitive learning can produce directional or non-directional place fields depending on whether the exploration of the rat is restricted or not. An alternative, less simple, supposition is that subsets of place cells fire as a function of position relative to a single reference point (or even a single cue; McNaughton et al., 1993) and that an attentional system makes different reference points active as a function of direction of movement (Wan et al., 1993). The presence of edge fields would be accommodated in our model by allowing sections of the environment's edge to form extended cues, rather than the point-like cues in the current model.

In the present model, connections are only ever switched on, and a new set of goal cells must be used whenever a new goal is encountered (we also made the simplifying assumption that a completely different subset of place cells is used for each new environment). Thus after learning many or crowded environments, the system could saturate, running out of available cells or connections. Whether we should use this to predict a finite capacity for learning, or wait for a consensus on the form of long-term synaptic depression to model, is an open question.

We interpreted values of $\langle C_{in} \rangle$ greater than 1.0 (for which PCs initially have more than one peak) as resulting from interference from other learned environments. The higher $\langle C_{in} \rangle$, the more learning is required before reasonable navigation can be performed (see Figure 10). It is interesting to note that, for the case of $\langle C_{in} \rangle = 1.0$ performance with very little exploration is so good (see Figure 11). Given $C_{in} = 1$ for all PCs and SCs, and a perfectly selected set of sensory input (so that EC receptive fields cover the environment uniformly) it would be possible for perfect navigation to occur using connection modification only at the goal cells. In fact in our model, with $C_{in} = 1$ for all PCs and SCs and no learning at PCs or SCs, the mean escape latency is still 2.62 ± 0.75 s (with 16 extra-maze cues), which, although not good, is much better than random search. This underlines the importance of (a) the choice of sensory input representation, which we have done by hand in this model, and (b) our supposed function of DG and the CA3 recurrent collaterals in avoiding interference between different environmental representations.

5.4. Comparison With Other Models; Supervised, Unsupervised, and Reinforcement Learning

The unsupervised competitive learning used builds stable firing fields very quickly compared to other neural

network simulations of place field construction, such as the competitive learning of Rumelhart and Zipser (1986) used by Sharp (1991), or error back propagation (BP; Rumelhart, Hinton, & Williams, 1986) used by Shapiro and Hetherington (1993). It also does not rely on a "teaching input" telling the network where the rat is during learning (necessary for BP), which begs the question of why the rat should build up a hippocampal representation of space if another part of its brain already knows its location at all points in space. Muller et al. (1991) suggest that the synaptic strengths of CA3 recurrent collaterals could encode the distances between place fields after much exploration, although how this information could be used in any computation is not clear.

Hetherington and Shapiro (1993) propose a model of navigation in which the rat is moved to a goal from all possible starting positions 1000 times, and learning is by BP (given teaching input of the rat's location at each step). Recurrent connections learn to reproduce the sequences of locations from a given start position to a given goal position, and can do so in the absence of any sensory, inertial, or proprioceptive input whatsoever. The one-shot learning of the population vector (goal cell) representation, using a reinforcement signal only whenever a goal is encountered during exploration (or search), seems more biologically plausible and learns faster, although our model does require sensory input to navigate.

Temporal difference reinforcement learning schemes for navigation (see Barto and Sutton, 1981; Dayan, 1991) can produce optimal trajectories, even in cluttered environments, but require almost exhaustive exploration of the environment. Rough translation of the results in Dayan (1991) to an environment of size similar to our model's implies that on the order of half an hour of running to the goal would be required. Whenever a goal or obstacle is added to the environment, learning must begin again from scratch. Some latent learning is incorporated in Dayan (1991) that would reduce learning time somewhat if a goal were added, but not if an obstacle were added. Rat's navigation is better characterised by fast learning and good, but not necessarily optimal, trajectories.

Our model requires about 1 min (depending on the value of $\langle C_{in} \rangle$) to build an environmental representation that can be used to code for the positions of goals (or obstacles) the first time that a goal is encountered. This has an obvious speed advantage over the above approaches in general, and in particular deals well with the case of addition (or movement) of goals or obstacles after exploration because the environmental representation and the goal location are separate. Another important aspect of our model of navigation is that it has a directional output (north, east, etc.) that could be used by a motor system after translation to egocentric direction (left, right, etc.), the only extra information

required being that present in the head-direction cells. Models of navigation that provide a sequence of place cell firing to determine a trajectory (e.g., Hetherington & Shapiro, 1993; and possibly Muller et al., 1991) still need a mechanism to make the conversion into a signal capable of controlling navigation (e.g., a system that knows the directional relationships between all place fields).

Other models of hippocampal navigation include conceptual models that refer directly to the action of individual cells, but have not been simulated as neuronal systems (e.g. McNaughton et al., 1993; O'Keefe, 1991); and computer simulations of conceptual models that cannot be related to the action of individual cells (e.g., Worden, 1992; Schmajuk & Thieme, 1992; note that although the "neural" equations in the latter resemble models of a neuron's membrane potential, they in fact refer to conceptual quantities such as a "place" or a "view"). Wan et al. (1993) make an interesting attempt to implement a conceptual model as a simulation of a computer model. Although some of these models provide useful insights into the possible biological mechanism present, we do not have space to review them here.

6. CONCLUSIONS

We have considered the action of cells in the hippocampus in terms of the typical spatial firing rate map of a place cell, which resembles a radial basis function. We have presented a model of hippocampal function in which place cell firing is built up from tuning curve sensory inputs from cues around the rat's environment. Viewed as basis functions, the set of place cell firing rate maps could be used to endow cells downstream with firing rate maps approximating any desired function over the 2-D environment of the rat.

In particular we have used this framework to suggest an explicit mechanism by which a "hippocampal cognitive map" (O'Keefe & Nadel, 1978) could be read, enabling spatial navigation, and confirmed the workability of this scheme by a neuronal simulation. Our success in modelling hippocampal cells with realistic spatial correlates of firing, and in obtaining navigation (see Figure 12), enables us to begin to tie together electrophysiological and behavioural data in one model. As far as we are aware this is the first simulation of the hippocampal system including both biologically plausible cell firing and connection modification, and behaviour, that is, actual navigational trajectories from sensory input. The mechanism used explores one potential computational use of the ubiquitous EEG θ rhythm and the apparent phase relationship of place cell firing with θ (O'Keefe & Recce, 1993).

Limitations:

1. The simulated rat experiences only one environment (avoiding the problem of environmental recogni-

tion), and no computational use is made of the dentate gyrus or the recurrent collaterals in field CA3 (see Section 4).

2. We have only simulated "open-field" environments, in which place cell firing is not strongly modulated by direction; we have not examined mazes with narrow arms, in which place cell firing is directional.
3. All cues are point-like, distinguishable, fairly evenly distributed, and fixed (we have not yet examined the effect of the movement or removal of individual cues on firing fields or navigation) and neither goal nor obstacles are used as cues. This leaves open the problem of distinguishing individual cues and choosing the particular form of sensory input (see Section 5.3). It is not clear how our cues map onto the real cues that are controlled by the experimenter and the sensory information from the environment that is not.
4. The actual model of *navigation* (i.e., deciding where to go given the approximate directions of goals or obstacles) is limited to approaching a goal directly (or while being repelled by a small number of point-like obstacles), for example, trajectories are not modified with reference to environmental constraints such as being on a maze with narrow arms.
5. The simulated rat moves at a constant, rather fast, speed and cannot stop and start.

Advantages:

1. Very fast learning of approximate trajectories to the goal via one-shot learning of a goal's position, and rapid construction of place fields (cf. Wilson & McNaughton, 1993).
2. "Latent learning" (during exploration irrespective of the presence of goals) and the ability to cross unexplored areas en route to the goal, that is, taking short cuts.
3. The output of the simulated system is directional (the population vector of a set of goal cells continuously directs the rat).
4. We have developed a simple model of synaptic modification and local inhibition, thought to be present in hippocampal circuits considered, that leads to rapid construction of cell firing fields that compare well with extracellular recording.

Predictions (note that partial data exist concerning 2–4):

1. A specific phase relationship of the firing of entorhinal cells to the relative position of the cues to which they respond. This is not crucial however; it could be that the proposed phase coding of place cell firing is caused by a mechanism intrinsic to the place cells.
2. The firing fields of cells nearer to the sensory input end of the model (i.e., ECs vs. PCs, or PCs vs. SCs) are more sensitive to changes in the position of individual cues.
3. Navigational performance depends on time spent in

the environment, and on time spent exploring prior to encountering the goal.

4. The formation of spatially correlated firing rate maps that cover large parts of the environment in subicular cells (see also Burgess et al., 1993).
5. The existence of goal cells, one synapse downstream of the subiculum, whose firing fields cover the entire environment, which provide a population vector directing movement.
6. If the rat cannot look around in different directions when at the goal location, its navigation to the goal will be less accurate.

REFERENCES

- Amaral, D. G., & Witter, M. P. (1989). The three-dimensional organization of the hippocampal formation: A review of anatomical data. *Neuroscience*, **31**, 571–591.
- Ballard, D. J. (1986). Cortical connections and parallel processing: Structure and function. *Behavioral and Brain Sciences*, **9**, 67–90.
- Barnes, C. A., McNaughton, B. L., Mizumori, S. J. Y., Leonard, B. W., & Lin, L.-H. (1990). Comparison of spatial and temporal characteristics of neuronal activity in sequential stages of hippocampal processing. *Progress in Brain Research*, **83**, 287–300.
- Barto, A. G., & Sutton, R. S. (1981). Landmark learning: an illustration of associative search. *Biological Cybernetics*, **42**, 1–8.
- Bliss, T. V. P., & Lomo, T. (1972). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *Journal of Physiology (London)*, **232**, 331–356.
- Bostock, E. M., Taube, J. S., & Muller, R. U. (1988). The effects of head orientation on the firing of hippocampal place cells. *Society of Neuroscience Abstracts*, **14**, 127.
- Bugmann, G. (1991). Summation and multiplication: Two distinct operation domains of leaky integrate-and-fire neurons. *Network*, **2**, 489–509.
- Burgess, N., O'Keefe, J., & Recce, M. (1993). Using hippocampal 'place cells' for navigation, exploiting phase coding. In S. J. Hanson, C. L. Giles, & J. D. Cowan (Eds.), *Advances in neural information processing systems 5*. San Mateo, CA: Morgan Kaufmann.
- Burgess, N., O'Keefe, J., & Recce, M. (1994). Toward a mechanism for navigation by the rat hippocampus. *Proceedings of the 2nd Annual Computation and Neural Systems Conference: CNS'93*, Washington DC, 31st July–6th August, 1993. New York: Kluwer Academic Publishers, in press.
- Dayan, P. (1991). Navigating through temporal difference. In R. P. Lippmann, E. J. Moody, & D. S. Touretzky (Eds.), *Neural information processing systems 3* (pp. 464–470). San Mateo, CA: Morgan Kaufmann.
- Fox, S. E., Wolfson, S., & Ranck, J. B. (1986). Hippocampal theta-rhythm and the firing of neurons in walking and urethane anesthetized rats. *Experimental Brain Research*, **62**, 495–508.
- Georgopoulos, A. P., Kettner, R. E., & Schwartz, A. B. (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. *Journal of Neuroscience*, **8**, 2928–2937.
- Hetherington, P. A., & Shapiro, M. L. (1993). A simple network model simulates hippocampal place fields: II. Computing goal-directed trajectories and memory fields. *Behavioral Neuroscience*, **107**, 434–443.
- Keith, J. R., & McVety, K. M. (1988). Latent place learning in a novel environment and the influence of prior training in rats. *Psychobiology*, **16**, 146–151.
- McNaughton, B. L., Barnes, C. A., & O'Keefe, J. (1983). The contribution of position, direction and velocity to single unit activity

- in the hippocampus of freely-moving rats. *Experimental Brain Research*, **52**, 41–49.
- McNaughton, B. L., Knierim, J. J., & Wilson, M. A. (1993). Vector encoding and the vestibular foundations of spatial cognition: A neurophysiological and computational hypothesis. In M. Gazzaniga (Ed.), *The cognitive neurosciences*. Boston: MIT Press.
- Morris, R. G. M., Garrud, P., Rawlins, J. N. P., & O'Keefe, J. (1982). Place navigation impaired in rats with hippocampal lesions. *Nature*, **297**, 681–683.
- Muller, R. U., Kubie, J. L., & Ranck, J. B., Jr. (1987). Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *Journal of Neuroscience*, **7**, 1935–1950.
- Muller, R. U., Kubie, J. L., & Saypoff, R. (1991). The hippocampus as a cognitive graph. *Hippocampus*, **1**, 243–246.
- Muller, R. U., Kubie, J. L., Bostock, E. M., Taube, J. S., & Quirk, G. J. (1991). Spatial firing correlates of neurons in the hippocampal formation of freely moving rats. In J. Paillard (Ed.), *Brain and space* (pp. 296–333). Oxford: Oxford University Press.
- Murray, E. A., Davidson, M., Gaffan, D., Olton, D. S., & Suomi, S. (1989). Effects of fornix transection and cingulate cortical ablation on spatial memory in rhesus monkeys. *Experimental Brain Research*, **74**, 173–186.
- O'Keefe, J. (1979). A review of the hippocampal place cells. *Progress in Neurobiology*, **13**, 419–439.
- O'Keefe, J. (1991). The hippocampal cognitive map and navigational strategies. In J. Paillard (Ed.), *Brain and space* (pp. 273–295). Oxford: Oxford University Press.
- O'Keefe, J., & Conway, D. H. (1978). Hippocampal place units in the freely moving rat: Why they fire where they fire. *Experimental Brain Research*, **31**, 573–590.
- O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely moving rat. *Brain Research*, **34**, 171–175.
- O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford: Clarendon Press.
- O'Keefe, J., Nadel, L., Keightley, S., & Kill, D. (1975). Fornix lesions selectively abolish place learning in the rat. *Experimental Neurology*, **48**, 152–166.
- O'Keefe, J., & Recce, M. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, **3**, 317–330.
- Olton, D. S., Walker, J. A., & Gage, H. (1978). Hippocampal connections and spatial discrimination. *Brain Research*, **139**, 295–308.
- Ono, T., Nakamura, K., Fukuda, M., & Tamura, R. (1991). Place recognition responses of neurons in monkey hippocampus. *Neuroscience Letters*, **121**, 194–198.
- Parkinson, J. K., Murray, E. A., & Mishkin, M. (1988). A selective mnemonic role for the hippocampus in monkeys: Memory for the location of objects. *Journal of Neuroscience*, **8**, 4159–4167.
- Pavlidis, C., & Winson, J. (1989). Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. *Journal of Neuroscience*, **9**, 2907–2918.
- Powell, M. J. D. (1985). Radial basis functions for multivariable interpolation: A review. In J. C. Mason & M. G. Cox (Eds.), *Algorithms for approximation* (pp. 143–167). Oxford: Oxford University Press.
- Quirk, G. J., Muller, R. U., Kubie, J. L., & Ranck, J. B. (1992). The positional firing properties of medial entorhinal neurons: Description and comparison with hippocampal place cells. *Journal of Neuroscience*, **12**, 1945–1963.
- Recce, M., & O'Keefe, J. (1989). The tetrode: A new technique for multi-unit extra-cellular recording. *Society of Neuroscience Abstracts*, **15**, 1250.
- Recce, M., Speakman, A., & O'Keefe, J. (1991). Place fields of single hippocampal cells are smaller and more spatially localized than you thought. *Society of Neuroscience Abstracts*, **17**, 484.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing* (Vol. 1, pp. 318–362). Cambridge, MA: MIT Press.
- Rumelhart, D. E., & Zipser, D. (1986). Feature discovery by competitive learning. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing* (Vol. 1, pp. 151–193). Cambridge, MA: MIT Press.
- Schmajuk, N. A., & Thieme, A. D. (1992). Purposive behavior and cognitive mapping: A neural network model. *Biological Cybernetics*, **67**, 165–174.
- Schwartzkroin, P. A., Scharfman, H. E., & Sloviter, R. S. (1990). Similarities in circuitry between Ammon's horn and dentate gyrus: Local interactions and parallel processing. *Progress in Brain Research*, **83**, 269–286.
- Shapiro, M. L., & Hetherington, P. A. (1993). A simple network model simulates hippocampal place fields: I. Parametric analyses and physiological predictions. *Behavioral Neuroscience*, **107**, 34–50.
- Sharp, P. E. (1991). Computer simulation of hippocampal place cells. *Psychobiology*, **19**, 103–115.
- Stewart, M., Quirk, G. J., Barry, M., & Fox, S. E. (1992). Firing relations of medial entorhinal neurons to the hippocampal theta rhythm in urethane anesthetized and walking rats. *Experimental Brain Research*, **90**, 21–28.
- Suzuki, W. A., Zola-Morgan, S., Squire, L. R., & Amaral, D. G. (1993). Lesions of the perirhinal and parahippocampal cortices in the monkey produce long-lasting memory impairment in the visual and tactual modalities. *Journal of Neuroscience*, **13**, 2430–2451.
- Taube, J. S., Muller, R. U., & Ranck, J. B. Jr. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *Journal of Neuroscience*, **10**, 420–435.
- Thompson, L. T., & Best, P. J. (1989). Place cells and silent cells in the hippocampus of freely-behaving rats. *Journal of Neuroscience*, **9**, 2382–2390.
- Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, **55**, 189–208.
- Vanderwolf, C. H. (1969). Hippocampal electrical activity and voluntary movement in the rat. *Electroencephalography and Clinical Neurophysiology*, **26**, 407–418.
- Wan, H. S., Touretzky, D. S., & Redish, A. D. (1993). Towards a computational theory of rat navigation. In M. Mozer, P. Smolensky, D. S. Touretzky, J. L. Elman, & A. Weigend (Eds.), *Proceedings of the 1993 Connectionist Models Summer School*, Hillsdale, NJ: Erlbaum Associates.
- Wiener, S. I., Paul, C. A., & Eichenbaum, H. (1989). Spatial and behavioral correlates of hippocampal neuronal activity. *Journal of Neuroscience*, **9**, 2737–2763.
- Wilson, M. A., & McNaughton, B. L. (1993). Dynamics of the hippocampal ensemble code for space. *Science*, **261**, 1055–1058.
- Winson, J. (1978). Loss of hippocampal theta rhythm results in spatial memory deficit in the rat. *Science*, **201**, 160–163.
- Worden, R. (1992). Navigation by fragment fitting: A theory of hippocampal function. *Hippocampus*, **2**, 165–187.

APPENDIX A: SUMMARY OF ACTIVATION AND CONNECTION MODIFICATION RULES

The rat moves at the end of each cycle of the EEG θ rhythm. The frequency of the θ rhythm is taken to be 10 Hz, and each cycle is divided into three phases so that 1 time step in the model is equivalent to 1/30 s. The activation value of a cell represents the number of spikes fired in one time step.

A.1. Activation

A.1.1. *Sensory Cells.* Each sensory cell responds to one cue, with an activity (0, 1, 2, or 3) that is an approximate tuning curve function of the distance of the cue from the rat [see eqn (1)]. These activation values change only as a result of the rat moving, at the end of each θ cycle.

A.1.2. *Entorhinal Cells.* An EC receives input from two sensory cells, i and j say, that fire in response to the distance of two cues, say A and B . An EC can have a nonzero activation for only one time step (or phase) in each θ cycle; for this time step its activation is $[n, n_j/2]$, where n_i is the activation of sensory cell i and $[y]$ is the integer part of y . The phase at which it is active depends on the angle (from the rat's direction of travel) of the two cues A and B (see Figure 9).

A.1.3. *Place and Subicular Cells.* PCs and SCs are arranged in groups as in competitive learning (Rumelhart & Zipser, 1986). Consider a cell i in one of these two layers (layer l , say) at time step t , then its input from the layer of cells projecting to it (layer $l - 1$) is:

$$h_i^l(t) = \frac{1}{C_{sh}(1 + m_i^l(t))} \sum_{j=1}^{N_{l-1}} c_j w_{ji} a_j^{l-1}(t), \quad (A.1)$$

where $c_j = 1$ if a connection to it exists from cell j in layer $l - 1$ and $c_j = 0$ otherwise, w_{ji} is the connection weight, $a_j^{l-1}(t)$ is the activation of cell j , N_{l-1} is the number of cells in layer $l - 1$, $m_i^l(t)$ is the number of connections to the cell that have been switched on by time step t , and C_{sh} is a constant.

The cells within each group are placed in order (i_1, i_2, i_3, \dots) with respect to the size of their input $h_i^l(t)$, such that $h_{i_1}^l(t) \geq h_{i_2}^l(t) \geq h_{i_3}^l(t), \dots$. The four cells with the largest inputs in each group, that is, cells i_k where $1 \leq k \leq 4$, have activation

$$a_{i_k}^l(t) = \min(h_{i_k}^l(t), 5 - k),$$

and all other cells ($i_k, k > 4$) in the group, have activation 0. For PCs there are 50 cells per group and $C_{sh} = 1$, for SCs there are 25 cells per group and $C_{sh} = \frac{1}{2}$.

A.1.4. *Goal Cells.* The activation of goal cell i at time step t is given by:

$$a_i^G(t) = \left[\frac{25}{(1 + m_i^G(t))} \sum_{j=1}^{N_S} w_{ji} a_j^S(t) \right],$$

where N_S is the number of SCs, $a_j^S(t)$ is the activation of the j th SC, $m_i^G(t)$ is the number of connection weights that have been switched on to the goal cell by time step t , w_{ji} is the connection weight from the j th SC to the goal cell, and $[y]$ is the integer part of y .

A.2. Connections

A.2.1. *Initialisation.* All connection strengths take only values 0 or 1. Sensory-entorhinal connections have a fixed value of 1. The connectivity c between ECs and PCs, and between PCs and SCs is 0.5: for each cell i in the PC or SC layer (layer l , say) c_{ji} [see eqn (A.1)] is set to 1 for cN_{l-1} values of j between 1 and N_{l-1} , $c_{ji} = 0$ for all other values of j . The weight w_{ji} of each existing connection (i.e., when $c_{ji} = 1$) has a probability of $\langle C_{in} \rangle / cN_{l-1}$ of being initialised to 1 rather than 0. The connectivity between SCs and goal cells is 1.0 and all connection weights are initialised to 0.

A.2.2. *Modification.* EC-PC and PC-SC connections are set to 1 whenever the pre- and postconnection cells both fire at their maximum rate (4) in the same time step. Connections from SC to goal cells are only modified at the late phase of a θ cycle. At this phase in each θ cycle, the connection from an SC to a goal cell representing a particular goal and a particular direction is set to 1 whenever the rat is at the appropriate goal, facing in the appropriate direction, and the activity of the SC is nonzero.