

Predictions derived from modelling the hippocampal role in navigation

Neil Burgess, Andrew Jackson, Tom Hartley, John O'Keefe

Institute of Cognitive Neuroscience and Department of Anatomy, University College London, 17 Queen Square, London WC1N 3AR, UK

Received: 15 July 1999 / Accepted in revised form: 20 March 2000

Abstract. A computational model of the lesion and single unit data from navigation in rats is reviewed. The model uses external (visual) and internal (odometric) information from the environment to drive the firing of simulated hippocampal place cells. Constraints on the functional form of these inputs are drawn from experiments using an environment of modifiable shape. The place cell representation is used to guide navigation via the creation of a representation of goal location via Hebbian modification of synaptic strengths. The model includes consideration of the phase of firing of place cells with respect to the theta rhythm of hippocampal EEG. A series of predictions for behavioural and single-unit data in rats are derived from the input and output representations of the model.

1 Introduction

Our ability to perceive and remember spatial locations in the real world is central to much of our everyday behaviour. The exact nature of our representation of spatial location will affect our ability in a range of tasks such as remembering where we parked the car to taking a short-cut on the way to the shops. Neuropsychological studies (e.g. Habib and Sirigu 1987; Maguire et al. 1996; Bohbot et al. 1998) and, more recently, functional brain imaging (e.g. Aguirre and D'Esposito 1997; Maguire et al. 1997, 1998) have implicated the hippocampus and surrounding cortical areas in the medial temporal lobes in this type of spatial memory. However, detailed knowledge of the nature of the neuronal representation of spatial location requires information at the level of single cells. Indeed, the observation of 'place cells' in the hippocampus of freely moving rats (O'Keefe and Dostrovsky 1971) relates directly to the neural representation of the rat's spatial location. In this article we briefly review the data from single-unit recordings that

relates to the neural basis of navigation. We then describe a computational model consistent with these data, and examine the predictions that it makes regarding single-unit and behavioural data in rats.

2 Neural correlates of spatial processing

'Place cells' are neurons recorded in the hippocampus of freely moving rats that tend to fire at a high rate only when the rat is in a particular portion of its environment, independently of local sensory cues such as the odour of the floor covering (O'Keefe and Dostrovsky 1971; O'Keefe 1976). Since different place cells respond in different parts of the environment, collectively their firing rates encode the current location of the rat (Fig. 1; Wilson and McNaughton 1993). 'Head-direction' cells have been found near to the hippocampus in the dorsal pre-subiculum (Taube et al. 1990) and also in the anterior thalamic nuclei and mammillary bodies (e.g. Taube 1998). These cells code for the direction of the rat's head, regardless of its location within the environment. This representation of direction was predicted to accompany the place representation, under the hypothesis that the hippocampus is a system for spatial navigation (O'Keefe and Nadel 1978).

The hippocampal formation is well positioned to process multimodal information from all sensory cortices, although, when available, visual information tends to exert a dominant influence on place cell firing compared to the other sensory modalities. Visual stimuli at or beyond the edge of the rat's reachable environment are sufficient to control the overall orientation of the place (O'Keefe and Nadel 1978; Muller and Kubie 1987; O'Keefe and Speakman 1987) and head-direction (Taube et al. 1990) representations. Rotation of these stimuli cause rotation of the receptive fields of place and head-direction cells about the center of a symmetrical environment. In congenitally blind rats, objects at the edge of the environment also control the orientation of the place cell representation (Save et al. 1998). However, objects placed within the environment do not control the orientation of the place cell representation (Cressant




Fig. 1. The firing rate maps of 35 simultaneously recorded place cells, i.e. 35 'place fields' showing firing rate as a function of the rat's location in the *square box* shown in the top right corner (adapted from O'Keefe et al. 1998). The fields are arranged topographically for

display purposes only; there is no clear topographical relationship between the locations of the place cells in the hippocampus and the locations of their place fields in the environment. Notice that, collectively, the place fields cover the environment

et al. 1997). The representations of place and head direction appear to be linked in that, when both head direction cells and place cell have been recorded simultaneously, the orientations of the two representations appear to be locked together (Knierim et al. 1995).

Experiments have been conducted in which place cell firing rates are recorded in environments of varying shape and size (O'Keefe and Burgess 1996). These show that the shape and location of a place field is determined by combining information about the distances of the rat from the walls that bound its environment in each allocentric direction. In simple rectangular environments, the data can be well explained by considering only dis-

tances in the four directions normal to the walls. Figure 2 shows the data recorded from a place cell in four rectangular environments, and the qualitative fit to the data that can be made from a thresholded sum of four inputs that have Gaussian tuning curve responses to the presence of a wall at a particular distance from the rat along the four allocentric directions North, South, East and West. These findings suggest that the inputs driving hippocampal place cell firing are tuned to respond to walls at a particular bearing. One of the implications of this finding is that distant visual cues play an important role in anchoring the rat's sense of allocentric direction, and this in turn affects the overall




Fig. 2. *Left* The place fields of a place cell recorded in four *rectangular boxes*. *Right* A qualitative model of the place cell's firing as the thresholded sum of four sensory input cells, each tuned to respond whenever there is a wall a particular distance from the rat along one of the four allocentric directions normal to a wall. The Gaussian tuning functions decrease with increasing distance of peak response, as shown in Eq. (1). Adapted from O'Keefe and Burgess (1996)

orientation of the place fields by determining which wall is perceived as to the North etc.

The modality of inputs conveying information about the distance of the walls is not clear. Some authors have argued that these are internal signals (i.e. vestibular and proprioceptive) which are occasionally reset by reference to visual (McNaughton et al. 1996) or tactile (McNaughton 1996) information. However, it seems more likely that the system is driven by a mixture of whichever input modalities are afforded by a particular environment (e.g. Hill and Best 1981). These inputs would include visual, olfactory and tactile information as well as internal signals, all of which are available in the hippocampus.

The spatial role of the hippocampus and surrounding areas is also indicated by the fact that lesion of the rat's hippocampus or subiculum impair its navigational ability, specifically in tasks requiring an internal representation of space such as returning to an unmarked goal location from novel starting positions (e.g. Morris et al. 1982; Barnes 1988; Jarrard 1993).

2.1 Phase coding

In the above we have considered the representation of spatial information by neuronal firing rates. However, place cell firing appears to reflect the location of the rat using a temporal code as well as increasing its firing rate whenever the rat enters a circumscribed part of its environment (the 'place field'). Whenever the rat is doing something that involves displacement motion of its head, the hippocampal electroencephalograph (EEG) shows a sinusoidal oscillation termed the theta rhythm. The timing with which a place cell fires as the rat runs through the place field on a linear track exhibits a systematic relationship to the phase of theta, such that each burst of spikes occurs at a successively earlier phase (O'Keefe and Recce 1993). Perhaps the most interesting observation related to this phase coding is that the phase of firing correlates more strongly with the location of the

rat ($r = 0.66$) than with the time since entering the place field ($r = 0.42$). Whether or not such a phase coding exists when the rat explores freely in an open environment is less clear. It was predicted (Burgess et al. 1993) that, in open field environments, place cells firing at a 'late' phase would tend to have place fields centred ahead of the rat, and those firing at an earlier phase would tend to have place fields centred behind the rat (see Burgess et al. 1994; Skaggs et al. 1996 for some evidence supporting this).

2.2 The model

The experimental observation of place cells together with the data linking hippocampal lesions to impairments in navigation raise two immediate questions: How does the place specificity of place cells arise, and how do they contribute to behaviour? Over the past few years we have developed a neural network model that provides a possible solution to these two questions (Burgess et al. 1994; Burgess and O'Keefe 1996; O'Keefe and Burgess 1996). We have also tested this model by implementing it





Fig. 3. *Top* Hardware setup. *Bottom* Robot's eye view. The detected horizontal dark-light edge points are shown in *white*, and a *black arrow* marks the (x, y) position in the image that is returned by the visual processing. The 'North' wall can be identified (*on the right*) by its dark upper half. Adapted from Burgess et al. (1997)

on a mobile robot: forcing it to use real-world inputs, and seeing if it can indeed direct the navigation of a robot.

2.3 The neural network

The sensory inputs to the simulated hippocampus are encoded in a rectangular array of cells, organized such that each row of cells codes for the distance to a wall in a particular direction with each cell tuned to respond maximally at a particular distance (Fig. 5). We note that, in convex environments, identifying the walls on the basis of their allocentric direction from the rat conveniently solves the ‘binding’ or ‘correspondence’ problem of how information regarding a particular stimulus (a wall here) is channelled to a particular set of cells as the robot moves about. It is possible that the lack of influence on place fields of an object within the environment is due to the fact that its allocentric direction from the rat can vary by more than 180° , allowing the binding problem to rear its head once more as information regarding that object would not arrive on one constant set of channels.

Sensory cells respond to the distance of walls in particular directions according to the form of the independent place field components identified in O’Keefe and Burgess (1996), e.g. cell i in the row coding for distances from the West wall has firing rate:

$$\frac{A \exp[-(x - d_i)^2 / 2\sigma^2(d_i)]}{\sqrt{2\pi\sigma^2(d_i)}} \quad (1)$$

where x is the distance from the wall, d_i is the distance at which the cell responds maximally, the amplitude $A = 500$. The width of the tuning curve increases with the distance of peak response as $\sigma(x) = \sigma_0(L^2 + x^2)/L^2$. This reflects the decreasing reliability of the estimate of x at large distances; see later for a different form for this increase (in the simulations here $\sigma_0 = 10$ cm and $L = 30$ cm). In principle, there will be sensory cells tuned to respond to walls at all distances in all allocentric directions. However, our previous experiment (O’Keefe and Burgess 1996) indicated that consideration of the four orthogonal compass directions was sufficient for explanation of data recorded in rectangular boxes aligned with these directions. We discuss a more general model in Sect. 4.1.

The connections in the model are simply taken to be ‘on’ or ‘off’, and each cell is modelled as a linear threshold unit: firing at a rate proportional to the amount by which its net input exceeds a threshold. The simulation is performed in time steps during which all firing rates and connection weights are updated. Each step corresponds to one cycle of the theta rhythm (about 0.1 s) and is divided into two, corresponding to the early and late phases of the theta cycle (see below). Each cell in the entorhinal layer receives hard-wired connections from two sensory cells related to two orthogonal walls (Fig. 5). The connections from the entorhinal layer to the place cell layer include an element of unsupervised competitive learning (Burgess et al. 1994).




Fig. 4. *Top* A simple snapshot model of place cells and navigation. The place cell activity at a goal location is stored by taking a snapshot of place cell activity via Hebbian modification of connections to a goal cell. Moving so as to maximize the firing of the goal cell returns the rat to the goal location. *Filled circles* are active place cells; *empty circles* are inactive place cells and *filled squares* mark potentiated synapses. *Bottom* A population vector model of place cells and navigation. The firing rates of four goal cells are illustrated. Each cell has a peak rate displaced from the central goal location in an allocentric direction (North, South, East or West). The population vector of displacement directions weighted by firing rates indicated the direction of the rat from the goal. The net firing rate indicates the proximity of the goal. *G* Goal location. Adapted from Burgess and O’Keefe (1997)

The behaviour of place fields depends solely on which connections to the corresponding place cell have been turned on. A place field may maintain a fixed distance from two orthogonal walls, or may depend on many more than two inputs, all peaked at a fixed distance from a walls of the environment. Accordingly, some place fields will change in amplitude and shape when the environment is changed in size or shape consistent with experimental data (O’Keefe and Burgess 1996). Because entorhinal cells receive only two orthogonal inputs, their receptive fields will all remain at a fixed distance from two walls and will not change shape or amplitude during changes in the shape and size of the environment (consistent with the known data on entorhinal cells Quirk et al. 1992).

The basic idea behind how the model uses place cell firing for navigation can be illustrated by the following

very simple model. When the rat encounters a location in its environment that is associated with reward, a 'goal cell' downstream of the place cells is strongly excited by the attributes of the goal. This causes a one-shot Hebbian increment in the synaptic connections to the goal cell from the place cells that are active at that location. 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. Thus the activation of each goal cell will code for the proximity of a goal location, and could be used in a gradient-ascent style search for the goal: the rat could return to the goal location simply by moving so as to increase the firing rate of the goal cell (Fig. 4).

The model simulated here actually uses a more complicated model of learning of the goal location. This depends on one-shot Hebbian association of the place cells active at the goal location to a set of goal cells, in such a way that the subsequent goal cell activations form a 'population vector' (Georgopoulos et al. 1988) coding for the direction of the goal (Burgess et al. 1994; Burgess and O'Keefe 1996). This more complicated model takes advantage of the coding of location by the phase of the EEG theta rhythm at which spikes are fired, described below. It has advantages over the simple model such as enabling rats to take short-cuts towards the goal, and does not require the rat to hunt around to determine the direction in which to move. Finally, we note that the goal cell population vector is an allocentric direction (e.g. Northwest), and must be translated into an egocentric direction (e.g. left) before being used to guide the navigation of the rat. This transformation is straightforward given that the rat knows its own orientation, and might be expected to occur in the basal ganglia (Brown and Sharp 1995), or in the posterior parietal cortex (Burgess et al. 1999).

2.4 Phase coding and a population vector representation of goal location

The two questions regarding the formation and use of place cell firing in navigating rats include the questions of the formation and use of the phase-coded nature of their firing. The model makes use of the phase coding type of in open-field environments that we predicted on the basis of the phase coding on linear tracks: that place cells firing at a late phase of the EEG theta rhythm should tend to have place fields centred ahead of the rat while those firing at an early phase should tend to have place fields centred behind the rat.

Rather than the simple model described above, in which the proximity of a goal location is represented by the firing rate of one goal cell, we decided that each goal location should be represented by a set of goal cells, each associated with a particular goal and with a particular head direction (North, East etc). In this way, the direction to a goal location could be represented by the population vector of a set of goal cells, i.e. the vector

sum of each goal cell's associated direction weighted by its firing rate is a vector indicating the direction of the rat from the goal.

We assume that goal cells have inputs from the head direction and reward systems such that the connections from place cells to a goal cell can be modified whenever that rat encounters the corresponding goal and faces in the appropriate direction. The modification of these connections occurs at the late phase of the theta cycle and switches on connections from place cells active at that time. We further assume that whenever a rat encounters a goal location it turns around to face in several different compass directions. Application of the above learning rule causes a set of goal cells to provide a population vector encoding direction to a previously encountered goal. This follows from the fact that the goal cell associated with North will receive strong connections from place cells whose place fields tend to be centred to the North of the goal (i.e. place cells active at a late phase when the rat is facing North), and similarly for the goal cells associated with other directions. Thus, whenever the rat is North of the goal, the goal cell associated with North will fire more strongly than that associated with South, giving a net Northward component to the population vector (see Figure 5).

Although this model of goal learning is entirely speculative, it is consistent with evidence that the ease of inducing long-term potentiation of synapses (LTP) is modulated by the phase of theta at which stimulation occurs (Pavrides et al. 1988), and that rats' performance in the water maze is impaired when they are prevented from looking around at the goal location (Arolfo et al. 1994) and improved when they are allowed to (Keith and McVety 1988).

We assume that place cells in the model show phase coding as a result of phase coding of the sensory cells in the input to the model, as follows. The phase of firing of sensory cells depends on the current egocentric angle from the rat of the allocentric direction in which they respond to a wall at a particular distance (simply applying the rule used by Burgess et al. 1994 for discrete visual stimuli to the walls of the box). The phase of firing varies linearly from the latest phase for cells responding to a wall ahead of the rat to the earliest phase for those responding to a wall behind the rat. Each entorhinal cell receives inputs from two sensory cells and fires at a phase determined by the mean egocentric direction of the two walls from the rat (i.e. at the early phase if the mean direction is ahead, the late phase if behind). Each place cell receives inputs from many entorhinal cells, its phase of firing is simply taken to be the phase at which the majority of input firing arrives during each theta cycle. This mechanism produces a phase shift because sensory cells that respond maximally near to a given wall have a higher peak response than those that respond maximally far from the wall. Thus, as the rat moves through a place field, the sensory input from the wall in front of the rat (arriving at a late phase) tends to increase relative to the sensory input from the wall behind the rat.

2.5 Physical implementation

To check that the model can cope with the real-world problems of unreliable sensory and odometric information we have used a mobile robot to provide its inputs (and to be guided by its output). Fortunately, it is possible to provide a realistic simulation of these types of input since the rat's visual and odometric systems appear to be relatively unsophisticated. Rats have wide angle vision but do not necessarily segment stimuli into objects or extract much sensory information beyond the location or motion of the stimulus (e.g. Dean 1990). The ability of rodents to path integrate is also limited. For example, hamsters that are required to return to their point of departure using only path integration err significantly after an L shaped route of only 1 m per side, or after five active or two passive rotations in the dark (Etienne et al. 1996).

The hardware we used was a Khepera miniature robot (Fig. 3) with on-board video and a ring of short-range infra-red proximity detectors to provide artificial visual and haptic information. Two independently driven wheels allow movement around a rectangular environment formed by white walls and a dark floor (Fig. 3). Control of the robot's movement is imprecise: the speeds of the two independent wheel motors are set and occasionally monitored via shaft-encoders on the wheel axles, stopping when approximately the desired amount of turn has been achieved. To make our task easier, and to maintain close contact with the experimental situations in which the place cell data con-

straining the model was collected, the robot was tested in simple rectangular environments.

Visual estimates of the distances from the robot to the walls of the environment are used to drive the firing of 'sensory cells', entorhinal cells and thence place cells (Fig. 5; Burgess et al. 1997). The walls are identified by their allocentric direction from the robot. An estimate of the allocentric direction (orientation) of the robot is maintained by odometry and sightings of the North wall which is visually distinct from the other walls and serves to polarize the environment. When the robot encounters a goal location a reinforcement signal prompts one-shot Hebbian learning in connections from the place cells to a set of goal cells. The subsequent firing rates of these cells provide a continuous estimate of the direction and proximity of the goal location, enabling navigation, see above. Visual processing consists solely of filtering for horizontal dark-light edge points formed where a wall meets the floor, and finding the row (y) in the image containing the most dark-light edge points and the column (x) of the centroid of the edge points on that row (Fig. 3). The distance to the wall is estimated from y , and the bearing of the wall to the robot is estimated using x . This scheme does not work if the robot accidentally faces directly into a corner; however the estimated orientation does not drift fast enough for this to happen, see below. One wall (the North wall) is marked by a dark horizontal stripe along the top: its presence is detected by filtering for horizontal light-dark edge points. The infra-red proximity detectors detect the presence of a wall within about 4 cm and are used to prevent collisions.

At each step of the simulation the proximity detectors are read, the robot rotates on the spot to face in the estimated orientations N S E W capturing an image at each orientation. Each acquired image is used to estimate the distances to the wall and to provide on-going correction of the robot's estimated orientation (Fig. 3). If no wall was perceived by the proximity detectors the rat moves 3 cm forward in the desired direction. If a wall is perceived the robot moves 3 cm away from the wall, whose direction is estimated from the relative values of the proximity detectors. Since each step corresponds to 0.1 s or one theta cycle (see above) this implies a speed of 30 cm/s for the rat. (In practise each step takes around 3 s.) During exploration, each movement is made in a random direction within 30° of the previous direction (unless a wall is perceived). During navigation, each movement is made in the direction indicated by the output of the neural network model.




Fig. 5. The neural network. There are 60 sensory cells, 900 entorhinal cells, 900 place cells and 4 goal cells. Inputs from the sensors on the robot drive the firing of the sensory cells. Activation propagates through the model to form a representation of space in the place cell layer. Learning in the connections to the goal cells while at the goal location allows them to code for the direction and proximity of the goal location during subsequent movement. Adapted from Burgess et al. (1997)

3 Performance

Following the 300 steps of exploration (i.e. 30 s of rat-time) in a 50×50 cm square environment, the robot's navigation was tested in the same 50×50 cm environment or in a 50×75 cm environment. During exploration and navigation the robot performed well in maintaining estimates of the distance and direction of each wall relative to it. Put another way, relative to its

environment the robot shows good self-localization and maintenance of sense of direction. The relatively poor accuracy of the estimations of the distances to walls at each step can be judged from the spread of the points showing the robot's estimations of wall locations (Fig. 6), but these errors are not cumulative. The estimation of the angle to the wall is also rather imprecise, but is accurate enough to maintain a sense of direction accurate to the nearest 90° between sighting of the North wall.

In the 50×50 cm environment the robot successfully returned to unmarked reward locations having visited them once previously. The robot also shows generalization in returning to the goal from novel starting locations. Expanding the environment after the location of the goal had been learned caused the robot to search at a




Fig. 6. *Above* Navigation to a goal location (marked by *O*) in a square environment, following exploration of the environment. *Below* Navigation to the same goal in a rectangular environment (i.e. after deformation of the square). The figure shows the path of the robot, each *dash* representing 0.1 s of simulated time (or about 3 s of real time). *Dots* show the robot's estimate of the locations of the walls at each step. The robot received a (simulated) reward at the location marked *O* and was then replaced in the environment in two different locations and required to return to the goal location. Adapted from Burgess and O'Keefe (1997)

location intermediate to that indicated by each of the walls that have been pulled further apart (Fig. 6).

We also monitored the firing rates of the cells in the simulation, as a function of the robot's location. As expected, when the environment was increased in size along one axis most simulated place fields remain at a fixed distance from one of the two walls, although some become stretched and bimodal along that axis, see (Burgess et al. 1997). By contrast the entorhinal receptive fields are larger and always remain at a fixed distance from two of the walls. While reliable enough for maintaining a sense of direction, given frequent sightings of the walls of the environment, we note that odometry alone would not have been sufficient to support the firing of place cells over long periods of combined translation and rotation. However, experimental data from rats shows that, once established, both the place representation and the locus of searching can be maintained in the dark or in the absence of the environment's polarizing extra-maze stimuli. It is possible that uncontrolled auditory, olfactory and somatosensory cues may contribute to these findings.

4 Discussion and predictions

The model summarized above provides a detailed solution to the questions of how place cell firing could arise, and how it might be used, given the constraints of a particular physical manifestation in a particular (very simple) environment. We note that our robotic implementation only provides validation of the model's viability in these simple environments rather than cluttered real-world environments. In this section we compare the input side of the model with other approaches and generalize this aspect of the model to be applicable to environments other than simple rectangular boxes. This more general model makes several predictions which we examine briefly. We then consider the consequences of our choice of output representation. In consideration of both the model's inputs and outputs we attempt to make predictions regarding future experiments on rats in novel environments.

4.1 Sensory inputs

A major difference between the model presented here and most other models of navigation is our separation of information relating to place and direction. In the model, a place cell is tuned to fire in response to the presence of walls at a set of distances along particular allocentric directions from the rat. The set of distances and directions may be thought to define a 'receptive field' for the place cell, such that the firing rate is determined by the extent to which walls fall within the receptive field. Thus the rat's sense of direction controls the overall orientation of the set of place fields by defining the reference for allocentric directions, but where each cell fires depends on the distances along those directions at which walls are present.

Following Cressant et al. (1997) we suggest that distinctive features at or beyond the edge of the rat's immediate environment contribute to the rat's sense of direction, whereas intramaze features do not. Note that the features at the edge of the environment need not necessarily be visual (Save et al. (1998)), but that visual features tend to dominate other modalities when they are placed in conflict (Jeffery et al. (1997)). However, although distinct well-localized cues can drive the rat's sense of direction and hence play a role in determining the orientation of a place cell's 'receptive field', the receptive field itself responds only to extended features, like the walls of the box (or barriers within it). Thus narrow local features like the corner of a box will not influence place field shape and location, provided there are sufficient alternative directional cues (distant lights etc) for the corners not to affect the rat's sense of direction. In most other models of navigation (e.g. Bachelder and Waxman 1994), local features are directly used for both localization and orientation.

4.2 A general model: beyond rectangular environments

The model requires rather high-level 'sensory input cells' that are tuned to respond wherever there is a wall at a certain distance along a certain allocentric direction. In the case of the rectangular environments used in experiments discussed so far it is sufficient to consider only four directions (perpendicular to the walls) to provide a good account of the data. In this section we investigate a more detailed model of inputs with a smooth tuning to arbitrary distances and directions, so that as a population they cover the range of distances and directions at which the rat might experience a wall in more general environments. We use this model to provide more quantitative explanations of specific firing patterns, with the aim of producing more quantitative predictions of the firing pattern of individual place cells in specific novel environments.

We will term the particular form of sensory input cells explored in this section 'boundary vector cells', since each responds maximally when the boundary of the animal's environment is at a particular distance and allocentric angle. The firing rate of a particular boundary vector cell is determined by the overlap of the boundaries of the environment with its receptive field. Since a finite number of boundary vector cells is to cover the range of distances and directions the animal encounters, each must have a receptive field with angular and radial extent. Thus for a given cell i , peaked at polar coordinate (d_i, ϕ_i) from the rat's location, the receptive field is described by the product of two gaussian curves (cf. Eq. 1):

$$g_i(r, \theta) \propto \exp[-(r - d_i)^2 / 2\sigma_{\text{rad}}^2(d_i)] / \sqrt{2\pi\sigma_{\text{rad}}^2(d_i)} \\ \times \exp[-(\theta - \phi_i)^2 / 2\sigma_{\text{ang}}^2] / \sqrt{2\pi\sigma_{\text{ang}}^2} \quad (2)$$

where the angular width σ_{ang} is a constant but the radial width σ_{rad} increases linearly with increasing distance d_i .

This represents the accuracy of distance determination approximating by Weber's law.

To assess the effect of our assumptions about these cells as inputs to the place cell system we dispense with the entorhinal layer in the model described earlier (which effectively combined input cells into pairs before passing on their activations to place cells). Thus we consider each boundary vector cell as contributing directly to a thresholded sum of inputs that determines the place cell's firing rate. Note that within this framework, the firing rate of each place cell is solely determined by the subset of inputs to which it is connected, and possibly by the weights on those connections (if connection weights are not simply zero or one). In principle, by determining the subset of inputs which are involved in driving a particular place cell, we can predict the initial behaviour of that cell under arbitrary manipulations of environment size and shape. With reference to the effects of local features like corners we note that recent work (Lever et al. 1999) has shown that a place field in a square box will assume a consistent initial shape and location when the rat is transferred to a cylindrical box, as predicted by our model. Note that we do not discount effects of learning in the behaviour of place fields over the longer term. This is consistent with experiments showing that normal looking place fields are established almost immediately on the first visit to an environment (Hill 1978), and that this does not depend on NMDA mediated LTP, although their subsequent stability does (McHugh et al. 1996; Kentros et al. 1998).

Figure 7 illustrates the shape of a typical receptive field, and the firing field of the corresponding boundary vector cell. A given place cell's firing is determined by taking the thresholded sum of a set of these inputs. Figure 8 shows a selection of 30 simulated place fields generated in a rectangular environment, by assuming binary connections from a random set of 40 boundary vector cells and a threshold chosen to show the top 10% of firing in each case. We note that this simple model produces a realistic looking selection of place fields without any learning taking place. Note that we do not rule out learning effects; on the contrary, synaptic modification is implicated in a place field's subsequent stability or otherwise, and is required for the place cell to goal cell connections for the learning of goal locations.

The existence of boundary vector cells are a major prediction of the model, i.e. that there are neurons within, or upstream of, entorhinal cortex whose response is simply driven by the presence of a wall at a particular distance away from the rat along a particular allocentric direction. Furthermore, those cells responding maximally when a wall is nearby should have a higher peak firing rate and more sharply tuned response than those responding maximally when a wall is far off. We note that without systematic manipulation of the shape of the environment or introduction of a single movable wall, these cells might resemble place fields in that they would show spatially localized firing. Unlike place cells, their firing would depend only on the distance of the wall in the appropriate direction, and not on the other walls. The finding of subicular cells with

localised firing independent of environmental shape (Sharp 1997), and the existence of projections from subiculum to entorhinal cortex indicates this area as a good place to look first.

An issue that the simple model presented here, using the four directions N, S, E, W in rectangular environments, cannot resolve is whether such an input cell should be tuned to respond maximally when the surface


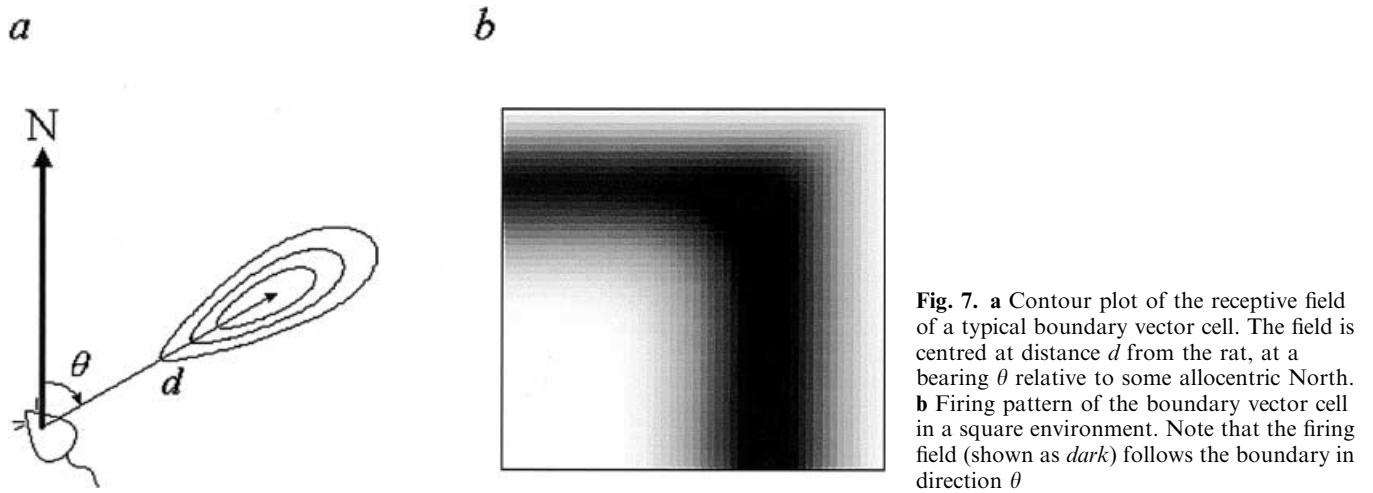


Fig. 8. Thirty simulated place fields. The firing of each place cell is modelled as the sum of 40 randomly chosen inputs from a set of 200 (tuned to respond at 20 distances in 10 allocentric directions),

thresholded to show the top 10% of firing. Their realistic appearance (cf. Fig. 1) implies that learning is not needed to explain place cell formation in a novel environment




Fig. 9. A simulation of the place cell shown in Fig. 2. The simulated place cell firing rate is the thresholded sum of weighted inputs from the 200 possible boundary vector cells. These cells respond to walls at arbitrary distances and directions with the Gaussian response function shown in Eq. (2). The connection weights are constrained to be positive and are chosen to fit the place fields shown in Fig. 2

of the wall is perpendicular to the direction along which it responds to distance (angle θ above). The current model does not include such a dependence on the angle of incidence with the wall. This may be an oversight, however, since the normal direction is intuitively the natural direction along which to measure a wall's proximity: being the one unique direction defined by a wall and, in most situations, the direction along which direction can be estimated most accurately.

In future modelling work we intend to reverse-engineer place fields recorded in many environments of differing shape by determining their constituent inputs, i.e. finding the set of sensory cells that, summed and thresholded, could produce the place field in all environments. This would allow us to make a quantitative prediction regarding the place field following further environmental manipulations, see (Jackson et al. 1999). An example of this process is shown in Figs. 9 and 10. Figure 9 shows a simulated place cell formed from the thresholded weighted sum of boundary vector cell inputs, with connection weights chosen to provide a least-squares fit of the data in Fig. 2. These weights are constrained to be greater than or equal to zero for biological plausibility (assuming excitatory inputs from boundary vector cells onto place cells). Figure 10 shows the prediction of the cell's place fields in the two square boxes if a barrier had been inserted. Unfortunately experimental data from that cell in this situation is not available. However this type of prediction (a second place field after introduction of a barrier) has recently been observed (Lever et al. 1999 and unpublished data)!




Fig. 10. Prediction of the effect a central barrier would have had on the place cell in Fig. 2. The simulation of this place cell (see Fig. 9) was tested in the large and small square environments shown in Fig. 2, with the addition of a central barrier. The simulation predicts that, for this cell, secondary place fields would be formed just below the barrier in addition to the fields just below the North wall

4.2.1 Phase coding. For the model to work as specified, sensory input cells should show a phase coding relative to the hippocampal theta rhythm according to the egocentric angle of the normal to the wall from the rat (or the angle to the nearest part of the wall when the wall is obscured in the normal direction). This is a very specific prediction relating to the model of phase coding used here.

4.3 Hippocampal outputs

In terms of the robot's behaviour, expanding the environment along one axis effectively stretches out the goal cell representation along that axis, but still results in a unimodal search pattern located between the loci indicated by fixed distances from each of the walls. These experiments have not yet been performed on rats, but the predicted search behaviour in the expanded environment runs contrary to that implied by simple extension of an experiment in which two cylinders indicating the goal location were moved further apart (Collett et al. 1986). However, the relative indifference of PCs to objects placed within an environment (Cressant et al. 1997) compared to the walls of the environment (O'Keefe and Burgess 1996), and recent lesion data (Pearce et al. 1998) indicate that navigating to places defined with respect to a single proximal cue is not mediated by the hippocampus.

The search pattern generated from the hippocampal representation of space depends on the storage and

output mechanism that makes use of it. Our particular model of this mechanism (i.e. the goal cell population vector) leads to the above behaviour. Whether or not these results predict the actual experimental performance of rats reflects directly on the validity of this mechanism. Using a different output mechanism, such as the simple model in Fig. 4 (top), would have different consequences for behaviour. In the model, contracting an environment by a large enough factor can cause place fields near to, and primarily associated with, opposing walls to cross over. This can cause the locations of peak firing of opposing goal cells to cross over, which has a disastrous effect on behaviour, with the robot searching only at the edges of the environment. Perhaps for this reason, such crossing over of place fields has not been observed, with the firing of one or other seemingly inhibited in this situation (e.g. Gothard et al. 1996).

The specific mechanism used to guide the model's behaviour makes predictions for single unit recording studies. The most clear prediction is that of sets of 'goal cells' immediately downstream of the place cells. How would we recognize a goal cell in a single-unit recording study? The clearest signature of a goal cell is that it develops a large receptive field peaked near to the goal location when the rat encounters the goal for the first time. Unfortunately, the chance of recording from the set of goal cells associated with a given goal at the moment it is discovered is rather small. Following discovery of the goal, goal cells would resemble place cells with very large receptive fields. The fact that head-direction signals must modulate the modification of synapses onto goal cells implies that the subiculum, pre-subiculum or nucleus accumbens (e.g. Brown and Sharp 1995) might be a good place to look for goal cells.

We have discussed the extension of the input side of the model to more complex environments. Extension of the goal cells to complex environments in a way that would support robust navigation is beyond the scope of this paper, but see Trullier and Meyer (2000, this issue).

Acknowledgements. We gratefully acknowledge many useful discussions with Colin Lever. The first half of this paper reviews work presented in Burgess et al. (1994, 1997) and in O'Keefe and Burgess (1996). N.B. is supported by a Royal Society University Research Fellowship, J.O.K. and T.H. are supported by the Medical Research Council and A.J. by a Wellcome prize studentship.

References

- Aguirre GK, D'Esposito M (1997) Environmental knowledge is subserved by separable dorsal/ventral neural areas. *J Neurosci* 17: 2512–2518
- Aroffo 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
- Bachelder IA, Waxman AM (1994) Mobile robot visual mapping and localization: A view based neurocomputational architecture that emulates hippocampal place learning. *Neural Netw* 7: 1083–1099
- Barnes CA (1988) Spatial-learning and memory processes – the search for their neurobiological mechanisms in the rat. *Trends Neurosci* 11: 163–169
- Bohbot VD, Kalina M, Stepankova K, Spackova N, Petrides M, Nadel N (1998) Spatial memory deficits in patients with lesions to the right hippocampus and to the right parahippocampal cortex. *Neuropsychologia* 36: 1217–1238
- 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 (1996) Neuronal computations underlying the firing of place cells and their role in navigation. *Hippocampus* 6: 749–762
- Burgess N, O'Keefe J, Recce M (1993) Using hippocampal 'place cells' for navigation, exploiting phase coding. In: Hanson SJ, Giles CL, Cowan JD (eds) *Advances in neural information processing systems*, vol 5. Morgan Kaufmann, San Mateo, Calif., pp 929–936
- Burgess N, Recce M, O'Keefe J (1994) A model of hippocampal function. *Neural Netw* 7: 1065–1081
- Burgess N, Donnett JG, Jeffery K, O'Keefe J (1997) Robotic and neuronal simulation of the hippocampus and rat navigation. *Phil Trans R Soc Lond B* 352: 1535–1543
- Burgess N, Jeffery KJ, O'Keefe J (1999) Integrating hippocampal and parietal functions: a spatial point of view. In: Burgess N, Jeffery KJ, O'Keefe J (eds) *The hippocampal and parietal foundations of spatial cognition*. Oxford University Press, Oxford, pp 3–29
- Collett TS, Cartwright BA, Smith BA (1986) Landmark learning and visuo-spatial memories in gerbils. *J Comp Physiol A* 158: 835–851
- Cressant A, Muller RU, Poucet B (1997) Failure of centrally placed objects to control the firing fields of hippocampal place cells. *J Neurosci* 17: 2531–2542
- Dean P (1990) Sensory cortex: visual perceptual functions. In: Kolb B, Tees RC (eds) *The cerebral cortex of the rat*. MIT Press, Cambridge, Mass., pp 275–307
- Etienne AS, Maurer R, Seguinot V (1996) Path integration in mammals and its interaction with visual landmarks. *J Exp Biol* 199: 201–209
- 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) Dynamics of mismatch correction in the hippocampal ensemble code for space: interaction between path integration and environmental cues. *J Neurosci* 16: 8027–8040
- Habib M, Sirigu A (1987) Pure topographical disorientation: a definition and anatomical basis. *Cortex* 23: 73–85
- Hill AJ (1978) First occurrence of hippocampal spatial firing in a new environment. *Exp Neurol* 62: 282–297
- Hill AJ, Best PJ (1981) Effects of deafness and blindness on the spatial correlates of hippocampal unit-activity in the rat. *Exp Neurol* 74: 204–217
- Jackson A, Hartley T, O'Keefe J, Burgess N (1999) A predictive geometric model of place field shape and location. *Soc Neurosci Abstr* 24: 556.1
- Jarrard LE (1993) On the role of the hippocampus in learning and memory in the rat. *Behav Neural Biol* 60: 9–26
- Jeffery K, Donnett JG, Burgess N, O'Keefe J (1997) Directional control of hippocampal place fields. *Exp Brain Res* 117: 131–142
- Keith JR, McVety KM (1988) Latent place learning in a novel environment and the influence of prior training in rats. *Psychobiology* 16: 146–151
- Kentros C, Hargreaves E, Hawkins RD, Kandel ER, Shapiro M, Muller RU (1998) Abolition of long-term stability of new hippocampal place cell maps by NMDA receptor blockade. *Science* 280: 2121–2126
- Knierim JJ, Kudrimoti HS, McNaughton BL (1995) Place cells, head direction cells, and the learning of landmark stability. *J Neurosci* 15: 1648–1659

- Lever C, Cacucci F, Burgess N, O'Keefe J (1999) Squaring the circle: place cell firing patterns in environments which differ only geometrically are not unpredictable. *Soc Neurosci Abstr* 24: 556.2
- Maguire EA, Burke T, Phillips J, Staunton H (1996) Topographical disorientation following unilateral temporal lobe lesions in humans. *Neuropsychologia* 34: 993–1001
- Maguire EA, Frackowiak RSJ, Frith CD (1997) Recalling routes around London: activation of the right hippocampus in taxi drivers. *J Neurosci* 17: 7103–7110
- Maguire EA, Burgess N, Donnett JG, Frith CD, Frackowiak RSJ, O'Keefe J (1998) Knowing where and getting there: a human navigation network. *Science* 280: 921–924
- McHugh TJ, Blum KI, Tsien JZ, Tonegawa S, Wilson MA (1996) Impaired hippocampal representation of space in CA1-specific NMDAR1 knockout mice. *Cell* 87: 1339–1349
- McNaughton BL (1996) Cognitive cartography. *Nature* 381: 368–369
- McNaughton BL, Barnes CA, Gerrard JL, Gothard K, 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
- 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
- O'Keefe J (1976) Place units in the hippocampus of the freely moving rat. *Exp Neurol* 51: 78–109
- O'Keefe J, Burgess N (1996) Geometric determinants of the place fields of hippocampal neurones. *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, Dostrovsky J (1971) The hippocampus as a spatial map: preliminary evidence from unit activity in the freely moving rat. *Brain Res* 34: 171–175
- O'Keefe J, Nadel L (1978) *The hippocampus as a cognitive map*. Oxford University Press, Oxford
- 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, Burgess N, Donnett JG, Maguire EA, Jeffery KJ (1998) Place cells, navigational accuracy, and the human hippocampus. *Philos Trans R Soc B* 353: 1333–1340
- Pavlidis 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
- Pearce JM, Roberts ADL, Good M (1998) Hippocampal lesions disrupt navigation based on cognitive maps but not heading vectors. *Nature* 396: 75–77
- 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
- Save E, Cressant A, ThinusBlanc C, Poucet B (1998) Spatial firing of hippocampal place cells in blind rats. *J Neurosci* 18: 1818–1826
- Sharp PE (1997) Subicular cells generate similar spatial firing patterns in two geometrically and visually distinctive environments: comparison with hippocampal place cells. *Behav Brain Res* 85: 71–92
- 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
- Taube JS (1998) Head direction cells and the neuropsychological basis for a sense of direction. *Prog Neurobiol* 55: 225–256
- 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
- Trullier O, Meyer J-A (2000) Animat navigation using a cognitive graph. *Biol Cybern* 83: 271–285
- Wilson MA, McNaughton BL (1993) Dynamics of the hippocampal ensemble code for space. *Science* 261: 1055–1058