

# HIPPOCAMPUS, SPATIAL MODELS

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

RUNNING HEAD: Spatial Models of the Hippocampus

Correspondence: Neil Burgess

Phone: +44 (0)20 7679 1147

Fax: +44 (0)20 7679 1145

# 1 INTRODUCTION

The hippocampus is the most-studied part of the brain, attracting interest due to its position many synapses removed from sensory transducers or motor-effectors, its role in human amnesia and Alzheimer's disease, and the discovery of long term potentiation (LTP, see SYNAPTIC PLASTICITY) and of spatially coded cell firing. 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. It contains one projection cell type, whose cell bodies are confined to a single layer, and receives inputs from all sensory systems and association areas. See Fig. 1 for more details.

-----  
Figure 1 about here  
-----

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; ASSOCIATIVE MEMORIES McNaughton & Nadel, 1990; and many others, for collected works see e.g. Gluck, 1996; Burgess et al., 1999 and Burgess et al., 2001 for a review). However the precise contribution of the hippocampus to memory, as opposed to the contribution of the nearby structures, remains controversial. Here we specifically consider neuronal models of spatial processing in the rat hippocampus: the domain in which the least controversial experimental data are available. Below we introduce data on the spatial correlates of hippocampal cell firing, and the idea of the hippocampus as a spatial map, and describe some models of hippocampal place cells and

navigation.

## 1.1 Electrophysiology

Single unit recordings in freely moving rats have revealed 'place cells' (PCs) in fields CA3 and CA1 of the hippocampus, so called because their firing is restricted to small portions of the rat's environment (the corresponding 'place fields'). There is little topographic organization of PCs relating their positions in CA3 or CA1 to the positions of their firing fields. 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. 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), anterior thalamus and mammillary bodies. 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 HEAD-DIRECTION SYSTEM, and Zhang (1996) and Sharp, Blair and Brown (in Gluck, 1996) for models.

The electroencephalogram (EEG) recorded in the hippocampus is the largest electrical signal in the rat brain. One form of the EEG, called the theta ( $\theta$ ) rhythm, is an oscillation of 7-12 *Hz*. O'Keefe & Nadel (1978) suggested that in the rat the  $\theta$  rhythm coincides with displacement movements. PC firing has been found to have a systematic phase relationship to  $\theta$  discovered by O'Keefe and Recce in 1993 (see Burgess and O'Keefe in Gluck, 1996): 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  $\theta$  cycle. Consistent with these data, PCs 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 in open field environments,

see Burgess and O'Keefe (in Gluck, 1996).

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 location of any other cue. Unfortunately there are no reports to date of cells that code for the destination of a rat's current trajectory.

## **1.2 Path Integration**

An animal may estimate its current position relative to some starting position purely on the basis of internal signals (e.g. vestibular or proprioceptive) relating to its movements in the intervening period. Such a process is often referred to as 'path integration' (PI). Many animals appear to be able to use PI to return to a home location in the absence of external stimuli. Interestingly, once a 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 (see also the role of PI in the HEAD-DIRECTION SYSTEM). Experiments indicate that PC firing can be supported by any one of visual, auditory, olfactory, tactile or internal information, as available.

## **1.3 Cognitive Maps**

Cognitive maps, meaning mental representations of the spatial layout of an environment, 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 synopsis see O'Keefe, 1991). The 'map' was taken to be a Euclidean description of the environment in an 'allocentric' 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  $\theta$  rhythm during displacement movement, (iii) deficits in performance of spatial tasks, including the Morris water maze and the Olton 8-arm maze, following 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). Note however that the goal-independence of place cell firing indicates that they only form part of a cognitive map: some read-out mechanism being required to guide behavior.

O'Keefe and Nadel (1978) proposed that, while the hippocampal cognitive map was clearly tied to external cues, some form of PI might support its intrinsic distance and direction metric. In their original formulation, a PC could be activated by two independent means. First, direct activation by the sensory inputs available to an animal in a particular location. Second, activation of the set of PCs corresponding to one location, coupled with inputs indicating the rat's performance of a movement translating and rotating it by a certain amount, would lead to the activation of the set of PCs corresponding to the new location. Mismatches between the two would provide the signal for exploration, which would bring the two sets of information into correspondence by strengthening some sensory inputs and weakening others (O'Keefe and Nadel, 1978; pp 220-230).

## 2 MODELS OF PLACE CELL FIRING

### 2.1 Sensory inputs

In this section we describe models of how the spatial firing of place cells develops and is maintained as the rat moves around an environment. Following an earlier, mathematical, model of PC firing (by Zipser in 1985), Sharp used a simple network with an input layer and two layers of cells governed by 'competitive learning' dynamics (see COMPETITIVE LEARNING), see Sharp, Blair and Brown (in Gluck, 1996). In this model two types of input (or 'neo-cortical') cell 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 angles relative to the rat's head. During exploration, competitive learning 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. 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. The place fields in this model are robust to the removal of a subset of the environmental cues. However, the model takes some time to develop realistic place fields, whereas experiments indicate that they are present as soon as they can be measured.

## **2.2 Attractors and path-integration**

Touretzky and Redish (in Burgess et al., 1999) proposed a model in which place cells form a coherent representation of location on the basis of estimates of the rat's location from PI and local-view information. The model investigates the interaction of frames of reference supported by different mechanisms and in different brain regions, assuming that resetting the PI system depends on the hippocampus. In a similar approach, Guazzelli et al. (2001) proposed a feed-forward competitive learning model that develops a place cell representation by combining PI and sensory inputs, simulating in some detail the effects of darkness, and the deletion or movement of extra-maze cues. See also Arbib (in Burgess et al., 1996).

Samsonovich and McNaughton (1997) went much further in placing the hippocampus at the heart of a PI system. They proposed that region CA3 of the hippocampus forms continuous attractors (Zhang, 1996) such that sets of place fields form pre-configured 'charts' as follows. The recurrent connections between PCs fix the relative locations of place fields within a chart, and also ensure stable and coherent patterns of PC activity (i.e. activity consistent with the rat being in a single location). This system serves as the neural basis of a PI system driven by hardwired motion related signals that shift PC activity so as to reflect the change in location of

the rat corresponding to its movements. A particular chart becomes associated to the sensory stimuli in a particular environment so as to make a correspondence between locations on the chart and locations in reality. This model predicts that hippocampal lesions will impair PI, although the experimental evidence for this is weak.

Kali and Dayan (2000) showed that continuous attractors could be formed by Hebbian learning in the recurrent connections during exploration. However, to create an attractor of equal depth across an un-evenly sampled environment required that learning be mediated by novelty. Hasselmo, Wyble and Wallenstein (in Gluck, 1996) suggest a mechanism for this. The CA3 representation reflects the expected state of the world, being influenced by the associations learned by the recurrent collaterals, while the CA1 representation reflects direct cortical input. Novelty is detected as a mismatch between the two representations, and mediates learning in CA3 by triggering the release of acetylcholine from the medial septum.

### **2.3 What inputs support PC's spatially tuned firing?**

Recording the same PCs in boxes of varying shape and size provides insight into the environment determinants of place fields. In these experiments the location of peak firing tends to maintain a fixed distance from the nearest walls, and a symmetrical, unimodal, place field in a small box may be elongated or multimodal in a larger box. These results are qualitatively fitted by a simple model in which PC firing is a thresholded sum of up to four 'boundary vector cell' (BVC) inputs, each tuned to respond maximally whenever there is a wall a given distance away along a given allocentric direction. The tuning to distance is Gaussian, with a width that increases with the distance of the peak response from the wall. (We note that the above attractor models can only show this behavior to the extent that they behavior is determined by feed-forward BVC inputs.) A random selection of hardwired BVC inputs is sufficient to model the characteristics of populations of place fields, and choosing BVCs to fit a given cell's firing in one set of environments enables prediction of its firing in a novel environment (Hartley et al., 2000).

While the model works well for the initial firing of place cells in an environment, a slow experience-dependent divergence (or 'remapping') of the representations of environments of different shape (Lever, Wills, Cacucci, Burgess, O'Keefe, *Nature*, 2002 in press) indicate an additional role for synaptic plasticity. Because the place cell representations in the models of Kali and Dayan (2000) and Gauzzelli et al. (2001) can be incrementally modulated by learning, they can begin to address the data showing varying degrees of remapping in different experiments, although it is not yet clear what changes will be necessary to provide an accurate model of the dynamics of these data.

### 3 NAVIGATION

How could the hippocampus 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. 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 seems 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 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, see Foster et al. (2000). Below we describe some recent models that have related navigation to the action of individual cells in the hippocampus.

#### 3.1 Using the CA3 recurrent collaterals

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 and Stead (in Gluck, 1996). 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', the synaptic strength of a connection between two CA3 PCs can come to depend on the proximity of their place fields. This is also the condition for the formation of a continuous attractor representation, see above. After brief exploration the synaptic strengths represent distances along the paths taken by the rat. The model proposes that the rat navigates by moving through the place fields of the cells most strongly connected to the cells with fields at the current and destination positions.

Blum and Abbott (1996) propose a related model in which the temporal asymmetry of LTP (synaptic strengthening can occur when presynaptic activity precedes postsynaptic activity, see TEMPORAL DYNAMICS OF BIOLOGICAL SYNAPSES) is invoked to strengthen recurrent connections from one place cell to another if they fire in sequence along the rat's trajectory during exploration. If synaptic modification is also weighted by how soon after the pre- and postsynaptic activity the goal was reached, then the effect of the current collaterals is to shift the location represented by place cell firing from the rat's current location towards the goal. This model proposes that the rat navigates by moving from the current location (e.g. read from CA1) to the shifted location in CA3.

Neither of these models make clear how a direction of motion is actually generated, or if it would be able to generate a short-cut or detour. To build up a true distance metric in complex environments would take a long time, and might be best achieved by reinforcement learning, see Foster et al. (2000). Clear experimental support for learned asymmetric connections between place cells comes from observation of that place fields tend to become elongated backwards along a path during the first few times that a rat runs along it. Asymmetric recurrent connections has also been invoked to explain the phase precession effect. In these models the spread of activation from cells with fields early in a learned trajectory to cells with fields further along the trajectory will, later in the theta cycle, cause a place cell to fire before reaching the location on the path where its original (externally driven) place field was located. However, recent evidence indicates that blocking LTP blocks the development of asymmetric place fields but not prevent the phase precession effect.

### 3.2 Local View Model

In 1989, McNaughton 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. It is also not clear whether this scheme is computationally feasible, see Sharp et al. (in Gluck, 1996) for a discussion of the limitations of simple associative nets in a related situation.

### 3.3 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  $0^\circ$  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 phasors: taking the  $\theta$  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. Thus, 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. Advantages of this system include the ability to perform shortcuts while disadvantages include the sensitivity of the slope (i.e. small movements of cues could lead to reversal of the coordinate system).

## 4 PLACE CELL FIRING & NAVIGATION

The population vector model (see Burgess and O'Keefe in Gluck, 1996) 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. It assumes that the output stage of the hippocampal system is groups of cells that represent the distance and direction of previously encountered goal locations from the rat as it moves around the environment, and that the input to the hippocampus is a set of 'sensory cells' with tuning curve responses to the distance of cues from the rat. Each goal location is represented by the firing rates of a group of 'goal cells' as a population vector (i.e. the vector sum of cells' preferred directions weighted by their firing rates, see REACHING: CODING IN MOTOR CORTEX), and is used to guide the rat back to a goal location (see Fig. 2).

The network operates in a feedforward manner. During exploration, a representation of current location is learned in the intermediate layers, entorhinal cells (ECs), PCs and subicular cells (SCs) map the sensory input to the population vector output. A type of competitive learning governs the dynamics of the PCs and SCs, similar to that used by Sharp (see Sharp et al. in Gluck, 1996). Latent learning during exploration is expressed as the development of large firing fields in SCs which avoids the locality of information access problem (so that goal cell firing fields cover the whole environment, see below). Upon encountering a goal location, learning by modification of connections to goal cells result in each goal cell having a conical firing rate map whose peak is displaced from the goal position in a particular absolute direction

(the 'preferred direction' for that cell), creating the appropriate population vector, see below.

-----

Figure 2 about here

-----

The model relies on the phase of firing of SCs (relative to  $\theta$ ) being such that those firing late in a cycle have place fields centered ahead of the rat. This is achieved 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. When the rat is at the goal the goal cell with preferred direction closest to the rat's heading direction receives a strong input, allowing connections to it to be switched on. This signal arrives at a late phase of the  $\theta$  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 goal cells representing each direction.

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 obstacle cells from that of the goal cells. Advantages for this model are that reasonable trajectories, including short-cuts, are performed after one visit to the goal following brief exploration, and its latent learning. The representation of directions is allocentric (e.g. North, South etc), the necessary translation into left-right body turns is assumed to occur in parietal cortex (taking into account the current heading direction and the locations of obstacles), see also Burgess et al., (2001), Recce and Harris (in Gluck, 1996), Arbib (in Burgess et al., 1999).

Brown and Sharp (see Sharp et al. in Gluck, 1996) proposed a similar feed-forward model of place cell firing and navigation. In their model the output representation (in the nucleus

accumbens) is of the egocentric directions (body-turns) that lead to the goal. The association from place cells to turn cells is built up over many runs to a goal. As with Blum and Abbott's model, synaptic modification must be weighted by how soon the goal was reached after pre- and postsynaptic activity. This model would not show latent learning, and navigation would be strongly affected if stereotyped routes were used during learning. Foster et al., (2000) suggest that, in addition to a fast learning mechanism related to that used by Burgess and O'Keefe (in Gluck, 1996), a slower process of reinforcement learning might build up an explicit metric representation of the environment over the course of several trials.

## **5 DISCUSSION**

We have reviewed several simple neuronal models of how the hippocampus takes in sensory information from environmental cues, turns it into a place cell representation of space, and uses this to support a spatial memory for where interesting things are located. Several of these models also consider the role of internal signals and recurrent connections in these processes. Together, these models represent some of the clearest examples of neuronal level explanations of cognitive behavior. As noted above, how different environments are distinguished so that where things are in each environment can be remembered remains to be well understood, with data on remapping and the roles of the DG and the CA3 recurrent collaterals in providing an associative memory likely to play a part (see Marr, 1971; McNaughton and Nadel, 1990; Gluck, 1996; Burgess et al., 2001; and ASSOCIATIVE MEMORIES). A related issue is the problem of navigation in complex environments: how can local maps be patched together to guide behavior over long distances, see the 'world graph' model of Arbib (in Burgess et al., 1999).

Having made progress in understanding the hippocampal role in rat navigation, extending these models to include its role in monkey and human behavior presents an exciting challenge. Recordings in monkey hippocampus that show neurons responding to the performance of actions in places (T. Ono and colleagues) or to the monkey looking in a particular place (see

Rolls in Burgess et al., 1999) give an indication of how hippocampal function might generalise from rats to monkeys. In humans, there is evidence that the right hippocampus plays a role in navigation similar to that in the rat (e.g. Burgess et al., 1999; 2001). However bilateral hippocampal damage in humans causes a general impairment in episodic memory. One idea relating the spatial and mnemonic roles is that a (non-verbal) episodic memory system could be formed in the right hippocampus by the addition of the human sense of linear time to the rat's spatial system, and, in the left hippocampus, the inputs have been supplemented by verbal information, producing memory for narratives (O'Keefe and Nadel, 1978). These issues, and the relationship between the hippocampus and the parietal cortex (generally considered the primary locus of spatial processing in primates, see POSTERIOR PARIETAL CORTEX - CODING OF GAZE AND REACH) are explored further in Burgess et al. (1999). Finally, the need to impose an egocentric point of view on the allocentric representations in long term memory provides a starting point for modelling the role of the hippocampus and the head-direction system in episodic retrieval (Burgess et al., 2001, see also Recce and Harris in Gluck, 1996).

## References


- Blum, K.I. and Abbott, L. F., 1996, A model of spatial map formation in the hippocampus of the rat *Neural Computation*, 8:85-93.
- \* Burgess, N., Jeffery, K.J. and O'Keefe, J. (Eds.), 1999, *The Hippocampal and Parietal Foundations of Spatial Cognition*. Oxford University Press.
- \* Burgess, N., Becker, S., King, J.A., and O'Keefe, J., 2001 Memory for events and their spatial context: models and experiments. *Philos. Trans. R. Soc., London B* 356: 1493-1503.
- \* Gluck, M.A., (Ed.) 1996, Special issue on computational models of the hippocampus *Hippocampus*, 6: 565-762.

- Guazzelli, A., Bota, M. and Arbib, M.A., 2001, Competitive Hebbian learning and the hippocampal place cell system: Modeling the interaction of visual and path integration cues *Hippocampus*, 11: 216-239.
- Foster, D.J., Morris, R.G., and Dayan, P., 2000, A model of hippocampally dependent navigation, using the temporal difference learning rule. *Hippocampus*, 10: 1-16.
- Hartley, T., Burgess, N., Lever, C., Cacucci, F., and O'Keefe, J., 2000, Modeling place fields in terms of the cortical inputs to the hippocampus. *Hippocampus* 10: 369-379.
- Kali, S. and Dayan, P., 2000, The involvement of recurrent connections in area CA3 in establishing the properties of place fields: a model. *J.Neurosci.* 20: 7463-7477.
- Marr, D., 1971, Simple memory: a theory for archicortex, *Phil. Trans. of the Royal Society B*, 262: 23-81.
- \* 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.
- 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.
- Samsonovich, A. and McNaughton, B.L., 1997, Path Integration and Cognitive Mapping in a Continuous Attractor Neural Network Model. *J. Neurosci* 17: 5900-5920.
- Zhang, K., 1996, Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. *J. Neurosci.* 16: 2112-2126.

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  $\theta$  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.

Figure 2: Left: simulated trajectories from 8 novel starting positions, to a goal encountered after 30 seconds of exploration (at  $60\text{cm/s}$ ) in a  $135 \times 135\text{cm}^2$  environment; the rat is shown to scale. Right: simulated navigation between two goals with an 'obstacle' in between, cues are marked by '+', goals by 'x', and obstacles by '/'. Adapted from N. Burgess, M. Recce and J. O'Keefe, 1994, *Neural Networks* 7: 1065-1081.

A)



B)

