

Robotic and Neuronal Simulation of Hippocampal Navigation

Neil Burgess, James G Donnett and John O'Keefe

Department of Anatomy
University College London
London WC1E 6BT, U.K.

n.burgess@ucl.ac.uk, j.donnnett@ucl.ac.uk, j.okeefe@ucl.ac.uk

7/4/97

Abstract

A robotic platform controlled by a neuronal simulation is presented for studying the mechanisms underlying the firing of hippocampal place cells and how they support navigation. The sensors on the robot provide visual, odometric and short-range proximity information which are combined to estimate the distance of the walls of the enclosure from the robot and its current heading direction. These inputs drive a neuronal simulation of the hippocampus in which the robot's location is represented as the firing of place cells. If a goal location is encountered, connection strengths are incremented between active place cells and a set of 'goal cells' which guide subsequent navigation (Burgess et al., 1994), allowing the robot to return to goal locations. The system shows good agreement with actual place cell firing, and makes predictions regarding the firing of cells in the subiculum and the locus of search of rats after deformation of their environment.

1 Introduction

Recordings from single neurons in the hippocampus and neighbouring brain regions of freely moving rats provide an insight into the neuronal basis of mammalian navigation. Understanding the working of these neurons offers the tantalising prospect of developing control algorithms that directly emulate mammalian navigational abilities. The firing rates of *place cells* (PCs) in regions CA3 and CA1 of the hippocampus mostly resemble tuning curve functions of the rat's location in an environment, and contain sufficient information to localise the rat (Wilson & McNaughton, 1993). The firing rates of *head-direction cells* (HDCs) in the postsubiculum (Taube et al., 1990), and elsewhere, resemble tuning curve functions of the direction of the rat's head, regardless of its location within the environment. Further evidence relating the hippocampal region to navigation and topographical memory comes from the effects of lesions in rats (see e.g. Morris et al., 1982; O'Keefe & Nadel, 1978) and neuropsychological (Habib & Sirigu, 1987; Maguire et al., 1996a) and functional imaging (Maguire et al., 1996b; Maguire et al., 1997) studies in humans.

Recent evidence has begun to reveal the nature of the signals underlying the firing of these spatially tuned cells. 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 of space. However objects placed within the environment do not show this control (Cressant et al., 1996). Interoceptive inputs relating to self-motion (possibly vestibular) also influence the overall orientation of place (Sharp et al., 1995; Jeffery et al., 1997) and head-direction (Blair & Sharp, 1996) representations. In uniform rectangular walled environments place cells' receptive fields ('place fields') 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).

These remarkable properties of cell firing are derived from relatively unsophisticated sensory inputs. Rats have wide angle vision (320 – 360° 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 (e.g. 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 integrate’) is also limited. For example, hamsters err significantly in returning to the start location after an L shaped route of 1m per side or after 5 active rotations or 2 passive rotations in the dark, see Etienne et al., 1996. Thus, while useful for maintaining a sense of direction (as in the model presented here), path integration is clearly not sufficient to support the firing of place cells over long periods of motion. 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 polarising extra-maze stimuli (O’Keefe & Speakman, 1987). Uncontrolled auditory, olfactory and somatosensory cues may contribute to these findings.

We present a computational model of the hippocampal system implemented on a mobile robot to ensure the realism of the assumed sensory inputs and to enable true evaluation of the model’s navigational ability. The behaviour of neurons in the model is compared to single unit recordings from the rat hippocampus. Navigation is driven by place cell firing in a manner similar to that in Burgess et al. (1994), and is compared to data showing that rats can return to an unmarked goal location from novel starting positions (but not after hippocampal lesions, Morris et al., 1982) and that gerbils search in two separate locations after two cues indicating a single reward site are pulled apart (Collett et al., 1986).

2 The model

The 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 Fig. 2. Walls are identified on the basis of their allocentric direction from the rat, solving the ‘binding’ or ‘correspondence’ problem of how information regarding a particular wall is channeled to a particular set of cells as the robot moves about. The estimate of the allocentric direction (orientation) of the rat is maintained by odometry and sightings of the North wall which is visually distinct from the other walls and serves to polarise 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).

2.1 Physical implementation

A Khepera miniature robot with on-board video and a ring of short-range infra-red proximity detectors provides artificial visual and haptic information. Two independently driven wheels (with shaft-encoders) allow movement around a rectangular environment formed by white walls and a dark floor. 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 Fig. 1B. 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 4cm. Their function might be compared to that of a rat’s whiskers.

Movement is controlled by setting the speeds of the wheel motors and occasionally monitoring the shaft-encoders (this is not precise so that 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, the robot rotates on the spot to face in the estimated orientations N S E W 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 rat moves 3cm forward in the desired direction. If a wall is perceived the robot moves 3cm away from the wall, whose direction is estimated from the relative values of the proximity detectors. Each step corresponds to 0.1s (one theta cycle, see below) implying a speed of 30cm/s for the rat, but actually takes around 3s (processing on a SUN Ultra). 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.

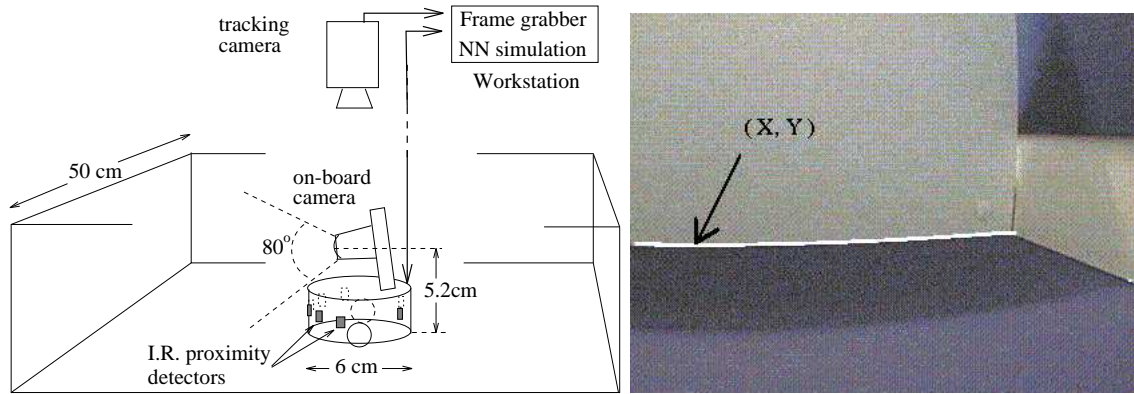


Figure 1: Left: hardware setup. Right: robot's eye view, detected horizontal dark-light edge points are shown in white, and a black arrow marks the (x, y) position in the image that is returned by the visual processing. The 'North' wall can be identified (on the right) by its dark upper half.

2.2 The neural network

The visual inputs to the hippocampus are represented by a rectangular array of cells organised 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 Fig. 2. The tuning of these 'sensory' cells follows the form of the independent place field components identified in O'Keefe & 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 (estimated visually), d_i is the distance at which the cell responds maximally, 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 would 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 in O'Keefe & Burgess, 1996, since the environments used here are about half the size).

All connections in the model take value 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 2 sensory cells related to 2 orthogonal walls, see Figure 2. 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 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% of PCs with the greatest input are active, and connections between maximally active ECs and the 4 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 walls 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

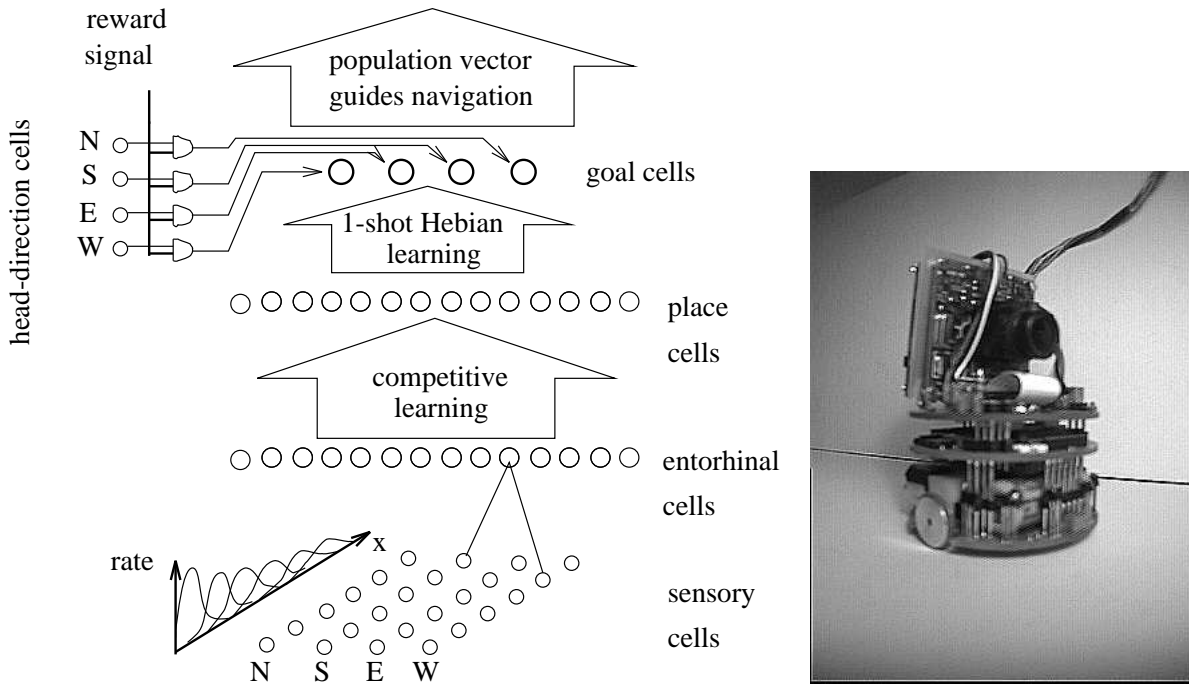


Figure 2: Left: schematic diagram of the neuronal simulation. There are 60 sensory cells, 900 ECs, 900 PCs and 4 goal cells. Inputs from the sensors on the robot drive the firing of the sensory cells. Activation propagates through the model to form a representation in space in the place cell layer. Learning 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. Right: the Khepera mobile robot.

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 3 A and B.

In fact we use a more complicated model of learning of the goal location 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’ (Georgopolous 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 e.g. 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 relationship to the phase of the theta rhythm of the EEG (O’Keefe & Recce, 1993; the theta rhythm is a large 7-10Hz oscillation in the extra-cellular potential recorded throughout the hippocampus). An implication of this phase relationship is that PCs active at a ‘late’ phase tend to have place fields that are centered ahead of the rat while those firing at an ‘early’ phase tend to have place fields centered behind the rat (Burgess et al., 1994). In the simulation, each time-step (0.1s) is divided into two intervals corresponding to the early and late phases of a 10Hz 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 for both the rat’s head-direction and for the sensory attributes (e.g., food, water, etc.) of a particular goal, see Figure 2. These connections deliver a ‘learn now’ type reinforcement signal to a goal cell whenever the rat is at the appropriate goal location and facing in the appropriate direction. If this learning signal arrives at a ‘late’ phase of the EEG, or if synaptic plasticity is restricted to this phase (see e.g. Pavlides et al., 1988, for the relationship of long-




Figure 3: 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 LTP when a goal cell is excited by the attributes of a particular goal location. Filled circles are active place cells; empty circles are inactive place cells and filled squares mark 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 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.




Figure 4: (left) Exploration of a square environment. Dashes show the actual path of the robot, each dash representing 0.1s of simulated time (or about 3s 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-localisation. (right) 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.




Figure 5: Navigation in a square environment to a goal location (marked by ‘O’) before (left) and immediately following expansion of the environment into a long rectangle (right).

term potentiation of synapses to the theta rhythm), then the goal cell associated with, e.g. the direction North, will form active connections from place cells with receptive fields centered to the North of the location of the goal. 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 (e.g. 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 3C). The firing of these goal cells drives navigation of the robot, enabling it to return to a previously visited but unmarked goal location.

3 Performance

The robot was tested in two rectangular environments of size $50 \times 50\text{cm}$ and $50 \times 75\text{cm}$. Its movements were tracked by an overhead camera and tracking system that detected two LEDs on the robot. Figure 4 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-localisation and maintenance of sense of direction. Figure 4 also shows the performance of the robot in returning to an unmarked reward location having visited it once previously. The robot also shows generalisation in returning to the goal from novel starting locations.

Figure 5 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 6. This compares well with observed data in which the most common pattern was for place fields to maintain fixed positions relative to a wall, whilst 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 effectively stretches out the goal cell representation




Figure 6: 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 following 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.

along that axis, but still results in a unimodal search pattern located between the loci indicated by fixed distances from each of the walls.

3.1 Discussion

This model has developed from the work in 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 of Redish & Touretzky (1996). However their aim is more towards integrating suggested functional roles for many brain regions, whereas ours is more directed at the details of the neuronal implementation. For example we concentrate on the origins of the sensory inputs underlying place cell firing and on how the firing of place cell could drive navigation (i.e. specifying how a vector voting hypothesis could be implemented rather than simply postulating its existence).

We note 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). We suppose 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 response of ECs and PCs to changes in environmental shape (Quirk et al., 1992), indicating that this 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 and results in a unimodal locus of search midway between the locations corresponding to maintaining a 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). A second prediction concerns the existence of the 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.

Acknowledgements

NB is supported by a Royal Society University Research Fellowship, JOK and JD are supported by a programme grant from the MRC.

References

- Blair HT, Sharp PE (1996) 'Visual and vestibular influences on head-direction cells in the anterior thalamus of the rat', *Behavioral Neuroscience* **110**, 643-660.
- Benhamou S, Seguinot V (1995) 'How to find ones way in the labyrinth of path integration models', *J. Theo. Biol.* **174**, 463-466.
- Burgess N, O'Keefe J (1995) 'Modelling Spatial Navigation by the Rat Hippocampus', *Int. J. of Neural Systems* **7** (supplementary issue) 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* **5** Eds: Hanson S J, Giles C L & Cowan J D, 929-936, Morgan Kaufmann, San Mateo, CA.
- Burgess N, Recce M, O'Keefe J (1994) 'A model of hippocampal function', *Neural Networks* **7**, 1065-1081.
- Collett TS, Cartwright BA, Smith BA (1986) 'Landmark learning and visuo-spatial memories in gerbils', *J. Comp. Physiol. A* **158**, 835-851.
- Cressant A, Poucet B, Muller RU (1996) 'The spatial arrangement of 3-dimensional objects inside the arena influences their control over the positional firing of place cells' *Soc. Neurosci. Abstr.* **22**, 911.
- Dean P (1990) 'Sensory cortex: visual perceptual functions' *The Cerebral Cortex of the Rat* B Kolb & RC Tees (eds). MIT Press: Cambridge, MA.
- 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.
- 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.
- Jeffery K, Donnett JG, Burgess N, O'Keefe J (1997) 'Directional control of hippocampal place fields', *Exp. Brain Res.* submitted.
- Maguire EA, Burke T, Phillips J, Staunton H (1996a) 'Topographical disorientation following unilateral temporal lobe lesions in humans' *Neuropsychologia* **34**, 993-1001.
- Maguire EA, Frackowiak RSJ, Frith CD (1996b) 'Learning to find your way - a role for the human hippocampal region' *Proc. Royal Society B* **263**, 1745-1750.
- Maguire EA, Burgess N, Donnett JG, O'Keefe J, Frith CD (1997) 'Knowing where things are: Parahippocampal involvement in the encoding of object location in large-scale space', *J. Cog. Neurosci.*, in press.

- McNaughton BL, Nadel L (1990) 'Hebb-Marr networks and the neurobiological representation of action in space', in: *Neuroscience and Connectionist Theory* MA Gluck & DE Rumelhart (eds.), Hillsdale NJ: Lawrence Erlbaum Assoc. pp 1-63.
- 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, Burgess N (1996) 'Geometric Determinants of the Place Fields of Hippocampal Neurons' *Nature* **381** 425-428.
- O'Keefe J, Nadel L (1978) *The Hippocampus as a Cognitive Map* UK: Oxford University Press.
- O'Keefe J, Speakman A (1987) 'Single unit activity in the rat hippocampus during a spatial memory task' *Exp. Brain Res.* **68**, 1-27.
- Pavrides C, Greenstein YJ, Grudman M, Winson J (1988) 'Long-term potentiation in the dentate gyrus is induced preferentially on the positive phase of θ -rhythm', *Brain Res.* **439** 383-387.
- 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. Neuroscience* **12** 1945-1963.
- Redish AD, Touretzky, DS (1997) 'Navigating with landmarks: computing goal locations from place codes', in *Symbolic Visual Learning* (eds. Ikeuchi, K. & Veloso, M.) (Oxford University Press, in the press).
- Sharp PE, Blair HT, Etkin D, Tzanetos DB (1995) 'Influences of vestibular and visual-motion information on the spatial firing patterns of hippocampal place cells' *J. Neurosci.* **15**, 173-189.
- 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.
- Tolman EC (1948) 'Cognitive maps in rats and men', *Psychol. Rev.* **55**, 189-208.
- Wilson MA, McNaughton BL (1993) 'Dynamics of the hippocampal ensemble code for space', *Science* **261** 1055-1058.