the monkey produce long-lasting memory impairment in the visual and tactual modalities, J. Neuroscience, 13: 2430-2451.

* Traub, R.D., Miles, R., Muller, R.U. and Gulyas, A.I., 1992, Functional organization of the hippocampal CA3 region: implications for epilepsy, brain waves and spatial behavior, Network, 3: 465-488.

Treves, A. and Rolls, E.T., 1992, Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network, Hippocampus, 2: 189-200.

et al. (1994). Finally, a model should be able to distinguish between many different environments, and 'remember' where things are in each. There is evidence that distinct place cell representations are formed to code for different environments. Incorporating this into a model leads us back to consideration of the role of the DG and the CA3 recurrent collaterals in providing an associative memory (see Marr, 1971; McNaughton and Nadel, 1990; Treves and Rolls, 1992; see also ASSOCIATIVE MEMORIES). Exciting future possibilities include application of these models to related tasks, allowing consideration of experiments involving monkeys and humans, and to the navigation of autonomous mobile robots.

Acknowledgements:

Thanks to Kate Jeffery and Jim Donnett for carefully reading this manuscript. NB was supported by a Joint Councils Initiative on Human Computer Interaction, and as a Royal Society University Research Fellow.

References

Barnes, C.A., McNaughton, B.L., Mizumori, S.J.Y., Leonard B.W. and Lin, L.-H., 1990, Comparison of spatial and temporal characteristics of neuronal activity in sequential stages of hippocampal processing, Progress in Brain Research, 83: 287-300.

Barto, A.G. and Sutton, R.S., 1981, Landmark learning: an illustration of associative search, Biological Cybernetics, 42: 1-8.

Burgess, N., Recce, M. and O'Keefe, J., 1994, A model of hippocampal function, *Neural Networks*, 7, in press.

Dayan, P., 1990, Navigating through temporal difference, Neural Information Processing Systems 3, San Mateo, CA: Morgan Kaufmann, pp 464-470.

Georgopoulos, A.P., Kettner. R.E. and 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. Neuroscience, 8: 2928-2937.

Marr, D., 1971, Simple memory: a theory for archicortex, Phil. Trans. of the Royal Society B, 262: 23-81.

McNaughton, B.L., Knierim, J.J. and Wilson, M.A., 1994, Vector encoding and the vestibular foundations of spatial cognition: a neurophysiological and computational hypothesis, in: The Cognitive Neurosciences, (M. Gazzaniga, Ed.) Boston: MIT Press.

* McNaughton, B.L. and Nadel, L., 1990, Hebb-Marr networks and the neurobiological representation of action in space, in: Neuroscience and Connectionist Theory, (M.A. Gluck and D.E. Rumelhart, Eds.), Hillsdale NJ: Lawrence Erlbaum Assoc. pp. 1-63.

* Muller, R.U., Kubie, J.L., Bostock, E.M., Taube, J.S. and Quirk, G.J., 1991, Spatial firing correlates of neurons in the hippocampal formation of freely moving rats, in: Brain and Space, (J. Paillard, Ed.), Oxford University Press pp. 296-333.

* O'Keefe, J., 1991, The hippocampal cognitive map and navigational strategies, in: Brain and Space, (J. Paillard, Ed.), Oxford University Press pp. 273-295.

* O'Keefe, J. and Nadel, L., 1978, The Hippocampus as a Cognitive Map, Oxford: Clarendon Press.

O'Keefe, J. and Recce, M., 1993, Phase relationship between hippocampal place units and the EEG theta rhythm, Hippocampus 3: 317-330.

Sharp, P.E., 1991, Computer simulation of hippocampal place cells, Psychobiology, 19: 103-115.

Suzuki, W.A., Zola-Morgan, S., Squire, L.R. and Amaral, D.G., 1993, Lesions of the perirhinal and parahippocampal cortices in

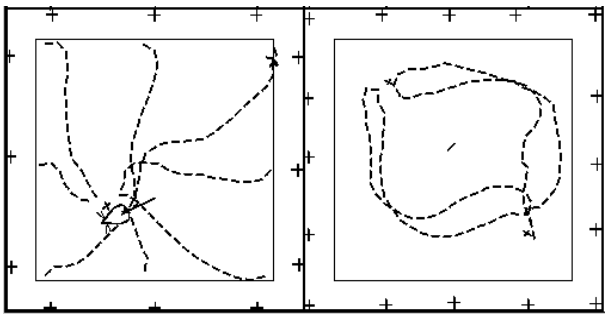

Figure 2: Left: simulated trajectories from 8 novel starting positions, to a goal encountered after 30 seconds of exploration (at $60cm/s$) in a $135 \times 135cm^2$ environment; the rat is shown to scale. Right: simulated navigation between two goals with an 'obstacle' in between, cues are marked by '+', goals by '×', and obstacles by '/'. Adapted from (Burgess et al., 1994).

in a cycle have place fields centered ahead of the rat. This is achieved, somewhat arbitrarily, by making the phase of firing of ECs depend on the angle between the rat's heading direction and the direction of the centroid of the corresponding pair of cues (if the centroid is ahead the cell fires at a late phase, if behind it fires early). This property propagates throughout the PC and SC layers. Less severe competition amongst SCs than PCs leads to larger firing fields, which avoids the locality of information access problem (so that goal cell firing fields cover the whole environment).

The rat's direction of motion is updated at the end of each θ cycle; during 'exploration' it is simply a random variable within 30° of the previous direction, during 'searching' it is determined by the goal cells. A set of 8 goal cells codes for each new goal location that is encountered (representing north, northeast, east, etc). When the rat is at the goal the goal cell with preferred direction closest to the rat's heading direction (given by head-direction cells) receives a strong input, allowing connections to it to be switched on. Information from the goal is only ever received at the goal location (cf. Barto and Sutton, 1981). This signal arrives at a late phase of the θ rhythm, and connections are switched on from those SCs active at that time (which tend to have firing rate maps that are peaked ahead of the rat). When a goal is encountered the rat looks around in all directions so that connections are switched on to all 8 goal cells.

The direction and proximity (represented by the net firing of the group of goal cells) of interesting objects is the output of the hippocampal 'map', and allows the simulated rat to navigate. A small number of obstacles can be avoided during navigation by subtracting the population vector of cells coding for an obstacle from the population vector for the goal. The model assumes that the role of the dentate gyrus (DG) and CA3 recurrent collateral system is to cause different subsets of PCs to be active in different environments (e.g. via mossy fiber 'detonator' synapses in CA3 and an associative network, see McNaughton and Nadel, 1990; Treves and Rolls, 1992). Advantages of this model over those discussed above are that reasonable trajectories, including short-cuts, are performed after one visit to the goal following brief exploration (during which the goal need not be present), i.e. the model learns very quickly, and shows latent learning. Future work includes modelling the discrimination of many different environments.

5 DISCUSSION

There has been much development of spatial models of hippocampal processing in recent years. We are now close to being able to construct a simple neuronal model of the hippocampus that takes in sensory information from environmental cues, turns it into a place cell representation of space, and uses that spatial representation to create a representation of the direction in which interesting things are located, relative to it, as it moves about an environment. Such a model would allow systematic examination of the role of individual cells in the hippocampus in terms of their effect on (spatial) behavior, and provide direction for experiments.

Section 4 describes a possible starting point for this type of model, following the ideas of Sharp (1990) and O'Keefe (1991). Inertial/directional navigation systems should also be integrated with the hippocampal place system; as in McNaughton

fields can be rotated by moving the cue between sessions. However this is not the case if (from the very first session) the rat is disoriented between each session.

It is also proposed that each PC codes for a single cue being at a particular angle (provided by the H system) and distance. The rat could then navigate by comparing the current vectors $\vec{p}_i(t)$ to individual cues with the memory $\vec{m}_i$ of their respective values from the goal location. The vector back to the goal is: $\vec{p}_i(t) - \vec{m}_i$ for any cue i. This has the advantage that navigation would be possible using only one cue, but is hard to reconcile with the robustness of place cell firing to removal of individual cues.

3.4 Centroid Model

O'Keefe (1991) proposed a navigational mechanism in which environments are characterized by two parameters: the centroid and the slope of the positions of environmental cues, which can be used as the origin and 0° direction of a polar coordinate framework. The firing of a PC could represent the average position of a small number of cues (their mini-centroid), while head direction cells could represent the translation vector between pairs of cues. The environmental centroid and slope are then found by averaging the mini-centroids and slopes. It was proposed that single cells could represent two-dimensional vectors as phasor: taking the θ rhythm of the EEG as a clock cycle, the amplitude of firing would code for proximity, and the phase of firing within a clock cycle would code for angle. In this formalism, summing the output of several neurons results in vector addition, and subtraction is equivalent to a phase inversion followed by addition. Thus, the summed PC activity could provide a vector $\vec{v}(t)$ continually pointing to the centroid of the environment, so that, if the 'goal' was encountered at time t_g, storing $\vec{v}(t_g)$ (outside the hippocampus) would enable the translation vector $\vec{v}(t) - \vec{v}(t_g)$ from the rat to the goal to be calculated whenever the rat wanted to go to the goal.

There are some points in favor of this model: (i) it provides a powerful alternative set of mechanisms to those usually considered by 'connectionists'; (ii) it gives a plausible interpretation of PC firing; (iii) if a PC's phase of firing with respect to θ was given by $360° - 2\alpha$, where α is the angle between the rat's direction of travel and the direction of the mini-centroid, then the model would be consistent with the phase shift data in O'Keefe & Recce, 1993; (iv) the rat would be able to take short-cuts. Disadvantages include the sensitivity of a unique reference direction to movements or occlusion of cues.

4 PLACE CELL FIRING & NAVIGATION

The population vector model (Burgess et al., 1994) is implemented at the neuronal level but also generates actual movement trajectories for the simulated rat, and aims to surmount some of the difficulties discussed above. The input to the model is a set of sensory cells with tuning curve responses to the distance of cues from the rat. The output is able to guide the rat back to a goal that has been visited one or more times in a single environment (see Fig. 2). The output representation chosen is analogous to that found in primate motor cortex (see Georgopoulos et al., 1988): a 'population vector' (i.e. the vector sum of cells' preferred directions weighted by their firing rates). In the model it is instantiated by a set of 'goal' cells such that each cell has a conical firing rate map whose peak is displaced from the goal position in a particular absolute direction (the 'preferred direction' for that cell). The intermediate layers, entorhinal cells (ECs), PCs and subicular cells (SCs) map the sensory input (from 4 to 16 intramaze and extramaze cues) to the population vector output. A type of competitive learning governs the dynamics of the PCs and SCs, similar to that in Sharp (1991).

Each EC receives connections from a pair of sensory cells, although not all possible pairs are represented. The choice of cues and distances at which the cues in a pair are represented results in ECs with large spatial firing fields that tend to be centered near to the centroid of the pair of cues. The model relies on the phase firing of SCs (relative to θ) being such that those firing late

navigation to the action of individual cells in the hippocampus.

3.1 Graph Model

A role in navigation was proposed for the CA3 recurrent collaterals (the axonal projections by which each CA3 PC contacts approximately 5% of the other CA3 PCs), by Muller, Kubie and Saypoff (see Traub et al., 1992). Given a model of LTP in which pre- and post-synaptic firing within a short time interval leads to a small increase in synaptic 'strength', then it is possible to construct a model of CA3 in which PC firing as the rat moves around leads to the synaptic strength of a connection between two PCs depending on the proximity of their place fields. If two PCs have closely overlapping fields they will both fire near-coincidentally relatively often, whereas if they have well separated place fields they will very rarely fire simultaneously.

After brief exploration the synaptic strengths represent distances along the paths taken by the rat: however, to build up a true distance metric would take a long time, as with reinforcement learning approaches. Another problem is that of accessing the distance information (i.e. getting projection cells to fire so that information could be passed out of region CA3); it also falls short of being a model of navigation in that it does not provide a direction to move in.

3.2 Local View Model

McNaughton (1989) proposed that the hippocampus functions as an associative memory as follows. As the rat explores, it learns to associate each local view and movement made with the local view from the place visited as a result of the movement. Thus routes through an environment can be stored as a chain of local view/movement associations, see McNaughton and Nadel (1990). The model is supported by the fact that, in some situations, PC firing depends on the rat's direction (and therefore its 'local view'). Some major problems with this theory are (i) simple route-following strategies appear to be the kind of navigation of which hippocampectomized rats are capable (see O'Keefe and Nadel, 1978); (ii) it is difficult to know which particular route will lead to the desired goal: solving this problem leads one back to the reinforcement learning approach; (iii) the model is not capable of more sophisticated navigation such as taking short-cuts; (iv) in open fields, PC firing does not seem to depend on direction.

3.3 Inertial/Directional Systems

The existence of a directional system in rodents is indicated by the head-direction cells and by experiments in which (i) a goal location seems to be stored in terms of a distance and absolute direction from one symmetrical cue, and (ii) inertial homing (i.e. an internal sense of direction) is used for navigation rather than external cues (e.g. in complete darkness). It is also interesting to note that, once the rat has got its orientation from the array of cues, PCs can continue to fire in the correct places after all the salient cues in an environment have been removed, or after the lights have been switched off.

Various aspects of this system have been modelled (see McNaughton et al., 1994). A model of 'dead-reckoning' describes how the rodent's current heading direction could be calculated as it moves about. Briefly, it is proposed that the combinations of all heading directions ($H(t)$) and all possible angular displacements ($\Delta H(t)$) over a short time step (τ) are associated with the resulting heading directions ($H(t+\tau)$) by an associative neural network. Evidence for this includes cells (in posterior parietal cortex) with responses that depend both on the rat's orientation ($H(t)$) and on the sense of its angular velocity ($\Delta H(t)/\tau$) .

The hippocampus is invoked to correct cumulative errors in direction by associating the 'local views' at each location with the appropriate heading direction (H). However, they suggest that if the local view – H associations cannot be learned, then the heading direction system is the fundamental determinant of the orientation of the PCs' spatial map. Evidence comes from experiments in a symmetrical environment with one polarizing cue: in normal circumstances when the rat is familiar with the environment the position of place

tion of any other cue. Unfortunately there are no reports to date of the existence of the postulated destination cells.

1.2 Cognitive Maps

Cognitive maps were first introduced by Tolman to explain place learning in rats, including e.g. their ability to take short-cuts (see COGNITIVE MAPS). An alternative view, suggested by Hull, is that navigation is achieved by following a list of stimulus-response-stimulus steps. O'Keefe and Nadel (1978) proposed that independent neural systems exist in the brain to support a 'taxon' system for route navigation and a 'locale' system for map-based navigation (for a current synopsis see O'Keefe, 1991). The 'map' was taken to be a Euclidean description of the environment in a coordinate system based on the world and not on some part of the animal's body. They proposed that the locale navigation system resides in the hippocampus, based on: (i) the firing properties of hippocampal PCs, (ii) the presence of θ rhythm during displacement movement, (iii) spatial tasks, including the Morris water maze and the Olton 8-arm maze, showing clear deficits in performance after hippocampal lesions, and (iv) the interpretation of the amnesic syndrome as the loss of episodic memory (memory for specific events set in a spatio-temporal context).

2 PLACE CELL FIRING

In this section we describe a model of how the spatial firing of place cells might develop as the rat explores an environment. Following an earlier, mathematical, model of PC firing (by Zipser in 1985), Sharp (1991) used a simple network with an input layer and two layers of cells governed by 'competitive learning' dynamics (a simple model of the effect of Hebbian synaptic increment combined with inhibitory interneurons, see COMPETITIVE LEARNING). In this model there are two types of input (or 'cortical') cell that respond to cues placed around the environment: a type 1 cell 'fires' whenever a particular cue is at a given distance from the rat; type 2 cells also respond to a particular cue being at a given distance, but only if the cue is within a certain range of angle from the rat's head.

As the simulated rat moves around its environment, competitive learning in the synaptic weights (initially random) leads to unsupervised clustering of the input vectors: a PC learns to fire in a portion of the environment in which the inputs (i.e. the distances and angles of cues) are similar. Cells in the same competitive group will fire in different places, cells from different groups can have overlapping firing fields. The model generates realistic-looking place fields after 64 minutes of simulated exploration. Interestingly, if the simulated rat's exploration is restricted to movements consistent with being on an 8-arm maze, then PC firing tends to be much more strongly correlated with the orientation of the rat (as well as its location) than in the case of unrestricted exploration, which fits well with the experimental data.

3 NAVIGATION

In this section we look at models addressing how the hippocampus could be used to enable navigation. The simplest map-based strategies are based on defining a surface over the whole environment, on which gradient ascent leads to the goal, see REINFORCEMENT LEARNING (see also Barto & Sutton [1981] for a simple related mechanism which learns to navigate given that it can 'smell' how far away the goal is). These tend to have the problem that, to build up the surface, the goal must be reached very many times, from different points in the environment. A new surface must be computed if the goal is moved, and multiple goals, as in the 8-arm maze, cannot be handled. Learning in these models is much slower and more goal dependent than in rats, and they are unable to perform 'latent learning' (e.g. in rats, exploration in the absence of any goals improves subsequent navigation). Interestingly, the performance of these models improves somewhat when a spatially diffuse representation (like place fields), rather than a punctate representation, is built up during exploration (Dayan, 1990). Below we describe some recent models that have related

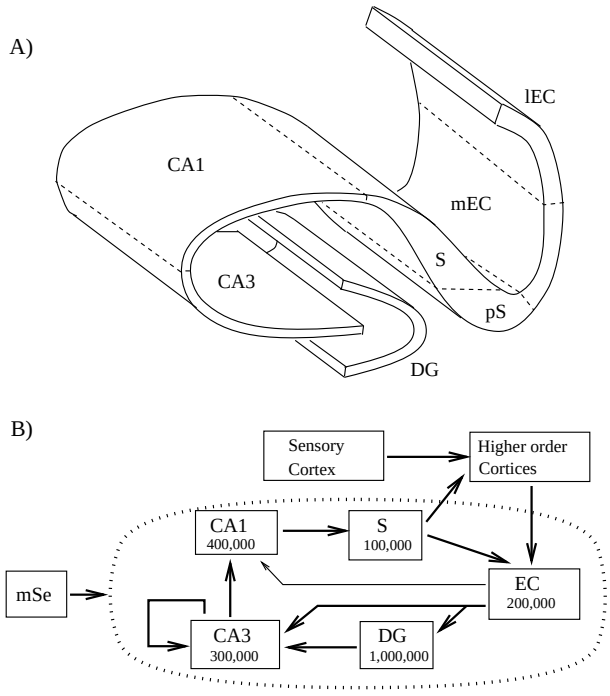

Figure 1: The hippocampus is formed from sheets of cells. A) A schematic section cut perpendicular to the longitudinal axis of the hippocampus. EC–entorhinal cortex (mEC–medial EC, lEC–lateral EC); S–subiculum; pS–pre- and para-subiculum; DG–dentate gyrus. B) The major projections between subfields (mSe–medial septum), and approximate numbers for the major cell type in each subfield (i.e. pyramidal cells, except for the DG in which it is granule cells) in the rat. The human hippocampus contains one order of magnitude more cells. In the DG-CA3 projection a single mossy fiber projects from each granule cell, making very large synapses onto only 14 or so pyramidal cells. All the other projections have large divergence and convergence (many thousands to one), and involve the type of synapse in which 'Hebbian' LTP has been observed. A variety of interneurons provide feed-forward and feed-back inhibition. Cells in the mSe project into DG, CA3 and (less strongly) CA1, playing a role in producing the θ rhythm of the hippocampal EEG. Cells in CA3 and CA1 also project out to the lateral septum via the fornix. Adapted from B.L. McNaughton, 1989, in: Neural Connections, Mental Computation (L. Nadel, L.A. Cooper, P. Culicover and R.M. Harnish, Eds.), MIT Press.

of place cells fire in different environments. The firing properties of PCs can be manipulated by changing the rat's environment: rotating the major cues in an environment can cause the place fields to rotate; increasing the environment's size can cause some of the place fields to expand correspondingly. In environments in which direction of movement is restricted (e.g. mazes with narrow arms) PC firing rates appear to depend on the rat's direction of travel as well as its location.

Cells in the entorhinal cortex (the main cortical input to the hippocampus, see Fig. 1) also have spatially correlated firing, but tend to have larger, less well defined place fields than those in CA3 or CA1. Cells whose primary behavioral correlate is 'head direction' have also been found, in the (dorsal) pre-subiculum (see Fig. 1) They fire when the rat points its head in a specific direction relative to the cues in the environment, and independently of its location. See Muller et al. (1991) and Barnes et al. (1990) for reviews of these data.

The electroencephalogram (EEG) recorded in the hippocampus is the largest electrical signal in the brain. One form of the EEG, called the theta (θ) rhythm, is a sinusoidal oscillation of 7-12 Hz. O'Keefe & Nadel (1978) suggested that in the rat the θ rhythm coincides with displacement movements, e.g. walking. Recently, PC firing has been found to have a systematic phase relationship to θ (O'Keefe & Recce, 1993): when a rat on a linear track runs through a place field, the PC tends to fire groups of spikes, with each successive group occurring at an earlier phase of the θ cycle. Consistent with these data, initial analysis of the phase of firing of PCs as a rat runs in an open field environment indicates that cells firing at a late phase tend to have place fields centered ahead of the rat, whereas those firing at an early phase tend to have place fields centered behind the rat (Burgess et al., 1994).

There are two features of PC firing that raise immediate problems for their use in navigation: (i) information about a place in an environment (i.e. the firing of the corresponding PCs) can only be accessed locally (by actually visiting that place); (ii) place fields appear to be no more affected by the location of the goal (which is obviously essential for navigation) than by the loca-

To be published in: The Handbook of Brain Theory and Neural Networks, M. A. Arbib (Ed.), Bradford Books/ MIT Press, 1995. ©MIT press.

HIPPOCAMPUS – SPATIAL MODELS

Neil Burgess, Michael Recce & John O'Keefe
Dept. of Anatomy, University College London,
London WC1E 6BT, U.K.
e-mail: n.burgess@ucl.ac.uk

1 INTRODUCTION

The hippocampus is one of the most studied areas of the brain. It has attracted interest due to its position many synapses removed from sensory transducers or motor-effectors, the putative role of hippocampal damage in human amnesia, the discovery of long term potentiation (LTP, see SYNAPTIC PLASTICITY) and the discovery that cell firing is spatially coded. It has also been studied for its role in neurological disorders including epilepsy, schizophrenia and Alzheimer's disease. Bilateral damage to the hippocampus and nearby structures in patient H.M., as treatment for epilepsy, produced a profound retrograde and anterograde amnesia, prompting extensive cross-species research to uncover the specific memory deficit that results from hippocampal damage (the most prominent of which, in the rat, appears to be a deficit in spatial navigation).

In short, the hippocampus has become the primary region in the mammalian brain for the study of the synaptic basis of memory and learning. Structurally, it is the simplest form of cortex (Figure 1). It contains one projection cell type, confined to a single layer (compared with the large number of cell types and six principal layers of the neocortex), and receives inputs from all sensory systems and association areas. See Fig. 1 for more details.

Attempts to model the hippocampus differ both in the level of anatomical detail and in the functionality which they seek to reproduce. Marr (1971) proposed a theory for how the hippocampus could function as an associative memory, from which have grown many extensions, usually focussing on the role of the CA3 recurrent collaterals (see Fig. 1), see ASSOCIATIVE MEMORIES, McNaughton & Nadel (1990) and Treves & Rolls (1992). However, the discovery that much of the deficit in short term memory tasks for monkeys is due to lesion of the nearby perirhinal cortex (see Suzuki et al., 1993) is prompting a revision of the idea of the hippocampus as a general store for declarative memories. Here we specifically consider neuronal models of spatial processing in the rat hippocampus: the domain in which the least controversial experimental data are available. Models at this level, which are based on repeatable data, are the most likely to provide insight into hippocampal function with respect to higher levels of description and wider functionality. Below we introduce data on the spatial correlates of hippocampal cell firing and the seminal idea of the hippocampus as a spatial map, and describe some models of hippocampal place cells and navigation.

1.1 Single Unit Recordings

Single unit recordings in freely moving rats have revealed place cells (PCs) in fields CA3 and CA1 of the hippocampus whose firing is restricted to small portions of the rat's environment (the corresponding 'place fields'). More recently PCs have been identified in monkeys. There appears to be little topographic organization of PCs relating their positions in CA3 or CA1 to the positions of their firing fields, and different subsets