

What can the hippocampal representation of environmental geometry tell us about Hebbian learning?

Colin Lever¹, Neil Burgess^{1,2}, Francesca Cacucci¹, Tom Hartley^{1,2}, John O'Keefe^{1,2}

¹Department of Anatomy and Developmental Biology,

²Institute of Cognitive Neuroscience, University College London, Gower Street, London WC1E 6BT, UK

Received: 8 March 2002 / Accepted: 13 June 2002

Abstract. The importance of the hippocampus in spatial representation is well established. It is suggested that the rodent hippocampal network should provide an optimal substrate for the study of unsupervised Hebbian learning. We focus on the firing characteristics of hippocampal place cells in morphologically different environments. A hard-wired quantitative geometric model of individual place fields is reviewed and presented as the framework in which to understand the additional effects of synaptic plasticity. Existent models employing Hebbian learning are also reviewed. New information is presented regarding the dynamics of place field plasticity over short and long time scales in experiments using barriers and differently shaped walled environments. It is argued that aspects of the temporal dynamics of stability and plasticity in the hippocampal place cell representation both indicate modifications to, and inform the nature of, the synaptic plasticity in place cell models. Our results identify a potential neural basis for long-term incidental learning of environments and provide strong constraints for the way the unsupervised learning in cell assemblies envisaged by Hebb might occur within the hippocampus.

Key words: Hippocampus, place cell, remapping, space, neural network

1 Introduction

Hebb's (1949) postulate regarding the creation of cell assemblies has come to be seen as the pre-eminent model of learning in neural systems. Here we examine the evidence for it in single-unit recordings from the hippocampi of freely moving rats. As we shall see, this paradigm provides a particularly appropriate testing ground for several reasons.

Correspondence to: C. Lever or N. Burgess
(e-mail: colin.lever@ucl.ac.uk; n.burgess@ucl.ac.uk,
Tel.: +44-20-76793388 or 1147, Fax: +44-20-76791306 or 1145)

The current pre-eminence of Hebb's postulate derives at least in part from the discovery of long term potentiation (LTP) (Bliss and Lomo 1973) in the hippocampus. This seems to provide a biological instantiation of Hebb's rule in that roughly simultaneous pre-synaptic activity and post-synaptic activity (or at least depolarisation) produce a long-lasting increase in the efficacy of the synaptic connections.

By coincidence, at about the same time, one of the most striking behavioural correlates of neural activity was discovered, also in the hippocampus. The firing of 'place cells' in fields CA1 and CA3 of the hippocampus of freely moving rats appears to represent the current location of the animal: each one firing whenever the animal enters a restricted portion of its environment (O'Keefe and Dostrovsky 1971; O'Keefe 1976). The integrity of the hippocampus, and of some of the processes linked to LTP, have been shown to be required for spatial behaviours such as finding the hidden platform in a water maze (e.g. Morris et al. 1982; Davis et al. 1992; Steele and Morris 1999).

There is thus solid ground to suppose that LTP and place cells in the hippocampus form components of the neural basis of spatial behaviour of the rat. In addition, as O'Keefe and Nadel (1978) pointed out, Hebb's graduate seminar at McGill emphasised incidental or latent learning as one of the few areas that pointed out the limitations of the behaviourist approach. Indeed, Hebb's postulate concerns exactly this type of 'unsupervised' learning. Appropriately, latent learning has been demonstrated in aspects of rodent spatial behaviour (e.g. Blodgett 1929; Keith and McVety 1988; Harley 1979; Tolman 1932, 1948). Thus examination of changes to the hippocampal representation of space during the rat's exploration of its environment should provide a good opportunity to observe Hebbian learning at the level of single cells in an ecologically valid situation.

We start by reviewing a model of the properties of place cells that does not invoke synaptic or cellular plasticity (O'Keefe and Burgess 1996; Burgess and O'Keefe 1996; Burgess et al. 2000; Hartley et al. 2000). We argue that quantitative definition of the basic system must be the first

step upon which an understanding of the additional effects of plasticity can be built. We also briefly discuss some of the existing models of the effects of synaptic plasticity on place cell firing. The main body of the article then examines in detail recent experimental data concerning the time course of plasticity in the place cell representation of environmental geometry under two experimental manipulations. The first concerns short term changes following the introduction of a barrier into the environment, extending the findings of (Muller and Kubie 1987), while the second concerns further analysis of the time course of plasticity in the representation of environments of different shape (Lever et al. 2002). These data provide constraints on any mechanisms of synaptic plasticity that might be at play in the place cell system.

2 Computational models of place cell firing

2.1 Introduction to computational models

In this section, we briefly review a model of the firing of place cells, specifically their spatial receptive fields or ‘place fields’, that requires no learning at all. We show that this adequately handles much of the data regarding the basic characteristics of place cell firing. As such, we argue that this model provides an appropriate framework into which to incorporate experimental data regarding the effects of learning, such as that presented in the main section of this paper. To provide a broader context for the consideration of place cell plasticity, we also briefly review some of the existent models of place cell firing that depend on synaptic plasticity. In this section we also introduce the additional concepts of pattern completion and pattern separation that have been respectively attributed to the recurrent collaterals in CA3 and the dentate gyrus by many authors.

2.2 A simple geometric model of place cell firing, without learning

A major motivation for not invoking some mechanism of learning is that, on a rat’s very first exposure to an environment, many cells already have place fields in that environment (Hill 1978; Lever Cacucci, and O’Keefe unpublished; Wilson and McNaughton 1993). The type of simple feed-forward model of place cell responses discussed here originates from the work of Zipser (1985). In Zipser’s model, sensory details of the environment feed-forward to ‘landmark detectors’ and hence to place cells. Each landmark detector is tuned to detect a particular aspect of the sensory scene (a ‘location parameter’) and simply performs a match between the stored state of the location parameter and its currently perceived state. A place cell’s activity then corresponds to a threshold sum of the strengths of the matches of the landmark detectors that connect to it. Location parameters included the retinal angle between two landmarks, on the assumption that place field size should scale proportionally with the size of the environment (in fact

the relevant experiment shows a much lower scaling factor Muller and Kubie 1987).

Constraints on the functional form of the sensory input to place cells can be derived by systematically varying the shape and size of the rat’s environment while recording from the same place cells. In an experiment of this form, O’Keefe and Burgess (1996) showed a lawful pattern of firing across environments such that the location of peak firing often maintained a fixed distance to the nearest two walls, while the place field sometimes became stretched or bimodal in the larger environments. The pattern of firing was consistent with a thresholded linear sum of inputs tuned to respond to the presence of a boundary at a given distance along a given allocentric direction (Fig. 1). These inputs were labeled ‘boundary vector cells’ (BVCs) (Hartley et al. 2000; Burgess et al. 2000). Two aspects of the BVCs should be noted. First, the directions along which the postulated BVCs are oriented (ϕ_i) are independent of the orientation of the

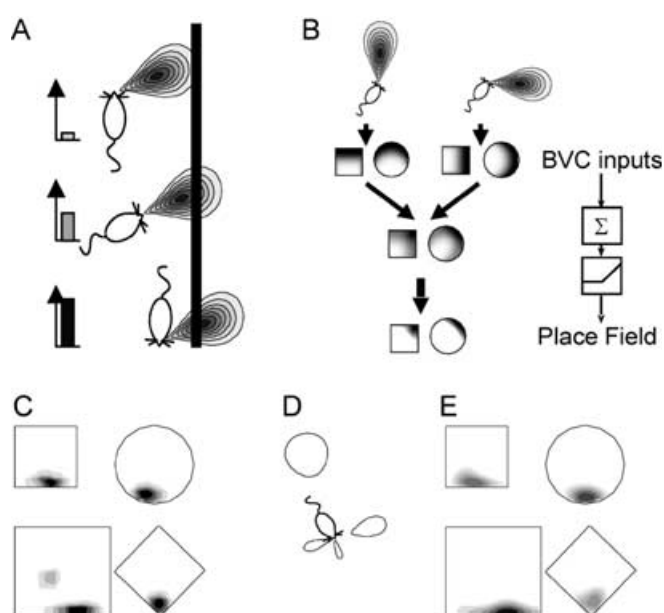


Fig. 1A–E. A geometrical model of place fields. Place cells receive inputs from boundary vector cells (BVCs) which are tuned to respond to walls at a particular distance and compass direction (the greater the distance tuning, the broader the tuning of the response). A place field is modelled as a thresholded summation of two or more BVCs. (A) Example of a BVC. The firing rate of the BVC (indicated by the shaded bars left) depends on the proximity of a boundary in a particular allocentric direction. The receptive field of the BVC (indicated by the shaded teardrop) maintains a fixed orientation with respect to the environment (in this case east-north-east) regardless of the rat’s facing direction. The presence of a boundary in this region leads to increased firing. (B) A given BVC has a firing field that follows the walls of a given environment (square and circle shown here) at a particular distance. Two BVCs are shown, one with northern field (left), one with eastern field (right). Summating, then thresholding, these two BVC inputs by a place cell would result in restricted north-eastern place fields in the square and circle for that place cell. As an example, the experimental data from a place cell recorded in four different environments shown in C) can be modelled as a cell with four BVC inputs, whose receptive fields are illustrated in D). The simulated place fields are shown in E). Adapted from Hartley et al. (2000)

rat and probably dependent on the head-direction system (e.g. Taube 1990, 1998) and the various orientation cues around the environment. Second, the sharpness of tuning of a BVC's response is affected by the distance to which it is tuned (d_i), with sharper tuning to shorter distances. This means that boundaries near to a field will tend to provide the most powerful determinant of the field's subsequent location.

More specifically, a given BVC (i.e. BVC i) is tuned to respond to the presence of a boundary at a given bearing (d_i, ϕ_i). The response of BVC i to a boundary element at distance r subtending angle $\delta\theta$ is:

$$\delta f_i = g_i(r, \theta) \delta\theta,$$

where

$$g_i(r, \theta) = \frac{\exp\left(-\frac{(r - d_i)^2}{2\sigma_r^2(d_i)}\right)}{\sqrt{2\pi\sigma_r^2(d_i)}} \times \frac{\exp\left(-\frac{(\theta - \phi_i)^2}{2\sigma_a^2}\right)}{\sqrt{2\pi\sigma_a^2}}.$$

The width of the radial tuning increases with the preferred distance d_i , i.e. $\sigma_r(d_i) = \sigma_o(1 + d_i/\beta)$. The firing rate of a place cell with n BVC inputs is then:

$$F(\underline{x}) = AH \left(\sum_{i=1}^n \left(\int_0^{2\pi} g_i(r, \theta) d\theta \right) - T \right),$$

where T is the threshold, A determines the amplitude, H is the Heaviside function (i.e. $H(u) = u$ if $u > 0$; $H(u) = 0$ otherwise) and \underline{x} is the location of the rat which (together with the geometry of the environment) determines the locus $r(\theta)$ of the boundary from the rat.

Hartley et al. (2000) showed that the characteristic properties of populations of place cells could be

captured by the above model by fixing A and T and simply assuming that each place cell receives a random sample of BVCs. The place field properties modelled included the distributions of the numbers of fields and their shapes, sizes and orientations across the four environments used by Burgess and O'Keefe (1996). Beyond this, Hartley et al. (2000) showed that a place cell's firing pattern across different environments could be modelled by choosing the appropriate BVCs and threshold T . For most cells, a reasonable fit could be obtained using no more than four BVCs along orthogonal directions, corresponding to only six degrees of freedom in the model (their overall orientation, the four distances and the threshold). Such a model can then be used to predict the pattern of firing in an environment of novel shape (Fig. 2). The model is also consistent with the patterns of firing of place cells recorded on a linear track whose length is varied systematically (Gothard et al. 1996).

2.3 Learning in models of place cell firing

Models of place cell firing that focus on synaptic plasticity fall into three categories according to the principal reason for which plasticity is required: directionally modulated firing of place cells, asymmetric expansion of place fields and stability and remapping in the place field representation. We briefly discuss each of these categories but before doing so we provide a brief introduction to the concepts of continuous attractors and pattern completion, as applied to the recurrent collaterals in CA3, and of pattern separation as applied to the dentate gyrus. These concepts are important for interpretation of the computational and experimental studies of stability and remapping and are not included in the simple geometrical model discussed above.

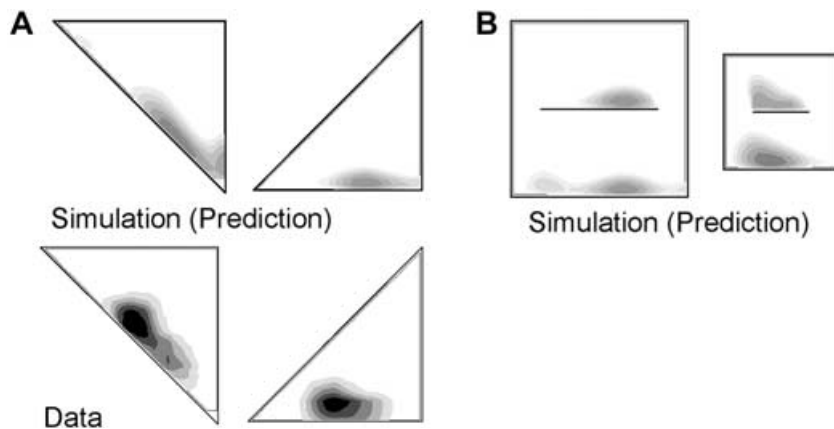


Fig. 2A,B. Predictions from the geometric place field model. Having fitted the data for a given place cell in several different environments we can use the same set of BVCs to predict the behaviour of that cell in any novel environment. In this example, data has been fitted for a cell based on its firing in the four environments shown in Fig. 1 (square, circle, large square, diamond). **A**) shows the predicted place fields for the cell fitted in Fig. 1 in a right angled triangular box in two different orientations relative to distal cues in the laboratory.

Experimental data from the same cell are shown for comparison. **B**) shows the predicted place field for two environments (standard square and large square) in which a barrier has been placed. The model predicts an additional second field will appear north of the barrier in both the large square (left) and standard square (right). Experimental data from this cell in the corresponding environments are shown in Figs. 3D and 4A (Cell 2). Adapted from Hartley et al. (2000)

2.3.1 Recurrent collaterals, pattern separation, pattern completion and continuous attractors. The recurrent collaterals in region CA3 and the dentate gyrus whose cells project into this region comprise the two anatomical features of the hippocampus most noticeably absent from the simple geometric model. These two features have been the focus of the hippocampus as an associative memory device since the seminal contribution of Marr (1971). Many authors have brought attention to the possibility that the extensive recurrent connections between pyramidal cells in CA3 could support an auto-associative memory (Marr 1971; McNaughton and Morris 1987; McNaughton and Nadel 1990; Treves and Rolls 1992; McClelland et al. 1995), see also (Kohonen 1972; Gardner-Medwin 1976; Hopfield 1985). In these systems, based on Hebbian learning in the recurrent connections, a partial cue can produce retrieval of an entire stored representation, a process referred to a 'pattern completion'. Interference between similar stored representations can be a problem in these systems such that performance is improved when non-overlapping representations are to be stored (e.g. Amit 1990). For this reason, it has also been proposed that the dentate gyrus serves to ensure that even similar cortical inputs to the hippocampus are stored as non-overlapping representations in CA3 (Marr 1971; McNaughton and Nadel 1990; Treves and Rolls 1992; McClelland et al. 1995). This is achieved by generating an intermediate, highly sparse, representation of the cortical input using the very large number of cells in the dentate gyrus (which contains an order of magnitude more neurons than either entorhinal cortex or CA3). This process is referred to as 'pattern separation'. Much evidence points to the presence of processes akin to pattern completion. For example, it has long been known that representation of only a subset of the original cues present during training may be sufficient to maintain normal place field firing (e.g. O'Keefe and Conway 1978; O'Keefe and Speakman 1987; Quirk et al. 1990). There is now evidence that this property depends on CA3 NMDA receptors (Nakazawa et al. 2002). The primary direct evidence for pattern separation concerns experiments in which the place cell representation appears to 're-map' entirely between reasonably similar environments. We consider the details of these data in the following sections.

Many computational models of remapping have concerned a class of models of the hippocampus that use the recurrent collaterals in region CA3 to support 'continuous attractor' states in the activity of CA3 place cells (Zhang 1996). In these models, each place cell is assumed to have a preferred location in the environment and to fire according to the rat's proximity to this location. The recurrent connections are arranged so that the strength of the connection between two place cells is a simple (increasing) function of the proximity of their preferred locations, a situation that might be created by unsupervised Hebbian learning during exploration (Muller et al. 1996; Muller and Stead 1996). This causes patterns of firing consistent with the rat being in a single location to form attractor states, and allows for smooth transitions between the representations of neighbouring locations.

2.3.2 Directionality of place cell firing. When rats move in an unconstrained manner through an open environment, the spatially specific firing of place cells is independent of the animal's orientation (Muller et al. 1994; O'Keefe 1976). However, when the rat is constrained to run on a linear track or a narrow-armed maze, responses become direction dependent. In other words, a given cell will fire when the rat runs through the place field in a specific direction but not when it runs through it in the reverse direction (McNaughton et al. 1983; Muller et al. 1994; O'Keefe and Recce 1993).

To model the directionality of place fields, Sharp (1991) extended the simple feed-forward model of Zipser (1985) by adding an element of 'competitive learning' (Rumelhart and Zipser 1986). Her model consists of a layer of inputs projects to an intermediate layer (entorhinal cortex) which projected to a layer of place cells. Inputs correspond to one of two types of information regarding environmental landmarks (representing their distances and angles relative to the rat's heading direction). To simulate competitive learning, neurons in the two processing layers are divided into groups dominated by lateral inhibition such that only one 'winner' can be active at a time. Hebbian learning is then applied to the initially random connection strengths such that connections to the active neuron in each group from active neurons in the preceding layer are strengthened. However, an important modification must be made to the Hebbian learning rule for the algorithm to work. The net strength of the connections to each neuron must be normalised (i.e. kept equal and constant over time), to ensure that several neurons in each group can 'win' according to the pattern of input to the group. This learning rule results in specific neurons coming to respond to a particular pattern of input, or to patterns similar to it.

The orientation dependence of half of the inputs in Sharp's model means that initial responses in the place cell layer are modulated by the rat's orientation as well as by its location. If the rat subsequently moves through a place field solely along restricted directions, then the response remains directionally modulated. However, if the rat subsequently moves through the place field in diverse directions, the cell learns to respond to a broader and broader selection of directions at that location, eventually showing directionally-independent firing. As a variant on this model, models with similar behaviour have been proposed in which direction-independence arises as a result of unsupervised Hebbian learning in recurrent collaterals in the place cell layer (Brunel and Troullier 1998; Kali and Dayan 2000). The main argument against the way all of these models incorporate learning is that, as far as we can tell, place cell firing is initially direction-independent, and subsequently becomes direction-dependent under conditions of constrained motion (O'Keefe unpublished observation).

2.3.3 Experience dependent increases in the firing field. The phenomenon of LTP shows an interesting variation from a simple dependence on coincident pre-

and post-synaptic activity, with preferential induction when presynaptic activity precedes postsynaptic activity (Gustaffson and Wigstrom 1986; Levy and Stewart 1983; Markram this volume). Perhaps the best evidence for short-term experience-dependent changes in place cell firing is related to this phenomenon and comes from experiments by Mehta et al. (1997, 2000). They found that the spatial distribution of a place cell's firing rate on a linear track becomes more asymmetrical during the first few runs of the rat along the track: tending to fire at a lower rate on entry to the place field than on leaving it. They suggest that the temporal asymmetry of LTP, acting on the CA3 to CA1 pathway, causes this effect by strengthening the connections from place cells firing earlier on the track to those firing later on the track on a given run. This would cause the cells firing later on the track to begin firing earlier on it after learning. Other similar models have implicated asymmetric Hebbian learning in the recurrent connections within CA3. Importantly, this experience-dependent asymmetry has been shown to be dependent on the NMDA receptor and thus linked to LTP (Ekstrom et al. 2001). However, one aspect of this phenomenon that differentiates it from the Hebbian idea of the long-term formation of cell assemblies is that it appears to reset each day, despite the occurrence of many runs during the day.

2.3.4 Stability or remapping in the place cell representation. Initial experiments in which place cells were recorded in environments of different shape (Muller and Kubie 1987) reported completely different patterns of firing, or 'remapping', in the two environments. A place cell active in one environment might fire in an unrelated location in the second environment or might be silent. Interestingly, this phenomenon appears to be at odds with the experiments on which the simple geometric model was based, showing systematic regularities in the place fields recorded in environments of different shape (O'Keefe and Burgess 1996). As discussed below, the remapping likely results from a process of experience-dependent plasticity over the two week or longer training period that was used in the experiments by Muller and Kubie (1987) but not by O'Keefe and Burgess (1996) (Sect. 4). Furthermore, this learned discrimination of the place cell representation of environments of different shape has been shown to last for several weeks (Lever et al. 2002). A final experimental finding pertinent to this discussion is the observation that the day to day stability of the place fields within a constant environment has been shown to depend on the NMDA receptor and thus linked to LTP (Kentros et al. 1998).

Samsonovich and McNaughton's (1997) model of the place cell representation of space showed that a large number of independent continuous attractor representations (or 'charts') could be supported by the recurrent CA3 network. In their model, as animals explore an environment, local view representations become bound by Hebbian learning to specific places in that environment's chart. They suggested that these charts were preconfigured (i.e. hard-wired) in CA3, with specific

charts becoming associated with given environments as and when necessary. Within this model, remapping corresponds to switching between different charts. Notice that the distances between place fields are fixed by the preconfigured recurrent connections, a condition seemingly at odds with the plastic and multimodal fields seen after environmental manipulations (O'Keefe and Burgess 1996; Gothard et al. 1996). For the chart model to accommodate the occurrence of a bimodal field under these circumstances, each lobe of the now bimodal place field must exist on a different chart with resetting or switching between charts occurring while the rat moves from one lobe to the other (McNaughton 1996). Alternatively, the model can be set up so that feed-forward input of the form suggested by the geometrical model dominates any recurrent input (Samsonovich and McNaughton 1997).

While Samsonovich and McNaughton's (1997) chart model can be shown to support separate (remapped) environmental representations, some interesting issues arise regarding the nature of the plasticity that would be required to allow these representations to develop during exploration rather than simply being somehow pre-configured. Kali and Dayan (2000) showed that, if serial exploration is simulated, simple Hebbian learning is not sufficient to form a well-behaved continuous attractor representation i.e. one in which all locations are equally stable. An additional mechanism is required such that synaptic plasticity is modulated by novelty, preventing the attractor for frequently visited locations to become deeper than for less well visited locations. A plausible mechanism for this involves novelty being detected by mismatch between a retrieved CA3 representation and a CA1 representation activated directly from entorhinal cortex. The level of novelty (mismatch) would then control the release of acetylcholine into the hippocampus from the medial septum which in turn would modulate synaptic plasticity (Hasselmo et al. 1996). Interestingly, Kali and Dayan (2000) also argued that novelty-mediated learning was sufficient to allow distinct representations to develop for different environments, partially distinct representations for related environments and similar representations for environments differing only geometrically (as in O'Keefe and Burgess 1996; Gothard et al. 1996). However, to get place fields to maintain fixed distances to environmental boundaries during geometrical manipulations to the environment (as opposed to maintaining a fixed ratio of distances between boundaries) requires feed-forward inputs of the BVC form to dominate over recurrent connections, as with Samsonovich and McNaughton's (1997) model.

An alternative analysis of the development of remapping was performed by Fuhs and Touretzky (2000). Of direct relevance to our own work described in Sect. 4 (Lever et al. 2002), Fuhs and Touretzky adopt a partial, or gradual, remapping approach to map separation suggested by one study (Tanila et al. 1997) but not others (e.g. Bostock et al. 1991). Their partial remapping is effected by the combined effects of pattern completion and pattern separation mechanisms. Phenomenologically, they focused on what may turn out to be a common expression

of remapping in individual place cells, that is, for it to continue to fire in the environment in which it fires most strongly, and to stop firing in the other. In contrast to the recurrent models previously discussed, they examined whether or not plasticity in the perforant path projection from entorhinal cortex to CA3 place cells could support this type of behaviour. Interestingly, changing connection weights according to the product of pre and post synaptic activation (i.e. 'Hebbian' learning), and or to their covariance, did not produce the required behaviour. Hebbian learning leads to a place cell that fires strongly in one environment and weakly in another, thus strengthening its firing in both environments. Using covariance learning, exposure to the second environment leads to loss of the place cell representation of the first environment (including cells with high and low firing rates). However, the BCM learning rule (Bienenstock et al. 1982) did produce the desired result, that is, strong firing remains stable and weak firing reduces with experience. The critical aspects of the BCM rule are that synaptic modification depends on pre and post synaptic activity (avoiding interference where different inputs are active in the two environments) and the direction of modification (increasing or decreasing) depends on the strength of the post synaptic activity.

Models implicating plasticity in providing a stable representation of an environment are supported by some recent studies. Basically, rats and mice which are old (Barnes et al. 1997) or whose candidate molecular plasticity-machinery (e.g. NMDA receptor) is compromised (Kentros et al. 1998; Rotenberg et al. 1996, 2000) can show different patterns of firing upon exposure to a previously-experienced environment (but see also McHugh et al. 1996). Although various models would interpret, say, an old rat's remapping of a familiar environment quite differently and identify assumed plasticity in concomitantly different projections (Barnes et al. 1997; Redish and Touretzky 1999), consensus is emerging which suggests that the place cell's daily stability in familiar environments is plasticity dependent. These models are also consistent with studies relating plasticity to pattern completion within an environmental representation (Nagazawa et al., 2002 unpublished data). There have been hints throughout the place cell literature that alterations in a familiar environment sometimes cause firing rate reduction in many cells in the altered environment, as though there were a learned template firing pattern. This was specifically studied in Fenton and Muller (2000) with clear results stating that in a circle with two cards on the internal wall, firing rates were significantly lower in the two manipulations where the cards were placed further apart and closer together than in the standard trained configuration.

We return to the subject of bimodal place fields and attractor charts in Sect. 3 "Experiments with barriers – Predictions and limitations of the geometric model". We return to stability and remapping in Sect. 4 "Long-term memory and incidental learning of environmental geometry". As we shall see, the simple geometric model needs to be extended to include plasticity.

3 Experiments with barriers – predictions and limitations of the geometric model

3.1 Introduction to barrier experiments

In this section we describe some experimental place cell data which highlight some of the strengths and limitations of the model described above. We stress at the outset that these preliminary data are not quantitative and were collected in rats with diverse experience, though sharing a common exposure to square and circular environments formed from a deformable walled environment (the "morph box", further described in Lever et al. 2002). We introduce this topic by noting that experimenting with barriers formed an important part of the seminal studies by Muller and Kubie (1987) in the cue controlled circular-walled environment. Muller and Kubie used barriers to suggest the "kinematic hypothesis" of place cell firing. Their key finding was that the placement of a barrier within the place field tended to severely reduce the amount of firing in the field. Figure 3A, B schematises this finding and shows a clear replication of this effect for a cell firing in the square, with further demonstration of its reproducibility. We emphasise that the powerful rate-reduction effect exerted by the barrier can occur when the barrier is not directly within the place field but at some distance from it (trials 1, 3, 6, 8) as well as when the barrier is in the field centre (first half of trial 9). Such effects on place fields can be understood within the framework of our model because the introduction of the barrier "occludes" more distant walls and thus prevents BVC inputs in those directions from contributing to firing at the original location of the field. We would not expect an overall decrease in firing rates (across a population of cells) as the introduction of the barrier will also produce additional firing in some cells.

Of particular interest here is the prediction that the insertion of a barrier will often *induce* a new second field at a predictable location. This happens because the barrier acts as a boundary and thus produces additional input to the place cells at the new location (Burgess et al. 2000; Hartley et al. 2000). Sometimes this additional input, when combined with the other inputs to the cell, will be sufficient to exceed its firing threshold and thus produce a second field (typically near to the barrier). To our knowledge, such effects had not been suggested before the development of our model. In the top two rows of Fig. 3C, one can see that the insertion of an east-west barrier can create a second field close to, and north of, the barrier, as if the barrier were acting as an additional south wall (Fig. 3C, bottom row). The model also predicts that some cells will not fire north of the barrier where the additional input is not sufficient to drive the cell beyond its firing threshold (Fig. 3C, top row) and that some previously silent cells will start to fire because the cell only receives sufficient input to fire when the barrier is added (not shown).

If place cells provide the rat's signal of its allocentric location within an environment, it might plausibly be reasoned that the hippocampal network will not tolerate too many cells firing in two distinct positions. A cell

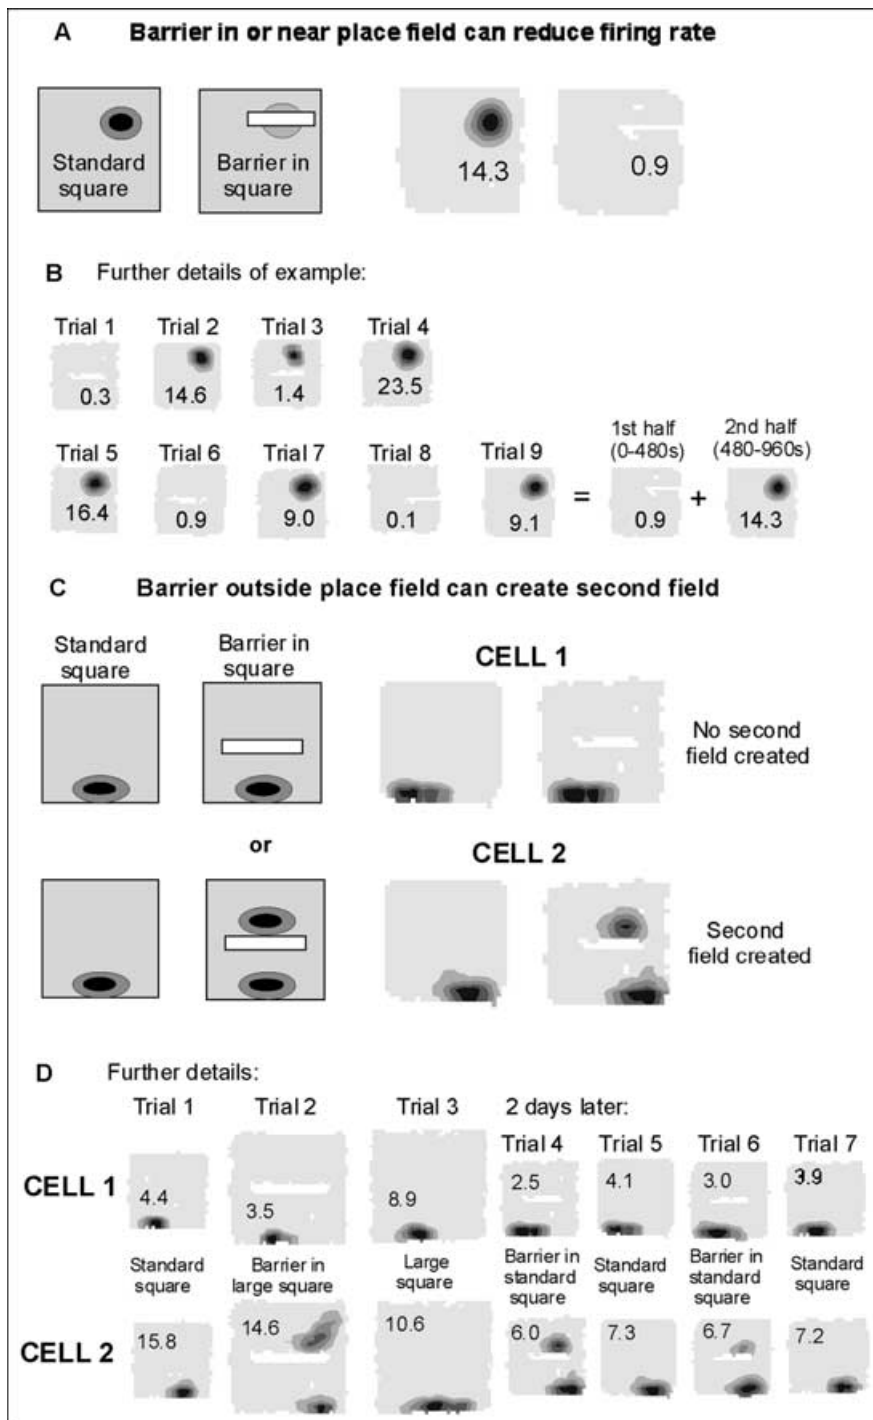


Fig. 3A–D. Two classes of place field firing changes induced by insertion of a barrier into a square-walled enclosure. **A)** Barrier in or near place field can severely reduce firing rate (top). This effect is seen schematically (top left) and in a real place cell (top right). **B)** Smoothed firing rate maps are shown for the cell in nine trials. In this and subsequent figures, methods of firing rate map construction are as described in Lever et al. (2002). The five shades of firing rate maps are each autoscaled to represent 20% of the peak rate which is indicated numerically (light grey through to black). Unvisited bins are shown in white. ‘Fields’ with a peak of less than 1.0 Hz are not shown (i.e. map only shows light-grey region denoting visited area, and white bins denoting unvisited areas). **C)** Barrier outside place field can create a second field in a position related to the position of the original place field. Schematic (top left) and real place cells (top right and **D)** are shown. In Cell 1, the barrier reduces firing somewhat but does not create a second field. No change is seen with repetitions. In Cell 2, the barrier creates a duplicate field which is stable across repetitions. Note that trials 2 and 3 are conducted in a larger square

continuing to fire as in the bottom row of Fig. 3C is highly ambiguous, and the network’s “correction” of such firing may well give us important clues about the system as a whole. The paradigm may open a window into the network, to see if, for instance, it obeys attractor dynamics. The next section describes some preliminary work illustrating this line of approach.

3.2 Stable barrier-related firing in place cells

The basic methodology behind the results described in Sects. 3 and 4 of this paper is the same. Briefly,

recording took place within a black-curtained, circular testing arena. Walled enclosures, such as the standard 62 cm sided, 50 cm high square box, or the larger 93 cm sided, 50 cm high square box, were placed inside the testing arena such that the centre of each box always had the same location in the arena. Between trials, rats are placed on a holding platform outside the curtained testing arena. We actively encouraged directional constancy with the use of stably positioned cue cards and procedures standardising the way that the rats were led into the testing arena. To encourage the rats to walk around and provide even coverage of the boxes’ two-

dimensional surface, sweetened rice was randomly thrown into the given box at about 30 s intervals. In this section, we describe results where, upon finding place cells firing along the south wall of the standard square, we tested whether a second field is created by the insertion of an east-west barrier within the square.

Cell 1 in Fig. 3D shows the simplest type of result, where the insertion of the barrier in the large and the small square does not induce a second field. We note that cells of this type should be relatively rare according to the geometric model, and finding a high proportion of these would constitute evidence against it. Cell 2 in Fig. 3D shows another type of result simultaneously recorded from the same animal. Here the barrier reliably induces a second field in the expected region in the large square (trial 2), and this effect is also seen when tested two days later in the small square (trial 4) where the effect is reproduced on a later trial on the same day (trial 6). The firing in the square without a barrier insertion is unaffected. We point out that the trials labeled 4–7 in Fig. 3D and the trials 1–4 in Fig. 3B are the same trials. Accordingly, in the same cell ensemble, we can simultaneously observe rate inhibition, second field induction, and no obvious change effected by the barrier's insertion. This rules out interpretations based on global changes. In contrast to Cell 1, the behaviour of Cell 2 corresponds well to the geometric model. In fact the geometrical model is particularly well supported by this example, as the barrier was made of a different substance (old wood) with a different texture, visual appearance, and smell to the rest of the environment.

For models in which place fields must have fixed relative locations within a given representation (Samsonovich and McNaughton 1997), cells that develop sub-fields following some environmental manipulation require elaborate explanation. One possibility is that each sub-field exists on a different chart, or that charts are somehow reset¹ between each manifestation of activity within a given sub-field. Thus it is important to ask if the firing in two firing locations for Cell 2 represents the outcome of two genuine “simultaneously” bimodal fields. It could be that the firing rate map averaged over the whole trial is misleading and that only one field is really operative for half the trial, then stops firing, while the other starts. One might even imagine that the two fields oscillate every two minutes, and so on. Of course, the fields cannot literally be examined simultaneously since the rat cannot be in two places at once. Given natural coverage of the environments, however, can it be shown that there is firing in both fields at the shortest measurable intervals?

In Fig. 4, we deal with these issues and show that the fields can be genuinely bimodal. Figure 4 presents data from Cell 2, Fig. 3D showing segments from Trial 2 in parts A, B, C and from Trial 4 in parts D, E. Parts A and D present firing rate maps from the first and second half

of the respective trials, showing substantial firing in both locations for each trial half. Parts B, C, and E present firing rate maps for short equal time slices, 7 or 10 s, and show firing in both locations at the shortest natural intervals. To show genuinely bimodal firing, one needs to show firing first in one field, then the other, and then again in the first field, as quickly as the rat can go between them. In both parts B and C, we show two examples of three alternating firing episodes, both occurring within 25 seconds. Part E shows many such alternating firing episodes from Trial 4. For this cell, firing in the two locations occurs as close together in time as one could hope to observe in a freely moving rat.

3.3 Plasticity in barrier-related firing

We now consider plastic changes in the place cell representation following barrier insertion. Obviously, these changes will not be consistent with the simple geometrical model which is restricted to firing patterns that remain static over time given an unchanging environment.

Figure 5 shows a place cell (Cell 3) whose response to the barrier changes rapidly within a single 9 minute trial. Its original field is against the south wall (Fig. 5A, Trial 1). Introduction of the barrier quickly induces a second, duplicate field above the barrier (Trial 2, 0–180 s). The original field is gradually eliminated during Trial 2 (0–180 s, 180–360 s, 360–540 s) and does not return in the subsequent trial (Trial 3, 0–270 s, 270–540 s). Removal of the barrier in the latter part of Trial 3 (540–900 s) reinstates the original field configuration which continues stably through the next control trial (Trial 4). The time slices from Trial 2 (Fig. 5B) show the kind of bimodal firing pattern seen in Cell 2. The time slices from Trial 3 (Fig. 5C) show that firing is now restricted to the barrier field, despite frequent visits to the original wall field (e.g. in 6/10 time slices from Trial 3).

Cell 4 (Fig. 5D) is another example of a short-lived field. In this case the second, duplicate field, is transitory. This is a cell which fired consistently in the southwest of the square (Fig. 5D, Trials 1–4). However, closer examination of the temporal firing over Trial 2 showed that the cell fired in the first 160 seconds in the expected duplicate field location above the barrier (albeit at a reduced rate of 2.1 Hz) but changed this pattern such that it did not fire in this location thereafter. The absence of firing above the barrier was then maintained in a subsequent trial (Trial 3) and normal firing resumed in the standard square (Trial 4).

These examples suggest an interesting line of approach by using barrier-type experiments to examine relatively rapid incidental learning in the hippocampal network. It is unclear whether such rapid plasticity is stable in the long-term. The next section describes an experiment in which we have used square and circular enclosures to demonstrate plasticity in the representation of these environments that is both transferable and stable in the long term.

¹ One suggestion (McNaughton 1996) was that path integrative inputs are reset by an event such as the rat bumping into a wall.




Fig. 4A–E. Stable bimodal place cell firing in the barrier-in-square environment. Bimodal place field firing of Cell 2 (Fig. 3D) is seen in both the original, and barrier-induced, duplicate locations, in both Trial 2 in the large square (**A, B, C**) and Trial 4 in the small square (**D, E**), even over the shortest feasible time intervals. **A**) shows firing rate maps for the first and second halves of Trial 2. **B**) and **C**) show small time slices from the first half (**B**), and second half (**C**), of Trial 2. Firing rate maps in **B**) and **C**) are arranged as ‘equations’: firing in successive, very brief (e.g. 10 s) time slices is shown in the two or more firing rate maps on the right hand side of the equal sign. The firing rate map on the left of the equal sign represents the firing over the whole period summing these time slices. In effect, with firing rate maps of small time slices, light grey regions form lines showing the trajectory of the rat, medium-grey through to black regions along the trajectory denote above-threshold firing fields, and the white region denotes the large, unvisited area. Arrows show the broad direction of the rat’s movement. This style of presentation is also used in Fig. 4E and Fig. 5B & C. **D**) shows firing rate maps for the first and second halves of Trial 4. **E**) shows small time slices from first half of trial 4

4 Long-term memory and incidental learning of environmental geometry

4.1 Rationale and description of a study (Lever et al. 2002) showing incremental remapping

In this section we summarise and comment on a recent study (Lever et al. 2002) that tested place cells in geometrically different environments. We also present further details, where appropriate, which were not given in that study. The rationale for this study is as follows. Previous work in this laboratory had shown similarities in place fields across various rectangular walled environments, such as different sized squares and differently oriented rectangles (O’Keefe and Burgess 1996). Previous studies in the Muller-Kubie laboratory involving the comparison of place cell firing in rectangular and circular environments had shown clear evidence of

dissimilar representations between these two types of environments (Muller and Kubie 1987; Quirk et al. 1992). In contrast to these studies, our experience (Lever et al. 2002) showed that place fields were found to be “homotopic” when tested in circles and squares; i.e. in corresponding locations in both shapes. Various environmental manipulations, such as wall translation, wall removal, and reconfiguration of the walls into shapes other than circles and squares, showed that the similarity of the firing patterns were determined by the box walls, not by the identical square and circle room locations. We reasoned that the apparent contradiction with the results of Muller and Kubie (1987) might be due to the fact that their animals had received considerable pre-training whereas ours were naïve. To test if experience was a critical factor, we recorded from a new group of animals, for up to three weeks, on successive days from first exposures. In some cases we followed indi-




Fig. 5A–D. Plasticity seen in firing patterns in two place cells (**Cell 3: A, B, C; Cell 4: D**) after induction of a new second field by the insertion of a barrier into the square. **A)** Plasticity in place Cell 3. Left-hand column shows firing rate maps for Trials 1 to 4, each map averaged over the whole trial. Firing rate maps for 1st, 2nd, and last third of trial 2 are shown to right of whole Trial 2 map. Firing rate maps for 1st and 2nd half of the barrier-in-place portion of Trial 3, and portion of Trial 3 where the barrier has been removed, are shown to right of whole Trial 3 map. Note that the southern, original field gradually declines in strength (Trial 2: 0–180 s, 180–360 s, 360–540 s segments) and then disappears reliably (Trial 3: 0–270 s, 270–540 s segments). When the barrier is removed (Trial 3: 540 s) the original field returns (Trial 3: 540–900 s segment, and Trial 4). **B)** Bimodal firing in Cell 3 occurs over the shortest feasible time intervals. A 50 s segment of bimodal firing from Cell 3 in the early part of Trial 2,

segmented further into five 10 second time slices, is shown. **C)** The unimodal firing seen when the barrier is present in Trial 3 is not due to lack of sampling of the location of the original place field. A 200 s segment of firing from Cell 3 in Trial 3 showing place field in the new, northern region only, segmented further into ten 20 second time slices. The difference in firing between **B** and **C** is attributed to rapid hippocampal plasticity. **D)** Plasticity in place Cell 4. Trials 1 to 4 show firing rate maps averaged over the whole trial. To the right are shown firing rate maps for the 1st third (0–160 s segment) and last two-thirds (160–480 s segment) of Trial 2. Note that initially, a new second field is created above the barrier in the predicted region, albeit at a reduced rate. This disappears after about 3 min. The absence of firing above the barrier is maintained in a subsequent trial (Trial 3), and normal firing resumes in the standard square (Trial 4)

vidual cells for over a week. The entire duration of the animals' experience in the two shaped environments was recorded (Lever et al. 2002).

Replicating our first experiment, place cell firing on initial exposures to the circles and squares was highly similar. Gradually, however, with increased experience in these environments, the place cell firing patterns became divergent across the two shapes (but not between environments of the same shape). In other words, the evolution of this shape-specific remapping over time could be observed. In later trials, we found that many cells fired in shape-specific patterns. For instance, a cell might fire in one shape only (monotopic) or in different locations in the two shapes (heterotopic: e.g. in the centre of the circle, but in the north-west of the square). Two further aspects of the phenomenon were explored, namely, transfer and long-term stability, in an attempt to relate these findings to spatial learning and memory. First, it was found that the cells' geometrically-tuned responses showed good generalisation from circles and squares made of one kind of material to those of another. Second, after delays of about a month, the firing patterns across shapes remained highly divergent, suggesting that remapping was permanent. We interpret these results as identifying a potential neural basis for hippocampal long-term memory of environments.

It is important to re-emphasise that such learning occurs independently of explicit reward. There is nothing in the rice-throwing procedure used to encourage active exploration of the boxes which would reinforce the development of different representations of the two differently-shaped environments. The animals are *not* trained to differentiate the square and the circle. Accordingly, we believe the paradigm provides a good example of incidental learning and memory. This may be particularly useful as this unreinforced type of learning is often emphasised in theories of hippocampal function (Cohen and Eichenbaum 1993; Morris and Frey 1997; O'Keefe and Nadel 1978).

4.2 Individual cells – can we identify different hippocampal mechanisms involved in incremental remapping?

This section describes data from two individual place cells in enough detail to suggest that there may be different hippocampal mechanisms involved in incremental remapping. First, however, we need to address the question: is the incremental remapping seen in Lever et al. (2002) effected by specifically hippocampal synaptic plasticity? While we cannot be certain that the learning site(s) is/are hippocampal, it is the most compelling possibility. An interesting study which partly motivated Lever et al. (2002) found that cells in the superficial layers of entorhinal cortex, which comprise the dominant projection to the hippocampus, showed inter-shape similarity in exactly those circumstances (circles and squares similar to ours) in which their hippocampal cells showed divergence (Quirk et al. 1992). This permits some confidence that hippocampal processes are responsible for incremental remapping. Nev-

ertheless, future work in our laboratory will address the issue of learning sites.

We now turn to remapping in individual cells and discuss two examples of remapping place cells while focusing on the transition from similar to divergent firing patterns. The first cell (Fig. 6) becomes gradually heterotopic by developing a second field in the square, and then losing the original homotopic field (Fig. 6A). By day 21 the cell had established a heterotopic firing pattern with a northern field in the circle and a southeastern field in the square environment. This pattern remained stable during subsequent testing and generalised to shapes of different material (not shown). During the transition period both fields were evident in the square environment (see day 20, Trials D and F). These trials are examined in detail in Fig. 6B, C. At the beginning of each trial the cell fired in the homotopic location, followed by firing in both locations, and ending at the divergent location alone (see Lever et al. 2002 for alternative time slicing). The cell shows a kind of “two-steps forward, one-step back” process. Although varying interpretations are possible, the data appear to suggest that two processes with different time scales are involved. We take up the implications of this after consideration of the second remapping place cell.

Rather than changing field position, this cell ceases firing in the square, while continuing to fire in the circle (Fig. 7A). The top row shows data recorded in the circular environment transformed into the square form (see Lever et al. 2002 for details) to permit a direct comparison with the corresponding square trials (bottom row). From day 16 onwards, the cell fires in the circle only (top row, D16–D20). We have previously emphasised the gradual decline in firing in the square on days 13 to 15. Firing rate peaks in the square are similar to those in the circle on day 13, about half those in the circle on day 14, and less than a quarter of those in the circle on day 15 (Lever et al. 2002, Fig. 3d). Figure 7B presents further details from a crucial phase of this transition period, comprising the last two trials of day 14, and all six trials of day 15. What we wish to draw attention to here is the remarkably clear evidence for a “two-steps forward, one-step back” process during this phase. In Fig. 7B, firing rate maps are shown for the first, middle, and last third of each trial. Thus we can consider the *within-trial* and *across-trial* dynamics. The across-trial data show a clear incremental decrease in firing rates in the square. The striking feature of the within-trial dynamics is that firing is consistently *much higher at the beginning of each square trial* than at its end, while there is no such relationship in the circle where firing is roughly constant. This pattern is reminiscent of the cell in Fig. 6. For both cells, the firing at the beginning of the square trial is similar to that in the circle, then becomes more dissimilar as the square trial proceeds. The degree of divergence reached at the end of a square trial is not obtained at the beginning of the next. Note that firing in the first 60 seconds in the square, on trial days 14f and 15b but not on day 15d, f (Fig. 7B, insets to right), is clearly comparable to that seen in the circle.

Finally, it perhaps needs clarifying that we could not see evidence of this type of firing pattern change (i.e. both the across- and within-trial rate decrease) in the

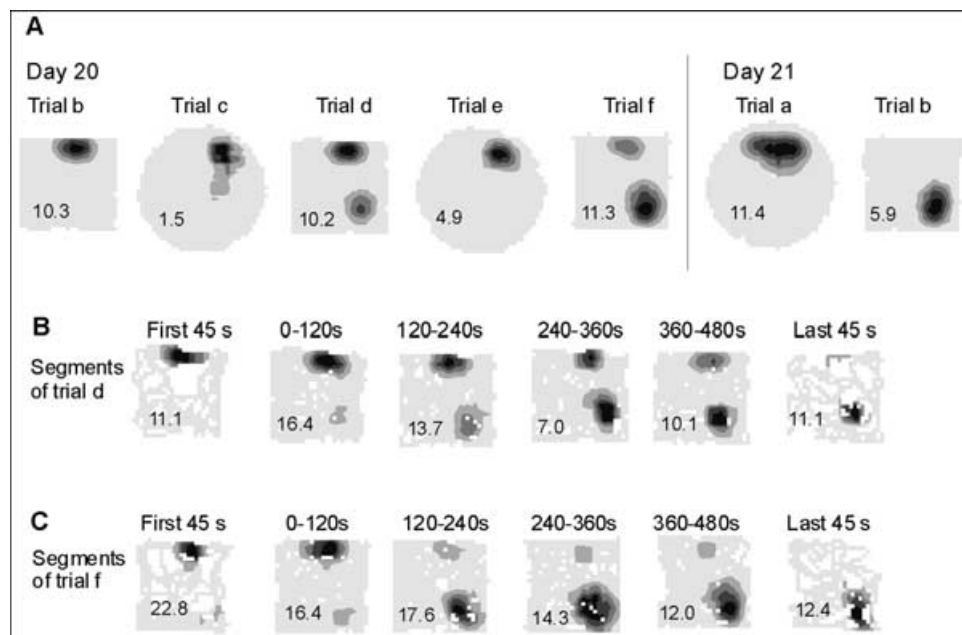


Fig. 6A–C. Example of a cell with initially homotopic firing pattern (similar position in both shapes) that gradually develops a heterotopic pattern (different position in each shape). **A)** Firing rate maps from seven consecutive trials beginning and ending with square trials (from Trial b of day 20 to Trial b of day 21 of main experiment in Lever et al. 2002) showing the evolution of the heterotopic pattern. **B)** and **C)** Firing rate maps of smaller time segments taken from trials d (B)

and f (C) in the square. First and last maps on each row show first and last 45 s of each trial respectively. Middle four maps show each trial segmented into equal quarters of 2 min each. These temporal sequences reveal the dynamics of the processes underlying remapping in a single cell. Note that the pattern divergence occurs both within trials (rapidly) and between trials (more slowly). Trial times were 10 min in the circle, and 8 min in the square (also applies to Fig. 7)

other cells recorded at the same time. We re-emphasise here that the precise nature and time course of the remapping is individual to each cell.

4.3 Shedding and recruitment in the network

The above consideration of changes in the contribution of single cells to the hippocampal representation should not blind us to the possibility of additional processes of changes in the hippocampal representation across cells. If we consider that the “active subset” of cells in a network that fire in an environment represent that environment (Muller 1996; McNaughton and Nadel 1990) then we must appreciate that this active subset changes. In other words, the hippocampal network may both shed cells from, and recruit cells to, the active subset.

Definitive evidence for these processes is difficult to obtain but we suspect that both of these possibilities, particularly recruitment, do occur. The occurrence of recruitment is consistent with our findings in Lever et al. (2002) that more cells per day were recorded later, rather than earlier, in the time-series from each of the three animals. Note that if some cells initially firing in both environments later become silent in one of them (for which we have good evidence) in a situation of *no recruitment* at all, then it follows that later on there will be fewer cells in absolute terms representing each environment. This seems unlikely given that we continue to observe plenty of active cells near to the tetrode throughout several weeks of continuous re-

coding. Pattern differentiation involving cells becoming silent thus implies recruitment and some sort of normalisation process through which the total number of active cells in an environment remains approximately constant. These issues are vital for understanding the network dynamics in a more than superficial manner. What about the data? The demonstration of both shedding and recruitment face the same technical problem. It is hard to provide convincing evidence that a particular cell is within the sensitive range of the electrode, but is not firing. This problem may be helped by recording in sleep or anaesthesia, when normally silent cells may fire, before or after environmental experience (Best and Thompson 1989; Wilson and McNaughton 1993), or by more global imaging methods (Guzowski et al., 1999).

5 Discussion

5.1 Incidental or unsupervised Hebbian learning

Reward-mediated plasticity is clearly required by models of behavioural learning and experimental evidence for its occurrence is growing (e.g. Kilgard and Merzenich 1998; Schultz et al. 1997). Many models of mammalian learning in the rat and monkey involve using reward to shape the animal’s learning and behaviour. These models may be more tractable in terms of analysing and manipulating learning at the level of the animal’s behaviour. They fail to capture, however, the kind of learning that Hebb often

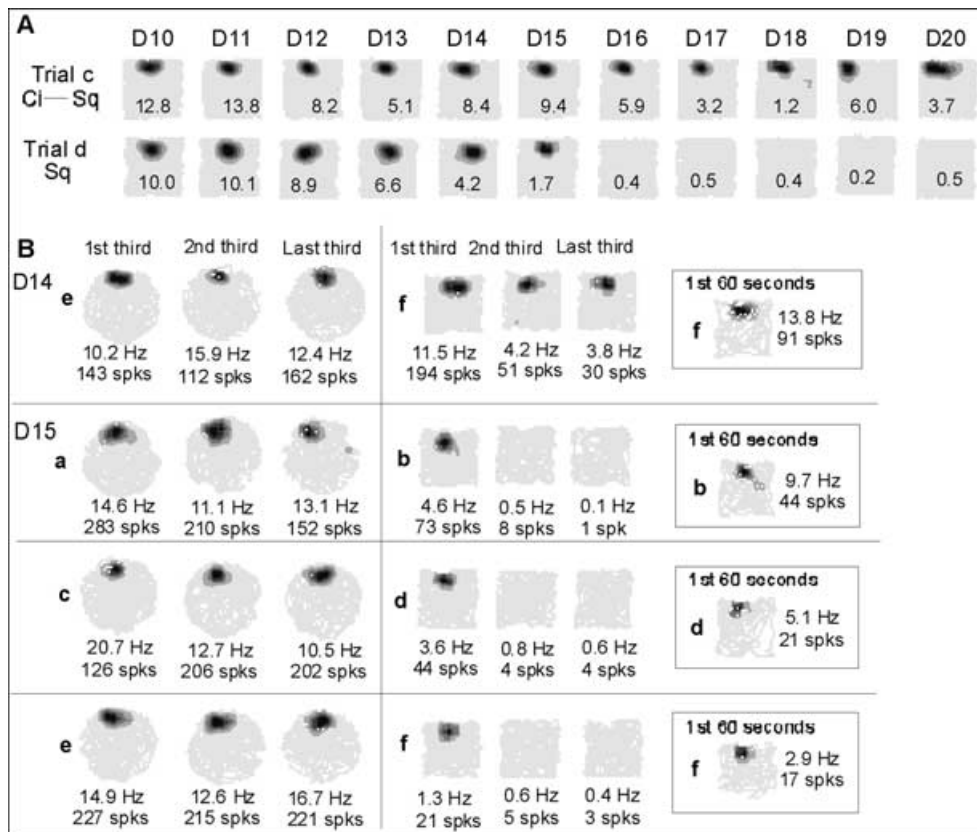


Fig. 7A–C. Example of a cell that initially fires in both shapes in related locations and gradually comes to fire only in one shape only (the circle). **A**) Firing rate maps from middle trials on each day from day 10 to day 20 of main experiment in Lever et al. (2002). Top row shows transformed-circle trials (i.e. after topological transformation of the circle data into square data), for direct comparison with firing in the square (bottom row). After day 15, the cell fires in circle only (top row, D16–D20). **B**) Firing rate maps of smaller time segments taken from the eight consecutive trials from Trial e on day 14 to Trial f on

day 15. Firing rate maps are shown for the 1st, 2nd, and last third of each trial. Right-hand panels show the first 60 seconds of each trial in the square. Note that the firing is always highest in the first third of each square trial, and then declines quite rapidly within the trial. Indeed, a disproportionate amount of firing takes place in the first 60 s. Note too that this early-phase firing also declines in strength over trials, from the last trial of day 14 to the last trial of day 15. Firing in the circle shows none of these patterns, and is basically stable

emphasised and indeed that which is often equated with learning from experience in the popular mind: the automatic acquisition of knowledge. Here we have presented evidence at the level of cells and assemblies for the unsupervised learning described by Hebb. We have described two types of plasticity in place cells associated with environmental changes that are good candidates for such Hebbian learning. First, the rapid changes sometimes induced by insertion of a barrier. Second, the slow, experience-dependent, incremental divergence of the representations of environments of different shapes over trials that are sometimes accompanied by more rapid but less permanent changes within a trial. An important issue to be explored is the degree to which the rapid plasticity seen in barrier experiments is stable in the long term, as we suspect.

There are other candidates for Hebbian learning paradigms involving place cells. Perhaps the most obvious candidate is the experience dependent development of asymmetry of place fields (Ekstrom et al. 2001; Mehta et al. 1997, 2000). Although clearly demonstrated to be dependent on the NMDA receptor, this plasticity

seems to be rather short-lived (resetting daily despite repetitions) compared to the creation of long-lasting cell assemblies envisaged by Hebb. Also, it is perhaps less easy to see evolutionary selection pressure in the behavioural uses for asymmetric or expanding field shape than, say, map divergence (e.g. Bostock et al. 1991; Jeffery 2000; Kentros et al. 1998; Lever et al. 2002; Sharp 1997; but see Blum and Abbott 1996).

5.2 What kinds of plasticity are involved in incremental remapping?

Can our data be related to physiological models of hippocampal plasticity? Figures 6 and 7 suggest a possible role for depression (short term and long term) as well as LTP, and more speculatively perhaps, the potential for both processes at the same synapse as in the BCM learning rule used by Fuhs and Touretzky (2000). Recent plasticity studies have suggested that both depression and enhancement can occur depending upon activity frequency and timing (Martin et al. 2000), as

suggested by the BCM rule, and that this frequency-response curve can be altered in favour of LTD in mouse mutants (Bach et al. 1995). It might be interesting to compare these mutant mice against mutants with impairments restricted to the processes related to LTP alone in a gradual remapping paradigm.

What does the time-course of remapping tell us? Although there can be several interpretations of the intra- and inter-trial processes in Figs. 6 and 7, they are sufficiently striking to pose an important test for any remapping model to reproduce them in some of the cells of a simulated network. Can the “two steps forward, one step back” remapping of both cells be interpreted in the same way?

One possibility is that comparable mechanisms for Hebbian learning exist in both long- and short-term forms. Indeed, a short-term (i.e. rapidly decaying) potentiation seems to invariably accompany the occurrence of the more famous long-term potentiation in experimental studies (e.g. Bliss and Collingridge 1993; McNaughton 1982). These parallel forms of plasticity have been reflected in various computational models of short- and long-term memory processes as ‘fast’ and ‘slow’ connection weights (e.g. Gardner-Medwin 1989; Hinton and Plaut 1987; Burgess and Hitch 1999). Thus, similar changes to a connection weight may occur both as a large amplitude but rapidly decaying change, and as a small amplitude permanent adjustment.

An interesting aspect of the place cell changes shown in Figs. 6 and 7 seems to point strongly towards a parallel ‘fast and slow’ interpretation. That is, whatever process causes the slow divergence in the representation of each environment across trials, *the same process* also appears to occur within each trial. However, the changes within each trial proceed relatively quickly, while only a small and incremental residue of the within-trial changes remain over the longer term. Thus the process of divergence at the beginning of the next trial has advanced only a modest amount from that at the beginning of the previous trial despite the rapid advance made during the trial.

A different level of interpretation could include a hierarchical process of environmental recognition, in which the balance between the processes of pattern completion and pattern separation alters over time. Thus, on entry to an environment, it might be behaviourally useful to first classify the general type of environment and then to become successively more specific. Within an auto-associative memory this might correspond to relaxation of the pattern of activation to the basin of attraction for the representations of all similar environments followed by settling into the representation corresponding to a specific environment. Alternatively, within CA3, there might be some dynamic alteration to the relative influence of the inputs from the dentate gyrus (supporting pattern separation) compared to the recurrent collaterals (supporting pattern completion).

5.3 Relationship to hippocampal network models

The geometric model captures much of the place cell data in a single static environment. Using barriers, we

can see that bimodal place fields can be created in accordance with the model. We should perhaps stress that our data (not presented here) also clearly shows that such bimodality, indicative of local or partial remapping (Muller and Kubie 1987; Muller 1996), occurs within the same representation where the active subset of cells representing the square does not change i.e. there is no complex remapping (Muller 1996).

However, the dynamic field changes we have described require synaptic plasticity. The rapid remapping shown by Bostock et al. (1991) might be consistent with addition of a learned colour preference to the BVCs of the geometric model. In addition, the reduction in firing rates following novel cue card manipulations (Fenton and Muller 2000) might be consistent with a strengthening of the, initially hardwired, inputs to place cells that are active in a much visited environment (Burgess and Hartley 2002). In the more general terms of modeling remapping by incorporating plasticity into feed-forward models, Fuhs and Tourtzky’s (2000) model looks pertinent to some of the data presented here, such as in Figs. 5 and 7. However, this model only deals with the evolution of monotopy, i.e. how a cell becomes silent in one of the two environments, and would need to also account for the processes whereby homotopic patterns are replaced by heterotopic patterns.

In terms of auto-associative models, the Charts model (Samsonovich and McNaughton 1997) does not deal well with bimodal firing such as seen in Fig. 4, and predicts only instantaneous, discontinuous remappings between different environments. Such remapping should be seen in the time it takes to minimally sample a standard environment (1–3 minutes). Although Bostock et al. (1991) observed rapid remapping (in many but not all their animals) after changing the colour of the cue card, even this form of remapping was not fully expressed until the next day. Clearly, obligatorily instantaneous remapping is also inconsistent with our recent data. A revised attractor model (e.g. Kali and Dayan 2000) could possibly capture important aspects of the plasticity we have seen although these would still predict place fields that maintained their relative locations rather than the absolute distances to boundaries shown by O’Keefe and Burgess (1996), (see earlier discussion of these models).

How conservative can we be in adducing mechanisms needed to explain the full spectrum of data on stability and remapping? Can a sliding threshold based on the magnitude of difference between environments explain Bostock et al.’s (1991) remapping in terms of a very fast version of gradual remapping? It will be important to describe the temporal dynamics and other aspects of remapping in more detail in order to appreciate any distinctions that may exist between map separation effected by rapid remappings based on several multimodal environmental differences (as in Kentros et al. 1998) and slow remapping effected through repeated experience (Lever et al. 2002). Both may result in an equally thorough pattern divergence. Clearly, as well as environments being “sufficiently different” to induce remapping

(Muller 1996), experience is also important. Remapping may proceed even with small differences, so long as they are perceived to be stable differences.

6 Conclusions

We have argued that study of the hippocampal representation of environmental geometry in freely moving rats should provide one of the best paradigms within which to observe the effects of Hebbian learning at the level of single cells in vivo. The slow, long-term and incremental plasticity of this representation observed across environments of different shape (Lever et al. 2002) appears to be at least consistent with Hebb's postulate. These changes also appear to be consistent with the incidental or unsupervised nature of the type of learning stressed by Hebb (1949), and additionally indicate constraints on the nature and time-course of its exact implementation. Other forms of plasticity in the place cell representation, such as some of the rapid changes caused by insertion of a barrier (Sect. 3), or changes to the sensory features of an environment (e.g. Bostock et al. 1991) may indicate the presence of long-term changes occurring over much shorter time-scales. Interestingly, closer analysis of the dynamics of the slow incremental shape-based remapping also reveals the action of a faster but less enduring form of plasticity. Taken together with our quantitative understanding of the basic feed-forward organisation of the place cell system (Hartley et al. 2000), these data provide a powerful test-bed for investigation of the mechanisms of plasticity at work within a cognitive representation.

Acknowledgements. This work was supported by the Medical Research Council of the United Kingdom.

References

- Amit DJ (1989) Modeling brain function. The world of attractor neural networks. Cambridge University Press
- Bach ME, Hawkins RD, Osman M, Kandel ER, Mayford M (1995) Impairment of spatial but not contextual memory in CaMKII mutant mice with a selective loss of hippocampal LTP in the range of the theta frequency. *Cell* 81: 905–915
- Barnes CA, Suster MS, Shen J, McNaughton BL (1997) Multistability of cognitive maps in the hippocampus of old rats. *Nature* 388: 272–275
- Best PJ, Thompson LT (1989) Persistence, reticence, and opportunism of place-field activity in hippocampal neurons. *Psychobiol* 17: 230–235
- Bienenstock EL, Cooper LN, Munro PW (1982) Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J Neurosci* 2: 32–48
- Bliss TVP, Collingridge GL (1993) A synaptic model of memory: long-term potentiation in the hippocampus. *Nature* 361: 31–39
- Bliss TV, Lomo T (1973) Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J Physiol Lond* 232: 331–356
- Blodgett HC (1929) The effect of the introduction of reward upon the maze performance of rats. *Univ Calif Publ Psychol* 4: 113–134
- Blum KI, Abbott LF (1996) A model of spatial map formation in the hippocampus of the rat. *Neural Comput* 8: 85–93
- Bostock E, Muller RU, Kubie JL (1991) Experience-dependent modifications of hippocampal place cell firing. *Hippocampus* 1: 193–205
- Brunel N, Trullier O (1998) Plasticity of directional place fields in a model of rodent CA3. *Hippocampus* 8: 651–665
- Burgess N, Hartley T (2002) Orientational and geometric determinants of place and head-direction. *Neural Inf Process Syst* 14 (in press)
- Burgess N, Hitch GJ (1999) Memory for serial order: a network model of the phonological loop and its timing. *Psych Rev* 104: 551–581
- Burgess N, Jackson A, Hartley T, O'Keefe J (2000) Predictions derived from modeling the hippocampal role in navigation. *Biol Cybern* 83: 301–312
- Burgess N, O'Keefe J (1996) Neuronal computations underlying the firing of place cells and their role in navigation. *Hippocampus* 6: 749–762
- Cohen NJ, Eichenbaum H (1993) Memory, amnesia, and the hippocampal system. MIT Press, Cambridge, Mass
- Davis S, Butcher SP, Morris RG (1992) The NMDA receptor antagonist D-2-amino-5-phosphonopentanoate (D-AP5) impairs spatial learning and LTP in vivo at intracerebral concentrations comparable to those that block LTP in vitro. *J Neurosci* 12: 21–34
- Ekstrom AD, Meltzer J, McNaughton BL, Barnes CA (2001) NMDA receptor antagonism blocks experience-dependent expansion of hippocampal "place fields" *Neuron* 31: 631–638
- Fenton AA, Muller RU (2000) Conjoint control of hippocampal place cell firing by two visual stimuli. *J Gen Physiol* 116: 191–209
- Fuhs MC, Touretzky DS (2000) Synaptic learning models of map separation in the hippocampus. *Neurocomput* 32: 379–384
- Gardner-Medwin AR (1976) The recall of events through the learning of associations between their parts. *Proc Roy Soc Lond B Biol Sci* 194: 375–402
- Gardner-Medwin AR (1989) Doubly modifiable synapses: a model of short and long term auto-associative memory. *Proc Roy Soc London B Biol Sci* 238: 137–154
- Gothard KM, Skaggs WE, McNaughton BL (1996) Dynamics of mismatch correction in the hippocampal ensemble code for space: interaction between path integration and environmental cues. *J Neurosci* 16: 8027–8040
- Gustafsson B, Wigstrom H (1986) Hippocampal long-lasting potentiation produced by pairing single volleys and brief conditioning tetani evoked in separate afferent. *J Neurosci* 6: 1575–1582
- Guzowski JF, McNaughton BL, Barnes CA, Worley PF (1999) Environment-specific expression of the immediate-early gene Arc in hippocampal neuronal ensembles. *Nature Neurosci* 2: 1120–1124
- Harley CW (1979) Arm choices in a sunburst maze: effects of hippocampotomy in the rat. *Physiology and Behavior* 23: 283–290
- Hartley T, Burgess N, Lever C, Cacucci F, O'Keefe J (2000) Modeling place fields in terms of the cortical inputs to the hippocampus. *Hippocampus* 10: 369–379
- Hasselmo ME, Wyble BP, Wallenstein GV (1996) Encoding and retrieval of episodic memories: role of cholinergic and GABAergic modulation in the hippocampus. *Hippocampus* 6: 693–708
- Hebb DO (1949) The organization of behavior. Wiley, New York
- Hill AJ (1978) First occurrence of hippocampal spatial firing in a new environment. *Exp Neurol* 62: 282–297
- Hinton GE, Plaut DC (1987) Using fast weights to deblur old memories. Proceedings of the Ninth Annual Conference of the Cognitive Science Society, Seattle, WA
- Hollup SA, Molden S, Donnett JG, Moser MB, Moser EI (2001) Accumulation of hippocampal place fields at the goal location in an annular watermaze task. *J Neurosci* 21: 1635–1644

- Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci USA* 79: 2554–2558
- Jeffery KJ (2000) Plasticity of the hippocampal cellular representation of place. In: Holscher (ed) *Neuronal Mechanisms of Memory Formation* Cambridge University Press, Cambridge
- Jensen O, Lisman JE (1996) Hippocampal CA3 region predicts memory sequences: accounting for the phase precession of place cells. *Learn Mem* 3: 279–287
- Kali S, 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
- Kentros C, Hargreaves E, Hawkins RD, Kandel ER, Shapiro M, Muller RV (1998) Abolition of long-term stability of new hippocampal place cell maps by NMDA receptor blockade. *Science* 280: 2121–2126
- Keith JR, McVety KM (1988) Latent place learning in a novel environment and the influences of prior training in rats. *Psychobiol* 16: 146–151
- Kilgard MP, Merzenich MM (1998) Cortical map reorganization enabled by nucleus basalis activity. *Science* 279: 1714–1718
- Kohonen T (1972) Correlation matrix memories. *IEEE Trans Comp C-21*: 353–359
- Lever C, Wills T, Cacucci F, Burgess N, O'Keefe J (2002) Long-term plasticity in the hippocampal place-cell representation of environmental geometry. *Nature* 416: 90–94
- Levy WB, Steward O (1983) Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Neurosci* 8: 791–797
- Marr D (1971) Simple memory: a theory for archicortex. *Philos Trans R Soc Lond B Biol Sci* 262: 23–81
- Martin SJ, Grimwood PD, Morris RGM (2000) Synaptic plasticity and memory: an evaluation of the hypothesis. *Annu Rev Neurosci* 23: 649–711
- McClelland JL, McNaughton BL, O'Reilly RC (1995) Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol Rev* 102: 419–457
- McHugh TJ, Blum KI, Tsien JZ, Tonegawa S, Wilson MA (1996) Impaired hippocampal representation of space in CA1-specific NMDAR1 knockout mice. *Cell* 87: 1339–1349
- McNaughton BL (1982) Long-term synaptic enhancement and short-term post-tetanic potentiation in the rat fascia dentata act through different mechanisms. *J Physiology* 324: 249–262
- McNaughton BL (1996) Cognitive cartography. *Nature* 381: 368–369
- McNaughton BL, Barnes CA, O'Keefe J (1983) The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Exp Brain Res* 52: 41–49
- McNaughton BL, Barnes CA, Gerrard JL, Gothard K, Jung MW, Knierim JJ, Kudrimoti H, Qin Y, Skaggs WE, Suster M, Weaver KL (1996) Deciphering the hippocampal polyglot: the hippocampus as a path integration system. *J Exp Biol* 199(Pt 1): 173–185
- McNaughton BL, Morris RG (1987) Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in Neurosci* 10: 408–415
- McNaughton BL, Nadel L (1990) Hebb-Marr networks and the neurobiological representation of action in space. In: Gluck MA, Rumelhart DE (eds) *Neurosci and connectionist theory*, Lawrence Erlbaum Assoc, Hillsdale, NJ pp 1–63
- Mehta MR, Barnes CA, McNaughton BL (1997) Experience-dependent, asymmetric expansion of hippocampal place fields. *Proc Natl Acad Sci USA* 94: 8918–8921
- Mehta MR, Quirk MC, Wilson MA (2000) Experience-dependent asymmetric shape of hippocampal receptive fields. *Neuron* 25: 707–715
- Morris RG, Garrud P, Rawlins JN, O'Keefe J (1982) Place navigation impaired in rats with hippocampal lesions. *Nature* 297: 681–683
- Morris RG, Frey U (1997) Hippocampal synaptic plasticity: role in spatial learning or the automatic recording of attended experience? *Philos Trans R Soc Lond B Biol Sci* 352: 1489–1503
- Muller RU (1996) A quarter of a century of place cells. *Neuron* 17: 813–822
- 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
- Muller RU, Bostock E, Taube JS, Kubie JL (1994) On the directional firing properties of hippocampal place cells. *J Neurosci* 14: 7235–7251
- Muller RU, Stead M, Pach J (1996) The hippocampus as a cognitive graph. *J Gen Physiol* 107: 663–694
- Muller RU, Stead M (1996) Hippocampal place cells connected by Hebbian synapses can solve spatial problems. *Hippocampus* 6: 709–719
- Nakazawa K, Quirk MC, Chitwood RA, Watanabe M, Yeckel MF, Sun LD, Kato A, Carr CA, Johnston D, Wilson MA, Tonegawa S (2002) Requirement for hippocampal CA3 NMDA receptors in associative memory recall. *Science* 297: 211–218
- 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 neurons. *Nature* 381: 425–428
- O'Keefe J, Conway DH (1978) Hippocampal place units in the freely moving rat: why they fire where they fire. *Exp Brain Res* 31: 573–590
- O'Keefe J, Dostrovsky J (1971) The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res* 34: 171–175
- O'Keefe J, Nadel L (1978) *The hippocampus as a cognitive map*. Clarendon Press, Oxford
- O'Keefe J, Speakman A (1987) Single unit activity in the rat hippocampus during a spatial memory task. *Exp Brain Res* 68: 1–27
- O'Keefe J, Recce ML (1993) Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3: 317–330
- Quirk GJ, Muller RU, Kubie JL (1990) The firing of hippocampal place cells in the dark depends on the rat's recent experience. *J Neurosci* 10: 2008–2017
- Quirk GJ, Muller RU, Kubie JL, Ranck JB (1992) The positional firing properties of medial entorhinal neurons: description and comparison with hippocampal place cells. *J Neurosci* 12: 1945–1963
- Redish AD, Touretzky DS (1999) Separating hippocampal maps. In: Burgess N, Jeffery KJ, O'Keefe J (eds) *The hippocampal and parietal foundations of spatial cognition*. Oxford University Press, pp 203–219
- Rotenberg A, Mayford M, Hawkins RD, Kandel ER, Muller RU (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
- Rotenberg A, Abel T, Hawkins RD, Kandel ER, Muller RU (2000) Parallel instabilities of long-term potentiation, place cells, and learning caused by decreased Protein Kinase A activity. *J Neurosci* 20: 8096–8102
- Rumelhart DE, Zipser D (1986) Feature discovery by competitive learning. In: Rumelhart DE, McClelland JL (eds) *Parallel distributed programming Vol 1: foundations*, MIT Press, pp 151–193
- Samsonovich A, McNaughton BL (1997) Path integration and cognitive mapping in a continuous attractor neural network model. *J Neurosci* 17: 5900–5920
- Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. *Science* 275: 1593–1599
- Sharp PE (1991) Computer simulation of hippocampal place cells. *Psychobiol* 19: 103–115

- Sharp PE (1997) Subicular cells generate similar spatial firing patterns in two geometrically and visually distinctive environments: comparison with hippocampal place cells. *Behav Brain Res* 85: 71–92
- Steele RJ, Morris RGM (1999) Delay-dependent impairment of a matching-to-place task with chronic and intrahippocampal infusion of the NMDA-antagonist D-AP5. *Hippocampus* 9: 118–136
- Tanila H, Shapiro M, Gallagher M, Eichenbaum H (1997) Brain aging: changes in the nature of information coding by the hippocampus. *J Neurosci* 17: 5155–5166
- Taube JS, Muller RU, Ranck-JB J (1990) Head-direction cells recorded from the postsubiculum in freely moving rats. I: Description and quantitative analysis. *J Neurosci* 10: 420–435
- Taube JS (1998) Head direction cells and the neuropsychological basis for a sense of direction. *Prog Neurobiol* 55: 225–256
- Tolman EC (1932) *Purposive behavior in animals and men*. Century, New York
- Tolman EC (1948) Cognitive maps in rats and men. *Psychol Rev* 55: 189–208
- Treves A, Rolls ET (1992) Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. *Hippocampus* 2: 189–199
- Tsodyks MV, Skaggs WE, Sejnowski TJ, McNaughton BL (1996) Population dynamics and theta rhythm phase precession of hippocampal place cell firing: a spiking neuron model. *Hippocampus* 6: 271–280
- Wilson MA, McNaughton BL (1993) Dynamics of the hippocampal ensemble code for space. *Science* 261: 1055–1058
- Zhang K (1996) Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. *J Neurosci* 16: 2112–2126
- Zipser D (1985) A computational model of hippocampal place fields. *Behav Neurosci* 99: 1006–1018