

Learning in a Geometric Model of Place Cell Firing

Caswell Barry^{1,2,3*} and Neil Burgess^{1,2}

ABSTRACT: Following Hartley et al. (Hartley et al. (2000) *Hippocampus* 10:369–379), we present a simple feed-forward model of place cell (PC) firing predicated on neocortical information regarding the environmental geometry surrounding the animal. Incorporating the idea of boundaries with distinct sensory qualities, we show that synaptic plasticity mediated by a BCM-like rule (Bienenstock et al. (1982) *J Neurosci* 2:32–48) produces PCs that encode position relative to specific extended landmarks. In an unchanging environment the model is shown to undergo an initial phase of learning, resulting in the formation of stable place fields. In familiar environments, perturbation of environmental cues produces graded changes in the firing rate and position of place fields. Model simulations are compared favorably with three sets of experimental data: (1) Results published by Barry et al. (Barry et al. (2006) *Rev Neurosci* 17:71–97) showing the slow disappearance of duplicate place fields produced when a barrier is placed into a familiar environment. (2) Rivard et al.'s (Rivard et al. (2004) *J Gen Physiol* 124:9–25) study showing a graded response in PC firing such that fields near to a centrally placed object encode space relative to the object, whereas more distant fields respond to the surrounding environment. (3) Fenton et al.'s (Fenton et al. (2000a) *J Gen Physiol* 116:191–209) observation that inconsistent rotation of cue cards produces parametric changes in place field positions. The merits of the model are discussed in terms of its extensibility and biological plausibility. © 2007 Wiley-Liss, Inc.

KEY WORDS: hippocampus; navigation; neural network; computational model; spatial memory

INTRODUCTION

Pyramidal cells in the rodent hippocampus exhibit spatially localized firing (O'Keefe and Dostrovsky, 1971), presumably on the basis of the neocortical inputs to the hippocampal formation. Place cells (PCs) fire whenever the animal enters a specific area of its environment, the "place field," and do so independently of heading (Muller et al., 1994). Location-specific firing appears rapidly after an animal enters a new environment (Hill, 1978; Wilson and McNaughton, 1993) and is reinstated on subsequent visits, even months later (Best and Thompson, 1984). Subsequently discovered in other animals (Rolls et al., 1989; Ono et al., 1991; Bingman et al., 2006; Ulanovsky and Moss, 2007), including humans

(Ekstrom et al., 2003), a small population of PCs is sufficient to accurately encode an animal's position (Wilson and McNaughton, 1993).

How is PC firing controlled? Early experiments showed that firing was not the product of a single sensory modality (Best and Thompson, 1989) and was robust to removal of subsets of cues (O'Keefe and Conway, 1978; Save et al., 2000; Nakazawa et al., 2002). Subtle perturbation of an animal's environment, such as lengthening along one or both dimensions (O'Keefe and Burgess, 1996) or changing the relative position of cues (Fenton et al., 2000a), produced parametric changes in place fields. Further, the representation of identical environments connected by a corridor is initially similar (Skaggs and McNaughton, 1998) despite the existence of a route between them. Synaptic plasticity is also important, in the presence of NMDA block (Kentros et al., 1998) or protein synthesis inhibitors (Agnihotri et al., 2004) place fields fail to stabilize, and several experimenters have observed the accumulation of small plastic changes in PC firing (Lever et al., 2002b; Barry et al., 2006). Together these findings suggest that place fields are defined, at least in part, by their position relative to environmental cues, a relationship that is learnt or refined during exploration.

These observations are addressed by models that describe place fields as a function of the geometric position of cues around an animal (Zipser, 1986; Sharp, 1991; Burgess et al., 1994; O'Keefe and Burgess, 1996; Hartley et al., 2000). In particular, Hartley et al.'s (2000) model posits the existence of boundary vector cells (BVCs), a class of cells that respond to the presence of any barrier, such as a wall or a drop, at a given distance and allocentric direction. Though not explicit in the model, the position of BVCs upstream of the hippocampus, together with their multimodal, highly processed geometric response, suggest a neocortical location for these cells; indeed Hartley et al. (2000) proposed that they might be found in the entorhinal cortex, and some evidence for BVCs has been found in the subiculum (Barry et al., 2006). The thresholded sum of a population of these cells provides a good analogue of PC firing and makes accurate predications as to how place fields respond to geometric manipulation of an environment. Predictions include: the duplication of place fields by addition of a barrier (Lever et al., 2002a; Barry et al., 2006); that some fields will follow the shape of existing barriers (e.g. crescent fields in a circular enclosure) (Muller et al., 1987); and specific predictions as to how individual fields would change

¹Institute of Cognitive Neuroscience, University College London, London, United Kingdom; ²Department of Anatomy and Developmental Biology, University College London, London, United Kingdom; ³Department of Psychology, University College London, London, United Kingdom

Grant sponsors: MRC, BBSRC.

*Correspondence to: Caswell Barry, Institute of Cognitive Neuroscience, 17 Queen Square, London WC1N 3AR, United Kingdom. E-mail: caswell.barry@ucl.ac.uk or Neil Burgess, Institute of Cognitive Neuroscience, 17 Queen Square, London WC1N 3AR, United Kingdom.

E-mail: n.burgess@ucl.ac.uk

Accepted for publication 1 May 2007

DOI 10.1002/hipo.20324

Published online 27 June 2007 in Wiley InterScience (www.interscience.wiley.com).




FIGURE 1. (A) Each BVC responds optimally when a barrier is encountered at a preferred distance and allocentric direction from the animal. The BVC in this example shows an increasingly vigorous response as the rat moves closer to a wall, causing the wall to overlap with the BVCs receptive field. (B) In the model, firing of a population of BVCs afferent to a PC is summed, and thresholded to calculate the cell's firing. Using this approach the PC's firing can be estimated for any arrangement of barriers. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

when the shape of an enclosure was changed (Hartley et al., 2000). The BVC framework also predicts human search behavior when returning to a previously visited location and the effects upon it of manipulations to environmental geometry (Hartley et al., 2004). The model continues to provide insight into contemporary findings. For example, it explains the observation that place fields were not affected by the substitution of one item for another but were changed by shifts in the relative position of those objects (Burgess and Hartley, 2002; Lenck-Santini et al., 2005).

The original BVC model, however, simply contained hard-wired connections of equal strength from BVCs to PCs; no attempt was made to capture experiential changes in PC firing. Here we describe an extension of the BVC model allowing experience-dependent plasticity in the connections from BVCs to PCs, and allowing experience-dependent modification of the response to a boundary according to its sensory qualities. See Barry et al. (2006) for preliminary simulations with the BCM rule (Bienenstock et al., 1982; Fuhs and Touretzky, 2000), and Burgess and Hartley (2002) for preliminary simulations of the effect of allowing sensitivity to a boundary's sensory qualities. We show that the new model produces PCs that encode position relative to specific objects; an area west of a black cue card, for example. We apply this model to data we recorded in collaboration with K. Jeffery (Barry et al., 2006), as well as results reported by Rivard et al. (2004) and Fenton et al. (2000a). In general, the model shows good accordance with these independent experimental findings. In particular, it captures the experience-dependent disappearance of one or other of the duplicate place fields produced by the placement of an additional barrier into an enclosure (Barry et al., 2006). It also provides a mechanistic explanation of Rivard et al.'s observation that some PCs encode position in room centered coordinates while others fire relative to centrally placed objects (Rivard et al., 2004). Finally it captures the graded distortion in place field position produced by inconsistent rotation of cue cards together or apart (Fenton

et al., 2000a); unlike previous models (Fenton et al., 2000b; Burgess and Hartley, 2002; Touretzky et al., 2005), it does so at the level of single neurons without making prescriptive assumptions about the cues that PCs or their inputs will respond to.

MODEL FRAMEWORK

Following Hartley et al. (2000), PC firing was described as a function of the relative location of barriers around an animal. Specifically, the activity of PCs was considered to be the thresholded sum of inputs received from putative BVCs. The firing of BVCs themselves was a Gaussian tuned response to the presence of extended barriers (e.g. walls) around an animal's location, each BVC firing optimally when a barrier was encountered at a specific distance and allocentric direction. PC firing is then simply a thresholded sum of the firing rates of all the BVCs connected to it (Fig. 1).

The main departure from Hartley et al. (2000), see also Barry et al. (2006), was that BVCs were allowed to learn, such that they would come to respond more or less strongly to barriers with distinct sensory qualities (e.g. different colored cue cards or a free-standing barrier vs. the external wall of an enclosure). To capture the effect of learnt sensitivity to sensory qualities learning was described as occurring to sets of BVCs. All BVCs in each set shared the same distance and directional tuning but each responded solely to barriers of a specific type (Fig. 2A). For exam-

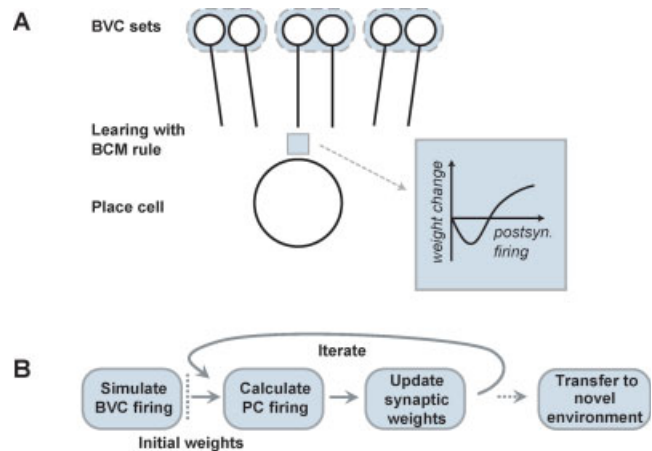


FIGURE 2. BVC sets consist of multiple BVCs with the same preferred firing direction and distance. Each BVC in a set responds to boundaries with different sensory qualities (e.g. one might respond to a cue card and another to the enclosure walls). All BVCs in a set initially drive the upstream PC with the same efficacy—the connections have the same weight. Through learning with the BCM rule, weights are updated and so a BVC set can come to drive the PC more strongly in response to certain barriers (e.g. the PC may be driven more strongly by a white cue card than by a black one). The sign of weight change produced by the BCM rule depends on the postsynaptic cell's firing rate—firing above a threshold leads to LTP, below the threshold produces LTD. Learning is applied iteratively, PC firing is calculated and weight changes applied, PC firing is recalculated, and so on. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

ple, in a black environment with a white cue card fixed to the wall, each BVC set would consist of two BVCs ($n = 2$), one of which would respond to the white cue card, the other to the black environment wall. Prior to learning, all the connections from the BVCs in a set to PCs have the same weight, but subsequent iterative application of a BCM-like learning rule (Bienenstock et al., 1982) allowed weights to vary (Fig. 2B). This schema is equivalent to allowing individual BVCs to change the efficacy with which they respond to barriers with distinct sensory properties.

More precisely, the contribution to the firing of BVC ik (tuned to distance d_i , angle ϕ_i , and responding to barrier type k) from a segment of boundary of type k at distance r in allocentric direction θ , subtending an angle $\delta\theta$ at the rat, is given by:

$$\delta f_{ik} = g_{ik}(r(\theta), \theta)\delta\theta \tag{1}$$

where

$$g_{ik}(r(\theta), \theta) \propto \frac{\exp[-(r(\theta) - d_i)^2 / 2\sigma_{\text{rad}}^2(d_i)]}{\sqrt{2\pi\sigma_{\text{rad}}^2(d_i)}} \times \frac{\exp[-(\theta - \phi_i)^2 / 2\sigma_{\text{ang}}^2]}{\sqrt{2\pi\sigma_{\text{ang}}^2}} \tag{2}$$

The angular width σ_{ang} is constant, but radial tuning width increases linearly with the preferred tuning distance: $\sigma_{\text{rad}}(d_i) = ((d_i/\beta) + 1) \times \sigma_0$. In the simulations presented here $\sigma_{\text{ang}} = 0.2$ radians, $\beta = 1,830$ mm, and $\sigma_0 = 122$ mm. The preferred firing distance (d_i) for each BVC set was selected randomly from the following values: 81.0, 169.0, 265.0, 369.0, 482.5, 606.5, and 741.0 mm. It can be seen that BVCs with shorter preferred firing distance, and hence narrower tuning curves, are more densely represented. The preferred firing direction (ϕ_i) for each BVC set was selected randomly from the continuous range $0-2\pi$.

For each location \underline{x} in the environment, the contribution of all boundaries of type k to the firing of BVC ik is determined by integrating Eq. (1) over θ :

$$f_{ik}(\underline{x}) = \int_0^{2\pi} g_{ik}(r(\theta), \theta) d\theta \tag{3}$$

Although a barrier not of type k does not directly contribute to the firing of a BVC that responds only to barrier type k , its presence in the environment can nonetheless occlude other barriers that would otherwise contribute to the cell's firing.

The firing of PC j at location \underline{x} (referred to as $F_j(\underline{x})$) is then proportional to the thresholded, weighted sum of the N BVC sets that connect to it, each BVC set containing n BVCs responding to n different boundary types:

$$F_j(\underline{x}) = \Theta \left(A \sum_{i=1, N} \sum_{k=1, n} f_{ik}(\underline{x})w_{jik} - T \right) \tag{4}$$

where the threshold T and coefficient A are constants, w_{jik} is the variable connection weight between BVC ik and PC j , and Θ is

the Heaviside function (i.e. $\Theta(x) = x$ if $x > 0$; $\Theta(x) = 0$ otherwise). Constant A simply allows all non-zero connection weights w_{jik} to take initial values of 1, for consistency with Hartley et al.'s (2000) model. In all simulations $A = 5,000$ and $T = 12$. Firing does not depend upon the rat's heading and the directional tuning of all BVCs is determined relative to the same allocentric reference frame (assumed to be provided by the head-direction system (Taube et al., 1990a,b). For example, a PC that received input from BVCs tuned to respond to nearby barriers to the South and to the East would fire when the animal was in the Southeast corner of its environment. The sharper tuning of BVCs with shorter preferred directions implies that boundaries near to the location of peak firing will tend to have more influence than boundaries further away. In all simulations PCs received inputs from 10 unique BVC sets ($N = 10$), the number of BVCs in each set varied according to the sensory complexity of the environment, but did not exceed three.

The BCM rule (Bienenstock et al., 1982; Fuhs and Touretzky, 2000) was used to update the weights of connections between BVCs and PCs (w_{jik}). Under the BCM rule sustained firing of the postsynaptic cell below a dynamic threshold leads to weakening of the connection from the presynaptic cell. Conversely, firing in the postsynaptic cell above the threshold leads to strengthening of the connection from the presynaptic cell. In both cases the magnitude of the change is dependent upon the rate of firing in the presynaptic cell, and if the presynaptic cell does not fire then no weight change occurs. The value of the threshold itself reflects the recent level of activity in the postsynaptic cell, e.g., if the postsynaptic cell is highly active the threshold will increase, making synaptic weight reduction more likely. Specifically, weight changes were implemented as follows:

$$\Delta w_{jik} = D(f_{ik}(\underline{x}) \Phi(F_j(\underline{x}), \xi)), \tag{5}$$

where $f_{ik}(\underline{x})$ is the firing of BVC ik at position \underline{x} ; $F_j(\underline{x})$ is the firing of PC j at position \underline{x} ; $\Phi(F, \xi) = \tanh(F - \xi)$; and D is a scaling factor set to 0.2 in all simulations. The threshold separating positive and negative weight changes is given by:

$$\xi = \left(\frac{\bar{F}_j}{F_0} \right)^p \bar{F}_j, \tag{6}$$

where \bar{F}_j is the mean firing rate of PC j across the current trial, F_0 is a positive constant, and p is a constant greater than 1. In all simulations $F_0 = 0.3$ and $p = 3$. All weights from BVCs in the 10 sets connecting to PC have an initial value of 1.0 and can vary with the following two constraints: connections whose weight is reduced to zero remain at zero; and weights were not allowed to exceed a maximum of 3.0.

SIMULATION: EXPERIENCE-DEPENDENT CHANGES IN PC FIRING IN A STATIC ENVIRONMENT

We start by simulating the effect of experience-dependent plasticity on place fields in an unchanging environment. The fir-

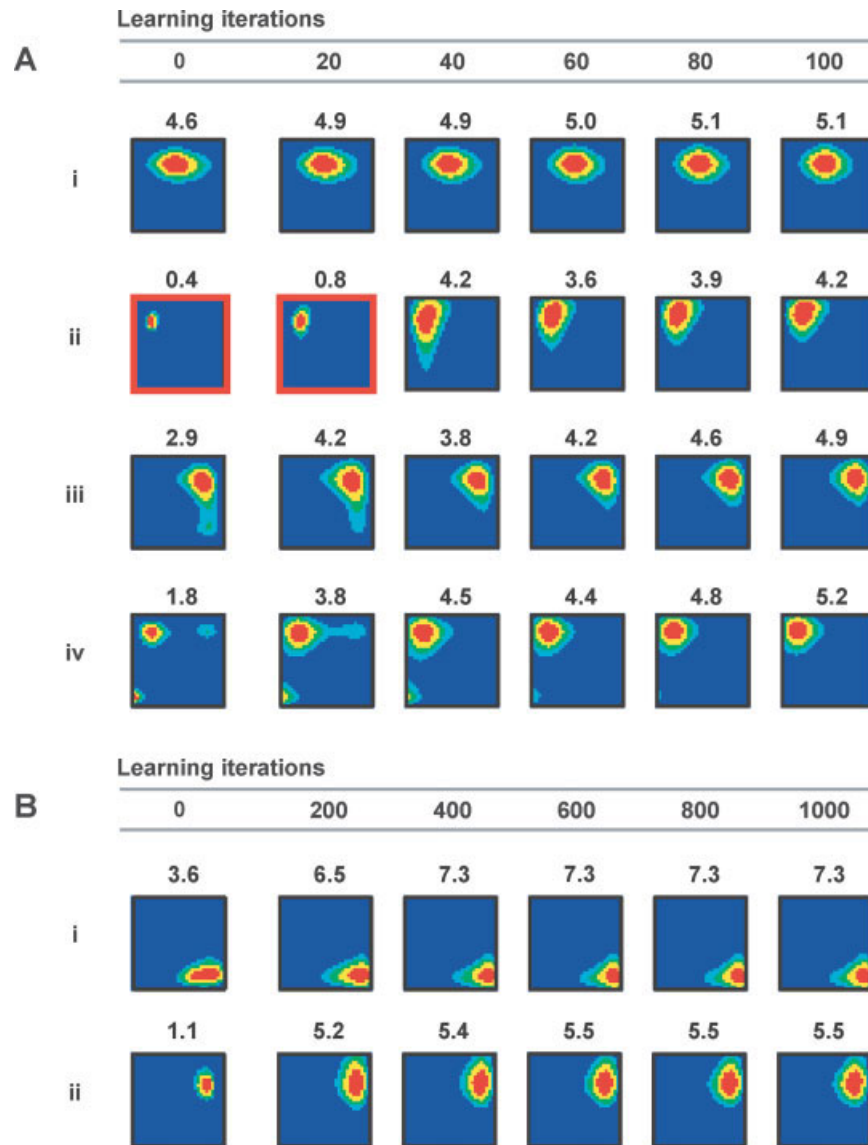


FIGURE 3. Simulated PC firing showing changes arising during learning in a static environment ($65 \times 65 \text{ cm}^2$). Peak firing rate is shown above each plot. Areas of high firing are indicated by red, low firing by dark blue. Each of the five colors indicate a firing range of 20%, thus dark blue is 0–20% of the peak firing rate, turquoise is 21–40%, etc. (A) Before learning (0 iterations) firing is spatially constrained, generally increases monotonically towards a single peak, and resembles the firing of experimentally recorded PCs. We observed that place fields changed in several ways during learning, individual cells often showed a mixture of

effects: Some fields showed very small changes often limited to a change in peak rate (i); A smaller proportion of cells became active [(ii) fields outline in red initially fired with a peak rate < 1 Hz and were considered to be off]; A similar number of fields “tightened up,” losing extended areas of low firing (iii); Fields which initially had multiple sub-peaks showed a similar effect, smaller peaks were normally removed (iv). (B) Changes initially accumulated quickly, but fields typically became stable before 200 iterations of learning. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

ing of 1,000 BVCs and 100 PCs were simulated at each node of a 2-cm square grid superimposed on a $65 \times 65 \text{ cm}^2$ enclosure, assuming that the rat visits all areas of the enclosure equally. The walls of the enclosure were treated as being indistinguishable, that is, a single BVC per set was used ($n = 1$). Learning was simulated by 100 iterations of Eq. (5). Prior to learning PC firing resembled that produced in Hartley et al.’s (2000) model; 38 cells fired at a peak rate of 1 Hz or above and so were considered to be “active,” place fields were discrete and generally had a sin-

gle peak with firing rate increasing monotonically towards that peak. A small number of cells had multiple peaks though in all cases one peak was clearly larger than the other (Fig. 3, first column).

After 100 iterations of learning an additional 14 cells had peak firing at or above 1 Hz and were considered to be active; none of the 38 cells that were originally active turned off. The majority of fields remained largely unchanged; mean peak firing rate increased only slightly from 2.3 to 2.8 Hz and fields

remained in the same position. There was, however, a tendency for fields with multiple peaks to become more coherent such that sub-peaks were lost, gradually shrinking until they fell below the firing threshold (Fig. 3A). A similar effect was seen in about 20% of cells with a single field, particularly those with large fields that included an extended area of low firing. The broad low tails of such fields were often reduced, leading to more circular fields with clearer edges. This increased coherence occurs because the BCM rule reduces the weight of BVC inputs if they drive PC firing outside of the main field. Conversely, a similar number of fields that initially had very low peak firing rates (<1 Hz), and so were considered to be “off,” showed a gradual increase in peak firing and field size with learning. This ramping-up of firing reflects the increasing efficacy of BVC inputs whose firing fields largely overlap with the place field. Experimental data provides a precedent for these observations; place fields in a novel environment are initially unstable and provide a poor estimate of the rat’s position, after 6–10 min fields become more stable (Wilson and McNaughton, 1993). Indeed, learning in the model initially progressed rapidly but was much reduced by the final 10 iterations. Inspection of longer periods of learning (1,000 iterations) confirmed that place fields generally reach a steady state within the first 100 iterations, changes observed after that point were small and diminishing (Fig. 3B). This is not surprising, the BCM rule is inherently stable; changes to synaptic efficacy generally move the mean firing rate of the postsynaptic cells towards the threshold ξ [Eq. (6)]. As the threshold is approached step size decreases and the cell settles to a steady state.

We now move to the question of how simulated place fields respond to the presence of a barrier inserted into an enclosure. Barry et al. (2006), see also Muller and Kubie (1987) and Lever et al. (2002a), described how placement of a barrier into a square environment initially caused some place fields to be duplicated, and subsequently how most duplicates were lost (Fig. 4A). Accordingly, we simulated the firing of the same 1,000 BVCs and 100 PCs in a 65×65 cm² enclosure with a 40-cm barrier placed perpendicular to and abutting the North wall. Two-hundred learning iterations were used to ensure that the system would reach a stable state. The centrally placed barrier was designated as being distinct from the external walls. That is, we allowed BVCs to learn to respond differently to these different surfaces ($n = 2$ BVCs per set). Note, although the barrier and enclosure used by Barry et al. (2006) were cut from similarly colored foam board they differed in two important ways: the enclosure had been used extensively for recordings and likely smelt very different to the recently prepared barrier; both sides of the barrier were accessible to the rats, whereas animals only ever had contact with one side of the external walls.

Results were superficially similar to the previous simulation: 37 PCs were initially active (peak firing ≥ 1 Hz), rising to 51 after learning. Mean peak rates were slightly higher, rising from 2.6 to 4.0 Hz. The key difference, however, was that initially 59% of active PCs had fields on both sides of the barrier; the other cells had unitary fields or at the most small subfields on the same side of the barrier. During learning we again saw the

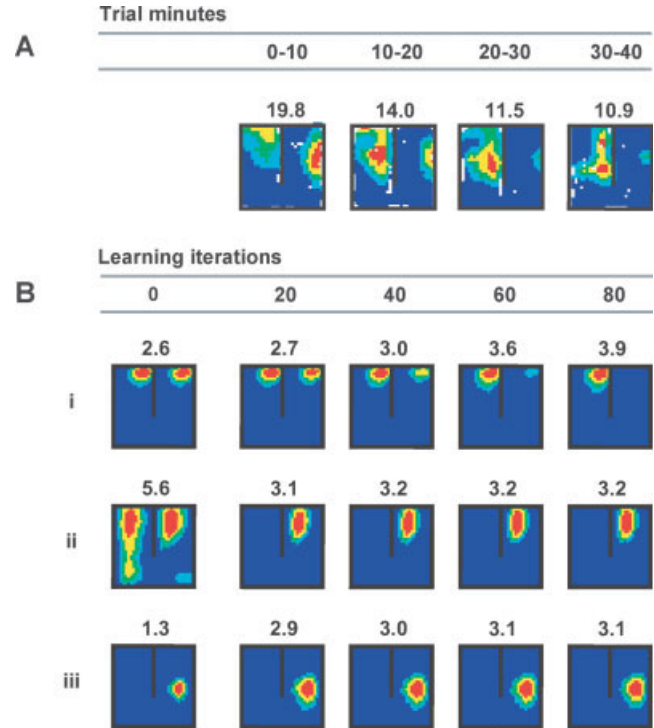


FIGURE 4. Changes in place fields accumulated during prolonged exposure to an environment partly divided by a barrier (65×65 cm² environment – 40-cm barrier), experimental, and simulated data. (A) Experimental data recorded during a 40-min trial – 10-min sections are shown. Insertion of the barrier caused a previously unitary place field to be duplicated on both sides of the barrier. With prolonged exposure the right hand field is gradually reduced and ultimately lost. (B) Simulated place fields. Presence of the centrally placed barrier initially produced a number of duplicate place fields [(i) and (ii)]. Repeated iterations of learning caused all duplicate place fields to be lost. The rate of reduction differed between cells. (iii) Like the experimental data, not all cells exhibited duplicate fields. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

gradual “tidying” of place fields, though this time the effect also extended to the duplicate fields (Fig. 4B). Hence, after 200 iterations all duplicate fields were lost with the rate of reduction varying between cells; most duplicates were lost within the first 40 iterations, while a few persisted for up to 100. In most cases the field that initially had the highest firing rate was maintained though this was not always the case; 11 cells maintained the Western field, and 11 maintained the Eastern field. By contrast, most duplicate fields were not removed when we simulated firing with the central barrier treated as being indistinguishable from the external walls, initially 25 duplicate fields were present and only three were removed after learning.

The results from the simulation accord well with the experimental data reported by Barry et al. (2006). In particular, they reported the gradual loss of most duplicate fields over one or more recording session of 40 min, and that the loss of such fields occurred at different times in different cells. An effect we did not see in the simulation was the partial reversal of learning that

occurred overnight between recording sessions. The experimental data showed that, on successive days of recording, place fields often reverted to a state intermediate to those found at the start and end of the previous recording session. Interestingly, a similar effect, described as being “two steps forwards, one step backwards” was noted by Lever et al. (2002b) during slow remapping in different shaped environments. This indicates an overnight, *partial* long-term consolidation of the changes that occurred during the day. Failure to observe this overnight decay in learnt changes is understandable, the simulation includes no decay term and all synaptic changes are assumed to be immediately stable. As such, each round of learning always picks up from the point reached in the previous iteration.

SIMULATION: EFFECT OF BARRIER MOVEMENT AFTER LEARNING

Rivard et al. (2004) reported a series of manipulations in which rats were run in a 76-cm-diameter grey circular environment while PCs were recorded. The environment was polarized by a 45° white cue card fixed to the wall and a 23.5-cm-wide Perspex barrier was present in the cylinder during training. Initially positioned in the South-West segment of the enclosure, during probe trials the barrier was either rotated by 45°, translocated to the North-East segment, or removed. The experimenters' principal finding was that place fields located close to the barrier in the training condition tended to be affected more strongly by movement (rotation or translocation) of the barrier than those located further away (Fig. 5A). More accurately, they showed a graded effect such that fields close to the barrier maintain their position relative to it during manipulations, while those more distant retained their position relative to the external wall. A similar, though more abrupt effect was seen in response to removal of the barrier, fields that were immediately adjacent to it during training became inactive when it was removed, again more distant fields were unaffected. The authors interpreted these results as evidence of two distinct sets of cells, one set responding to the barrier (object cells), the others responding to the external environment (PCs).

To better understand Rivard et al.'s result we simulated 200 iterations of learning for 100 PCs in a 76-cm-diameter circular environment with 23.5-cm barrier place in the South-West segment. The barrier was designated as being distinct from the walls, as was the 45° cue card fixed to the external wall (to the East). Thus, the environment included three distinct surfaces to which BVCs might learn to respond preferentially (i.e., $n = 3$ BVCs per set). Similar to previous simulations 42 cells were active initially (firing ≥ 1 Hz), this rose to 54 after 200 iterations. We also saw the same rapid “tidying” of place fields and at the same time duplicate fields produced by a few cells (13) on either side of the barrier were lost. After learning, the firing of the same set of cells was simulated in three probe environments comparable to the ones used by Rivard et al. (2004): one with the barrier rotated 45° clockwise, one with the barrier translocated to the North-East quadrant of the environment, and one

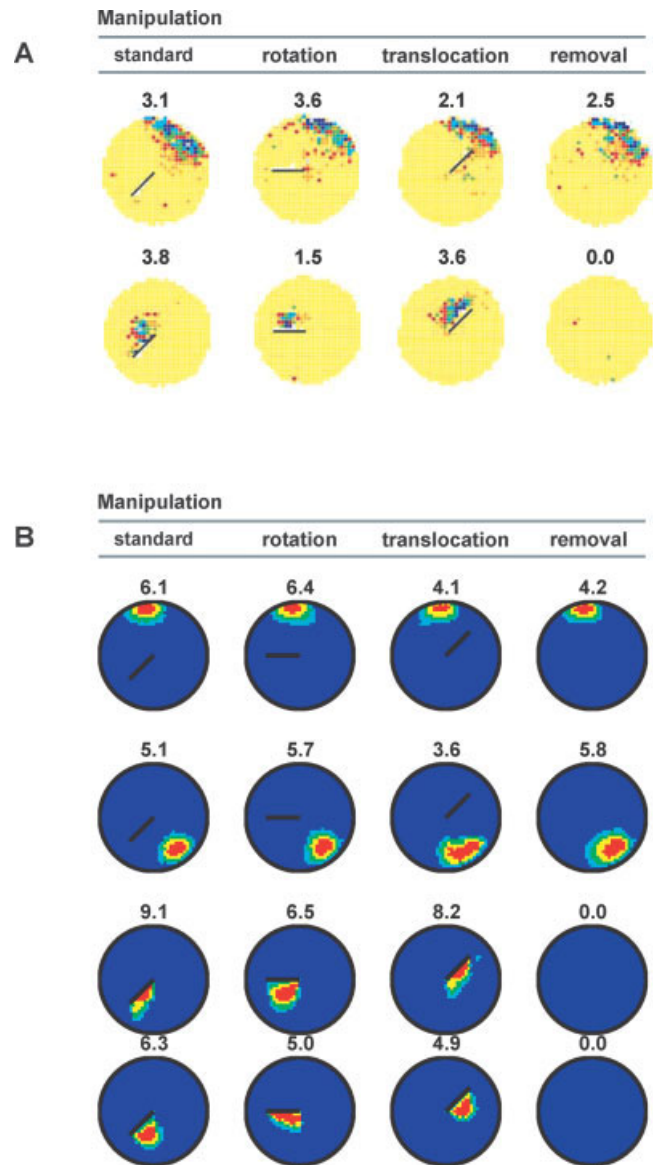


FIGURE 5. Changes in PC firing induced by manipulations made to a centrally placed barrier, experimental, and simulated data. (A) Experimental data, images adapted from Rivard et al. (2004). PCs were recorded in a 76-cm-diameter circular environment. Increased firing rate is indicated by darker colors, yellow indicates no firing. Value above each plot show mean infield firing rate. Animals were familiar with the standard condition in which a 23.5-cm barrier was positioned in the South-West section of the environment. Rotation, translocation, or removal of the barrier revealed a graded effect such that some cells, usually those most distant from the barrier, were largely unaffected by the manipulations. Fields closer to the barrier (bottom) were affected more strongly and appear to encode space relative to the barrier, in its absence the cell stops firing. (B) Simulated PC firing. After 200 iterations of learning in the standard condition cells were transferred to three probe environments with the barrier rotated 45° clockwise, translocated to the North-East section of the environments or removed. Like the experimental data, a differential effect was seen, cells distant to the barrier (top two cells) were largely unaffected by the change, whereas fields close the barrier (bottom) appeared to encode space relative to it. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

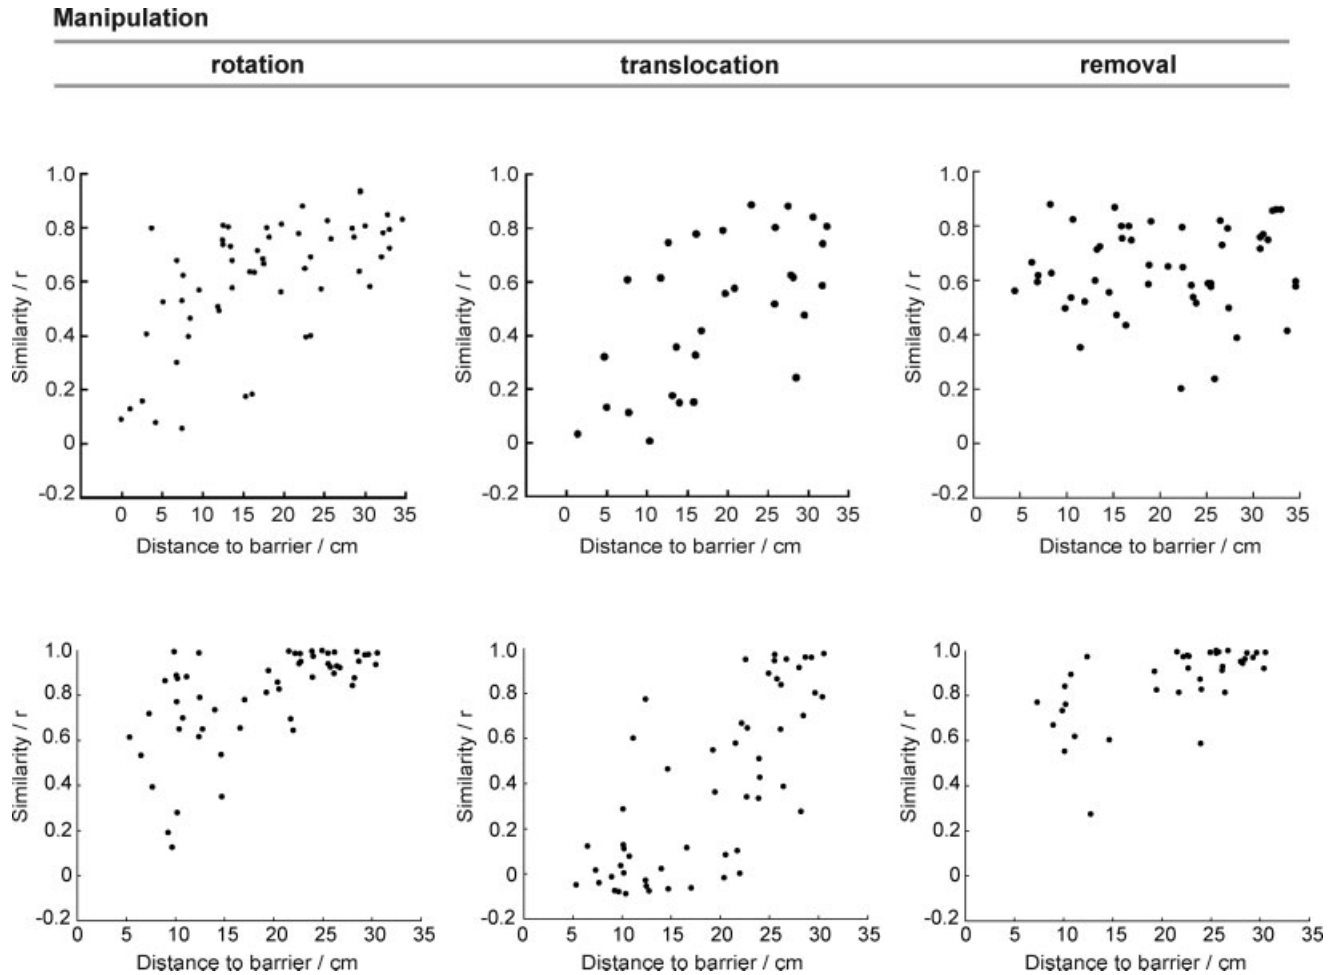


FIGURE 6. Similarity between standard condition and probe trials (barrier rotated, translocated, and removed) as a function of distance between field centroids and the barrier, experimental, and simulated data. Top row, experimental data, adapted from Rivard et al. (2004). In response to barrier rotation and translocation a graded effect is evident, similarity increases with distance from the

barrier. Removal of the barrier causes proximal fields to stop firing, distal fields are not affected. Bottom row, simulated data. After 200 iterations of learning in the standard condition, the relationship between similarity and distance to barrier is comparable to that seen in the experimental data.

with the barrier removed. We deal first with the effect of the barrier rotation on PC firing.

Rotating the barrier 45° clockwise had a variable effect on simulated PC firing, some fields appeared to be completely unchanged whereas others moved or changed their peak firing rate. All 54 cells that had been active before the barrier was rotated were active after the rotation, in addition a further two cells started to fire at or above the 1 Hz threshold. As a whole though, mean infield rate was slightly reduced by the change in barrier position, falling from 2.4 to 2.0 Hz. Interestingly, Rivard et al. (2004) reported a reduction in infield firing of similar proportions, 4.1–3.3 Hz, when the barrier was rotated. Visual observation of the simulated PC firing showed that fields distant to the barrier were more likely to be unaffected by the change, while those close to the barrier often changed position and firing rate (Fig. 5B, column 2).

To better understand how proximity to the barrier affected field stability we used the same strategy as Rivard et al. and, for

each active field, calculated the distance from its center of mass to the barrier. Additionally a measure of “similarity,” again adopted from Rivard et al., was used; for each active field the bin by bin correlation (spatial correlation) was calculated between firing rates before and after the barrier manipulation. The same similarity score was also calculated with firing in the probe trial rotated by 45° , this was done for clockwise and counter-clockwise rotations. Note, as it appears that Rivard et al. included bins with a firing rate of 0 Hz in both configurations in this calculation, we did the same thing.

The scatter plot of distance to barrier versus similarity for the simulation shows a strong relationship between the two, such that, similarity increases with distance from the barrier (Pearson’s correlation coefficient was $r = 0.70$) (Fig. 6, bottom row). This relationship mirrors the one found in the experimental data in which a correlation of $r = 0.61$ was reported (Fig. 6, top row). This relationship was reversed when firing in the probe trial was rotated 45° counter-clockwise before making the comparison,

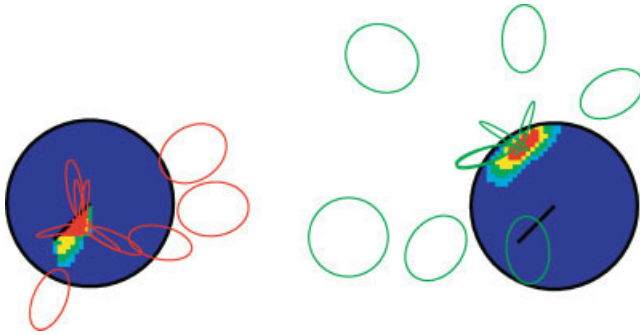


FIGURE 7. Simulated firing of two PCs showing the 10 BVC sets driving each cell. Firing is shown in an enclosure analogous to the “standard” configuration used by Rivard et al. In each case the receptive field of BVCs are portrayed as ellipses at a preferred firing distance and direction. Note BVC with large tuning distances are broader than those with short preferred firing distances, this is indicated by the relative size of the ellipses. Red BVCs originate from the cell on the left, green from the cell on the right. The field on the left is principally defined by four short range BVCs tuned to respond to the barrier, movement of the barrier (e.g. translocation) will likely result in a matching movement in the field. The cell on the right receives the majority of its input from four BVCs that are shown responding to the external wall of the environment (one of the BVCs is obscured by the others). Manipulations made to the barrier have limited effect on the firing of this cell. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

hence fields adjacent to the barrier showed the highest similarity scores ($r = -0.59$). Again this result is very similar to the one reported by Rivard et al. and suggests a graded response such that fields close to the barrier tend to rotate with it, while more distant fields remain static relative to the external walls. In agreement with this view, rotation of the fields by 45° clockwise before calculating similarity greatly reduces the correlation ($r = -0.28$ in simulated data, $r = -0.24$ in experimental data).

Rivard et al. reported similar results when the barrier was translocated to the North-East quadrant of the enclosure and also when it was removed. In both cases mean infield firing rate fell relative to the standard condition and a graded response was seen such that fields close to the barrier were affected most strongly. In the case of the translocation, local fields moved with the barrier, whereas distant fields maintained their position relative to the external walls. Removal of the barrier also produced a graded effect, albeit more abrupt; fields adjacent to the barrier were more likely to become inactive or show pronounced changes in firing rate.

Using the same set of BVCs and PCs described earlier, we simulated PC firing with the barrier translocated to the North-East and also with it removed. As with the 45° rotation we found excellent accord between simulated fields and experimental data. In both conditions, mean infield firing rate fell when compared to the standard condition (translocation: 2.4–1.5 Hz; removal: 2.4–1.7 Hz; Rivard et al. reported reductions from 4.3 to 3.5 Hz and from 4.5 to 3.6 Hz, respectively). Fields adjacent to the barrier were again most affected, in the case of the translocation, a plot of similarity versus distance to the barrier yielded a

correlation of $r = 0.73$ ($r = 0.64$ in the experimental data) (Fig. 6). This correlation was reversed when firing fields from the second environment were translocated before calculating similarity ($r = -0.48$, Rivard et al. found $r = -0.52$). When the barrier was removed 30% (16/54) of simulated cells that had been active turned off. The experimental data showed that fewer cells became inactive (14%), but the same bias between proximal and distal cells was evident: Rivard et al. found that rate change cells were on average 7.5 cm from the barrier, this rose to 20.4 cm for stable cells; in the simulation 63% of cells within 10 cm of the barrier turned off but only 24% of cells further from the barrier (Fig. 6).

So why do we see a graded effect such that fields adjacent to the barrier seem to encode space in a different coordinate system to those more distant to it? BVCs tuned to respond to proximal walls have tighter firing fields than those tuned to respond to more distant walls. Hence, at a given location, proximally tuned BVCs tend to contribute more to the firing of a PC than BVCs with longer tuning distances. Inspection of the simulated PCs show that fields adjacent to the barrier tend to be driven by BVCs that are firing as a result of the barrier; typically they have short preferred firing distances (d_i) and are orientated to encounter the barrier (Fig. 7). As a result of learning these BVCs come to drive the PC more strongly but also become more selective, responding preferentially to the barrier and not to other surfaces. If the barrier is moved the same BVCs continue to drive the PC, only now they fire in a different position in the cylinder. Conversely, in the absence of the barrier, the same BVCs will fire in response to the external wall but will do so in a different position and will drive the PC with reduced efficacy. The opposite is also true, place fields adjacent to external walls will tend to be controlled by BVCs that respond to those walls, and so will not be so affected by changes made to the barrier. The reduction in mean infield firing rate that we observed reflects a similar explanation. Learning conducted with the barrier in its standard position optimizes inputs from BVCs whose firing fields closely coincide with that of the postsynaptic PC. Any change made to the geometry of the enclosure will cause the BVCs' fields to move relative to one another and is likely to result in reduced drive to the PC.

A point worthy of note is that in response to each of the three manipulations, Rivard et al. noted a small number of cells that drastically changed their firing rate, effectively turning on or off. The authors construed these as conjunction cells, PCs that responded to the barrier but only when it was in a certain position relative to the enclosure. Other experimenters have seen similar effects, for example, placement of a barrier so that it bisects an existing place field can cause the cell to stop firing (Muller and Kubie, 1987). We saw similar effects in the simulations, for example, when the barrier was rotated two cells became active (peak firing ≥ 1 Hz), the space occupied by their fields having previously been occupied by the barrier (Fig. 8). In response to the translocation, two cells became active and a further nine showed either a large change in peak firing rate (more than 100% increase or reduction) or a pronounced change in the size of the place field. These effects are easily explained by the

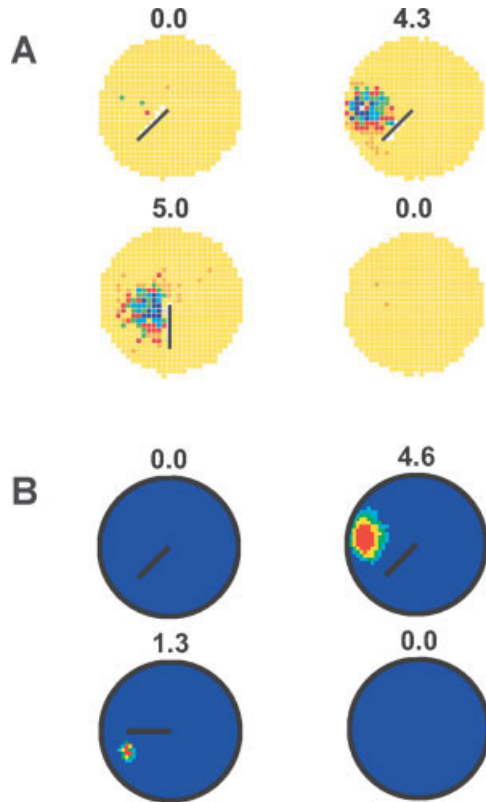


FIGURE 8. “Conjunction cells” that exhibited large changes in rate after manipulation of the barrier. a) Experimental data, adapted from Rivard et al. (2004). Two cells are shown, one per column, top row is the standard condition, bottom row is firing after making changes to the barrier position. In the standard condition the first cell (left) fired very few spikes and was considered to be off. After rotation of the barrier the cell developed a robust place field. The second cell (right) showed strong firing adjacent to the barrier in the standard condition but after removal of the barrier it became silent. b) Simulated place cell firing exhibiting similar effects to the experimental data, layout similar to (a). Rotation of the barrier caused the first cell (left) which had been silent, to develop a small place field. Removal of the barrier caused the second cell (right) to top firing. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

basic BVC model and do not require learning. Adding, removing, or moving a barrier changes the pattern of activity in a population of BVCs. The barrier provides new inputs for some cells, but can obscure more distant walls driving other cells. Hence, a PC downstream of these BVCs might be driven more strongly, exceeding its firing threshold, alternatively it may receive less input and become silent.

SIMULATION: EFFECT OF CUE ROTATION AFTER LEARNING

Fenton et al. (2000a) showed that movement of two cue cards together or apart caused place fields to move relative to one another and to the surrounding environment (Fig. 9A). The changes in place field location were spatially graded and locally

coherent; fields adjacent to a cue rotated with it, whereas centrally placed fields tended to be translocated in the same direction as the two cards. In addition, after removal of one cue the other was sufficient to maintain PC firing and to control the orientation of the ensemble. To better understand these effects we simulated the firing of 1,000 BVCs and 100 PC in a 76-cm-diameter circular environment analogous to the one used by Fenton et al. The “standard” enclosure was polarized by a two distinct cue cards fixed to the wall, each subtended an angle of 45° with their centers separated by 135° . The environment presented three distinct surfaces to which BVCs might learn to respond differentially: the enclosure walls, and the two cue cards. After 200 iterations of learning in the standard environment PC firing was assessed in response to three manipulations: cue cards rotated apart (separation increased to 160°); cue cards rotated together (separation decreased to 110°); and one cue card removed to leave only one card in the cylinder.

In addition to the basic model we follow Burgess and Hartley’s (2002) model of these data in assuming the rat’s sense of direction is perturbed by relative movement of the cues (see also Touretzky et al. (2005) for a very similar subsequent approach). Models of the rodent head direction system (Skaggs et al., 1995; Zhang, 1996) implicate learnt associations between visual cues

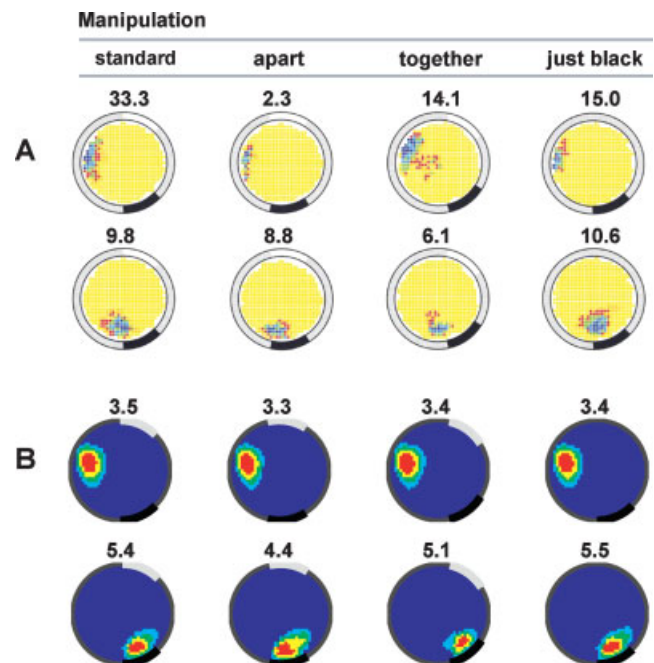


FIGURE 9. PC firing in a circular environment (76 cm diameter) polarized by two cue cards, experimental, and simulated data. Place fields are shown in a familiar environment and the same environment after the cue cards were moved apart, together, and after one card was removed. Experimental and simulated fields show subtle changes in field position and size in response to the manipulations. (A) Experimental data, images adapted from Fenton et al. (2000a). Two PCs are shown, numbers above each plot indicate peak rate, dark colors show high firing, yellow indicates zero firing. (B) Simulated firing. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

and head direction cells (Taube et al., 1990a,b) as a mechanism for maintaining the system's alignment with the environment. In Fenton et al.'s study, the rotation of the PC ensemble induced by consistent rotation of both cards, or by rotation of a single card, suggests that each card was sufficient to control the orientation of the head direction system. As such, we can ask what effect inconsistent rotation of the cards would have on the head direction system and in turn what implications that would have for PC firing. We assume that the influence of a distal orientation cue is proportionate to its proximity to the animal (Fenton et al., 2000b; Burgess and Hartley, 2002). Specifically, the angular deflection in the head direction system produced by a rotation of the cards apart by θ units would be:

$$H_{\text{def}} = \frac{\phi(d_B - d_W)}{2(d_B + d_W)} \quad (7)$$

where d_B and d_W are, respectively, the shortest distance from the current position to the black and white cards. For movement of the cards together θ is negative. H_{def} is measured clockwise from "East" which is defined as the point midway between the two cue cards. Hence, after movement of the cards apart by 25° (i.e. $\theta = 25^\circ$), an animal standing next to the white card would experience a deflection of 12.5° to its head direction system, an animal equidistant between the cards experiences no deflection (Fig. 11A). In the absence of one card, East has the same relationship to the remaining card as it did in the standard condition and no deflection is induced in the head direction system. Note, Fenton et al. reported that when both cards were rotated by 45° clockwise this resulted in a commensurate rotation of all place fields. We do not explicitly simulate this manipulation, but following the assumption made above, it is clear that rotation of the orienting cues would result in a matching deflection in the preferred tuning direction of all BVCs. As a result all place fields would be seen to rotate with the cues.

Learning in the standard environment produced the same changes in place fields noted earlier; initially 38 cells fired with a peak rate at or above 1 Hz, this rose to 56 after learning, place fields became tighter, and secondary fields were lost. Rotation of the cue cards apart or together affected PC firing in a similar way, like Fenton et al. we observed subtle changes in the position and firing rate of place fields (Fig. 9B). Fields close to one of the cue cards tended to rotate with the card, whereas fields equidistant to both cards tended to be translocated. All active cells remained "on" (peak firing ≥ 1 Hz) and none of the inactive cells became active. As with Rivard et al.'s study, Fenton et al. found a significant reduction in mean infield rate and centroid rate after the manipulations. For example, rotating the cards together caused the mean infield rate to fall from 5.0 to 3.9 Hz, a reduction of 21%, moving the cards together reduced firing by 15%. As expected, and for reasons outlined earlier, simulated PC firing showed a similar reduction, infield rate fell by 7% in response to rotation of the cards apart and by 4% when they were rotated together. The original experimenters also reported that movement of the cards produced small, nonsignificant reductions in field size of 1% (apart) and 3% (together). The

model showed a reduction of 2% in the apart condition and no change when the cards were moved together.

Following Fenton et al., we calculated the movement of the place field centroids between the standard and probe conditions. Inspection of these displacement vectors showed a good match with plots generated from experimental data (Fig. 10). Specifically, the angular and radial components of displacement show the same periodic modulation as the experimental data; fields equidistant to both cards show a large radial component to their displacement, while fields closer to one card show an angular component. Using a population of 1,000 PCs we calculated the mean vertical and horizontal displacement of centroids in each condition. Fenton et al. employed this same measure and showed that for both conditions vertical displacement was not different to zero (apart: 0.147 cm, together: 0.387 cm), that movement of the cards together produced a small horizontal shift in field centroids (1.80 cm), and that movement of the cards apart produced a larger negative shift (-4.84 cm). We found the same pattern, mean vertical movement was close to zero (apart: -0.15 cm, together: 0.16 cm), mean horizontal movement produced by moving the cards together was positive (3.56 cm), and moving the cards apart induced a negative shift of slightly larger magnitude (-3.66 cm). Individually these values are a good match to the experimental values, however, the difference in magnitude between the horizontal measures is greatly reduced in the model ("together" movement is 97% of the "apart" movement, in the experimental data it was 37% of the "apart" value). This discrepancy is not unique to our account, other models (Touretzky et al., 2005), including Fenton et al.'s own vector-field model of place field movement (Fenton et al., 2000b), provide a similar underestimate of this difference.

To characterize the contribution of the deflection to the head direction system, we recalculated the firing of the 1,000 PCs described earlier, without accounting for the effect of card rotation on the preferred firing direction of BVCs. In accordance with the previous simulations by Burgess and Hartley (2002), the principle difference induced by the head direction deflection is to accentuate the angular displacement of place fields produced by both manipulations (Figs. 11B,C). The difference is most obvious at the Western edge of the enclosure, before taking account of the distortion in the head direction system, fields in this area remained largely unchanged, with the distortion they show clear angular displacement; this seems to provide a better match to the experimental data (contrast Figs. 10 and 11). In the cards apart condition, taking account of the deflection in the head direction system increases mean horizontal field translocation by 39% (3.56 cm vs. 2.54 cm). Horizontal field movement in the cards together condition is similarly increased (-3.66 cm vs. -2.61 cm, a 40% increase). Mean vertical displacements are not affected and remain close to zero in either case.

The model provides a less-good account of changes in PC firing after removal of one of the cards. Using 100 PCs, we simulated removal of the white cue card. Simulated cards have arbitrary color, the model simply treats them as being visibly different to one another and to the walls of the enclosure. Hence, observations made after removal of the white card are valid for

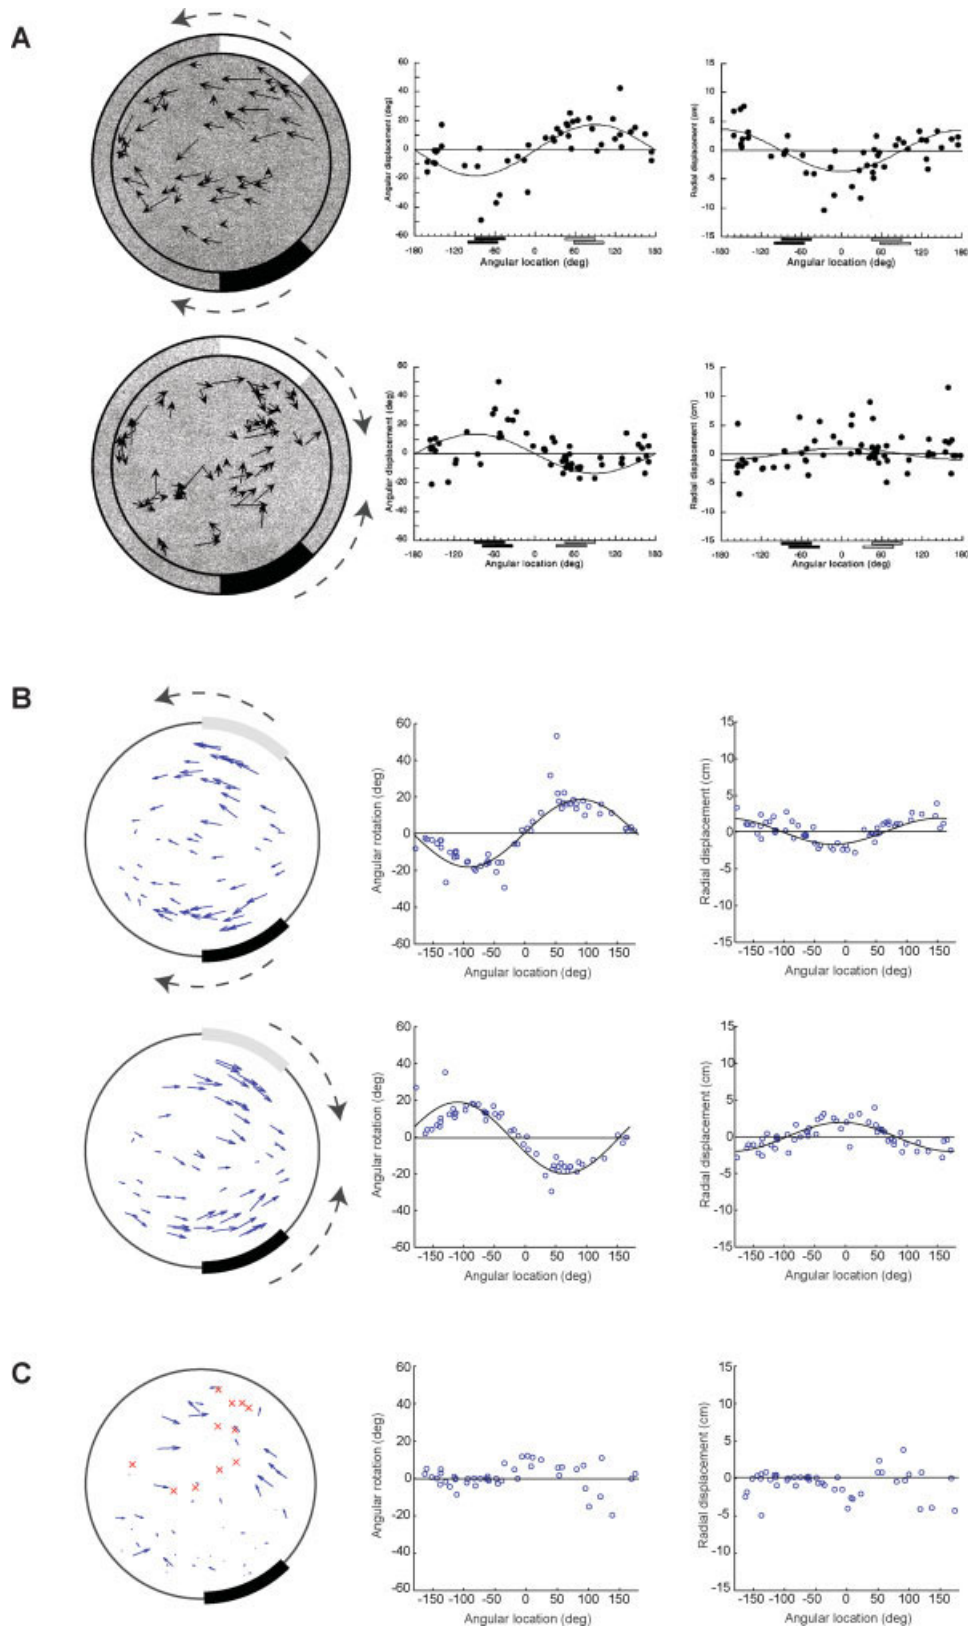


FIGURE 10. Movement of place field centroids produced by increasing and decreasing cue card separation, experimental, and simulated data. (A) Experimental data, images adapted from Fenton et al. (2000a). Left, movement of field centroids. Middle, angular displacement of centroids as a function of angular location in the standard condition (sine function fit by eye). Angular location is defined relative to the center of the enclosure such that 0° defines an axis running directly between the cards, movement clockwise from that axis is considered to be positive (e.g. “North” is at 90°). Angular displacement is defined relative to the same coordinate system (e.g. a field the

moves anticlockwise will have negative displacement). Right, radial displacement of centroids as a function of angular location in the standard condition (sine function fit by eye). Radial displacement is defined relative to the center of the environment such that movement directly away from the center is considered to be positive. (B) Simulated data, movement of centroids in probe conditions after 200 iterations of learning in the standard condition. Plots as above, sine functions fit to minimize squared error. (C) Simulation of removal of one cue card. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

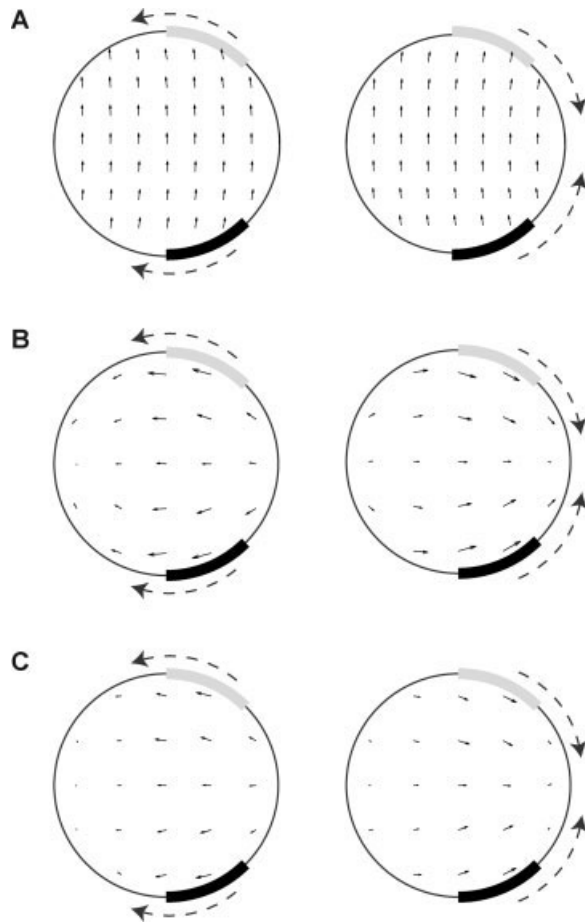


FIGURE 11. (A) Distortion in the head direction system induced by card movement. Arrows show perceived North – in the standard condition all arrows would point upwards. (B, C) Average movement of place field centroids produced by increasing (left) and decreasing (right) cue card separation, simulated data. Average displacement was calculated using a population of 1,000 PCs, mean displacement vector is shown at 21 locations within the enclosure. (B) Movements predicted by BVC model after taking account of distortions in the head direction system. (C) Movements predicted by pure BVC model alone.

removal of the black card. The key departure between simulated and experimental firing was that 11 out of 56 (20%) cells in the simulation became inactive (peak firing < 1 Hz), whereas none of the experimental cells turned off. All the cells that became inactive had fields located close to the missing card (Fig. 10C). Fenton et al. actually considered a cell to be active if it had a place field of at least nine contiguous pixels with above 0 Hz firing. Even using this less stringent criterion for the simulation, nine of 54 cells became inactive (17%). Observation of the inactive PCs confirmed that they all received some input from BVCs tuned to respond to the missing card. These BVCs were short range, oriented towards the card's position, and drove the PC more strongly in response to the cue card than the wall. A further departure between the simulation and experimental results is evident from the plot of displacement vectors (Fig. 10C); fields in the bottom half of the environment near the black card are,

like the experimental fields, largely unmoved. However, fields near the absent white card are rotated, their centroids being displaced towards the center of the missing card; this effect is not seen in the experimental data. Inspection of the BVCs showed that removal of the card reveals an additional surface for wall specific BVCs to respond to. Boundary vector fields that had previously been constrained by the edge of the card now spread into the region in front of it. In turn, this caused place fields driven by those cells to spread into the same space. Finally, simulated cells also showed a reduction in mean infield firing from 2.0 to 1.7 Hz (a 13% drop). Fenton et al. found no significant reduction in infield rate after removal of either card, nevertheless a trend similar to that seen in the simulation is evident, both experimental manipulations caused mean infield firing to fall by 10%.

DISCUSSION

The model presented here is consistent with a wide range of experimental findings. Hartley et al.'s (2000) original formulation provided an excellent account of the effect of geometric environmental manipulations on the firing of PCs. Now, by incorporating synaptic plasticity mediated by the BCM rule, the model is able to describe changes that occur to PCs during learning, such as the consolidation of place fields and removal of duplicate fields. Furthermore, in a familiar environment the model reproduces graded changes in place field position produced when cues are moved relative to one another. Independently, the model accounts for the reduction in peak and infield firing rates produced when previously stable cues are moved relative to each other. As such, we provide an alternative account of Rivard et al.'s finding without the need to evoke separate classes of PCs that respond in different reference frames. Similarly, we provide a neural level account of Fenton et al.'s vector-field model (2000b) that sought to explain experimental results from the same author (Fenton et al., 2000a). In addition, we believe our model to be preferable to others that have addressed Fenton et al.'s result (Burgess and Hartley, 2002; Touretzky et al., 2005) as it requires fewer assumptions specific to that experiment (namely not having to prescribe which sensory cues will drive BVCs).

Notably, the model failed to perform as well when cues were removed from the environment. Removal of a cue card in the simulation of Fenton et al.'s experiment caused place fields adjacent to the missing card to turn off and those nearby to move into the now vacant space. A related but less pronounced effect is evident in response to removal of the central barrier in the simulation of Rivard et al.'s finding; the experimenters reported that 14% of cells turned off but under the simulation 30% became inactive. Together these findings suggest that the model does not adapt in the same way as the brain when sensory drive to PCs is reduced. In vivo three possible mechanisms for adaptation seem plausible: dynamic reduction of inhibition in response to reduced excitation (Marr, 1971; McNaughton and Nadel,

1990; Kali and Dayan, 2000); maintenance of firing by path integration (Etienne et al., 1996; Gothard et al., 1996; McNaughton et al., 1996); or pattern completion of the remaining cues, probably mediated by CA3 (Marr, 1971; Kali and Dayan, 2000; Nakazawa et al., 2002; Rolls et al., 2002; Wills et al., 2004). Most likely all three play a role but nevertheless it is instructive that Touretzky et al. (2005) employed a recurrent network to model some aspects of Fenton's result, though not the cue removal manipulation. In principle, any of these mechanisms, each of which has previously been applied in computational models, could be utilized to expand the explanatory power of the current model.

Currently, the model deliberately avoids the question of remapping (Bostock et al., 1991). Remapping, at least in response to an entirely novel environment (Leutgeb et al., 2003, 2004), is best described as a wholesale reorganization of place fields, such that individual fields undergo position and rate changes. It stands in contrast to the slow, cumulative changes in PC activity discussed earlier (Lever et al., 2002b; Barry et al., 2006) and seems to be a distinct phenomenon. That is not to say the BVC framework is necessarily inapplicable to environments which induce remapping. Rivard et al., in addition to the manipulations simulated earlier, also showed that transferring the centrally placed barrier to an entirely new environment produced a graded response; fields adjacent to the barrier were preserved but more distant ones appeared to remap. In this case at least, remapping can be understood as a local effect not dissimilar to the changes produced by Rivard's other manipulations. It follows that BVCs that have "learnt" to respond most strongly to the barrier will continue to fire against it and so will maintain a population of place fields. On the other hand, BVCs that were driven by the grey walls and white cue cards in the standard enclosure will be driven less strongly, or not at all, by the new environment. Clearly the current model does not specify how a novel set of BVCs would be recruited to respond to the novel black cue card and white walls. Interestingly, Anderson and Jeffery (2003) have partially addressed this question, suggesting that contextual cues such as odour might "gate" the activity of BVCs. In effect, they described how sufficient change made to any aspect of the environment (e.g. floor texture, color, odor, etc.) could cause a novel set of BVCs to become active. At the level of PCs this rearrangement of BVC inputs would institute a change in place field position and firing rate.

A final pair of questions to be answered is, "do BVCs really exist?" and "how do they learn?" Burgess and O'Keefe (1996) initially suggested that BVCs might be located in the entorhinal cortex, the origin of most neocortical projections to the hippocampus. Subsequently, it has been suggested that BVCs might originate in the dorsal subiculum (Barry et al., 2006). Spatial cells from the subiculum, like BVCs, have stable elongated firing fields that often lie parallel to environmental barriers (Sharp, 1999). Interestingly, the PC firing driven by BVCs tuned to a specific boundary, as observed in bimodal place fields after expansion of an environment, appear to have two components—one reflecting sensory input from the boundary (be it visual, tactile, or acoustic) and the other reflecting additional path integra-

tive information from the boundary most recently visited (Gothard et al., 1996; O'Keefe and Burgess, 1996). The recent discovery of grid cells in the dorsolateral medial entorhinal cortex (Hafting et al., 2005), and the suggestion that they provide the path integration input to PCs (O'Keefe and Burgess, 2005; Sargolini et al., 2006) suggests a possible division between lateral entorhinal cortex providing the sensory component of BVCs and medial entorhinal cortex providing the path integrative component. Clearly, the multi-peaked, regular firing pattern of a single grid cell is unlike that of a putative BVC. However, grid firing, being stable across trials, is clearly sensory bound (Hafting et al., 2005) to the extent that geometric manipulation of an animal's environment can induce parametric changes in grid scale (Barry et al., in preparation). Thus, it is possible that sets of grid cells bound to environmental boundaries in particular directions could provide an input functionally equivalent to the path integrative component of a BVC. See Burgess et al. (2007) for further discussion.

Regardless of their location or composition, how do BVCs come to respond specifically to certain barriers in the environment? Our formulation saw each BVC represented a number of times, once for each of the different barrier types in the environment. Changes to the strength of synapses between BVCs and PCs resulted in place fields that showed specificity for certain surfaces. Such an arrangement is biologically improbable, a complex natural environment, for example, would require an arbitrarily large number of BVCs. Much more likely is the possibility that individual BVCs initially show low response specificity but, with experience, come to "prefer" certain barrier types to others. Thus, learning about the distinct nature of objects would occur upstream of the hippocampus, whereas learning about the configuration of those objects would occur within the hippocampus. What about the rate of learning? In the model all distinct surfaces were assumed to have the same saliency and so, given the same geometric conditions, BVCs would learn to distinguish objects equally fast. In the brain, learning is likely to be modulated by many factors including saliency of the surfaces present in the enclosure, and the animal's state of arousal. Considerations such as these probably account for Lenck-Santini et al.'s (2005) observation that, in their experiment, the firing of a small number of PCs was affected by movement of two intramaze landmarks, but none were affected by changing their identity. In that experiment a centrally placed object (which is not a very salient cue for all but the most proximal of firing fields) was replaced by another during two recording trials; likely leaving too little time for BVCs to establish distinct responses.

In summary, our findings suggest a plausible mechanism by which the hippocampus can learn to represent specific locales in terms of the surrounding cues. This has relevance not just for spatial memory and navigation in rodents but also shows how the firing of PCs can encode aspects of episodic memory. In essence, the firing of a set of PCs describes the position of the subject and of other objects, relative to each other and to the surrounding environment; the "what?" and the "where?" of episodic memory.

Acknowledgments

The place cell recordings referred to (see also Barry et al., 2006) were made by CB in Kate Jeffery's lab.

REFERENCES

- Agnihotri NT, Hawkins RD, Kandel ER, Kentros C. 2004. The long-term stability of new hippocampal place fields requires new protein synthesis. *Proc Natl Acad Sci USA* 101:3656–3661.
- Anderson MI, Jeffery KJ. 2003. Heterogeneous modulation of place cell firing by changes in context. *J Neurosci* 23:8827–8835.
- Barry C, Lever C, Hayman R, Hartley T, Burton S, O'Keefe J, et al. 2006. The boundary vector cell model of place cell firing and spatial memory. *Rev Neurosci* 17:71–97.
- Best PJ, Thompson LT. 1984. Hippocampal cells which have place field activity also show changes in activity during classical conditioning. *Soc Neurosci Abstr* 10[125].
- Best PJ, Thompson LT. 1989. Persistence, reticence, and opportunism of place-field activity in hippocampal neurons. *Psychobiology* 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.
- Bingman VP, Siegel JJ, Gagliardo A, Erichsen JT. 2006. Representing the richness of avian spatial cognition: Properties of a lateralized homing pigeon hippocampus. *Rev Neurosci* 17:17–28.
- Bostock E, Muller RU, Kubie JL. 1991. Experience-dependent modifications of hippocampal place cell firing. *Hippocampus* 1:193–206.
- Burgess N, Barry C, O'Keefe J. 2007. An oscillatory interference model of grid cell firing. *Hippocampus*. Appearing in the same issue as this paper.
- Burgess N, Hartley T. 2002. Orientational and geometric determinants of place and head-direction. In: Dietterich TG, Becker S, Ghahramani Z, editors. *Advances in Neural Information Processing Systems*, 14th ed. Cambridge, MA: MIT Press. pp 165–172.
- Burgess N, O'Keefe J. 1996. Neuronal computations underlying the firing of place cells and their role in navigation. *Hippocampus* 6:749–762.
- Burgess N, Recce M, O'Keefe J. 1994. A model of hippocampal function. *Neural Networks* 7:1065–1081.
- Ekstrom AD, Kahana MJ, Caplan JB, Fields TA, Isham EA, Newman EL, et al. 2003. Cellular networks underlying human spatial navigation. *Nature* 425:184–187.
- Etienne AS, Maurer R, Seguinot V. 1996. Path integration in mammals and its interaction with visual landmarks. *J Exp Biol* 199:201–209.
- Fenton AA, Csizmadia G, Muller RU. 2000a. Conjoint control of hippocampal place cell firing by two visual stimuli. I. The effects of moving the stimuli on firing field positions. *J Gen Physiol* 116:191–209.
- Fenton AA, Csizmadia G, Muller RU. 2000b. Conjoint control of hippocampal place cell firing by two visual stimuli. II. A vector-field theory that predicts modification of representation of the environment. *J Gen Physiol* 116:211–221.
- Fuhs MC, Touretzky DS. 2000. Synaptic learning models of map separation in the hippocampus. *Neurocomputing* 32:379–384.
- 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.
- Hafting T, Fyhn M, Moser M, Moser EI. 2005. Microstructure of a spatial map in the entorhinal cortex. *Nature* 436:801–806.
- Hartley T, Burgess N, Lever C, Cacucci F, O'Keefe J. 2000. Modelling place fields in terms of the cortical inputs to the hippocampus. *Hippocampus* 10:369–379.
- Hartley T, Trinkler I, Burgess N. 2004. Geometric determinants of human spatial memory. *Cognition* 94:39–75.
- Hill AJ. 1978. First occurrence of hippocampal spatial firing in a new environment. *Exp Brain Res* 62:282–297.
- 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.
- Lenck-Santini PP, Rivard B, Muller RU, Poucet B. 2005. Study of CA1 place cell activity and exploratory behaviour following spatial and nonspatial changes in the environment. *Hippocampus* 15:356–369.
- Leutgeb S, Kjelstrup KG, Treves A, Moser M, Moser EI. 2003. Differential representation of context in hippocampal areas of CA3 and CA1. *Soc Neurosci Abstr* 29:91.5.
- Leutgeb S, Leutgeb JK, Treves A, Moser MB, Moser EI. 2004. Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science* 305:1295–1298.
- Lever C, Burgess N, Cacucci F, Hartley T, O'Keefe J. 2002a. What can the hippocampal representation of environmental geometry tell us about Hebbian learning? *Biol Cybern* 87:356–372.
- Lever C, Wills TJ, Cacucci F, Burgess N, O'Keefe J. 2002b. Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature* 416:90–94.
- Marr D. 1971. Simple memory: A theory for archicortex. *Philos Trans R Soc Lond B Biol Sci* 262:23–81.
- McNaughton BL, Nadel L. 1990. Hebb-Marr networks and the neurobiological representation of action in space. In: *Neuroscience and connectionist theory*, ed. Gluck MA, Rumelhart DE. Hillsdale, NJ: Erlbaum.
- McNaughton BL, Barnes CA, Gerrard JL, Gothard K, Jung MW, Knierim JJ, et al. 1996. Deciphering the hippocampal polyglot: The hippocampus as a path integration system. *J Exp Biol* 199:173–185.
- 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, Kubie JL, Ranck JB. 1987. Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *J Neurosci* 7:1935–1950.
- Muller RU, Bostock E, Taube JS, Kubie JL. 1994. On the directional firing properties of hippocampal place cells. *J Neurosci* 14:7235–7251.
- Nakazawa K, Quirk MC, Chitwood RA, Watanabe M, Yeckel MF, Sun LD, et al. 2002. Requirement of hippocampal CA3 NMDA receptors in associative memory recall. *Science* 297:211–218.
- O'Keefe J, Burgess N. 1996. Geometric determinants of the place fields of hippocampal neurons. *Nature* 381:425–428.
- O'Keefe J, Burgess N. 2005. Dual phase and rate coding in hippocampal place cells: Theoretical significance and relationship to entorhinal grid cells. *Hippocampus* 15 pp. 853–866.
- 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.
- Ono T, Nakamura K, Fukuda M, Tamura R. 1991. Place recognition responses of neurons in monkey hippocampus. *Neurosci Lett* 121:194–198.

- Rivard B, Li Y, Lenck-Santini P, Poucet B, Muller RU. 2004. Representation of objects in space by two classes of hippocampal pyramidal cells. *J Gen Physiol* 124:9–25.
- Rolls ET, Miyashita Y, Cahusac PM, Kesner RP, Niki H, Feigenbaum JD, et al. 1989. Hippocampal neurons in the monkey with activity related to the place in which a stimulus is shown. *J Neurosci* 9:1835–1845.
- Rolls ET, Stringer SM, Trappenberg TP. 2002. A unified model of spatial and episodic memory. *Proc R Soc Lond Ser B Biol Sci* 269:1087–1093.
- Sargolini F, Fyhn M, Hafting T, McNaughton BL, Witter MP, Moser M, Moser EI. 2006. Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science* 312 pp. 758–762.
- Save E, Nerad L, Poucet B. 2000. Contribution of multiple sensory information to place field stability in hippocampal place cells. *Hippocampus* 10:64–76.
- Sharp PE. 1991. Computer simulations of hippocampal place cells. *Psychobiology* 19:103–115.
- Sharp PE. 1999. Subicular place cells expand or contract their spatial firing pattern to fit the size of the environment in an open field but not in the presence of barriers: Comparison with hippocampal place cells. *Behav Neurosci* 113:643–662.
- Skaggs WE, McNaughton BL. 1998. Spatial firing properties of hippocampal CA1 populations in an environment containing two visually identical regions. *J Neurosci* 18:8455–8466.
- Skaggs WE, Knierim JJ, Kudrimoti HS, McNaughton BL. 1995. A model of the neural basis of the rat's sense of direction. *Adv Neural Inf Process Syst* 7:173–180.
- Taube JS, Muller RU, Ranck JB. 1990a. Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *J Neurosci* 10:420–435.
- Taube JS, Muller RU, Ranck JB. 1990b. Head-direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations. *J Neurosci* 10:436–447.
- Touretzky DS, Weisman WE, Fuhs MC, Skaggs WE, Fenton AA, Muller RU. 2005. Deforming the hippocampal map. *Hippocampus* 15:41–55.
- Ulanovsky N, Moss CF. 2007. Hippocampal cellular network activity in freely moving echolocating bats. *Nat Neurosci* 10:224–233.
- Wills TJ, Cacucci F, Lever C, Burgess N, O'Keefe J. 2004. Abrupt shift in the hippocampal place cell representation from square-like to circle-like in a morph box. Poster.
- Wilson MA, McNaughton BL. 1993. Dynamics of the hippocampal 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. 1986. Biologically plausible models of place recognition and goal location. In: 1st ed. MIT Press. pp 432–470.