

Robotic and neuronal simulation of the hippocampus and rat navigation

NEIL BURGESS, JAMES G DONNETT, KATHRYN J. JEFFERY
AND JOHN O'KEEFE

Department of Anatomy, University College London, London WC1E 6BT, UK

SUMMARY

The properties of hippocampal place cells are reviewed, with particular attention to the nature of the internal and external signals that support their firing. A neuronal simulation of the firing of place cells in open-field environments of varying shape is presented. This simulation is coupled with an existing model of how place-cell firing can be used to drive navigation, and is tested by implementation as a miniature mobile robot. The sensors on the robot provide visual, odometric and short-range proximity data, which are combined to estimate the distance of the walls of the enclosure from the robot and the robot's current heading direction. These inputs drive the hippocampal simulation, in which the robot's location is represented as the firing of place cells. If a goal location is encountered, learning occurs in connections from the concurrently active place cells to a set of 'goal cells', which guide subsequent navigation, allowing the robot to return to an unmarked location. The system shows good agreement with actual place-cell firing, and makes predictions regarding the firing of cells in the subiculum, the effect of blocking long-term synaptic changes, and the locus of search of rats after deformation of their environment.

1. INTRODUCTION

The hippocampus has been implicated as the neural basis of mammalian navigation ever since the discovery of spatially tuned neurons (place cells) in the hippocampus of freely moving rats (O'Keefe & Dostrovsky 1971). The fact that each place cell (PC) tends 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, prompted the idea that they provide the neural representation of the location of the rat within its environment (O'Keefe & Nadel 1978). It has recently been shown that the firing of PCs does indeed contain sufficient information to localize the rat (Wilson & McNaughton 1993). Lesions of the rat's hippocampus 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 (see, for example, Morris *et al.* 1982; Jarrard 1993).

More recently, head-direction cells have been found near the hippocampus in the dorsal presubiculum (Taube *et al.* 1990) and elsewhere (Mizumori & Williams 1993; Taube 1995). These cells code for the direction of the rat's head, regardless of its location within the environment. The existence of cells of this type was predicted by O'Keefe & Nadel (1978). Although general amnesia is the primary symptom of lesions of the temporal lobes and hippocampus in humans (Scoville & Milner 1957), evidence is now beginning to emerge relating the human hippocampal region to navigation and topographical

memory in neuropsychological (Habib & Sirigu 1987; Maguire *et al.* 1996a) and functional imaging (Maguire *et al.* 1996b, 1997a,b) studies.

The behaviour of the spatially tuned neurons in and around the hippocampus of freely moving rats provides an insight into the neuronal basis of mammalian navigation. Understanding the working of these neurons offers the tantalizing prospect of developing control algorithms that directly emulate mammalian navigational abilities. In this paper the spatial properties of place-cell activity are reviewed, and a computational model of the hippocampus as a navigational system implemented on a mobile robot is presented. The use of a robot ensures the realism of the assumed sensory inputs and enables true evaluation of navigational ability. Navigation is driven by place-cell firing in a manner similar to that described by Burgess *et al.* (1994), and is compared to data showing that rats can return to an unmarked goal location from novel starting positions. The behavioural data that show two separate loci for gerbils' search after two cues indicating a single reward site are pulled apart (Collett *et al.* 1986) are also considered. The neurons in the model are compared to single-unit recordings from the corresponding brain regions, where possible.

2. WHAT INPUTS SUPPORT PLACE-CELL FIRING?

Recent experiments have begun to reveal the nature of the signals underlying the apparently mysterious

ability of place cells to restrict their firing to specific portions of an environment. 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 & Speakman 1987; Muller & Kubie 1987) and head-direction (Taube *et al.* 1990) representations. Rotation of these stimuli can be shown to cause rotation of the receptive fields of place and head-direction cells about the centre of a symmetrical environment. However, objects placed within the environment do not show this control (Cressant *et al.* 1997). Interoceptive (possibly vestibular) inputs relating to self-motion also influence the overall orientation of the place (Sharp *et al.* 1995; Knierim *et al.* 1995) and head-direction (Blair & Sharp 1996) representations.

Place cells' receptive fields ('place fields') recorded in a series of uniform, rectangular, walled environments appear to be composed of the thresholded sum of two or more separate Gaussian tuning curves, each peaked at a fixed distance from one wall of the environment (O'Keefe & Burgess 1996). (See figure 1.) Because the walls used in this experiment were indistinguishable, and often interchanged, they can be presumed to be disambiguated on the basis of their allocentric direction from the rat. Recent results (Jeffery *et al.* 1997) have proved consistent with this interpretation: altering the rat's sense of direction by rotating it very slowly (0.15 rpm) in the dark causes the fields to rotate so as to maintain a fixed distance from two walls identified as being in particular directions from the rat relative to its own (rotated) frame of reference. However, if the rat and the box are slowly rotated in the presence of polarizing visual cues in the experimental room, the place fields may rotate with the box, or may stay in fixed orientation with the room, demonstrating the additional effect of extramaze visual cues in determining the rat's sense of direction. Figure 2 shows a place field that rotated in the dark, but was not affected by rotation in the presence of the visual room cue: in this case it maintained a fixed distance from two walls identified as being in particular veridical allocentric directions.

(a) What types of input are available, and how precise are they?

These remarkable properties of place and head-direction cell firing are apparently derived from relatively unsophisticated sensory inputs. Rats have wide-angle vision (320–360° in the horizontal plane depending on head angle (Hughes 1977)) but do not necessarily segment stimuli into objects or extract much sensory information beyond the location or motion of the stimulus (see, for example, Dean 1990). The ability of rodents to maintain an estimate of their location and orientation by keeping a cumulative record of their own movements (referred to as 'path integration') is also limited. For example, hamsters err significantly in returning to the start location after an L-shaped route of 1 m per side or after five active rotations or two passive rotations in the dark; see Etienne *et al.* (1996) and figure 3. Thus, although it is useful for maintaining a sense of direction, path integration is clearly not sufficient to support the firing of place cells




Figure 1. Firing-rate maps showing the place fields of two hippocampal neurons in four rectangular boxes of varying size and shape (the peak firing rate is shown in white on each plot). The place fields maintain fixed distances from two or more walls of the environment. This occasionally leads to bimodal firing-rate maps. See O'Keefe & Burgess (1996).

over long periods of combined translation and rotation (see figure 3). However, once established, both the place representation and the locus of searching can be maintained in the dark (O'Keefe 1976; Quirk *et al.* 1990) or in the absence of the environment's polarizing extramaze stimuli (O'Keefe & Speakman 1987). Uncontrolled auditory, olfactory, and somatosensory cues may contribute to these findings.

(b) Phase coding

Place-cell activity shows interesting temporal as well as spatial properties. Place cells tend to fire in short bursts of between one and four action potentials. The timing of the firing of these bursts has a systematic relation to the concurrently recorded electroencephalogram (EEG). Whenever the rat is involved in




Figure 2. Effect of internal and external inputs on place-field location. Firing-rate maps as in figure 1, showing the effect of slowly rotating a rat with and without rotation of the environment and with the polarizing visual room cues either masked (the 'dark' condition) or visibly present (the 'light' condition). Left column: rotation of both rat and box in the dark, first 90° anticlockwise and then 90° clockwise, resulted in corresponding rotations of the field. Middle column: rotation of the rat alone in the dark by 180° rotated the field by 180° even though the box remained stationary. Right column: when the procedure shown in the left column was repeated in the light, the field failed to rotate with the rat and the box, but maintained a constant distance from the south and west walls (as defined in the reference frame of the room).

head-displacement movements or locomotion, the EEG exhibits a roughly sinusoidal oscillation of 6–11 Hz, called the 'theta' rhythm. As the rat runs through a place field on a linear track, the phase at which spikes are fired is not constant, but shifts in a systematic way (O'Keefe & Recce 1993). When the animal enters the field the firing occurs late in the cycle, but shifts to progressively earlier phases as the rat runs through the field.

3. THE MODEL

The sensory and motor aspects of the rat are simulated using a robot, see below. 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 (ECs) and thence place cells (PCs); see figure 5. The walls are identified by their allocentric direction from the robot. The 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 Burgess *et al.* 1994; Burgess & O'Keefe 1996).





Figure 3. Path integration and place-cell firing. (a) The mean return direction (continuous line) of hamsters that had followed an L-shaped path (broken line) from a point of departure (D) to the point P in darkness (adapted from Seguinot *et al.* (1993)). (b) Typical path of a rat searching for randomly scattered food pellets (37 s shown). The locations at which a place cell fired spikes are shown as black squares. The rat starts in the top right corner, initially enters the place field (indicated by a short string of spikes), leaves the place field and performs a long and circuitous trajectory before re-entering the field (indicated by the firing of many spikes) and eventually leaving for the bottom left corner. Note that the inaccuracy in path integration (a) implies that it would not be sufficient to support the place-specific firing of place cells (b) on its own.

(a) Physical implementation

The model is implemented using a Khepera miniature robot, with on-board video and a ring of short-range infrared 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 (see figure 4a). 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 (see figure 4b). The distance to the wall is estimated from y ; the bearing of the wall to the robot from 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 infrared proximity detectors detect the presence of a wall within about 4 cm. Their function might be compared to that of a rat’s whiskers.

Movement is controlled by setting the speeds of the two independent wheel motors and occasionally monitoring the shaft-encoders on the wheel axles, stopping when the desired amount of turn has been achieved. This control is not precise, so the odometry of the robot is not noticeably superior to that of a rodent (see above). Control of the robot proceeds in steps: the proximity detectors are read and the robot then rotates on the spot to face in the estimated orientations north, south, east and west, capturing an image at each orientation. After each rotation the acquired image is used to estimate the distance to the wall and to correct the robot’s estimated orientation to agree with the estimate of its angle to the wall. If no wall was perceived by the proximity detectors the robot 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. Each step corresponds to 0.1 s (one theta cycle; see below) implying a speed of 30 cm s^{-1} for the rat, but actually takes around 3 s (processing on a SUN Ultra computer). 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 goal cells (see below).

(b) *The neural network*

The visual inputs to the simulated hippocampus are represented by a rectangular array of cells organized such that each row of cells codes for the distance to a particular wall, with each cell tuned to respond maximally at a particular distance (see figure 5). Note that identifying the walls on the basis of their allocentric direction from the rat solves the ‘binding’ or ‘correspondence’ problem of how information regarding a particular wall is channelled to a particular set of cells as the robot moves about. This also suggests a reason for the lack of influence on place-fields of objects within the environment: if the rat moves around an object, its allocentric direction from the rat will vary

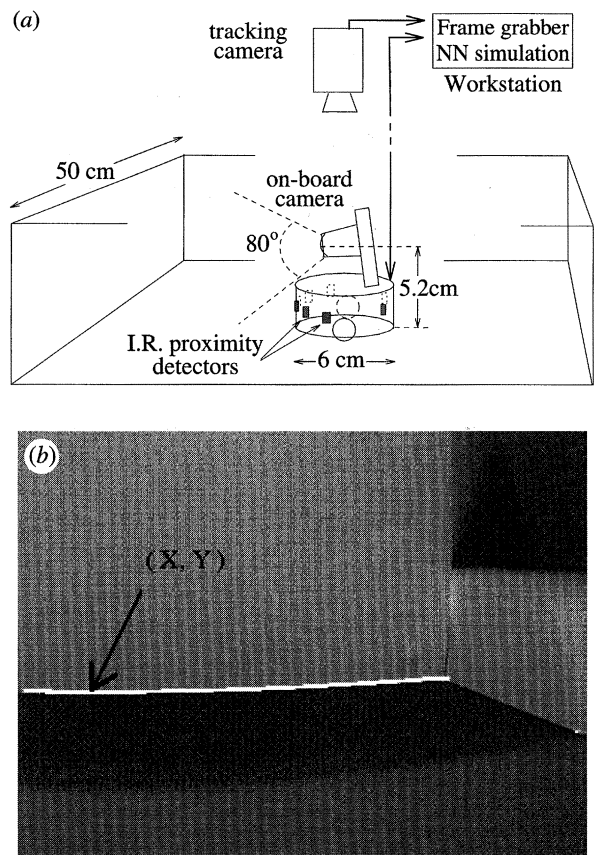



Figure 4. (a) Hardware set-up (see text for further details). (b) Robot’s-eye view. Detected horizontal dark–light edge points are shown in white; a black arrow marks the (x, y) position in the image that is returned by the visual processing. The north wall can be identified (on the right) by its dark upper half.

and information regarding it will not arrive on one constant set of channels.

The tuning of these ‘sensory’ cells follows the form of the independent place-field components identified by O’Keefe & Burgess (1996); for example, 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 (estimated visually), d_i is the distance at which the cell responds maximally, and the amplitude $A = 500$. The width of the response 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. If the rat uses the angle from vertical to the top or bottom of the wall to estimate its distance, then a fixed angular error will produce this form of error as a function of distance (to within a constant). The constants are assigned values $\sigma_0 = 10 \text{ cm}$ and $L = 30 \text{ cm}$ (half the values given by O’Keefe & Burgess (1996), because the environments used here are about half the size). All connections in the model take a



(b)




Figure 5. (a) Schematic diagram of the neuronal simulation. There are 60 sensory cells, 900 ECs, 900 PCs 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 in 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. (b) The Khepera mobile robot. See text for further details.

value of 0 or 1 ('on' or 'off'), and each cell fires at a rate proportional to the amount by which its net input exceeds a threshold. Each EC receives hard-wired connections from two sensory cells related to two orthogonal walls (see figure 5). The connections from

the EC layer to the PC layer include an element of learning. Initially, only one connection to each place cell is 'on', and a type of competitive learning turns on connections from a limited number of ECs with nearby receptive fields to the most active PCs at each time-step (0.1s) as the robot moves about its environment. Learning and activation in the PC layer occurs as follows. At each time step a threshold is set such that the 50 PCs with the greatest input are active, and connections between maximally active ECs and the four most active PCs are switched on. Each PC has a divisive threshold equal to the number of 'on' connections to it (preventing one PC always being the most active one and always receiving more 'on' connections; see Burgess *et al.* (1994)).

Depending on which connections to a PC have been turned on, its place field will maintain a fixed distance from two orthogonal walls, or reflect more than two inputs, all peaked at a fixed distance from a wall of the environment. Thus, some place fields will change in amplitude and shape when the environment is changed in size or shape. By contrast, the EC-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.

A simple model of navigation based on place-cell firing could work in the following way. When the rat encounters a 'goal' (i.e. a location in its environment that is associated with reward), a goal cell downstream from the place cells is strongly excited by the attributes of the goal. At the goal, a one-shot Hebbian increment is induced in the synaptic connections to the goal cell from the place cells that are active at the goal location. As the rat moves away from the goal location, the net activity of place cells with strong connections to the goal cell will be a monotonically decreasing fraction of the total place-cell activity. Consequently, the activation of each goal cell will code for the proximity of a goal location, and thus could be used as an evaluation function in a gradient-ascent-type search for the goal, i.e. the rat could return to the goal location simply by moving around so as to increase the firing rate of the appropriate goal cell (see figure 6*a,b*).

In fact, a more complicated model of learning of the goal location is used here, in which one-shot Hebbian association of the PCs active at the goal location to a set of 'goal cells' sets up a 'population vector' (Georgopoulos *et al.* 1988) that codes for the direction of the goal during subsequent navigation (see Burgess *et al.* 1994; Burgess & O'Keefe 1996). This has advantages over the simple model, such as enabling rats to take short cuts towards the goal (see, for example, Benhamou & Seguinot 1995; Tolman 1948), and does not require the rat to hunt around to determine the direction in which to move.

The population-vector model depends in part on the information carried by head-direction cells, and in part on the timing of PC firing showing the observed relation to the phase of the theta rhythm of the EEG (O'Keefe & Recce 1993). An implication of this phase relation is that PCs active at a 'late' phase tend to have place fields that are centred ahead of the rat whereas those firing at an 'early' phase tend to have place fields




Figure 6. (a) Simple model of place cells and navigation. A 'goal' cell stores a goal's location by taking a snapshot of place cell activity via long-term potentiation (LTP) when a goal cell is excited by the attributes of a particular goal location. Filled circles, active place cells; open circles, inactive place cells; filled squares, potentiated synapses. (b) The firing-rate map of the goal cell in the simple model during subsequent movements of the rat. This shows the cell's firing rate as a function of the location of the rat. It resembles an inverted cone, and codes for the proximity of the goal (G). (c) Population-vector model of place cells and navigation. Each goal location is represented by a group of goal cells. The firing of each cell indicates that the rat is displaced from the goal in a particular direction, such that the population vector of the group of cells represents the direction and proximity of the goal from the rat. The firing-rate maps of four cells corresponding to the directions north, south, east and west are shown. G marks the goal location. Adapted from Burgess & O'Keefe (1995).

centred behind the rat (Burgess *et al.* 1994). In the simulation, each time-step is divided into two intervals, corresponding to the early and late phases of a 10 Hz theta rhythm. The appropriate phase coding of PCs in the model results from the sensory cells responding to a wall ahead of the rat firing during the late phase, and those responding to a wall behind the rat firing at an early phase.

Briefly, each goal cell receives a projection coding both for the rat's head direction and for the sensory attributes (e.g. food, water, etc.) of a particular goal (see figure 5). These connections deliver a 'learn now' type of reinforcement signal to a goal cell whenever

the rat is at the appropriate goal location and facing in the appropriate direction. If this learning signal arrives at a 'late' phase of the EEG, or if synaptic plasticity is restricted to this phase (see, for example, Pavlides *et al.* (1988) for the relation of long-term potentiation of synapses to the theta rhythm), then the goal cell associated with, for example, the direction north will form active connections from place cells with receptive fields centred to the north of the location of the goal. As the rat looks around in different directions from the goal location, the connection weights to the set of goal cells are incremented such that each is associated with a particular allocentric direction, and will fire maximally at a location displaced from the goal in that direction. Thus the 'population vector', or vector sum of the directions associated with each goal cell weighted by their firing rates, estimates the direction of the rat from the goal (for example, whenever the rat is north of the goal, the goal cell associated with north will be firing more strongly than that associated with south; see figure 6c).

The firing of these goal cells drives navigation of the robot, enabling it to return to a previously visited but unmarked goal location. The goal-cell population vector is an allocentric direction (e.g. north-west) and must be translated into an egocentric direction (e.g. left) before being used. This transformation is simple given that the robot knows its own orientation, and might be expected to occur in the posterior parietal cortex or in the basal ganglia (see Brown & Sharp 1995).

4. PERFORMANCE

The robot was tested in two rectangular environments of size 50 cm × 50 cm and 50 cm × 75 cm. Its movements were tracked by an overhead camera and tracking system that detected two LEDs on the robot. Figure 7 shows the robot's exploration of a square environment. The robot performs 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. Figure 7 also shows the performance of the robot in returning to an unmarked reward location having visited it once previously. The robot also shows generalization in returning to the goal from novel starting locations.

Figure 8 shows the effect of expanding the environment after the location of the goal has been learned. When the environment is 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 figure 9). This compares well with observed data, in which the most common pattern was for place fields to maintain fixed positions relative to a wall, although some became stretched or bimodal (O'Keefe & Burgess 1996). By contrast, EC-receptive fields are larger and always remain at a fixed distance from two of the walls, consistent with the reported experimental data (Quirk *et al.* 1990). In terms of the robot's behaviour, expanding the environment along one axis

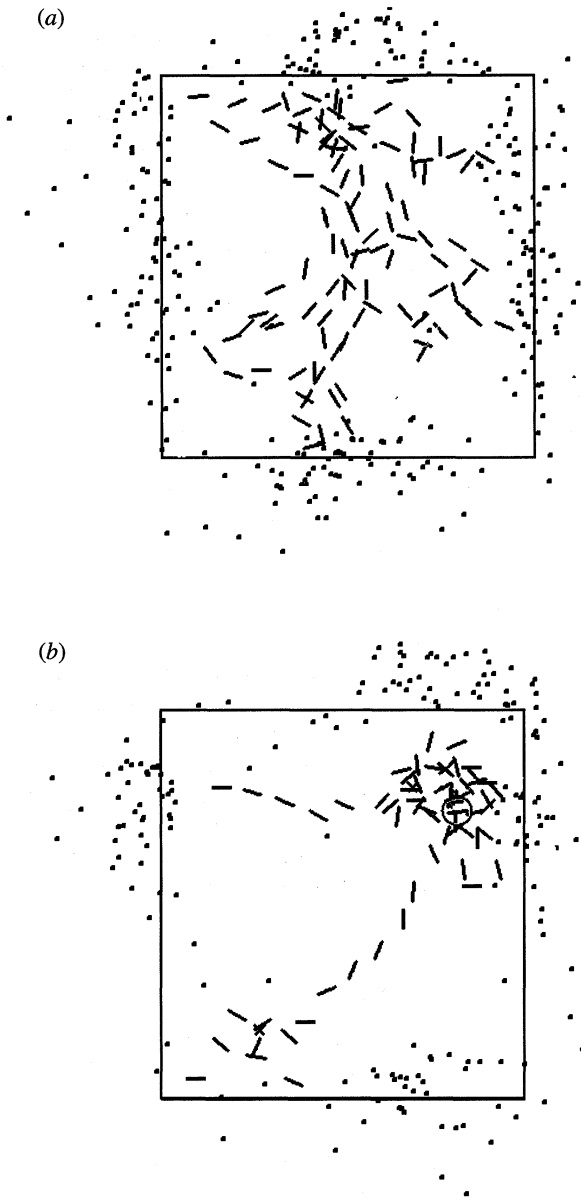


Figure 7. (a) Exploration of a square environment. Dashes show the actual path of the robot, each dash representing 0.1 s of simulated time (or about 3 s of real time). Dots show the robot's estimate of the locations of the walls at each step. Their proximity to the actual wall locations demonstrates the maintenance of good self-localization. (b) Navigation following exploration. 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. When searching, the robot follows the direction indicated by the goal cells at each time-step, and is successfully guided back to the goal location.

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.

The search pattern generated from the hippocampal representation of space depends on the storage and output mechanism that makes use of it. The particular model of this mechanism presented here (i.e. the goal-cell population vector) leads to the above behaviour. The way that phase coding is used in this model in the

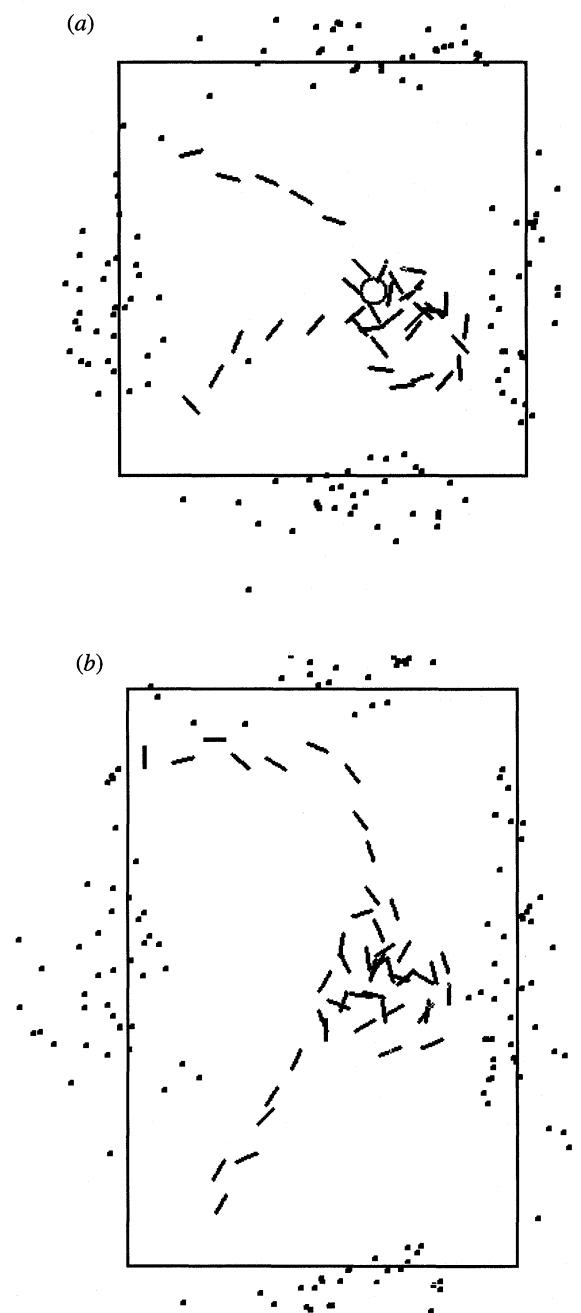


Figure 8. Navigation in a square environment to a goal location (marked by O) before (a) and immediately after (b) expansion of the environment into a long rectangle.

learning of connection weights to a goal cell leads to the north goal cell effectively being tied most strongly to the north wall, and similarly for the south goal cell and the south wall. Thus, the principal effect of expanding an environment is a separation of the peaks of each goal cell's firing-rate map (the locus of search remaining between them). By contrast, contracting an environment by a large enough factor can cause the locations of peak firing of opposing goal cells to cross over, and produces a more dramatic effect: the robot searches only at the edges of the environment. Whether or not these results predict the actual




Figure 9. Top row: firing rates for an entorhinal cell (EC, left) and two place cells (PCs, middle and right) calculated during exploration in a square environment. Bottom row: firing rates for the EC and PCs during exploration immediately after expansion of the square environment into a long rectangular environment. Note that the receptive field of the EC remains at a fixed distance from two of the walls (the top and left walls), whereas one of the PCs (middle) has a bimodal receptive field in the expanded environment.

experimental performance of rats reflects directly on the validity of this mechanism. Using a different output mechanism, such as the simple model in figure 6(a,b), would have different consequences for behaviour.

5. DISCUSSION

This model has developed from the work of Burgess *et al.* (1993, 1994) and O'Keefe & Burgess (1996). In terms of the overall aims and structure the work is similar to recent work by Touretzky & Redish (1996). However, their aim is more towards integrating suggested functional roles for many brain regions, whereas that of the present authors is more directed at the details of the neuronal implementation. For example, this model concentrates on the origins of the sensory inputs underlying place-cell firing and on how the firing of place cells could drive navigation (i.e. specifying how a vector-voting hypothesis could be implemented rather than simply postulating its existence).

It is noted that the representation of space in the entorhinal cell layer would be sufficient to enable navigation in a single environment if connected directly to the goal cell layer (presumed to be in the subiculum, immediately downstream of the place cells in region CA1 of the hippocampus). It is supposed that the role of the place cells in regions CA1 and CA3, and in particular of the long-range recurrent collaterals in CA3, is to support navigation in several distinct environments. These recurrent collaterals might enable the model to form an autoassociative memory for those place cells active in a given environment. Different subsets of place cells could then represent different

environments, with each subset forming a stable state of the autoassociator (see also McNaughton & Nadel 1990). The different responses of ECs and PCs to changes in environmental shape (Quirk *et al.* 1992) indicate that this response might play a role in environmental recognition. However, modifying the size and aspect ratio of a rectangular environment produced parametric changes in place fields (O'Keefe & Burgess 1996) rather than the discontinuous remapping that might be expected if each environmental shape was represented by the place cells as an orthogonal attractor state.

Expansion of an environment after the goal location has been learned stretches the goal-cell representation of goal location. With the current choice of output mechanism (see above) this results in a unimodal locus of search midway between the locations corresponding to maintaining fixed distances from each of the walls that have been pulled apart. 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 the result of Collett *et al.* (1986) in which two cylinders indicating the goal location were moved further apart. However, the relative indifference of PCs to objects placed within an environment (Cressant *et al.* 1997) compared with their response to the walls of the environment (O'Keefe & Burgess 1996) may indicate that this task is not mediated by the hippocampus in any case. A second prediction concerns the existence of goal cells, postulated to exist in the subiculum. It remains to be seen whether cells with the appropriate firing behaviour can be found in this part of the brain.

Finally, the likely effect on the model of preventing long-term changes of connection weights is noted. This would lead to an unstable association from ECs to PCs, leading to an unstable mapping of place fields onto place cells. This would be consistent with the finding of Rotenberg *et al.* (1996) in which place fields were recorded in mice genetically engineered not to show long-term potentiation of synapses in region CA1. The second effect of preventing long-term changes of connection weights would be the impermanence of the association from PCs to goal cells, and the loss of the population vector indicating the direction of the goal after the time interval over which any short-term potentiation had occurred. This is consistent with the behaviour of rats in a water maze after pharmacological blockade of long-term potentiation (see Morris & Frey, this volume).

A preliminary version of this paper can be found in: the Proceedings of the AISB workshop on *Spatial reasoning in mobile robots and animals*, Manchester, 1997. Technical Report Series, Department of Computer Science, Manchester University, ISSN 1361-6161, report no. UMCS-97-4-1. N.B. is supported by a Royal Society University Research Fellowship; J.O.K., J.D. and K.J. are supported by a programme grant from the MRC.

REFERENCES

- Blair, H. T. & Sharp, P. E. 1996 Visual and vestibular influences on head-direction cells in the anterior thalamus of the rat. *Behav. Neurosci.* **110**, 643-660.

- Benhamou, S. & Seguinot, V. 1995 How to find ones way in the labyrinth of path integration models. *J. Theor. Biol.* **174**, 463–466.
- Brown, M. A. & Sharp, P. E. 1995 Simulation of spatial learning in the Morris water maze by a neural network model of the hippocampal formation and nucleus accumbens. *Hippocampus* **5**, 171–188.
- Burgess, N. & O'Keefe, J. 1995 Modelling spatial navigation by the rat hippocampus. *Int. J. Neur. Syst.* **7**(suppl.), 87–94.
- 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 *Advances in neural information processing systems*, vol. 5 (ed. S. J. Hanson, C. L. Giles & J. D. Cowan), pp. 929–936. San Mateo, CA: Morgan Kaufmann.
- Burgess, N., Recce, M. & O'Keefe, J. 1994 A model of hippocampal function. *Neur. Networks* **7**, 1065–1081.
- Collett, T. S., Cartwright, B. A. & Smith, B. A. 1986 Landmark learning and visuo-spatial memories in gerbils. *J. Comp. Physiol. A* **158**, 835–851.
- Cressant, A., Muller, R. U. & 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 *The cerebral cortex of the rat* (ed. B. Kolb & R. C. Tees), pp. 275–307. Cambridge, MA: MIT Press.
- Etienne, A. S., Maurer, R. & Seguinot, V. 1996 Path integration in mammals and its interaction with visual landmarks. *J. Exp. Biol.* **199**, 201–209.
- 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. *J. Neurosci.* **8**, 2928–2937.
- Habib, M. & Sirigu, A. 1987 Pure topographical disorientation: a definition and anatomical basis. *Cortex* **23**, 73–85.
- Hughes, A. 1977 A schematic eye for the rat. *Visual Res.* **19**, 569–588.
- Jarrard, L. E. 1993 On the role of the hippocampus in learning and memory in the rat. *Behav. Neur. Biol.* **60**, 9–26.
- Jeffery, K., Donnett, J. G., Burgess, N. & O'Keefe, J. 1997 Directional control of hippocampal place fields. *Exp. Brain Res.* (In the press.)
- Knierim, J. J., Kudrimoti, H. S. & McNaughton, B. L. 1995 Hippocampal place fields, the internal compass, and the learning of landmark stability. *J. Neurosci.* **15**, 1648–1659.
- Maguire, E. A., Burke, T., Phillips, J. & Staunton, H. 1996a Topographical disorientation following unilateral temporal lobe lesions in humans. *Neuropsychologia* **34**, 993–1001.
- Maguire, E. A., Frackowiak, R. S. J. & Frith, C. D. 1996b Learning to find your way: a role for the human hippocampal region. *Proc. R. Soc. Lond. B* **263**, 1745–1750.
- Maguire, E. A., Burgess, N., Donnett, J. G., O'Keefe, J. & Frith, C. D. 1997a Knowing where things are: parahippocampal involvement in the encoding of object location in large-scale space. *J. Cogn. Neurosci.* (In the press.)
- Maguire, E. A., Burgess, N., Donnett, J. G., Frith, C. D. & O'Keefe, J. 1997 Differential involvement of right hippocampus, left hippocampus and inferior parietal cortex in human spatial navigation. *Soc. Neurosci. Abstr.* (In the press.)
- McNaughton, B. L. & Nadel, L. 1990 Hebb–Marr networks and the neurobiological representation of action in space. In *Neuroscience and connectionist theory* (ed. M. A. Gluck & D. E. Rumelhart), pp. 1–63. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Mizumori, S. J. Y. & Williams, J. D. 1993 Directionally sensitive mnemonic properties of neurons in the lateral dorsal nucleus of the thalamus of rats. *J. Neurosci.* **13**, 4015–4028.
- Morris, R. G. M., Garrard, 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. 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. & 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.
- 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.
- Pavlidis, C., Greenstein, Y. J., 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.
- Quirk, G. J., Muller, R. U. & Kubie, J. L. 1990 The firing of hippocampal place cells in the dark depends on the rat's experience. *J. Neurosci.* **10**, 2008–2017.
- 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. *J. Neurosci.* **12**, 1945–1963.
- Rotenberg, A., Mayford, M., Hawkins, R. D., Kandel, E. R. & Muller, R. U. 1996 Mice expressing activated CaMKII lack low frequency LTP and do not form stable place cells in the CA1 region of the hippocampus. *Cell* **87**, 1351–1361.
- Scoville, W. B. & Milner, B. 1957 Loss of recent memory after bilateral hippocampal lesions. *J. Neurol. Neurosurg. Psych.* **20**, 11–21.
- Sharp, P. E., Blair, H. T., Etkin, D. & Tzanetos, D. B. 1995 Influences of vestibular and visual-motion information on the spatial firing patterns of hippocampal place *J. Neurosci.* **15**, 173–189.
- Taube, J. S. 1995 Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. *J. Neurosci.* **15**, 70–86.
- Taube, J. S., Muller, R. U. & Ranck, J. B. 1990 Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *J. Neurosci.* **10**, 420–435.
- Tolman, E. C. 1948 Cognitive maps in rats and men. *Psychol. Rev.* **55**, 189–208.
- Touretzky, D. S. & Redish, A. D. 1996 Theory of rodent navigation based on interacting representations of space. *Hippocampus* **6**, 247–270.
- Wilson, M. A. & McNaughton, B. L. 1993 Dynamics of the hippocampal ensemble code for space. *Science* **261**, 1055–1058.