UNIVERSITY OF LONDON THESIS

Degree PhD Year 2007 Name of Author CASWELL JOHN BARRY

Terra Cognita: Representations of Space in the Rodent Hippocampus and Entorhinal Cortex

Caswell John Barry

Thesis Submitted to University College London for the Degree of Doctor of Philosophy in Neuroscience

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

I, Caswell John Barry, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm this has been indicated in the thesis.

Acknowledgments

The work for this thesis was conducted while I was funded by an MRC studentship awarded via UCL's Anatomy Department. I am indebted to both organisations for their support. Equally, I owe my thanks to the ICN and Psychology Department for providing me a place in which to work.

To Neil Burgess and Kate Jeffery, my supervisors, I can only offer my sincerest thanks for continual good natured guidance and inspiration in equal measures. Thank you both so much. Also my other colleagues past and present from the ICN and Bedford Way, in no particular order: Mike, Robin, Madeleine, Tom, John, Christian, Jack, Armelle, Ali, Chris, Lili, Dmitry, Kinga, and Lisa. Thank you for sharing your ideas, your time, and your friendship so freely. Also to Elvis, Fats, and your friends, thanks.

Mum and Dad, thank you for putting up with me all this time, nearly thirty years and still not respectable. I think it's fair to say that by the time you see this, the thesis will finally be finished.

Most importantly though I owe my absolute thanks to Jessica. You have loved, nurtured, and supported me throughout. Without you I would not even have started on this path. How do you do it?

Abstract

The rodent hippocampus and associated structures are implicated in spatial memory and navigation. A necessary requirement of these roles is the ability to integrate incoming sensory information with pre-existing knowledge about an environment. With this dichotomy in mind, sensory control over place cell firing in the hippocampus, and grid cell firing in the entorhinal cortex were investigated.

In the first experimental chapter a computational model of hippocampal place cell firing is presented. The model, a two layer feed-forwards network, describes place fields as a function of the distance and direction to boundaries surrounding an animal. It is shown that by incorporating the idea of boundaries with distinct sensory qualities, and by allowing synaptic weights to be updated by application of the BCM learning rule, the model is able to capture: (1) Experiential changes in place fields resulting from prolonged exposure to a static environment, and (2) Changes in place field position and firing rate induced by movement of cues and boundaries in a familiar environment. The model is shown to compare favourably with novel electrophysiological data collected for this purpose and with experimental findings published by other authors.

The second experimental chapter investigates the affects geometric manipulations of a familiar environment have on the firing of medial entorhinal grid cells. Novel data is presented that shows grid cell firing represents an experience-dependent interaction between sensory input and learnt expectation about the size of an animal's enclosure. It is also shown that this interaction evolves with time to resolve conflict between the two sources of information. Finally it is noted that grids from a single animal are aligned and have fixed relative sizes.

The final chapter discusses the data in terms of how the brain perceives and represents the world.

Publications Arising

The following publications were generated from work conducted as part of this thesis:

Barry, C. & Burgess, N. (In press). Learning in a geometric model of place cell firing. *Hippocampus*

Barry, C., Hayman, R., Burgess, N., & Jeffery, K. J. (2007). Experience-dependent rescaling of entorhinal grids. *Nature Neuroscience,* **10**, 682-684.

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. *Reviews in the Neurosciences* **17**[1-2], 71-97.

Table of Contents

Index of Figures

Index of Tables

1 Introduction

1.1 Preamble

"I'm not lost", goes the joke. *"I know exactly where I am... I just don't know where anything else is!"* And that, in a sense, is one of the key aspects of spatial memory research in a nutshell. To locate yourself in space 'here' must be defined relative to something else; a single point, a set of landmarks or even, in the case of path integration, a previously visited location. What's more, to recognise that location in the future, it is necessary to compare a memory of the arrangement of surrounding cues with their current arrangement. It is that final point that will be addressed in this thesis: How can the brain store, and indeed reinstate, information sufficient to encode location? Put another way, *'How does a rat know where everything else is?'*

Fortunately much work has already been done. A wealth of evidence suggests that the hippocampus and surrounding structures are critically involved in spatial memory (Scoville & Milner, 1957a; O'Keefe & Nadel, 1978; Morris, Garrud, Rawlins, & O'Keefe, 1982b). In particular, electrophysiological investigations of these areas have identified cells with spatially defined receptive fields, most notably hippocampal place cells (O'Keefe & Dostrovsky, 1971; O'Keefe, 1976) entorhinal grid cells (Fyhn, Molden, Witter, Moser, & Moser, 2004; Hafting, Fyhn, Moser, & Moser, 2005), and head direction cells (Taube, Muller, & Ranck, 1990b; Taube, Muller, & Ranck, 1990a). As such we can refine the previous question and ask instead the related questions, *'How is the spatial firing of these cells controlled?'* and *'What role does learning play in the development of this control?'*. This thesis employs two broad approaches to address these questions. In chapter three, a simple computational model of place cell firing is described. The model builds on the Boundary Vector Cell (BVC) model, first presented by O'Keefe and Burgess (1996) and developed more fully by Hartley *et al.* (2000). By incorporating the BCM learning rule (Bienenstock, Cooper, & Munro, 1982), the updated BVC model captures some of the experience-dependent changes observed in the place cell representation and, after a learning phase, exhibits place fields controlled by specific, salient environmental cues. The model's predictions are compared with experimental data collected by the author and also with results published by Fenton *et al.* (2000a) and Rivard *et al.* (2004). In chapter four, recordings of grid cells from the rodent dorso-lateral medial entorhinal cortex (dlMEC) are

described. In particular, the tessellated, triangular firing pattern, typical of grid cells is shown to exhibit a strong experience-dependent environmental influence; the grid's spatial scale varies parametrically with changes to the size and shape of a familiar environment. Finally in chapter five, the synergies between these studies are discussed with a view to answering the questions described above and other outstanding issues, such as, the relationship between grid cells, place cells and BVCs are examined. However, before reaching that point a brief overview of the relevant literature is presented below.

1.2 Nomenclature

Throughout this document the *hippocampal formation*, following Amaral and Witter (1995), is considered to include six distinct regions: dentate gyrus (DG); the Cornu Ammonis (CA) fields CA1, CA2 and CA3; subiculum; presubiculum; parasubiculum; and entorhinal cortex (EC). *'Hippocampus'*, where used, is considered to apply only to CA1, CA2 and CA3. Note, this nomenclature is predicated on the largely unidirectional connections that link these six areas and which distinguish them from neo-cortical structures (Amaral & Witter, 1995). As such it departs from the scheme applied elsewhere, for example by Scharfman *et al.* (2000), in which the hippocampal formation is considered to consist exclusively of allocortex; being the DG, CA fields, and subiculum; three layered structures consisting of a single cell layer between plexiform layers. The term *'place cell'* (O'Keefe, 1976) is used exclusively to apply to pyramidal cells of the CA1 and CA3 regions that exhibit spatially constrained firing. Similarly *'grid cell'* (Hafting et al., 2005) describes dlMEC cells that exhibit spatially stable firing arranged in a periodic, triangular grid that spans the environment. This definition likely includes stellate cells from layer II and pyramidal cells from layers III, V and VI. Other cell types will be defined as is necessary but in general will simply be referred to by region (e.g. 'cells in the subiculum that exhibit spatial firing').

1.3 Structure and morphology of the hippocampal formation

1.3.1 Overview

In the rat, the hippocampal formation represents a significant component of the central nervous system, occupying half the entire cortical volume and presenting a combined surface area nearly equivalent to the entire neocortex ($1.2cm^2$ vs. $1.5cm^2$) (Swanson, 1983; Amaral et al., 1995). In gross outline, the hippocampal formation resembles a pair of 'C's with the open face of the 'C' being most rostral, the entire structure being buried below the neocortical sheet (see Fig 1.1A). The long axis is referred to as the septo-temporal axis, with the septal pole being both rostral and dorsal to the temporal pole, for this reason septo-temporal is often used synonymously with dorso-ventral. A further axis, the transverse axis, is defined running medial to lateral, such that regions closest to the DG are considered to be proximal, those closest to the rhinal sulcus and postrhinal cortices are considered to be distal. Finally, the radial axis is defined orthogonally to both of these, such that surfaces close to the ventricle are considered to be deep, and those closer to the hippocampal fissure, superficial.

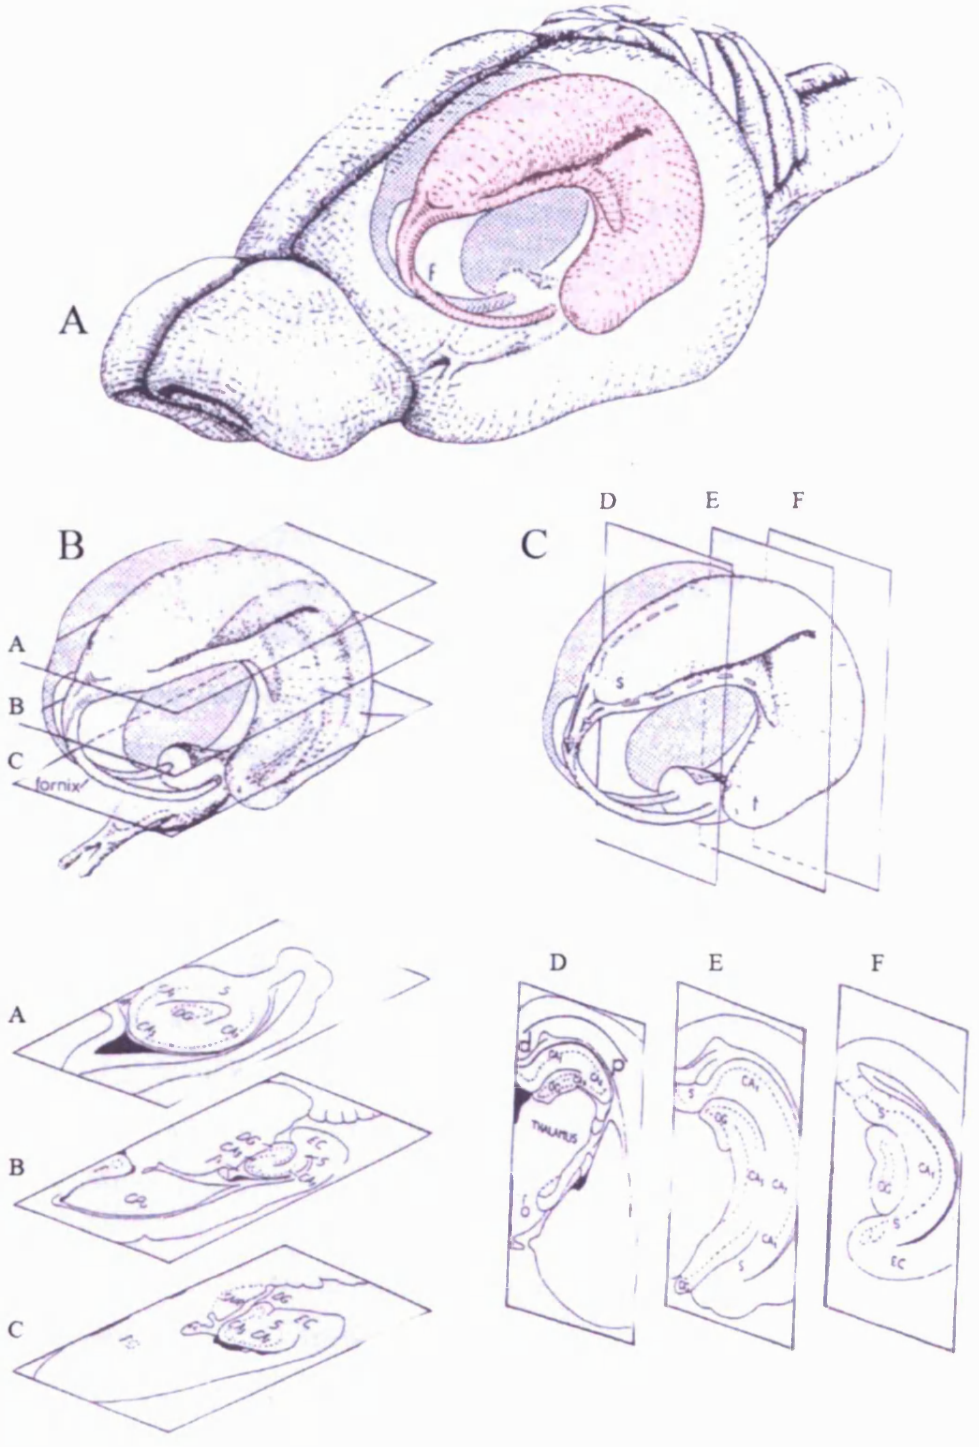

Fig 1.1. Rat hippocampus shown in horizontal and coronal section. **a**) The hippocampi are shown in position below the neocortical sheet. Left hippocampus is coloured red. The fornix (f) is indicated descending to its subcortical targets. **b**) Horizontal sections taken at three dorso-lateral levels along the hippocampus reveal its differing composition; EC is only evident at ventral levels. **c**) Coronal sections taken at three rostro-caudal levels. CA1 and CA3, principal fields of hippocampus. CPu, caudate and putamen. DG, dentate gyrus. Fi, fimbria. S, subiculum. Adapted from Amaral and Witter (1995)

The constituents of the hippocampal formation change considerably at various positions along the septo-temporal axis (Fig 1.1B&C). At the septal pole, for example, only the DG and CA fields are present, these are easily identified in slices by the distinctive interlocking 'C' shape formed by the cell layers in each. The subiculum appears about a third of the way along the axis with the presubiculum and parasubiculum located even more temporally. Finally, the EC is located ventrally, at the most caudal portion of the cortex. The major fibre bundle of the hippocampal formation is the fimbria-fornix. The fimbria, located at the lateral edge of the hippocampus, increases in thickness along the septal-temporal axis, being most well developed at the septal pole. It is formed from the alveus, a sheet of afferent and efferent fibres that originate from the hippocampus and subiculum and cover the deep surfaces of those structures. The fimbria departs the hippocampal formation at its septal extreme and descends to the basal forebrain, at this point it is referred to as the fornix. Immediately prior to the fornix a large number of fimbria fibres give rise to the ventral hippocampal commisure; the majority of the fibres project to hippocampal fields in the other hemisphere with a smaller number joining the contralateral fornix. The fornix itself branches several times, ultimately innervating multiple subcortical targets such as the septal nuclei, anterior thalamic nuclei, mammillary bodies, and hypothalamic areas (Gloor, 1997). A final point of note, is the angular bundle which carries entorhinal as well as para and presubicular fibres and is located between the entorhinal cortex and para and presubiculum. The angular bundle gives rise to the dorsal hippocampal commisure and is the main route by which fibres travel from the entorhinal cortex to other hippocampal fields.

Early studies of connectivity within the hippocampal formation emphasised the distinct, and well known tri-synaptic circuit (Fig 1.2A); a supposedly unidirectional course the flows from the entorhinal cortex to the dentate gyrus, then via CA3, onwards to CA1, finally returning to the entorhinal cortex through direct and indirect projections. We now know this view to be over simplistic as modern histological techniques have revealed additional intrinsic and extrinsic connectivity within the hippocampal formation. Still, the tri-synaptic circuit provides an obvious route to follow in order to study the hippocampal subfields, and it is to those that we now turn. The emphasis of the following subsections is on relevance. Anatomical detail is deliberately eschewed where it is not considered relevant to the models and data that

are the focus of this work. A complete review of hippocampal anatomy is provided by Amaral and Witter's excellent chapter in the Rat Nervous System (1995) also Amaral and Lavenex's work in the Hippocampus Book (2007).

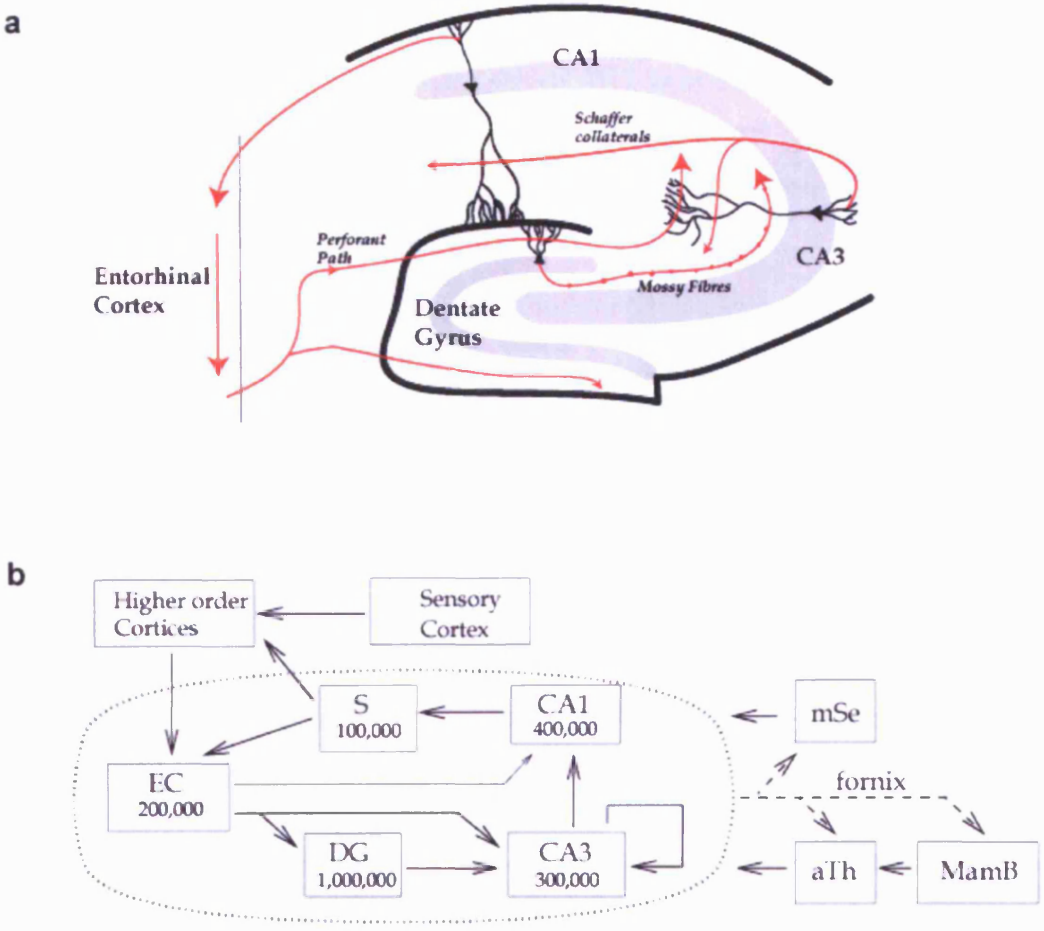

Fig 1.2. Two diagrammatic representations of hippocampal connectivity. **a**) A relatively simple view of the traditional tri-synaptic circuit. Projections from the EC reach the DG and CA3 via the perforant path. In turn the DG projects to CA3 via the mossy fibres. CA3 activity reaches CA1 through the Schaffer collaterals and the loop is 'closed' by projections back to the EC. **b**) A more complex but not complete sketch of connectivity, approximate numbers of principal cells are indicated next to each structure (S, subiculum; mSe, medial septum; aTh, anterior thalamic nuclei; MamB, mammillary bodies). Adapted from Burgess and O'Keefe (2002).

1.3.2 Entorhinal cortex

As stated previously, the rodent entorhinal cortex (EC) is a six layered structure located at the ventro-posterior extremity of the cerebral hemisphere. It is bordered laterally by the rhinal sulcus and medially by the presubiculum (Fig 1.3a&b). On the basis of morphology and connectivity, the region can be subdivided into two fields, lateral and medial entorhinal cortex (LEC and MEC), often referred to as the lateral

and medial entorhinal areas (LEA and MEA). Despite their names the MEC is, largely speaking, caudal to the LEC, with the two triangularly shaped regions being placed so as to form a trapezoid structure.

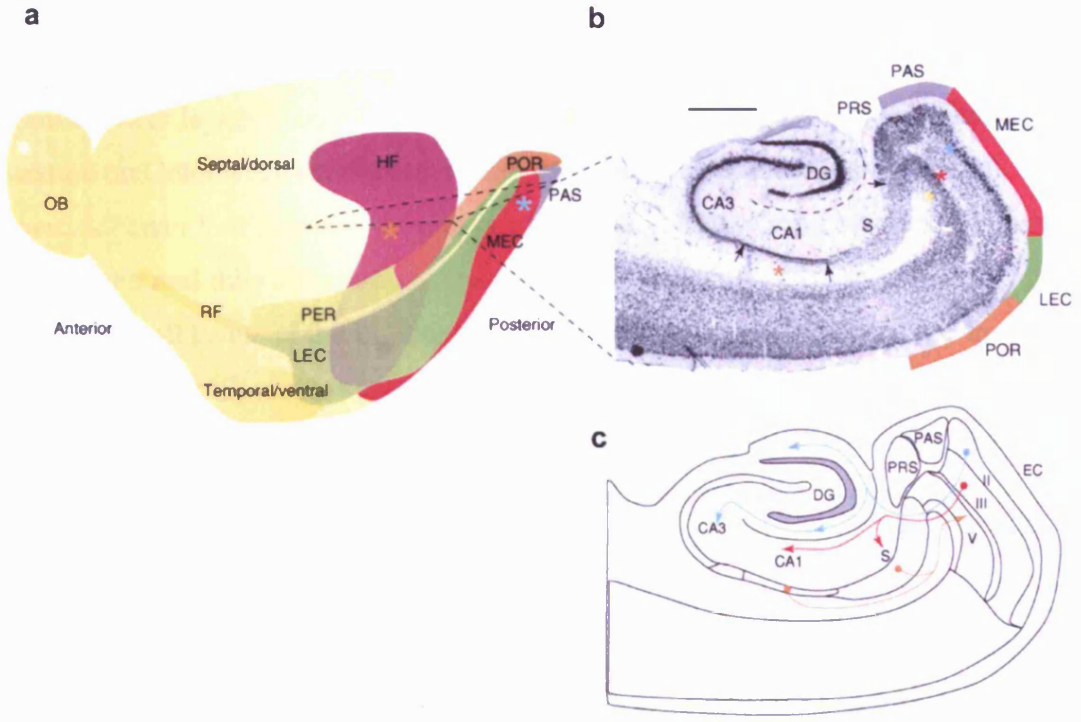

Fig 1.3. Position and connectivity of the entorhinal cortex. **a**) Left hemisphere of rat brain, the triangularly shaped LEC (green) and MEC (red) can be made out on the posterior surface bordered by perirhinal (PER) and postrhinal (POR) cortices. The hippocampal formation (HF) is indicated in purple. **b**) Horizontal section through the hippocampal formation's subfields; scale bar indicates 100um. The postrhinal cortex, LEC and MEC are indicated as before with the parasubiculum (PAS) and presubiculum (PRS) proximal to those structures. Below the cortical sheet, the interlocking 'C's of the CA fields and dentate gyrus (DG) are clearly visible. The less dense pyramidal cell layer of the subiculum (S) is also very clear. **c**) Main intrinsic EC connectivity. Layer II projects out to the DG and CA3 while layer III projects to CA1 and the subiculum. Reciprocal, point to point connects are returned from CA1 and the subiculum to layer V. Image adapted from Witter and Moser (2006)

In both the LEC and MEC, layers II, III, V and VI are the principal cell layers, leaving two plexiform layers, layers I and IV; layer IV, the lamina dessecans is less distinct in the LEC. Stellate cells, a form of pyramidal cell, comprise the majority of layer II and are the principal origin of the perforant path to the dentate gyrus and CA3. Pyramidal, multipolar, and horizontal cells are also found in layer II, though in lesser numbers (Fig 1.3c). In common with other cell layers in the EC, layer II produces multiple collaterals, which in this case innervate only superficial layers (Kohler, 1986; Kohler, 1988), and in the case of the MEC are known to provide excitatory connections

between stellate cells (Kumar, Jin, Buckmaster, & Huguenard, 2007). Pyramidal cells, the principal cell type of layer III, originate projections to CA1 and the subiculum and like layer II send collaterals to superficial layers, though mainly I and III (Fig 1.3c). These cells also send apical dendrites up to layer I where they likely synapse with the sparse inhabitants of that layer; GABA-ergic neurons, as well as horizontal and stellate cells (Witter, Groenewegen, Lopes da Silva, & Lohman, 1989). The organisation of layers V and VI is more complex and less well understood, relying to an extent on Golgi work conducted during the 1930s (e.g. (Lorente de No, 1933)). Superficial layer V is dominated by large pyramidal cells that send axons to the deep white matter and thus out to other cortical regions. The same cells also dispatch collaterals to all layers of the EC, though predominantly V and VI. However, in deep layer V and layer VI, various cell types send highly collateralised axons to influence a column of cells in layers I, II and III and also project out of the EC via white matter. Putative GABA-ergic interneurons are found in all layers though in increased concentration in superficial layers (Swanson, Kohler, & Bjorklund, 1987).

The superficial layers of the EC receive much of their extrinsic cortical input via the postrhinal and perirhinal cortices which in turn receive input from uni-modal and poly-modal association areas. In particular, the perirhinal cortex receives strong olfactory projection from the piriform cortex, while the postrhinal cortex receives a greater proportion of visual and visuo-spatial inputs from the retrosplenial, posterior parietal and cingulated cortices (Burwell, Witter, & Amaral, 1995; Burwell & Amaral, 1998b; Burwell & Amaral, 1998a). This distinction is maintained within the EC, with perirhinal cortex primarily projecting to the LEC, whereas postrhinal cortex prefers the MEC. In addition, the superficial EC receives direct neocortical projections from uni-modal, notably olfactory, and poly-modal association areas, and these too follow a similar scheme, such that the MEC receives predominantly visuo-spatial inputs (Burwell et al., 1998b). The deep layers of the EC also receive cortical afferents from a variety of areas including the medial prefrontal region and agranular insular cortex (Insausti, Herrero, & Witter, 1997). These deep inputs are generally viewed as modulatory, influencing the flow of information out of the hippocampal formation, and stand in contrast with superficial afferents which are assumed to convey information destined for the hippocampus. The EC also receives a number of

subcortical inputs from structures as diverse as the amygdaloid complex, the septum, thalamic nuclei and various structures in the hypothalamus.

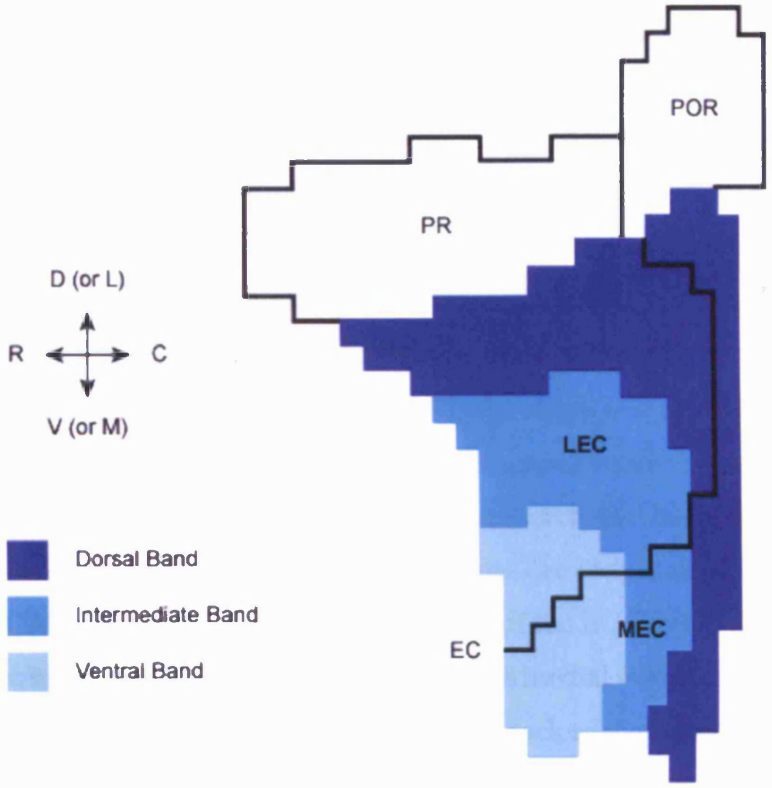

Fig 1.4. Unfolded map of the EC and neighbouring structures (PR, perirhinal cortex; POR, postrhinal cortex). The lateral EC (LEC) and medial EC (MEC) are shown overlaid by the gradient of topographical connections to the DG (shades of blue). Adapted from Burwell and Amaral (1998a)

The perforant path, so called because it passes through, or perforates the subiculum, represents the primary entorhinal input to the hippocampus. Originating mainly but not exclusively from layer II stellate cells and layer III pyramidal cells, it projects to all fields of hippocampus (Germroth, Schwerdtfeger, & Buhl, 1989; Kohler, 1985; Kohler, 1986; Kohler, 1988). Based on topographical considerations, it is possible to distinguish at least three axes that relate to the organisation of the perforant path. The most obvious being that layer II cells project to the DG and CA3, while layer III projects largely to CA1 and the subiculum (Fig 1.3c). A further distinction can be made based on whether projections originate in the MEC, LEC, or a putative intermediate zone that lies between them. At the level of the DG and CA3 this second axis is reflected in the radial termination site of perforant path fibres: projections from the LEC terminate in the outer third of the molecular layer and in similar superficial sites in the stratum lacunosum-moleculare; fibres originating in the MEC, on the other

hand, terminate in the middle third of the DG and at more superficial sites in CA3. Electrophysiological investigations support the intuitive conclusion that individual cells within these areas are likely to be influenced by projections from both the LEC and MEC (McNaughton & Barnes, 1977). In contrast, at the level of CA1 and the subiculum, the lateral and medial components of the perforant path terminate at different transverse positions. For example, the medial perforant path reaches proximal CA1 and distal subiculum, whereas lateral fibres terminate in distal CA1 and the adjacent, proximal section of the subiculum. A third axis lies across the MEC - LEC division such that cells located medially in the EC, project to temporal positions within the hippocampus and those located laterally project to the septal levels (Fig 1.4). To clarify, because of the triangular shape and misnomered identities of the MEC and LEC, medial and lateral bands drawn across the EC as a whole actually cut across both areas. A final, important consideration is that both the septo-temporal and distal-proximal arrangement, described above, are reciprocated by projections back from CA1 and the subiculum to deep layers of the EC. For example, medial-septal CA1 which receives perforant input from the medial extent of the MEC, the region in which grid cells are found, will project back to that same section of the MEC.

1.3.3 Dentate gyrus

The dentate gyrus (DG) is comprised of three layers: most superficial is the largely cell free molecular layer; below this is the granule cell layer containing densely packed granule cells; deep to both of these is the polymorphic layer, sometimes known as the hilus, which contains a variety of cell types.

The molecular layer is mainly occupied by dendrites of cells in the deeper layers, in addition to terminal arbours of entorhinal afferents. However, two inhibitory cell types worthy of note are the molecular layer perforant path-associated cell (MOPP) and axo-axonic cell. The latter descends to the granule cell layer forming symmetrical synapses on the axons of granule cells. Granule cells themselves are the principal cell type of the DG, generating afferents to CA3. Estimates of the number of these cells vary, though typically fall above 1×10^6 for a single hemisphere, making them one of the most numerous cell types in the hippocampal formation (Fig 1.2B) (West, Slomianka, & Gundersen, 1991). Unusually the number of granule cells continues to increase into adulthood with new cells being born in the subgranular zone.

Subsequently these neonate cells are preferentially recruited into spatial memory networks (Bayer, Yackel, & Puri, 1982; Kee, Teixeria, Wang, & Frankland, 2007). Intimately associated with granule cells are the inhibitory pyramidal basket cells, found along the deep surface of the granule cell layer, they form synapses with, and peri-cellular plexuses around, the cell bodies of granule cells. The most prominent cell type of the polymorphic layer is the mossy cell. These cells are glutamatergic, receive inputs from granule cell axons, and project back to the granule cell layer.

As described previously the primary extrinsic input to the DG contacts the molecular layer via the perforant path from the EC. Minor projections to this layer also arise in the pre and parasubiculum (Kohler, 1985). The DG also receives limited input from subcortical structures, the most notable being septal projections from the medial septal nucleus and the nucleus of the diagonal band of Broca. Septal fibres mainly terminate in the polymorphic layer, the majority being cholinergic with a smaller proportion being GABA-ergic. The DG also receives afferents from the brainstem, for example from the locus coeruleus and raphe nuclei, as well as from the hypothalamus.

In addition to intrinsic inhibitory inputs provided by a variety of basket cell types, the DG also exhibits two types of intrinsic excitatory connection. Granule cells, as well as originating mossy fibres bound for CA3, also form thinner collaterals that innervate the polymorphic layer forming synapses with mossy cells, pyramidal basket cells as well as other GABA-ergic cell types. In turn, mossy cells, and possibly other cell types from the polymorphic layer, project back to the molecular layer forming glutamatergic synapses with granule cells and basket cells. Curiously, the feed-back connections from mossy cells to granule cells reach much of the septo-temporal axis but are weakest at the level of origin and stronger at more distal levels along the axis (Amaral & Witter, 1989).

Mossy fibres, un-myelinated axons originating from granule cells, terminate in CA3 just above the pyramidal cell layer in the stratum lucidum; in proximal regions of CA3 mossy fibres also terminate within and below the pyramidal cell layer. Connections between the mossy fibres and proximal dendrites of CA3 pyramidal cells are unusual; comprising of complex en passant presynaptic terminals known as mossy fibre

expansions, coupled with postsynaptic branched spines called thorny excrescences. Not only are these terminals extremely large, up to 8μm, but each mossy fibre might make as many as 37 with a single pyramidal cell (Chicurel & Harris, 1992). Together with their size, the location of these synapses, close to the cell body may imply that mossy fibres exert a powerful influence over CA3 pyramidal cells. Conversely each mossy fibre contacts relatively few, about 15, pyramidal cells which indicates that each pyramidal cell receives input from around 70 granule cells.

1.3.4 CA3 & CA1

All of the Cornu Ammonis fields (CA1, CA2 and CA3) share a similar laminar structure. The principal cell layer being the pyramidal cell layer sandwiched between the deep, relatively cell free stratum oriens and superficial stratum radiatum. It is this latter structure, the stratum radiatum that contains CA3 to CA3 associational synapses and CA3 to CA1 Schaffer collateral synapses. In CA3 only, an additional layer, the stratum lucidium is interposed between the pyramidal cell layers and stratum radiatum. Finally, in all fields the perforant path terminates in the most superficial layer, the stratum lacunosum-moleculare. In the following description CA2 is largely ignored, this reflects both a historical uncertainty about its relationship to CA1 and CA3 (it can be seen as a terminal section of CA3) as well as its relatively small size compared to the other CA fields.

Pyramidal cells are the principal cell type of CA1 and CA3, comprising most of the cells in the pyramidal cell layer. From each cell basal dendrites descend into the stratum oriens while apical dendrites reach up into the stratum radiatum and stratum lacunosum-moleculare. CA1 pyramidal cells tend to be smaller than those from CA3 both in terms of somata and dendritic arbourisation. Along with the principal pyramidal cells, two types of interneuron, pyramidal basket cells and axo-axonic cells, have their cell bodies in or adjacent to the pyramidal cell layer. In both cases their dendrites span all levels of the field, with the axons of the basket cells innervating the soma and proximal dendrites of pyramidal cells, whereas axo-axonic cells form synapses onto the axons of pyramidal cells. Several other non-pyramidal cell types are found throughout CA1 and CA3, the majority being immunoreactive for GABA, and as such are presumed to be inhibitory interneurons. Indeed, a recent count identified

nearly 20 different types of interneuron distinguished on the basis of morphology and histology (Somogyi & Klausberger, 2005).

As described previously, CA3 and CA1 both receive extrinsic input from the EC via the temporoammonic path while CA3 is also driven by mossy fibres from the DG. However, the main source of input for both fields comes from within the hippocampus; CA3 generates highly collateralised axons that innervate all levels of the ipsilateral and contralateral CA fields. The CA3 to CA1 connections, known as Schaffer Collaterals are highly divergent and a single CA3 cell can project to as much as 65% of the septo-temporal extent of CA1 (Li, Somogyi, Ylinen, & Buzsaki, 1994). Connection density is not uniform though and a complex topography relates the septo-temporal and proximal-distal position of cells in CA3 to their termination site in CA1 (see (Ishizuka, Weber, & Amaral, 1990) for complete details). The CA3 to CA3 (associational connections) are similarly divergent and also exhibit a topographic organisation such that projection targets of cells in proximal portion of CA3 are less widely spread than those located more medially (Ishizuka et al., 1990). Additionally, in the rat but not in primates, CA3 cells produce commissures that innervate homologous fields in the contralateral hemisphere (Blackstad, 1956). Indeed, a single CA3 cell can generate collaterals destined for both the ipsilateral and contralateral hemispheres (Swanson, Sawchenko, & Cowan, 1980)

Both CA3 and CA1 also receive projections from the amygdala, though in the case of CA3 these are mainly limited to the temporal region. CA1 alone receives additional cortical input from perirhinal and postrhinal cortices, and sends return projections back to the former structure as well as to retrosplenial cortex. Similar to the DG, CA3 and to a lesser extent CA1 receives and returns subcortical input from the septum, specifically the medial septal nucleus and diagonal band of Broca. CA3 also receives some afferents from hypothalamic regions, whereas CA1 is targeted by thalamic afferents from regions such as the nucleus reunions. Finally, both structures receive input from brain stem nuclei such as the raphe nuclei.

Aside from the dense associational connections and Schaffer collaterals, the only other major target of CA3 is the lateral septal nucleus, to which it is reciprocally connected. Thus CA1 is essentially the main output of the CA fields. In addition to the

connectivity described in the previous paragraph, CA1 provides major input to the subiculum and EC. In the case of the latter structure the topography of these connections reciprocates the topography of direct EC to CA1 connects (described above). Subicular efferents are also topographically organised such that proximally positioned CA1 cells project to the distant, distal portion of the subiculum, whereas distal CA1 cells project to the proximal subiculum. CA1 axons exit CA1 at the level of the stratum oriens and terminate in the pyramidal cell layer and deep molecular layer of the subiculum.

1.3.5 Subiculum

In addition to CA1, the subiculum originates the other major projection back to the EC. It is also the source of the chief output from the hippocampal formation to subcortical structures and the major origin of the fornix. As such, the subiculum can be thought of as the final step in the hippocampal processing loop. The border between CA1 and the subiculum is marked by a loss of the stratum radiatum which is replaced by the subiculum molecular layer, and a concomitant widening of the pyramidal cell layer. The stratum oriens of CA1 is also lost at this point. The deep portion of the molecular layer is continuous with the stratum radiatum of CA1 and, as noted above, is the site at which CA1 afferents terminate. The most superficial section of the molecular layer, on the other hand, is continuous with the stratum lacunosum-moleculare of CA1 and receives afferents from the EC.

Similar to the hippocampus, the principal cell type of the subiculum is a large pyramidal cell that populates the pyramidal cell layer. Though morphologically similar, electrophysiological criteria distinguish at least two subpopulations, regular spiking cells and intrinsically bursting cells. Bursting cells are mainly found deep in the pyramidal layer, whereas regular spiking cells are more numerous superficially. Interestingly, it is believed that only the bursting cells project to the EC. Relatively little is known about subicular interneurons, though many smaller cell types are found in the pyramidal cell layer, some of which are immunoreactive for GABA.

As described previously, the subiculum is the major target of CA1 efferents. It also exhibits internal associational connections such that subicular pyramidal cells project to more temporal regions of the subiculum. Efferents from the subiculum target the

presubiculum, terminating largely in layer I with some deeper termination sites in the dorsal presubiculum. Similarly, projections to the parasubiculum also target layer I with some synapses in superficial layer II. Both sets of connections are topographically organised, such that the septal subiculum projects to dorsal and caudal regions of the para and presubiculum, while temporal subiculum targets ventral and rostral regions. Similar to CA1, subicular projections to the EC reciprocate the topography of EC to subicular connections. Hence the proximal subiculum projects to the LEC whereas more distal areas reach the MEC. The principal termination site in EC is layer V, though afferents also reach layer VI, as well as the more superficial layer III. The subiculum also generates input to the medial and ventral orbitofrontal cortices, prelimbic and infralimbic cortices, retrosplenial cortex as well as generating strong input to the perirhinal cortex. In addition to these connections, the subiculum forms substantial subcortical projections, the most extensive being those to the mammillary nuclei, septal nuclei, and nucleus accumbens.

1.3.6 Presubiculum and parasubiculum

The pre and parasubiculum do not fit quite so nicely into the notion of a 'hippocampal processing loop' as do the other structures of the hippocampal formation. Indeed it remains unclear whether they should be considered as an input, like the EC, or as an output similar to the subiculum. Amaral and Witter (1995) concluded that these two structures are most akin to an input as they project to superficial layers of the EC. However, they also concede that strong reciprocal connectivity with the anterior thalamic nuclei marks them out as a conduit for thalamic control over the hippocampal formation.

The presubiculum, placed distally to the subiculum, is sometimes subdivided into the presubiculum and postsubiculum, the latter structure being synonymous with 'dorsal presubiculum'. Together with the adjacent parasubiculum, the presubiculum is distinct from the allocortical hippocampus in that they are multilaminate like the EC, and as such six layers can be distinguished. In both structures layers II and III are densely populated by pyramidal cells, while deeper layers, which are continuous with the deep layers of the EC, are populated by a variety of different cell types.

The pre and parasubiculum both generate internal associational connections, as well as commissures to homotopic regions of the contralateral hemisphere. However, most striking is the dense interconnectivity with anterior thalamic nuclei exhibited by both structures. Thalamic connectivity is best described in the presubiculum, in that dorsal regions are known to receive input primarily from laterodorsal and anterodorsal nuclei, more ventral regions are reached by projections from latrodorsal and anteroventral nuclei. In addition to the subicular input mentioned previously, other cortical inputs arrive from the retrosplenial cortex and to a lesser extent visual area 18b. Interestingly, the presubiculum projects bilaterally to the parasubiculum. Subcortical inputs are received from the medial septal nuclei, the diagonal band of Broca, mammillary nuclei, and the brain stem.

Aside from the reciprocal thalamic connectivity the two major outputs of the pre and parasubiculum are to structures within the hippocampal formation. Both send substantial afferents to superficial layers of the EC; the presubiculum primarily to layer III of the MEC while the parasubiculum innervates layer II of the MEC and to a lesser extent the LEC. In addition, the presubiculum generates input to the DG, CA fields and subiculum. However in the parasubiculum the input to the DG is much more developed. Interestingly, Amaral and Lavenex (2007) suggest that these final two points may indicate a route by which thalamic input can influence the early stages of information processing in the hippocampal formation.

1.3.7 Summary

In synopsis, a simple view is that hippocampal connectivity describes a processing loop, albeit a loop with several sub-loops. Highly processed sensory information reaching the superficial layers of the EC is then routed into the hippocampus proper either via the dentate-CA3 fields or directly to CA1. Back projections from CA1 to the deep layers of the EC as well as similar projections from the subiculum close the loop. Finally, projections from the deep layers of the EC either return the output to other cortical regions or redirect it to superficial layers of the EC, presumably from where it re-enters the loop or modulates the entrance of novel information. However, not acknowledged in this reductionist model is the topographic organisation of connections seen, to a greater or lesser extent, within all fields of the hippocampal formation. Similarly downplayed are the substantial subcortical connections that, at

the very least, can be considered to modulate the flow of information around the 'loop'.

1.4 Place cells and their allies

Place cells, hippocampal pyramidal cells that exhibit spatially localised activity, were first identified by O'Keefe and Dostrovsky in 1971 (O'Keefe et al., 1971). Since the initial discovery, cells with spatially modulated firing have been found in almost all areas of the hippocampus and in some surrounding areas (e.g. dentate gyrus (Jung & McNaughton, 1993), entorhinal cortex (Quirk, Muller, Kubie, & Ranck, 1992; Hafting et al., 2005), subiculum (Sharp & Green, 1994)). Although early work was conducted on rats, similar cells have subsequently been found in other species of rodent (Tonegawa et al., 1996), bats (Ulanovsky & Moss, 2007), pigeons (Bingman, Siegel, Gagliardo, & Erichsen, 2006), monkeys (Rolls et al., 1989; Ono, Nakamura, Fukuda, & Tamura, 1991; Hori et al., 2003; Ludvig, Tang, Gohil, & Botero, 2004) and even humans (Ekstrom et al., 2003).

1.4.1 General properties of place cells

The striking quality of place cells is that they seem to provide a precise representation of an animal's position in its environment. The background firing rate of place cells is very low, effectively being zero. When an animal enters the receptive field (place field) of a cell, its firing rate rapidly increases, typically to a maximum between 5Hz and 15Hz (Fig 1.5). Interestingly, in a given environment, the majority of place cell are silent with only about a third of CA3 and CA1 cells being active (Thompson & Best, 1989; Guzowski, McNaughton, Barnes, & Worley, 1999; Leutgeb, Leutgeb, Treves, Moser, & Moser, 2004b).

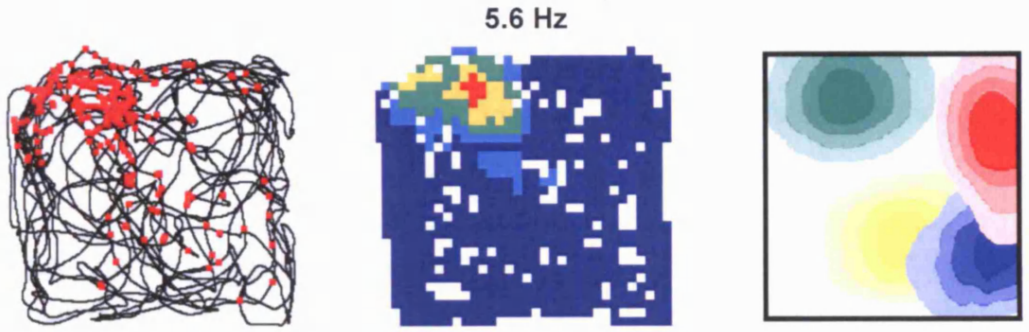

Fig 1.5. Example of the spatially constrained firing typical of a CA1 place cell. Data recorded by the author. Four minute recording made in a 70cm square enclosure. From left to right: **Left**) Rat's path is denoted by the continuous black line with action potentials indicated by the superimposed red dots; each dot indicates the animal's position when an action potential was fired. **Middle**) Rate map constructed from raw data shown to the left. The firing rate of the place cell was calculated for each bin - in this case the environment was divided into approximately 30 by 30 bins. Peak rate is indicated above the map and colours from 'hot' to 'cold' indicate firing rate as a percentage of the peak rate (dark blue 0-20%, light blue 21-40%, green 41-60%, yellow 61-80%, and red 81-100%). White indicates unvisited bins. **Right**) The firing of even a small population of place cells is sufficient to 'carpet' the enclosure.

In an open environment, activity is independent of the animal's orientation (O'Keefe, 1976). In essence, firing is best correlated with the position of an animal's head (Muller & Kubie, 1989), and can be used to infer it (Wilson & McNaughton, 1993). Conversely, environments or tasks that constrain an animal's path often result in place fields developing a strong directional component to their firing (McNaughton, Barnes, & O'Keefe, 1983; Markus, Qin, Leonard, McNaughton, & Barnes, 1995). Interestingly, the complementary representation of orientation, independent of location, is found in the 'head-direction' cells of the lateral mammillary bodies, anterior thalamus and presubiculum (Taube et al., 1990b; Taube et al., 1990a; Taube & Muller, 1998). These cells fire whenever the animal's head is pointing in a given direction, independent of the animal's location (Taube et al., 1990b; Burgess, Cacucci, Lever, & O'Keefe, 2005). More recently the conjunction of both cell types have been found in the dorsal presubiculum; theta-modulated place by direction cells (TPDs) exhibit spatially constrained firing that is dependent on orientation (Cacucci, Lever, Wills, Burgess, & O'Keefe, 2004).

Place fields seem to be randomly distributed across an environment, with perhaps a tendency to be smaller and more numerous near to the environmental boundaries (Hetherington & Shapiro, 1997). Place cell firing can be detected within the first few minutes of an animal's entrance into a novel environment, though is initially unstable and less robust than in a familiar space (Wilson et al., 1993). In invariant conditions,

place fields are extremely stable, in some cases persisting for months (Best & Thompson, 1984). During a single exposure to an environment, a place cell will continue to fire each time an animal moves through its field; although the cell's firing rate can be curiously variable (Muller, Kubie, & Ranck, 1987b).

1.4.2 Sensory control of place cell firing

So what causes place cells to fire where they do? Early studies showed that place fields are not defined by a single modality of sensory information, but will make use of whichever are available, including visual, tactile, olfactory and auditory inputs (Best & Thompson, 1989b; Save, Nerad, & Poucet, 2000). The representation of a given environment is robust to the removal of subsets of cues (O'Keefe J & Conway, 1978; Nakazawa et al., 2002). The absence of spatial cues from one modality tends to be compensated for by cues from other modalities (Save et al., 2000). For example, congenitally blind rats exhibit normal place fields (Save, Cressant, Thinus-Blanc, & Poucet, 1998). In the temporary absence of other cues, path integration (dead reckoning) (Etienne & Jeffery, 2004), can also support a place cell representation (O'Keefe, 1976; Quirk, Muller, & Kubie, 1990; Etienne, Maurer, & Seguinot, 1996). However, in the absence of external sensory information, the path integration process rapidly succumbs to cumulative error (Etienne et al., 1996). In terms of the hierarchy of influence of these various inputs, it seems that stable visual cues at or beyond the edge of the reachable environment exert a powerful control over the orientation of the place cell (O'Keefe J et al., 1978; Jeffery & O'Keefe, 1999) and head-direction cell (Taube et al., 1990a) representations. For example, rotation of a single polarising cue in a cylindrical environment causes a complementary rotation in the entire place cell ensemble (Muller & Kubie, 1987a). In the absence of reliable and stable (Jeffery et al., 1999) visual cues, similar effects can be produced by rotation of a sloped environment (Jeffery, Anand, & Anderson, 2005), and even of the rat itself (Jeffery et al., 1999). If rats are systematically disoriented between trials, the orientations of the place and head-direction representations can be made to rotate; failing to maintain a stable correspondence to the external world. Interestingly, in this situation, both representations rotate in synchrony with each other (Knierim, Kudrimoti, & McNaughton, 1998).

The above results pertained to cues that orient place cells in a rotationally symmetrical environment of constant dimensions. Changes to the shape or topography of an environment have also been used to investigate the nature of the inputs driving place cell activity. Comparing recordings from the same cells made in rectangular environments of varying shape and size, O'Keefe and Burgess (1996) noted that the location of peak firing typically maintained a constant position relative to the nearest walls. In addition to this, several fields were stretched along the axes of the environment, with some becoming bimodal in the larger environments. More recently, similar parametric responses were seen across a variety of different shaped environments (Lever, Cacucci, Burgess, & O'Keefe, 1999). O'Keefe and Burgess (1996), see also (Burgess & O'Keefe, 1996a), proposed that place cells received inputs that are tuned to respond to the presence of a barrier at a given distance along a given allocentric direction, with sharper tuning at shorter distances. Related, but more localised effects have also been observed by Muller (1987a). He showed that bisecting a place field with a barrier would often cause the cell to stop firing, while those with more distant fields remained unchanged. In short, it seems that impediments to movement, be they the walls of the environment, a free standing barrier or even a sheer drop at the edge of a platform, play a key role in defining place cell firing. In contrast to the influence of extended barriers, Cressant *et al.* (1997), showed that isolated objects within the environment failed to affect place cell firing. However, the same objects did influence place fields when moved to the edge of the environment, where they acted as orientation cues. Similarly, when placed in a line they functioned as an extended barrier.

1.4.3 Remapping

Pronouncing on sexually-explicit magazines, a judge once famously said "I can't define it, but I know what it is when I see it". Unfortunately, it seems that statement also applies to remapping (Knierim, 2003). For example, it is tempting to define remapping by way of contrast with the parametric changes described above. One might write that: sufficient perturbation of an animal's surroundings or motivational state can cause some or even all of the cells in a representation to radically change their behaviour (O'Keefe & Conway, 1978; O'Keefe & Speakman, 1987; Muller et al., 1987a; Markus et al., 1995). However, such a broad definition just highlights the

problem; remapping, as it is currently defined is likely to include a range of different phenomena, each mediated by distinct biological processes (Fig 1.6).

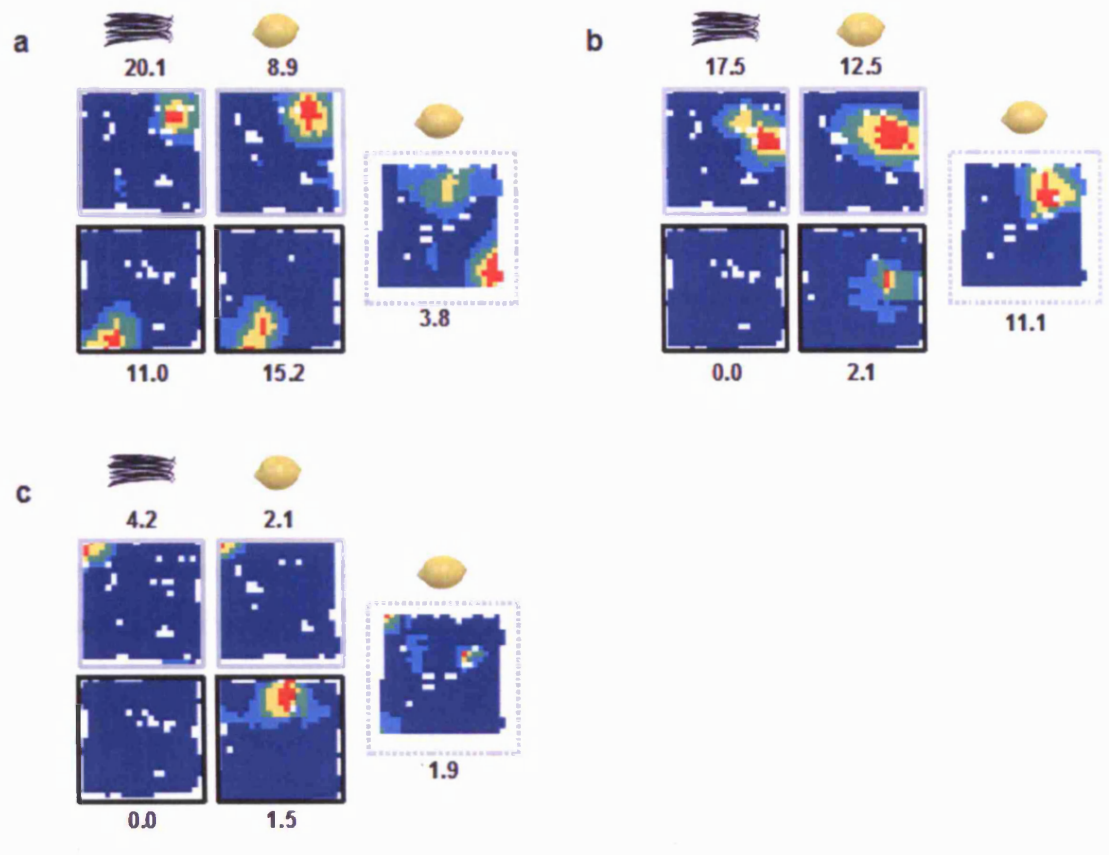

Fig 1.6. Remapping in CA1 place cells induced by olfactory and visual changes, a complex interplay of several different forms of remapping is evident. Data recorded by the author. Three simultaneously recorded cells are shown (a-c). Trials were run in a 70cm enclosure that allowed independent manipulation of colour and odour. Four configurations were possible: white-vanilla (top left), white-lemon (top right), black-vanilla (bottom left), black-lemon (bottom right). A further trial was run with all the lab lights turned out; thus the rat was introduced into the enclosure (configured as white-lemon) without being able to perceive its colour. All cells show clear remapping between the black and white environments; two distinct representations were apprent after the animal's first exposure to the enclosures, as such this seems to be an example of complex or complete remapping. Two of the cells (b&c) also show a configural response; they respond differently in the black-vanilla and black-lemon enclosures (Anderson & Jeffery, 2003). This effect developed slowly over the course of several recording session and might be thought of as slow rate remapping. Finally in the dark at least two of the cells (a&b) exhbit place fields different to any of those seen in the light trials. Again this effect developed slowly, early trials were always the same as the black-lemon trial or white-lemon trial. Hence this might be thought of as slow position remapping.

Historically, the concept of remapping has been shaped by Marr's (1971) ideas about pattern completion and separation, and O'Keefe and Nadel's (1978) restatement of Tolman's cognitive map (1948). As such, there seems to have been a tendency to see remapping purely as the process by which a unique hippocampal code is generated

for a new environment (Muller, 1996). This view was supported by work which demonstrated a complete reorganisation of place cell activity after changes were made to an animal's enclosure (Bostock, Muller, & Kubie, 1991). Also by functional considerations which saw the DG as being well placed to orthogonalise activity in the recurrent CA3 network (Marr, 1971; Rolls, 1989). At the time, it was recognised that remapping was not purely dictated by sensory input but also reflected the animal's internal state. In essence, the system had to 'decide' when it was in a novel environment and hence when to remap. This latter point was evident because in response to the same changes, remapping occurred at different times in different animals, and did not necessarily occur at all (Bostock et al., 1991; Muller et al., 1987a). In short, it made sense. When an animal entered a new space, a novel hippocampal code would be generated. Similarly, if a familiar environment was changed sufficiently it would not be recognised and remapping would occur. To an extent, this view is correct, however it fails to capture a group of related effects that have variously been called 'local remapping' (Knierim, 2003; Rivard, Li, Lenck-Santini, Poucet, & Muller, 2004; Lenck-Santini, Rivard, Muller, & Poucet, 2005), 'partial remapping' (Skaggs & McNaughton, 1998; Knierim et al., 1998; Yoganarasimha, Yu, & Knierim, 2006), 'slow remapping' (Lever, Wills, Cacucci, Burgess, & O'Keefe, 2002b; Barry et al., 2006), and 'rate remapping' (Hayman, Chakraborty, Anderson, & Jeffery, 2003; Leutgeb et al., 2004b).

So what can be done to distinguish these different forms of remapping? Fortunately, the recent discovery of grid cells (Hafting et al., 2005) upstream of the hippocampus (see 1.4.6) offers some hope of a solution and casts light on previous results. Work from Fyhn *et al.* (2007) indicates that 'complex remapping' (changes in field position and rate) in CA3 is accompanied by a coherent shift in the entire grid ensemble ('grid realignment'). In contrast, 'rate remapping' (changes in rate but not field position) is not accompanied by such a shift. Although data from CA1 were not reported, previous work from the same lab suggests that the activity of CA1 place cells is modulated to a greater extent by local cues (Leutgeb et al., 2004b). As such, a shift from complex remapping to rate remapping would not be so apparent in CA1. This final point is significant as most studies of remapping have been conducted with CA1 place cells. Indeed, it now seems that to an extent, the CA1 and CA3 representations might emerge independently of one another (Leutgeb, Leutgeb, Moser, & Moser,

2006) and be maintained by independent connectivity (Brun et al., 2002). Nevertheless, it is tempting to speculate that, at the level of CA1, some difference should be evident in the changes evoked by these two modes of remapping. A possible candidate is the rate at which the representations develop. In effect, there seems to be a divergence in the literature, some authors report that remapping of CA1 place cells appears rapidly and is fairly complete (e.g. (Bostock et al., 1991; Wills, Lever, Cacucci, Burgess, & O'Keefe, 2005); others find slower, experience dependent changes that can lead to intermediate, partially remapped states (e.g. (Lever et al., 2002b; Skaggs et al., 1998). Might it be that the former examples were marked by realignment of entorhinal grids while the latter arise from stepwise plastic changes, possibly occurring within the hippocampus? Clearly further experiments would be required to support such a view. Still a few circumspect results do suggest that quite different mechanisms drive fast-complete remapping as oppose to slow-partial remapping. For example, Wills *et al.* (2005) and Leutgeb *et al.* (2004a) recently published conflicting reports of how CA1 place cells respond to environments with different degrees of geometric similarity. Wills observed an abrupt transition, such that some environments were treated as circular and others were treated as square. Conversely, Leutgeb saw graded remapping that gradually transformed between the two extremes. The principal difference seems to be that Wills and colleagues initially obtained rapid and largely complete remapping by 'priming' their rats with environments that differed in colour and odour, as well as geometry. Hence, although Wills' result was portrayed as been driven by the recurrent dynamics of CA3, it is also consistent with a sudden realignment of grids. As such it might have been Leutgeb's failure to establish grid realignment at the outset, that led to them observing a different form of remapping in the probe enclosures.

What of the more gradual and partial changes in place cell activity reported by many authors; do these represent disparate outcomes produced in specific circumstances or are they mediated by a single process? For example, each of Skaggs and McNaughton (1998), Anderson and Jeffery (2003), Knierim (2002), as well as Lever *et al.* (2002b) have observed forms of partial remapping. However, these were variously driven by discrepancies in path integration, combinations of olfactory and visual cues, conflicting rotation of local and distal cues, and in the latter case geometric differences between enclosures. Although it is unclear how these effects relate to one another it

does seem that Lever's (2002b) slowly accumulating changes are a very different effect to the rapid and complete remapping witnessed by Bostock *et al.* (1991), Wills *et al.* (2005) and Fyhn *et al* (2007). Interestingly, a similar incremental modification of place fields was reported by Barry *et al.* (2006); in this case the change seems to have been driven by the presence of an additional barrier placed into the rat's enclosure. On the basis of these results, Barry *et al.* proposed a simple model of place cell firing which saw the changes in place fields arising as a result of plasticity at the level of inputs to place cells. As such, it is tempting to imagine that some of the other examples of partial remapping seen in the literature result from a similar process. That is, partial remapping might arise from the slow divergence of multiple representations driven by stepwise plastic changes upstream of place cells.

Clearly we are still some distance from a clear conception of remapping and there are many outstanding questions to be answered. For example, the behavioural relevance of the phenomenon remains unclear; with some authors reporting that goal directed navigation is abolished after remapping (Barnes, Suster, Shen, & McNaughton, 1997) and others claiming that it is preserved (Jeffery, Gilbert, Burton, & Strudwick, 2003). Conversely, Markus *et al.* (1995) showed that behaviour can apparently promote remapping. A further crucial problem seems to be that the functional mechanisms that mediate remapping are not understood. We do know that NMDA receptors (Kentros et al., 1998) and protein synthesis are necessary for a remapped representation to stabilise. However, the locus of the receptors and protein synthesis machinery targeted in these studies is unknown and as such it is difficult to interpret the results (Agnihotri, Hawkins, Kandel, & Kentros, 2004) (though see (Nakazawa et al., 2002). Based on our current understanding a tentative conclusion might be that 'remapping' actually represents several different processes that can occur independently or in tandem. At least one of these processes being a form of rapid pattern separation while one, or more of the others, is characterised by a slow discrimination between partially conflicting representations.

1.4.4 Models of place cell firing

Provoked by the discovery of place cells (O'Keefe et al., 1971; O'Keefe, 1976) and LTP (Bliss & Lomo, 1973), also the suggestive architecture of the hippocampus (Burwell et al., 1995) and its association with amnesia (Scoville & Milner, 1957b) and spatial

memory (Morris, Garrud, Rawlins, & O'Keefe, 1982a), many authors have proposed computational models of hippocampal function (e.g. (Marr, 1971; Hopfield, 1982; Zipser, 1985; Zipser, 1986; Samsonovich & McNaughton, 1997; Kali & Dayan, 2000; Hartley, Burgess, Lever, Cacucci, & O'Keefe, 2000; Burgess & Hartley, 2002). The text below provides a brief review of the models that directly influenced the development of Hartley *et al.*'s (2000) boundary vector cell (BVC) model. Hence, we primarily treat models that account for the representation of spatial location by place cells (e.g. (Zipser, 1985; Sharp, 1991). See (Trullier, Wiener, Berthoz, & Meyer, 1997; Redish, 1999; Burgess, 2007) for reviews covering models of spatial navigation (e.g. (Burgess & O'Keefe, 1996c; Zipser, 1986; Muller, Stead, & Pach, 1996) and of content addressable associative memory system (e.g. (Marr, 1971; Hopfield, 1982).

Before turning to the specifics of certain models it seems appropriate to highlight the general role that computational models occupy with respect to experimental neuroscience. As suggested by Burgess (2007), models couched in the explicit language of maths provide an unambiguous and precise definition of a theory or hypothesis. In this sense, computational models leave less space for misunderstanding and reinterpretation. More specifically though, a model accompanied by simulations provides a greater certainty that an idea is actually practicable than does a simple verbal explanation. Furthermore, a well constructed and relevant model can also suggest important experiments and manipulations to resolve outstanding theoretical questions.

The earliest computational model of place cell firing was developed by Zipser (1985). In his basic feed forwards model, Zipser described place fields as the thresholded sum of a set of landmark detectors. In turn, each landmark detector responded maximally when the retinal area subtended by specific distal landmarks matched a pre-stored state. Depending on the number of landmarks available and on the number of landmark detectors employed, place fields produced by the model produced a good match to experimental data. Fields were resistant to cue removal and exhibited motility in the face of cue manipulation. Zipser's approach, relating multiple egocentric views to a single allocentric position, is typical of a local view model and has been used widely. However, like other models that describe place cell firing

purely in terms of visual input it makes no predictions about the maintenance of place fields in the dark (Quirk et al., 1990).

Sharp's 1991 paper is interesting as it is one of the first to present a simulation incorporating learning and the notion of egocentric bearing to environmental landmarks. Sharp described a three layer network such that sensory inputs drove putative entorhinal cells that fed forwards to drive place cells. The sensory inputs took two forms, one type being sensitive to the distance of a cue from the rat and the other being jointly sensitive to distance and egocentric angle to the cue. Being a local view model, rats were only assumed to perceive cues within their current field of view. At each level of the network, simple competitive learning mediated by basic Hebbian plasticity followed by weight normalisation, served to bind specific 'views' to the firing of specific place cells. The model learnt to produce well defined place fields that were omni directional in the open field but which remained directional if exploration was confined to certain orientations. Crucially though, the model saw all place fields as initially being directional, becoming omni directional only after exploration; experimental data seems to suggest that the opposite effect occurs (Frank, Stanley, & Brown, 2004). At about the same time O'Keefe's (1991) 'centroid' model employed a similar approach to sensory input but did not utilise learning; the position of multiple landmarks were encoded in terms of egocentric position and distance from the rat.

Based on their observation that the position and size of place fields can be controlled by manipulation of environmental boundaries, O'Keefe and Burgess (1996) proposed a rudimentary model of place cell activity. The model suggested that place fields might be described as the summed threshold of two or more Gaussian tuning curves, peaked a specific distance from certain environmental boundaries. Three key points here are: that the Gaussian's were defined in allocentric coordinates; that they responded to objects that were treated as continuous surfaces; and finally that their width was proportional to their tuning distance (i.e. distance between the Gaussian's peak and the wall it responded to). Note that the model implicitly assumes that sensory information is transformed from egocentric to allocentric coordinates prior to reaching the hippocampus, for a similar approach in which the transform is explicit see (McNaughton, Knierim, & Wilson, 1994). Subsequently, Burgess *et al.* (2000)

presented numerical simulations showing that the model produced place fields consistent with those observed by O'Keefe and Burgess. Hartley *et al.*'s (2000) more complete formulation of the BVC model (see 1.4.5) essentially builds on the foundations of that work.

A derivation of the Hartley *et al.* model was presented by Burgess and Hartley (2002) to tackle Fenton *et al.*'s (2000a) observation that inconsistent rotation of cue cards produced a inhomogeneous parametric variations in place field position. The authors assumed that with experience, boundary vector cells would come to distinguish between different environmental boundaries. Thus, a BVC might respond only to a specific cue card or to the grey enclosure wall. No mechanism was proposed to account for this effect. A further assumption required that inconsistent rotation of the two cue cards introduced a local deflection into the head direction system, such that some directions were overrepresented relative to others. Barry *et al.* (2006) also built on the basic BVC model, describing how synaptic plasticity mediate by the BCM-learning rule (Bienenstock et al., 1982; Fuhs & Touretzky, 1999) could account for slow plastic changes seen in place fields.

Though not strictly related to the sensory-bound BVC model, a second set of models employing path integration are worthy of note because of their relevance to models of grid cell firing (Hafting et al., 2005). The first model to jointly incorporate the notion of path integration and sensory control of place cells was published by Wan *et al.*(1994); each place cell being driven by egocentric, allocentric and path integrative inputs. However, Samsonovich and McNaughton's (1997) subsequent 'chart' model of path integration has proved to be more influential, by simulating an actual mechanism for path integration. Samsonovich saw path integration as being conducted by a continuous attractor manifest in the recurrent CA3 network; an approach not dissimilar to Zhang's (1996) linear attractor model of the head direction system (see also (Skaggs, Knierim, Kudrimoti, & McNaughton, 1995). To function, the model has two main requirements: that the strength of connections between CA3 place cells decreases proportionate to the distance between their place fields; and that there exists a population of appropriately connected 'shifter cells' with spatially constrained firing that is also modulated by speed and heading (e.g. similar to Cacucci *et al.*'s (2004) TPD cells). The former of these requirements, in conjunction with appropriate

feedback inhibition, can lead to the creation of a stable 'bump' of activity in CA3 neurons. The final requirement is necessary if that 'bump' is to move and track the rat's position. Hence, when an animal enters an environment for the first time a novel combination of CA3 cells is activated, the 'active chart'. As the animal explores, place cell activity is initially driven by path integrative processes but through learning, available sensory cues are bound to the active chart such that it can be reactivated on subsequent visits to the same location. The massively recurrent CA3 network has also been applied to place cell models by Kali and Dayan (2000) as well as Brunel and Trullier (1998), though in both cases the emphasis on CA3 was of binding together multiple local views.

1.4.5 The Boundary Vector Cell model

The BVC model (Hartley et al., 2000) posits the existence of 'boundary vector cells' to explain the effects of geometrical manipulations of an environment on the firing of place cells. As a result, the model is able to make explicit, testable predictions of firing patterns for specific cells in a novel environment (Hartley et al., 2000). The same framework also predicts the search location of human subjects when remembering a location in a virtual reality environment subject to geometric manipulations (Hartley, Trinkler, & Burgess, 2004). Below I provide an outline of the BVC model as it was described by Hartley *et al.*(2000), this framework is developed further in chapter three.

The model describes place cell firing as a continuous function of the relative location of the barriers in and around the animal's environment. Place cell activity is driven by feed-forward connections from putative BVC inputs whose firing is determined by the presence of extended barriers (e.g. walls, large objects, an impassable drop). More specifically, each BVC fires optimally when a barrier is encountered at a defined distance and allocentric direction from the rat (Fig 1.7A), that is, each BVC has a Gaussian tuned response to the presence of a barrier, peaked at a preferred distance and direction. The firing of BVC i (tuned to preferred distance d_i and angle ϕ_i) to a boundary at distance r and direction θ, subtending an angle $\delta\theta$ at the rat, is given by:

$$\delta f_i = g_i(r,\theta)\delta\theta, \qquad (1)$$

where: $g_i(r,\theta) \propto \frac{\exp[-(r-d_i)^2/2\sigma_{rad}^2(d_i)]}{\sqrt{2\pi\sigma_{rad}^2(d_i)}} \times \frac{\exp[-(\theta-\phi_i)^2/2\sigma_{ang}^2]}{\sqrt{2\pi\sigma_{ang}^2}}$, (2)

the angular width σ_{ang} is constant, but radial tuning width increases linearly with the preferred tuning distance: $\sigma_{rad}(d_i) = d_i + c$. For each location $\underline{x}$ in the environment, the contribution of all boundaries to the firing of BVC i is determined by integrating equation 1 over θ. The firing of place cell j at that location $F_j(\underline{x})$ is then proportional to the thresholded linear sum of the N BVCs that happen to be connected to it (Fig 1.7B):

$$F_j(\underline{x}) = A\Theta\left(\sum_{i=1,N} f_i(\underline{x}) - T\right), \quad (3)$$

where the threshold T and coefficient A are constants, and Θ is the Heaviside function (i.e. $\Theta(x)=x$ if $x>0$; $\Theta(x)=0$ otherwise). Firing does not depend upon the rat's heading and the direction tuning of all BVCs is determined relative to the same allocentric reference frame (assumed to be provided by the head-direction system). For example, a place cell that received input from BVCs tuned to respond to nearby barriers to the North and to the East would fire when the 'rat' was in the Northeast corner of its environment. More generally, the sharper tuning of BVCs with short preferred directions implies that boundaries near to the location of peak firing will tend to have more influence than boundaries further away.

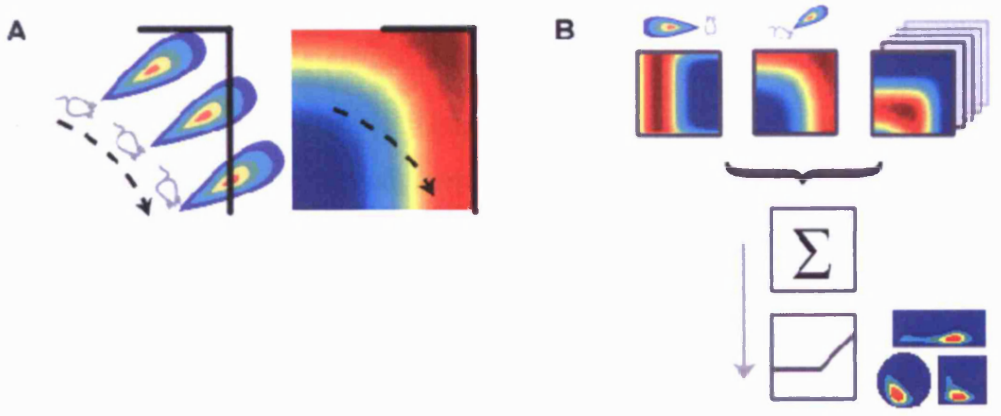

Fig 1.7. Basic properties of Hartley *et al.*'s (2000) BVC model. **a**) Each BVC fires optimally when a boundary is encountered at a certain allocentric direction and distance from the animal. In this example the rat approaches that eastern edge of its enclosure causing the BVC to fire more strongly as the wall moves into its receptive field. **b**) The model describes place fields as the summed and thresholded firing of several BVCs.

Simulated firing produced by the BVC model has similar properties to that obtained experimentally. In response to topographic changes of the environment, simulated cells responded like those observed by O'Keefe and Burgess (1996). The majority of cells maintain their position relative to nearby walls, while some stretch and become bimodal and a smaller number turn on or off, in the appropriate proportions. It also explains consistencies across more radical changes in shape (Hartley et al., 2000; Lever et al., 2002b). Furthermore, by determining the subset of BVCs which best fit a set of firing data from a given cell, it is possible to predict how that cell will respond to a novel change in the environment. The likely BVC nature of the input to place cells is also beginning to be recognised in other modelling studies (Kali et al., 2000; Touretzky et al., 2005).

1.4.6 Grid cells

Grid cells are relative newcomers to electrophysiology, alluded to by Fyhn *et al.* (2004) and Hargreaves *et al.* (2005), they were described formally by Hafting *et al.* (2005) in recordings made from layer II of the dorso-lateral medial entorhinal cortex (dlMEC). Like place cells, they show stable spatially constrained firing but with the peculiarity that each cell has multiple firing fields positioned in a grid defined by the vertices of tessellated, equilateral triangles (Fig 1.8A). Although the firing of layer II grid cells is agnostic of head direction, subsequent work showed that some grids from deeper

layers (III, V, and VI) do exhibit directional modulation (Sargolini et al., 2006). Interestingly, the directional tuning of these cells appears to be broader than that of presubicular head direction cells; seeming to approach a rectified cosine function.

In most cases grids are remarkably regular, to the extent that they can be accurately described in terms of three variables: orientation, the angle of the grid relative to an arbitrary axis; spacing, the distance between adjacent grid peaks; and offset, the position of the grid in two dimensional space (Fig 1.8B also see methods in section 4.2.3). It is striking that grids from the same animal, or more accurately the same hemisphere, apparently share the same orientation (Sargolini et al., 2006). Spacing, however, is topographically organised such that grids recorded from dorsal positions have a finer scale than those found more ventrally; Hafting *et al.* (2005) reported lengths from approximately 30cm to 50cm for recordings made over a range of 800μm. In contrast, offset is apparently randomly distributed, even for cells recorded from the same tetrode. This final point is significant as it implies that a relatively small population of grids from the same dorso-ventral position will effectively tile the environment. Note, under the assumption that peaks are indistinguishable, a population of same scale grids has limited ability to uniquely encode an animal's position. Consideration of additional grid sets from different dorso-ventral positions increasingly mitigates this problem (O'Keefe & Burgess, 2005; Blair, Welday, & Zhang, 2007).

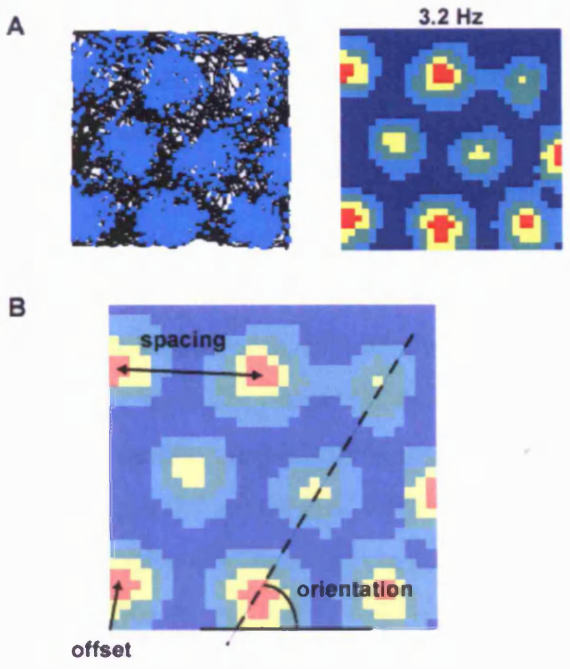

Fig 1.8. Example of the spatial firing of an entorhinal grid cell. **a**) Left, firing recorded from a 1m by 1m enclosure, 20 minutes worth of data shown. Rat's path is depicted by the continuous black line, action potentials (light blue) are superimposed. Right, rate map constructed from raw data to left. Peak rate is shown above the map, colour code as before. **b**) Three descriptive measures can be extracted from a regular grid: orientation, how the entire grid is rotated relative to an arbitrary horizontal axis; offset, displacement of the grid relative to an arbitrary point; and spacing, the distance between adjacent grid peaks.

Grids, like place cells, have stable firing correlates and are positioned with reference to environmental cues. For example, in a circular environment grid orientation is controlled by a single polarising visual cue. Grid firing also persists in the dark, likely being maintained by a combination of path integration and non-visual cues. Interestingly, grid scale is apparently unaffected by enclosure size and would seem to be a product of the system rather than the environment (Hafting et al., 2005). However, a recent paper from the Mosers' lab reported that in some cases grids exposed to a novel environment showed a small but significant increase in scale (Fyhn, Hafting, Treves, Moser, & Moser, 2007). It is not clear why or how this increase is induced or if it would persist with continued exposure to the novel environment. Importantly, the same paper demonstrated that environmental changes sufficient to instigate global remapping in CA3 were invariably accompanied by a coherent shift in the entire grid population.

So what is the purpose of these remarkable cells and how might the grids be produced? It seems likely the answer to both these questions is 'path integration'. A number of authors were quick to point out that the grids' regular periodicity, rapid appearance in novel enclosures, and independence of visual cues all support the view that they are generated by a system that uses self-motion cues to calculate position (O'Keefe et al., 2005; McNaughton, Battaglia, Jensen, Moser, & Moser, 2006). The subsequent discovery of direction and velocity modulated cells in layers III, V and VI (Sargolini et al., 2006) propelled this notion and suggested a specific mechanism by which self motion cues might be integrated. To date, four models of grid cell firing have been published, three of which explicitly see grids as the product of a path integrative network (Fuhs & Touretzky, 2006; McNaughton et al., 2006; Burgess, Barry, & O'Keefe, 2007). The fourth model is slightly unusual in that it assumes that mini-hexagonal grids are already present in the firing of theta cells and that the firing of grid cells results from interference between mini-grids (Blair et al., 2007); as such this model sidesteps the issue of how any grid is produced in the first place. Interestingly, two of these models, Burgess *et al.*'s and McNaughton *et al.*'s, can be seen as updated versions of models of place cell firing; specifically O'Keefe and Recce's (1993) oscillator model and Samsonovich and McNaughton's (1997) 'chart' model. Broadly speaking, these models including Blair's mini-grids, fall into two camps; those which are dependent upon recurrent connections (Fuhs et al., 2006; McNaughton et al., 2006) and those which propose that grids results from interference patterns generated between smaller-scale oscillators (Blair et al., 2007; Burgess, Barry, Jeffery, & O'Keefe,). Both approaches enjoy support from experimental data: functional, excitatory recurrent connections have recently been found in layer II of the MEC (Kumar et al., 2007); and it is known that layer II stellate cells exhibit intrinsic sub-threshold membrane oscillations (Alonso & Linas, 1989) the frequency of which varies in a topographic manner compatible with the arrangement of grid scale along the dorso-ventral axis (Giocomo, Zilli, Fransen, & Hasselmo, 2007).

Finally, what is the relationship of grid cells to place cells? Given their position a single synapse upstream of the hippocampus, it seems unlikely that grid cells do not contribute to the activity of place cells. Indeed it has been known for some time that place fields are in part defined by path integrative input (Gothard, Skaggs, & McNaughton, 1996; Gothard, Hoffman, Battaglia, & McNaughton, 2001; O'Keefe &

Burgess, 1996); grid cells seem a likely source of that information. However, are place fields solely the product of grid cell activity? A recent model showed that unitary place fields can be generated by summing and thresholding grids that have a common offset but differ in orientation or spacing (Soldstad, Moser, & Einevoll, 2006). Blair *et al.* (2007) went further and demonstrated that with the correct connectivity, grids could be used as a basis set to decompose complex shapes. Furthermore, Fyhn *et al.*'s paper showing that complex remapping in CA3 is coincident with grid shifts in the MEC also seems to support the view that the activity of place cells is to a great extent dictated by grid cells. Nevertheless, the picture is far from clear. The hippocampus receives input from many other brain regions besides the MEC, not least the LEC. In addition, the hippocampus not only receives input from the MEC but in the case of CA1 projects directly back to it. There are also experimental reasons to doubt the primacy of grid cells; preliminary work from the Mosers' lab suggest that when the hippocampus is temporarily inactivated, grid firing in layers II and III is lost (Bonnevie, Fyhn, Hafting, Moser, & Moser, 2006). Conversely, work from the same lab showed that despite bilateral lesions of the dorsal hippocampus, layer II MEC cells still exhibited stable spatial firing (Fyhn et al., 2004).

Considerations such as these have led O'Keefe and Burgess (2005) to suggest that place cell projections back to the dlMEC are necessary to maintain the alignment of grids with the environment. In essence, they see place cells as the route through which spatial information is linked to the grid's path integrator.

2 General Methods

2.1 Animals

Male Lister Hooded rats (250-400g and 3-6 months of age at implantation) were used in both studies. Prior to surgery animals were housed communally, four to a cage, with free access to water and food. After surgery the rats were transferred to individual Perspex cages (70cm long x 45cm wide x 30cm high) with free access to water and food sufficient to keep them at ~90% their free-feeding weight. All animals were maintained on a 12:12 light-dark cycle with two, one hour half-light periods simulating dawn and dusk. All work was conducted within the terms of appropriate Home Office Project and Personal licences.

2.2 Surgery and electrodes

Anaesthesia was induced and maintained with an isoflurane-oxygen mix (1.5-4l/min). Post-operative analgesia was provided by 0.1ml Vetergesic (buprenorphine - intramuscular injection) and for a further three days in the form of Vetergesic laced jelly (0.5ml buprenorphine twice a day). Prior to surgery animals received 0.1ml Baytril (enrofloxacin - subcutaneous injection) as a microbial prophylaxis and continued to receive Baytril in their water for seven days after surgery (4.0ml enrofloxacin/100ml water). During surgery Vaseline was used to cover the rats' eyes to prevent corneal damage and the animals' bodies were wrapped with bubble wrap to maintain body temperature.

Each animal received a single 'poor-lady' microdrive (Fig 2.1) (Axona Ltd., St. Albans, UK) loaded with four tetrodes (Recce & O'Keefe, 1989) of twisted 17-25μm HM-L coated platinum-iridium wire (90% - 10%) (California Fine Wire, USA). The 'poor-lady' design allows tetrodes to be advanced through the brain in steps as small as 25μm, though tetrodes cannot be moved independently of one another. Prior to surgery tetrodes were cut level with one another and were sterilised by immersion in ethanol. The animal's head was fixed in a stereotaxic frame with lambda and bregma in the horizontal plane, and the skull was revealed. For implants to CA1, a single 1mm trephine hole was made 3.8mm posterior and 2.5mm lateral to bregma, over the right hemisphere. The dura was retracted and the tetrodes were implanted, perpendicular to the ground, at a depth of 1.3 - 1.5mm relative to the intact dura. For entorhinal implants one or two 1mm trephine holes were made to locate the transverse sinus. Electrodes were implanted above the right dorso-lateral medial entorhinal cortex (MEC), 4.5mm lateral to the midline, 0.2 - 0.5mm anterior to the sinus, angled forwards in the sagittal plane at 8 - 10°, and to a depth of 1.0 - 2.5mm. Microdrives were secured to the skull using seven stainless steel bone screws and dental acrylic mixed 3:1 with Aureomycin (chlortetracycline hydrochloride) powder. The most anterior screw on the right side was used as a ground electrode.

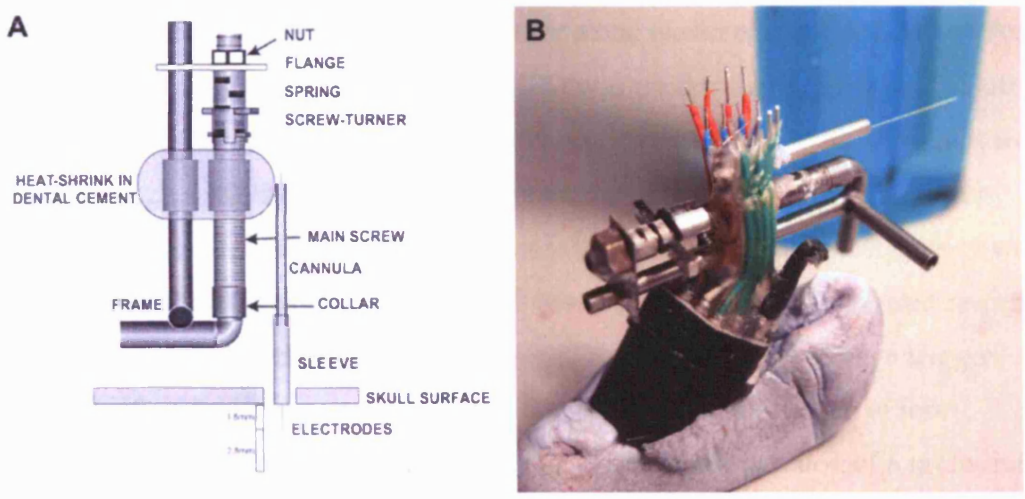

Fig 2.1. Design of the 'poor-lady' microdrive. During surgery the frame is attached to the animal's skull with bone screws and dental acrylic. A full turn of the screw advances the cannual and electrodes by 200μm.

2.3 Single unit recording

Screening was performed post-surgically after a one week recovery period. An Axona recording system (Axona Ltd., St. Albans, UK) was used to acquire the single unit and positional data. The local field potentials recorded from each of the 16 channels were passed through a RC-coupled, unity-gain operational amplifier mounted on the animal's head and fed to the recording system using lightweight wires. Each channel was amplified 20,000 to 40,000 times, bandpass filtered (500Hz - 7kHz) and recorded differentially against a channel on a separate tetrode. Spikes exceeding a trigger threshold (50 - 90mV) were sampled at 48kHz (50 samples from each of four channels) and time stamped with a 96kHz clock signal. The position of a rat in the apparatus was captured using an overhead video camera to record the position of the one or two LEDs on the animal's head-stage. The image was digitized and sampled at a rate of 50Hz to identify the rat's position. The animal's head direction was extracted using the relative position of the two LEDs, one large, one small, positioned 8cm apart at a known angle to the rat's head. In the case of the single LED head-stage, head direction was inferred from the animal's trajectory.

Prior to recording, tetrodes were advanced in 50 - 200μm steps. In the case of the hippocampal rats, electrodes were advanced until ripples, high frequency (~200Hz) oscillations, were seen on the EEG. At this point electrodes were moved more cautiously until hippocampal complex spike cells were detected on the oscilloscope. Electrophysiological markers for the MEC are less distinct and electrodes were advanced slowly (steps <=100μm) until multiple large-amplitude units were obtained. This process took place in the experimental room and animals remained in an elevated holding area. Animals were returned to their home cage for at least four hours between screening sessions.

2.4 Spike sorting & binning

Spike sorting was performed offline using a data analysis suite (Tint, Axona Ltd., St. Albans, UK). Action potentials were assigned to putative cells based on differential amplitude, waveform and temporal autocorrelation, criteria applied elsewhere to entorhinal grid cells (Hafting et al., 2005) and hippocampal place cells (Hayman et al., 2003) (Fig 2.2). The animal's position and concomitant spikes were binned into a 64 x 64 bin array covering the camera's field of view; each bin being the equivalent of 8 x 8 pixels, roughly 2 x 2cm. Unsmoothed rate maps were calculated by dividing the number of spikes assigned to a bin by the cumulative occupancy of the bin. Smoothed rate maps were constructed as follows, the firing rate for bin i was the number of spikes in a 5 x 5 kernel centred on i divided by the cumulative occupancy of the same bins.

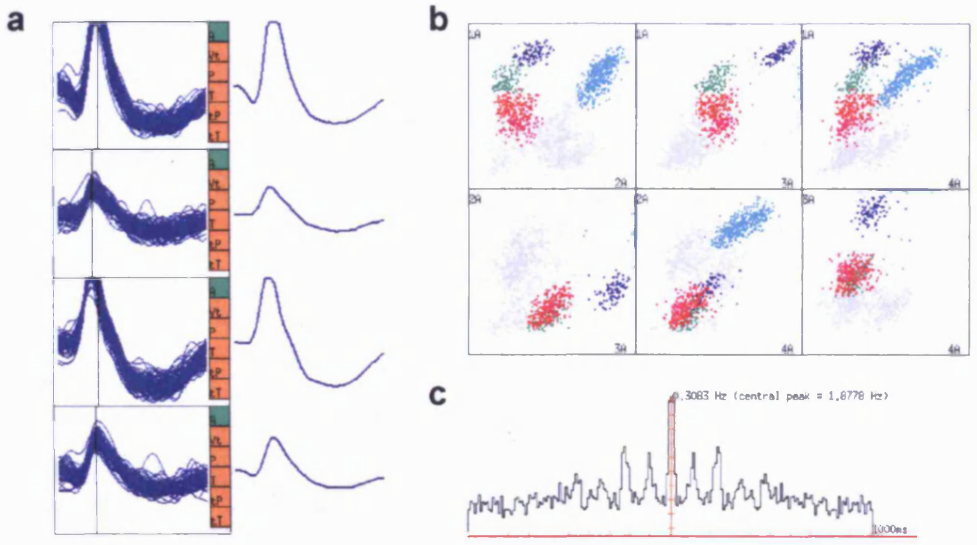

Fig 2.2. Example screen shots from Tint spike sorting software. a) Superimposed waveforms (left) corresponding to multiple action potentials captured on four spatially separated electrodes. Mean waveform for each channel is shown on the right. b) Peak to trough amplitude of each spike on each channel is plotted against its amplitude on each of the other channels. The dark blue 'cluster' corresponds to the waveforms in A. c) Temporal autocorrelation corresponding to the light blue cluster in B. The autocorrelation is calculated for each pair of spikes and indicates the probability of encountering spikes with a given lag, a range of -1000 to 1000ms is shown. Theta modulation of approximately 8Hz is clearly visible.

2.5 Histology

At the end of the experiment rats received an overdose of Euthatal (Sodium pentobarbital) and were transcardially perfused first with phosphate buffered saline and then with 4% paraformaldehyde (PFA) solution. The brains were removed and stored in 4% PFA for at least one week prior to sectioning. 40μm frozen sections were cut (sagittal sections for entorhinal implants, coronal for hippocampal), mounted on gelatine-coated glass slides and stained with cresyl violet. High resolution images were acquired using an Olympus microscope, Xli digital camera (XL Imaging Ltd.) and Xli-Cap image capture software. Individual images were imported into Photoshop CS2 for PC (Adobe Systems) and assembled into composites. The depth and layer at which cells were acquired was extrapolated by reference to the record of tetrode movements after taking account of brain shrinkage.

3 Learning in a Geometric Model of Place Cell Firing[1]

3.1 Introduction

Pyramidal cells in the rodent hippocampus exhibit spatially localised firing (O'Keefe et al., 1971). 'Place cells' (PCs) fire whenever the animal enters a specific area of its environment, the 'place field', and do so independently of heading (Muller, Bostock, Taube, & Kubie, 1994). Location-specific firing appears rapidly after an animal enters a new environment (Wilson et al., 1993; Hill, 1978) and is reinstated on subsequent visits, even months later (Best et al., 1984). Subsequently discovered in other animals (Bingman et al., 2006; Ulanovsky et al., 2007; Ono et al., 1991; Rolls et al., 1989), including humans (Ekstrom *et al.*, 2003), a small population of place cells is sufficient to accurately encode an animal's position (Wilson et al., 1993).

How is place cell firing controlled? Early experiments showed that firing was not the product of a single sensory modality (Best & Thompson, 1989a) and was robust to removal of subsets of cues (O'Keefe J et al., 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 et al., 1996) or changing the relative position of cues (Fenton, Csizmadia, & Muller, 2000a), produced parametric changes in place fields. Further, the representation of identical environments connected by a corridor is initially similar (Skaggs et al., 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 stabilise, and several experimenters have observed the accumulation of small plastic changes in place cell 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, Recce, & O'Keefe, 1994; O'Keefe et al., 1996; Hartley et al., 2000). In particular, Hartley *et al.*'s (2000) model posits the existence of boundary vector cells (BVCs), a class of

[1] The model and simulations described in this chapter appear in Barry & Burgess (2007). Place cell recordings conducted by the author are reported in Barry *et al.* (2006).

cells that respond to the presence of any barrier, such as a wall or a drop, at a given distance and allocentric direction. The thresholded sum of a population of these cells provides a good analogue of place cell 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, Burgess, Cacucci, Hartley, & O'Keefe, 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., 1987b); and specific predictions as to how individual fields would change when the shape of an enclosure was changed (Hartley *et al.*, 2000). The BVC framework also predicts human search behaviour 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 poison of those objects (Burgess et al., 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 place cell 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 et al., 1999), 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 place cells that encode position relative to specific objects; an area West of a black cue card, for example. We apply this model to data recorded in a square enclosure, with and without the presence of a central barrier(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 barriers into an enclosure. It also provides a mechanistic explanation of Rivard *et al.*'s observation that some place cells encode position in room centred coordinates, while others fire relative to centrally place 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, Csizmadia, & Muller, 2000b; Burgess et al., 2002; Touretzky et al., 2005), it does so at the level of single neurons without making assignments as to the specificity and efficacy of cues that place cells or their inputs respond to.

3.2 Model framework

Following Hartley *et al.* (2000), place cell firing was described as a function of the relative location of barriers around an animal. Specifically, the activity of place cells was considered to be the thresholded sum of inputs received from putative boundary vector cells (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. Place cell firing is then simply a thresholded sum of the firing rates of all the BVCs connected to it.

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 coloured 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, the n BVCs in each set shared the same distance and directional tuning but each responded solely to barriers of a specific type (Fig 3.1a). For example, in a black environment with a white cue card fixed to the wall, each BVC set would consist of two BVCs, one of which would respond to the white cue card, the other to the black environment wall. Note, the cue card is assumed to obscure the wall behind it. Prior to learning, all the connections from the BVCs in a set to place cells have the same weight, but subsequent iterative application of a BCM-like learning rule (Bienenstock *et al.*, 1982) allowed weights to vary (Fig 3.1b). This schema is equivalent to allowing individual BVCs to change the efficacy with which they respond to barriers with distinct sensory properties. The BCM rule was chosen in preference to a basic Hebb rule as it encapsulates both LTP and LTD, conveys stability to the network, and only initiates changes in synapse strength when both the pre and postsynaptic cells are active.

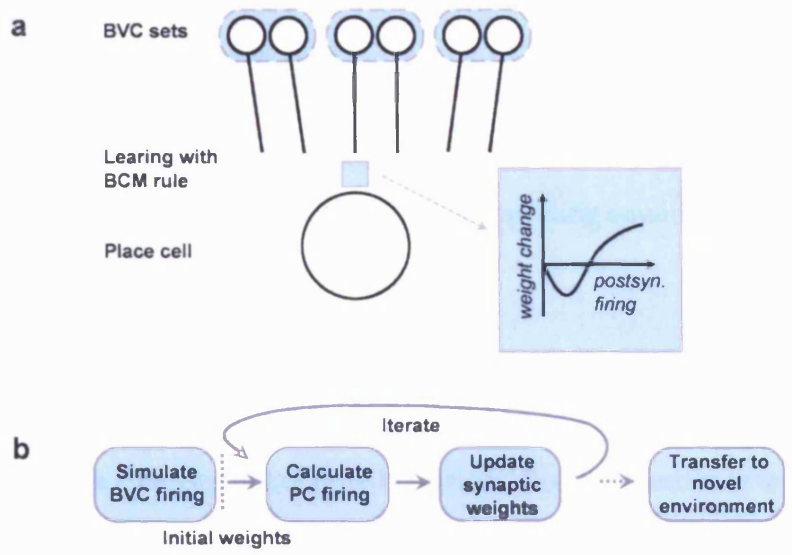

Fig 3.1. a) BVC sets consist of multiple BVCs with the same preferred firing direction and distance. Each BVC in a set responds to boundaries with different sensory qualities (e.g. one might respond to a cue card and another to the enclosure walls). All BVCs in a set initially drive the upstream place cell with the same efficacy; the connections have the same weight. **b**) Through learning with the BCM rule, weights are updated and so a BVC set can come to drive the place cell more strongly in response to certain barriers (e.g. the place cell 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, place cell firing is calculated and weight changes applied, place cell firing is recalculated and so on.

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:

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

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

The angular width σ_{ang} is constant, but radial tuning width increases linearly with the preferred tuning distance: $\sigma_{rad}(d_i) = ((d_i/\beta)+1) \times \sigma_0$. In the simulations presented here σ_{ang}=0.2 radians, β =1830mm, and σ_0=122mm. 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, 741.0mm. 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 up to 2π.

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 equation 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 place cell j at location $\underline{x}$ (referred to as $F_j(\underline{x})$) is then proportional to the thresholded, weighted sum of the N BVCs 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 place cell 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=5000 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., 1990b; Taube et al., 1990a)). For example, a place cell 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 place cells 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 et al., 1999) was used to update the weights of connections between BVCs and place cells (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:

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

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

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

Where $\overline{F}_j$ is the mean firing rate of place cell 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 place cell 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.

3.3 Simulation: experience-dependent changes in place cell firing in a static environment

We start by simulating the effect of experience-dependent plasticity on place fields in an unchanging environment. The firing of 1000 BVCs and 100 place cells were simulated at each node of a 2cm square grid superimposed on a 65x65cm square 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 equation 5. Prior to learning, place cell firing resembled that produced in Hartley *et al.*'s (2000) model; 38 cells fired at a peak rate of 1Hz or above and so were considered to be 'active', place fields were discrete and generally had a single 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 3.2 -first column).

After 100 iterations of learning an additional 14 cells had peak firing at or above 1Hz 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 2.3Hz to 2.8Hz 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 3.2a iv). A similar effect was seen in about 20% of cells with a single field, particular 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 (Fig 3.2a iii). This increased coherence occurs because the BCM rule reduces the weight of BVC inputs if they drive place cell firing outside of the main field. Conversely, a similar number of fields that initially had very low peak firing rates (<1Hz), and so were considered to be 'off', showed a gradual increase in peak firing and field size with learning (Fig 3.2a ii). 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 six to 10 minutes fields become more stable (Wilson et al., 1993). Indeed, learning in the model initially progressed rapidly but was much reduced by the final 10 iterations. Inspection of longer periods of learning (1000 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 3.2b). 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 (equation 6). As the threshold is approached, step size decreases and the cell settles to a steady state.

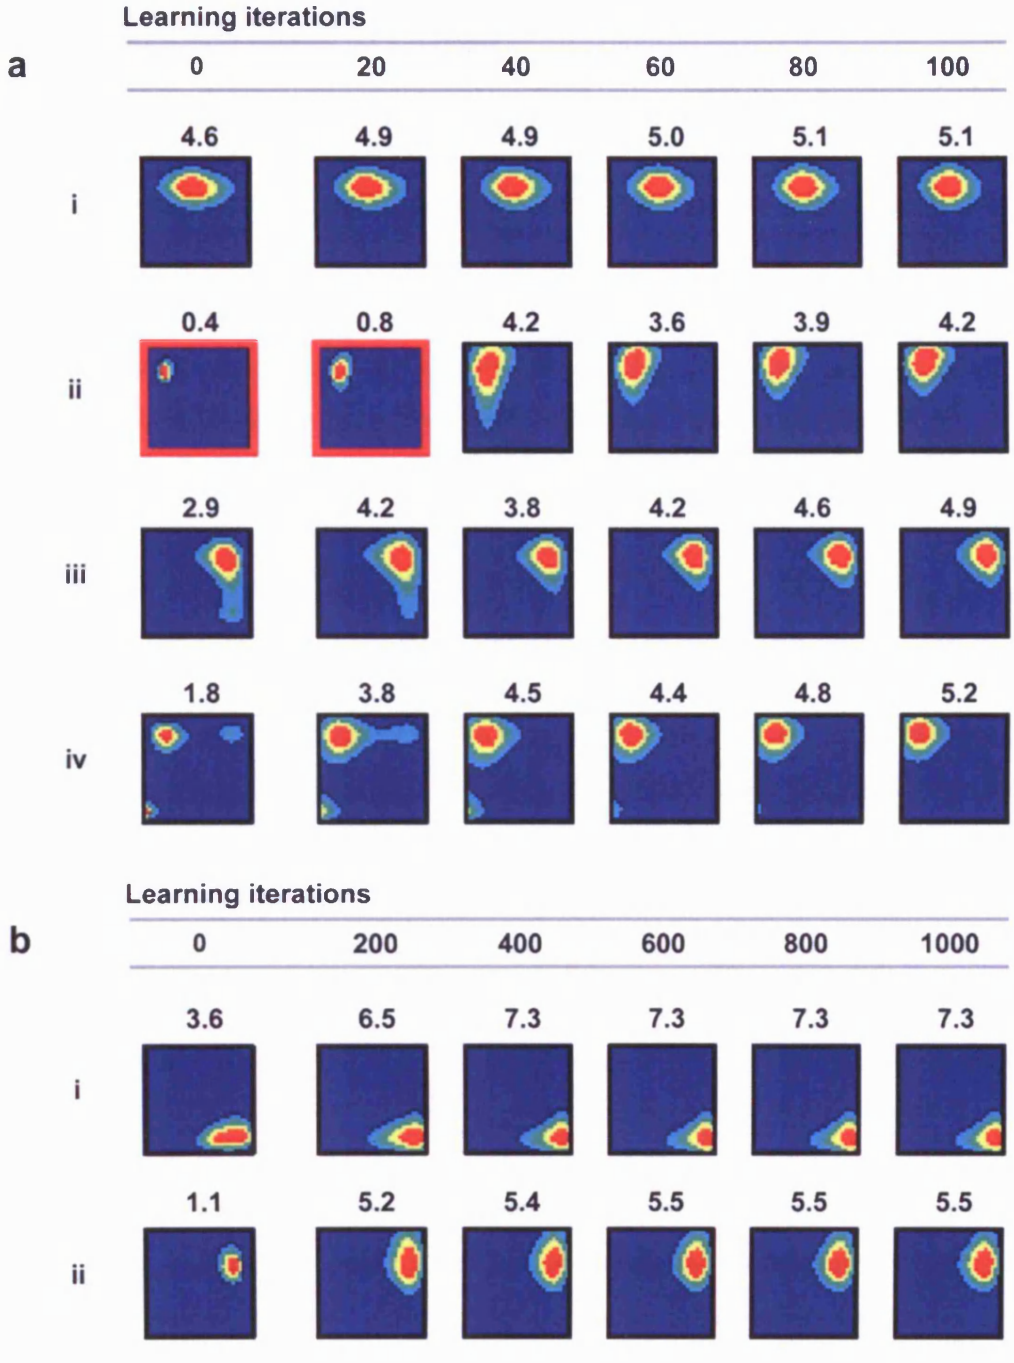

Fig 3.2. Simulated place cell firing showing changes arising during learning in a static environment (65cm x 65cm). Peak firing rate is shown above each plot in Hz. Areas of high firing are indicated by red, low firing by dark blue. Each of the five colours indicate a firing range of 20%, thus dark blue 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 place cells. 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 <1Hz and were considered to be off); A similar number of fields 'tightened up', losing extended areas of low firing (iii); Field 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.

3.4 Simulation & data: experience-dependent changes in place cell firing in response to insertion of a barrier

We now move to the question of how place fields respond to the presence of a barrier inserted into an enclosure. The clearest novel prediction of the basic BVC model (i.e. without learning) is the formation of double fields in response to the addition of an extended barrier into the environment. For instance, a BVC which fires strongly when the rat is a short distance South of the Northern boundary of the environment will also fire strongly when it is a short distance South of a newly-inserted East-West barrier. This implies that some of the place fields near to the North wall of the environment would be predicted to develop a second field in a similar location relative to the barrier. Another prediction is that place fields will tend to have fields oriented parallel to the walls of the environment, reflecting the firing pattern of individual BVCs.

As mentioned earlier, experimental data from several sources concord with these predictions. In terms of field shape, cells recorded in circular environments tend to have crescent fields (Muller et al., 1987b) and elongation of a rectangular environment produces elongated fields that run parallel to the direction of stretch (O'Keefe et al., 1996). Several experimenters have produced field doubling. Skaggs *et al.* (1998) recorded from two visually identical environments connected via a corridor. They observed a large number of cells with similar fields in both environments. Similarly, Lever *et al.* (1999) have observed field doubling in response to barrier insertion.

But what happens after prolonged exposure to a barrier? If, as has been suggested (O'Keefe et al., 1978), place cells really do provide an allocentric representation of space then it is not desirable for two areas in the same environment to be represented by the same cell. With this in mind it might be thought that processes exist to resolve contradictions such as these. Indeed, data from Lever, albeit only two cells, showed that during a 10 minute trial one of a pair of doubled fields created by a barrier faded away (Lever et al., 2002a). To better understand this effect we recorded place cells on consecutive days as an animal repeatedly explored an environment, with and without the presence of a barrier. Experimental data were then compared with simulations derived from the updated BVC model.

3.4.1 Experimental data

The barrier data described below were collected from a single animal[2]. The place fields presented should be considered as qualitative examples of the possible outcomes of experience-dependent plasticity. Data were captured and analysed using standard procedures and equipment described in chapter two. The animal received a single 16 channel microdrive to the right dorsal hippocampus. Post mortem histology subsequently confirmed that the tetrodes had reached the pyramidal cell layer of region CA1. The bulk of the recordings were made on eight consecutive days with additional manipulations performed after the end of the experiment. Where possible cells were followed across days. Recording sessions were conducted at the same time each day, allowing approximately 24 hours between sessions. Each session followed the same protocol, the first two trials saw the rat forage for rice in a 65cm square, walled environment for five minutes each. The environment was cut from dark grey foam board. Laboratory equipment and furniture (e.g. computers) provided directional cues. The animal was returned to its home box between trials and the box was rotated pseudo-randomly by 90^0, 180^0 or 270^0. A barrier, made of the same material as the box, was then positioned in the centre of the environment with one end abutting the North wall. The animal foraged continuously for 40 minutes with the barrier present and was then returned to its home box. The barrier was removed and a final two trials were conducted similarly to the first two.

In total we recorded 10 spatially defined cells (peak firing rate >=1Hz) before and after the addition of the North-South barrier. In response to initial insertion of the barrier three cells exhibited duplicate fields as predicted by the basic BVC model. A further three cells either turned on or moved to have fields adjacent to one side of the barrier. Two cells were unaffected, one turned off, and the final cell moved away from the barrier.

Cell 1, recorded on day 1, exhibited changes in firing similar to those reported by Lever *et al.* (2002a). Insertion of the barrier prompted the rapid appearance of an additional field adjacent to the west side of the barrier (Fig 3.3a, day1). Indeed, the cell fired on the rat's first pass through the new place field. The original field, adjacent to the east wall of the box, was initially more robust and fired more strongly than the

[2] Data collected by the author.

new field but during the 40 minute trial it gradual shrank away while the new field was maintained. The peak rate of the new field was constant at about 14Hz whereas the rate of the original field fell from 23.6Hz in the first 5 minutes to 4.9Hz during the last 5 minutes. When the barrier was removed towards the end of the session, the original field was reinstated and the new field was lost.

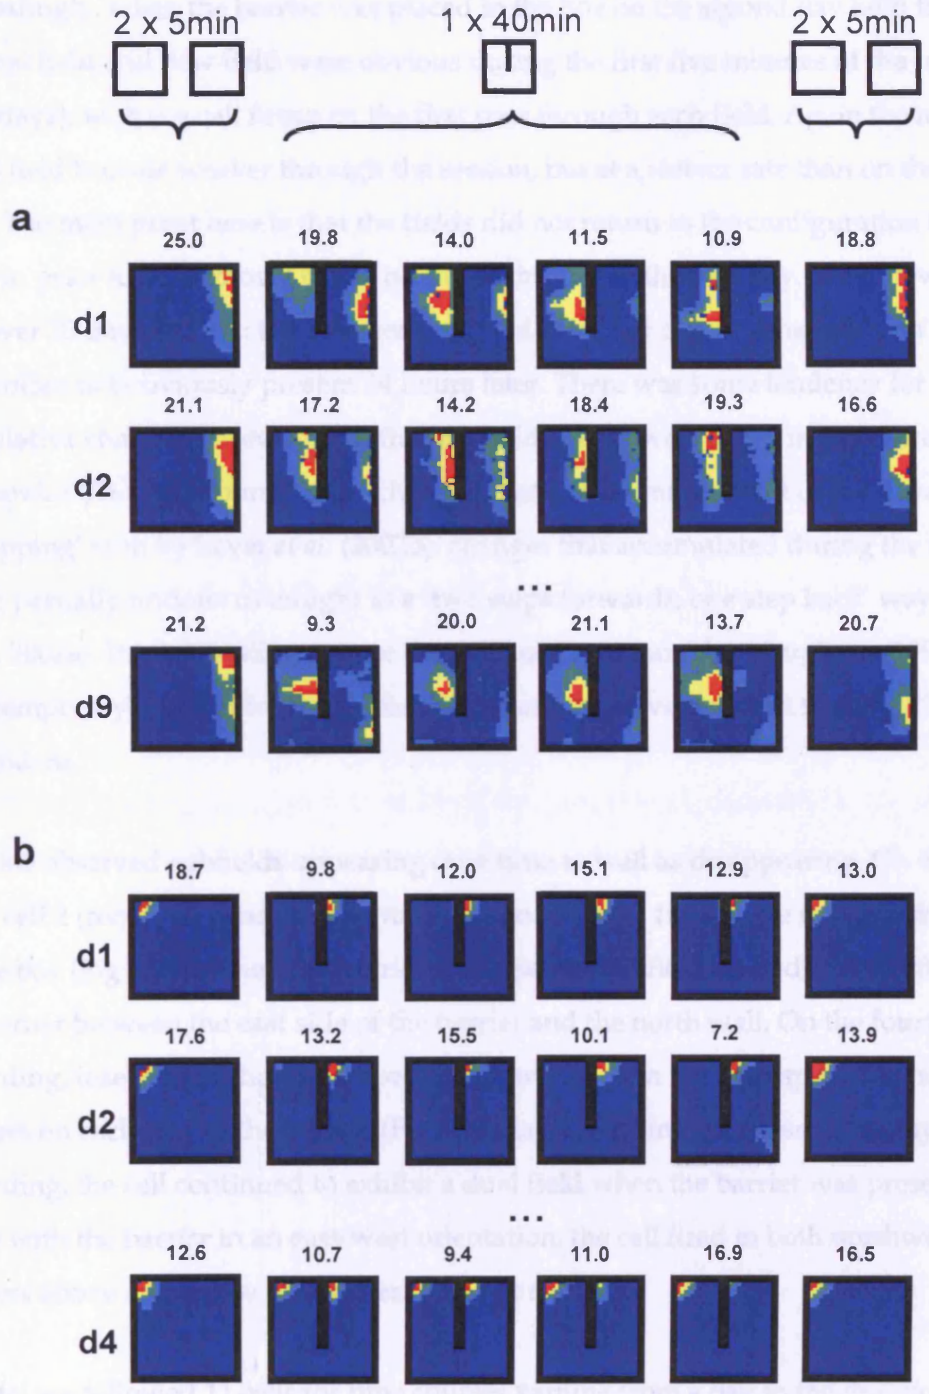

Fig 3.3. Firing rate maps from 2 place cells (a&b) simultaneously recorded over several days (one day per row) in the 65cm square enclosure with or without a 40cm North-South barrier. Each plot represents 10 minutes of exploration. The first column shows combined firing from 2, 5 minute trials recorded without a barrier. The next 4 columns show firing from a single continuous 40 minute trial (as 4 x 10 minute slices) in the presence of the barrier. The final column shows firing with the barrier removed. **a**) Cell 1 doubles its field when the barrier is introduced on day 1, while the eastern field is gradually lost both within each day and over days, see main text. **b**) Cell 2 fired only on one side of the barrier for the first 3 days, then developed a field on both sides on day 4 and continued to do so on subsequent days (not shown).

Interestingly, when the barrier was placed in the box on the second day both the original field and new field were obvious during the first five minutes of the trial (Fig 3.3a, day2), with the cell firing on the first pass through each field. Again the initial (east) field became weaker through the session, but at a slower rate than on the first day. The main point here is that the fields did not return to the configuration they were in prior to the removal of the barrier at the end of the first day. We followed this cell over 10 days in total: the changes observed in one or other of the fields on one day were often not obviously present 24 hours later. There was some tendency for cumulative change however, with the east field slowly weakening in the presence of the barrier. The time course of the changes seen here is reminiscent of the 'slow remapping' seen by Lever *et al.* (2002b): changes that accumulated during the day being partially undone overnight in a 'two steps forwards, one step back' way (Lever et al., 2002a). It is interesting to note that Mehta, Barnes and McNaughton (1997) also saw temporary changes in place field shape, and these were shown to be NMDAR dependent.

We also observed subfields appearing over time as well as disappearing. On the first 3 days cell 2 (recorded concurrently to cell 1 above) had a field in the northwest corner of the box (Fig 3.3b). When the barrier was inserted the field immediately shifted to the corner between the east side of the barrier and the north wall. On the fourth day of recording, insertion of the barrier provoked two fields in the corresponding northwest corners on each side of the barrier (Fig 3.3b, day 4). During six subsequent days of recording, the cell continued to exhibit a dual field when the barrier was present. Even with the barrier in an east-west orientation, the cell fired in both northwest corners above and below the barrier (data not shown).

In total we followed 11 cells for time courses varying from a day to the duration of the experiment. Clear it is difficult to draw firm conclusions from such a small data set. Nevertheless we can make some important observations that bear on the predictions of the BVC model, and our implementation of learning. First, the majority of the cells (six) exhibited some form of gradual plasticity. This was largely of the forms described above: duplication of a field, loss of a duplicate field or movement of field. Some cells showed all three forms at various stages. Second, there was no obvious relationship between changes observed concurrently in different cells. Third, cells we

followed for longer seemed more likely to exhibit change. Two cells we held for the course of the experiment demonstrated plasticity, whereas, cells recorded for a single day remained unchanged. The implication of this is that these plastic changes are not coherent across cells, occurring at different times in different cells (Lever *et al.* saw this same effect (2002b)).

3.4.2 Simulation

We calculated the firing of 1000 BVCs and 100 place cells in a 65x65cm enclosure with a 40cm 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 described above were cut from similarly coloured 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.

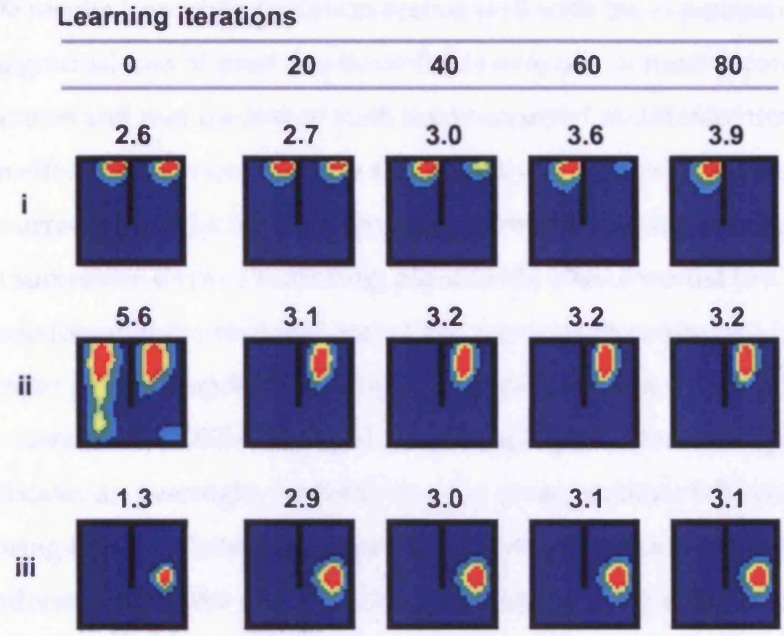

Fig 3.4. Changes in place fields accumulated during prolonged exposure to an environment partly divided by a barrier (65cm x 65cm environment – 40cm barrier), simulated data. Presence of the centrally placed barrier initially produced a number of duplicate place fields (i & ii). Repeated iterations of 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.

Results were superficially similar to the previous simulation: 37 place cells were initially active (peak firing >=1Hz), rising to 51 after learning. Mean peak rates were slightly higher, rising from 2.6Hz to 4.0Hz. The key difference, however, was that initially 59% of active place cells 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 gradual 'tidying' of place fields, though this time the effect also extended to the duplicate fields (Fig 3.4). 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. In particular, the gradual loss of most duplicate fields over one or more recording session of 40 minutes 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.

3.5 Simulation: effect of barrier movement after learning

Rivard *et al.* (2004) reported a series of manipulations in which rats were run in a 76cm diameter grey circular environment while place cells were recorded. The environment was polarised by a 45° white cue card fixed to the wall and a 23.5cm 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 3.5a). 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 ('place cells').

To better understand Rivard *et al.*'s result we simulated 200 iterations of learning for 100 place cells in a 76cm diameter circular environment with a 23.5cm 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 >=1Hz), 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 with the barrier removed. We deal first with the effect of the barrier rotation on place cell firing.

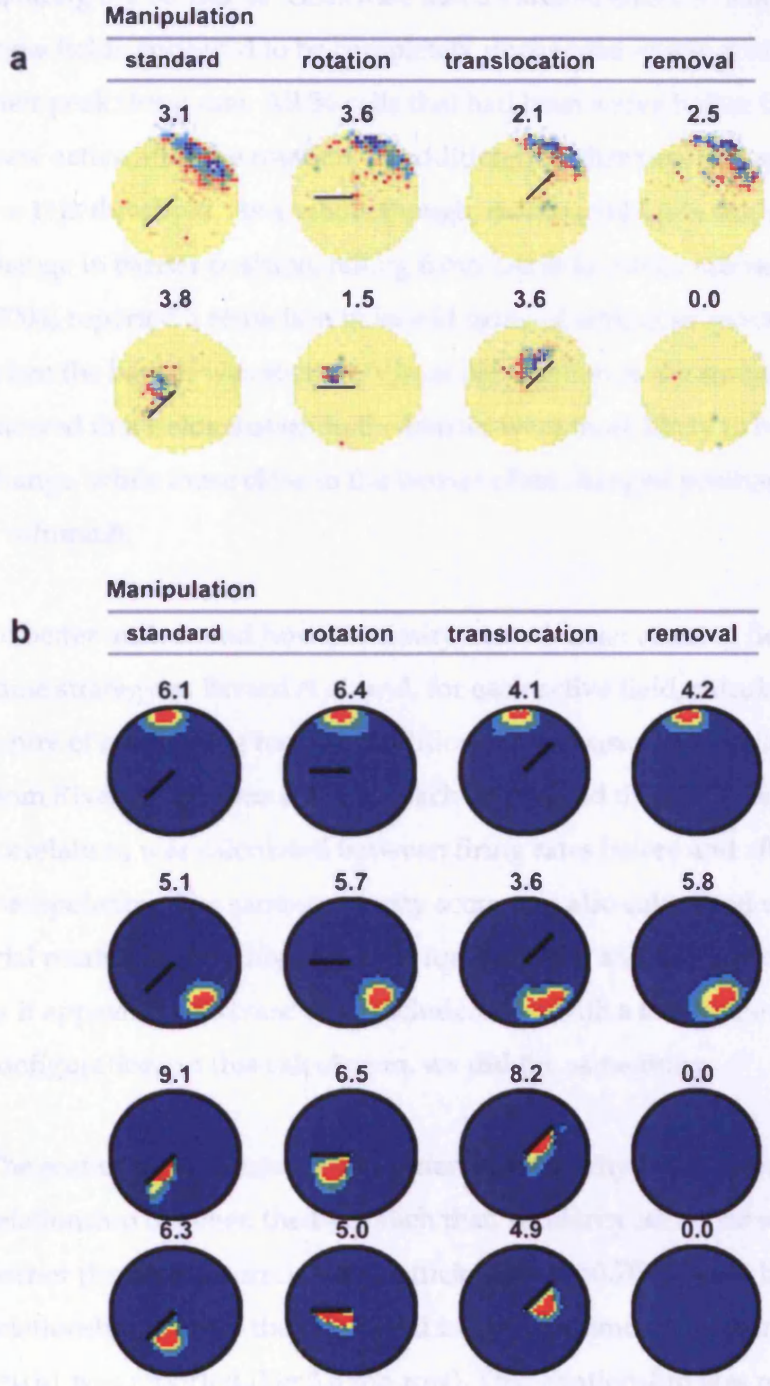

Fig 3.5. Changes in place cell firing induced by manipulations of a centrally placed barrier. **a**) Experimental data, images adapted from Rivard *et al.*(2004). Place cells were recorded in a 76cm diameter circular environment. Increased firing rate is indicated by darker colours, yellow indicates no firing. Values above each plot show mean infield firing rate in Hz. Animals were familiar with the standard condition in which a 23.5cm 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 place cell firing. After 200 iterations of learning in the standard condition cells were transferred to three probe environment with the barrier rotated 45° clockwise, translocated to the North-East section of the environment 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, while fields close the barrier (bottom) appeared to encode space relative to it.

Rotating the barrier 45° clockwise had a variable effect on simulated place cell 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 1Hz threshold. As a whole though, mean infield rate was slightly reduced by the change in barrier position, falling from 2.4Hz to 2.0Hz. Interestingly, Rivard *et al.* (2004) reported a reduction in infield firing of similar proportions, 4.1Hz to 3.3Hz, when the barrier was rotated. Visual observation of the simulated place cell 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 3.5b - column2).

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 centre 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 0Hz in both configurations in this calculation, we did the same thing.

The scatter plot of distance to barrier vs. 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 3.6 bottom row). This relationship mirrors the one found in the experimental data in which a correlation of r=0.61 was reported (Fig 3.6 top row). This relationship was reversed when firing in the probe trial was rotated 45° counter-clockwise before making the comparison, 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.

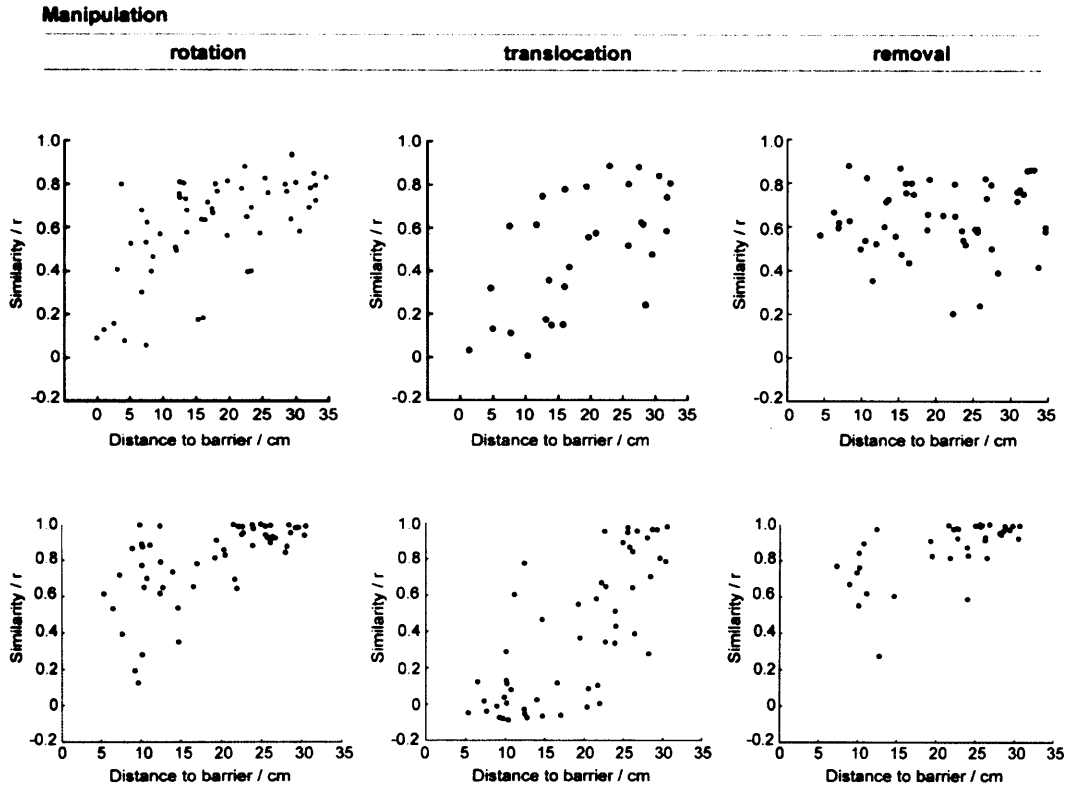

Fig 3.6. Similarity between standard condition and probe trials (barrier rotated, translocated and removed) as a function of distance between field centroids and the barrier, experimental and simulated data. Top row, experimental data, adapted from Rivard *et al.*(2004). In response to barrier rotation 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.

Using the same set of BVCs and place cells described above, we simulated place cell firing with the barrier translocated to the North-East and also with it removed. As with the 45° rotation we found excellent accordance between simulated fields and experimental data. In both conditions mean infield firing rate fell when compared to the standard condition (translocation: 2.4Hz to 1.5Hz; removal: 2.4Hz to 1.7Hz; Rivard *et al.* reported reductions from 4.3Hz to 3.5Hz and from 4.5Hz to 3.6Hz respectively). Fields adjacent to the barrier were again most affected, in the case of the translocation, a plot of similarity vs. distance to the barrier yielded a correlation of r=0.73 (r=0.64 in the experimental data) (Fig 3.6). 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 7.5cm from the barrier, this rose to 20.4cm for stable cells; in the simulation 63% of cells within 10cm of the barrier turned off but only 24% of cells further from the barrier (Fig 3.6).

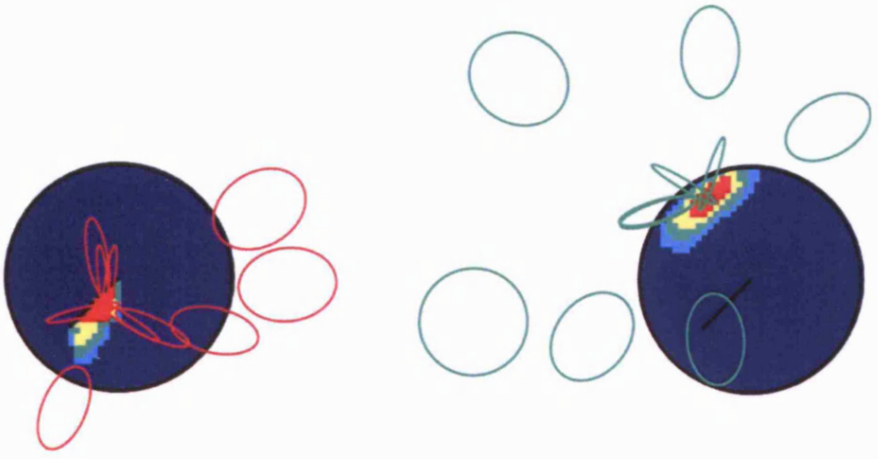

Fig 3.7. Simulated firing of two place cells showing the 10 BVC sets driving each cell. Firing is shown in an enclosure analogous to the 'standard' configuration used by Rivard *et al.* In each case the receptive field of BVCs are portrayed as ellipses at a preferred firing distance and direction. Note 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 with have limited effect on the firing of this cell.

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 place cell than BVCs with longer tuning distances. Inspection of the simulated place cells 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 3.7). As a result of learning these BVCs come to drive the place cell 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 place cell, 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 place cell 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 optimises inputs from BVCs whose firing fields closely coincide with that of the postsynaptic place cell. 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 place cell.

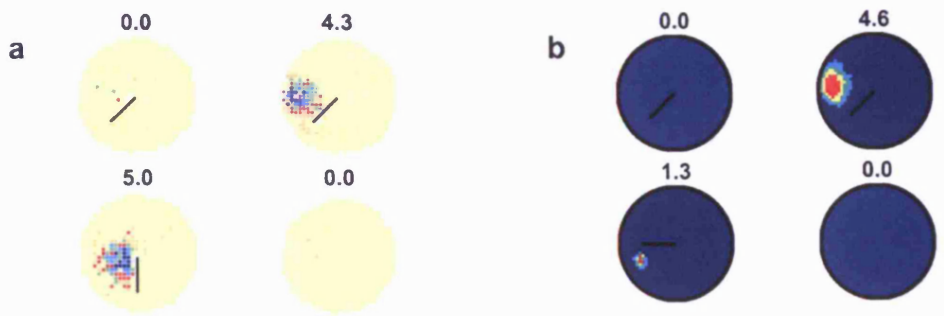

Fig 3.8. 'Conjunction cells' that exhibited large changes in rate after manipulation of the barrier. **a**) Experimental data, adapted from Rivard *et al.* (2004). Two cells are shown, one per column, top row is the standard condition, bottom row is firing after making changes to the barrier position. 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.

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 (Fig 3.8a). The authors construed these as conjunction cells, place cells 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 et al., 1987a). We saw similar effects in the simulations. When the barrier was rotated two cells became active (peak firing>=1Hz), the space occupied by their fields having previously been occupied by the barrier (Fig 3.8b). In response to the translocation two cells became active and a further nine showed either a large change in peak firing rate (more that 100% increase or reduction) or a pronounced change in the size of the place field. These effects are easily explained by the 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 place cell downstream of these BVCs might be driven more strongly, exceeding its firing threshold, alternatively it may receive less input and become silent. In contrast the graded control of place fields by adjacent cues seen *in vivo* is only matched by the

model after learning with the BCM rule. Prior to learning, that is, in the basic BVC model, many cells have duplicate fields that fire against the barrier in addition to the external walls of the enclosure. Similarly, others cells have large, elongated fields jointly controlled by the barrier and external walls. As such, movement of the barrier is likely to result in fields turning on or off.

3.6 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 3.9a). 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 place cell firing and to control the orientation of the ensemble. To better understand these effects we simulated the firing of 1000 BVCs and 100 place cell in a 76cm diameter circular environment analogous to the one used by Fenton *et al.* The 'standard' enclosure was polarised by a two distinct cue cards fixed to the wall, each subtended an angle of 45° with their centres 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, place cell 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 (Zhang, 1996; Skaggs et al., 1995) implicate learnt associations between visual cues and head direction cells (Taube et al., 1990b; Taube et al., 1990a) as a mechanism for maintaining the system's alignment with the environment. In Fenton *et al.*'s study, the rotation of the place cell 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 place cell firing. We assume that the influence of a distal orientation cue is proportionate to its proximity to the animal (Fenton et al., 2000b; Burgess et al., 2002). Specifically, the angular deflection in the head direction system produced by a rotation of the cards apart by ϕ units would be:

$$H_{def} = \frac{\phi(d_B - d_W)}{2(d_B + d_W)} \qquad (6)$$

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 ϕ 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. ϕ =25°) , 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 3.11a). 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.

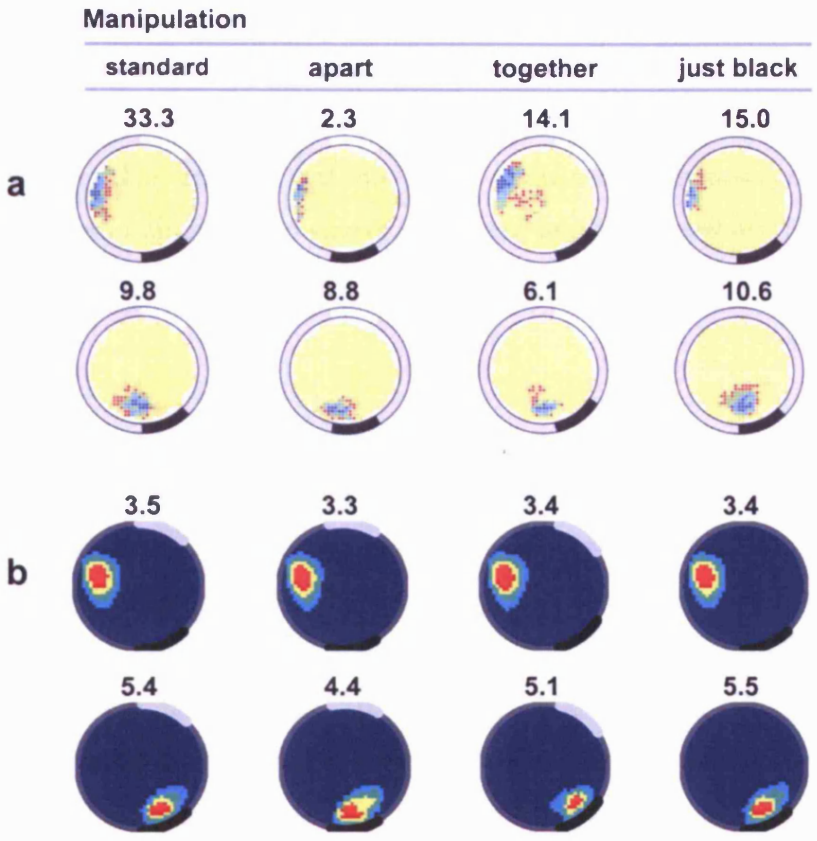

Fig 3.9. Place cell firing in a circular environment (76cm diameter) polarised by two cue cards, experimental and simulated data. Place fields are shown in a familiar environment and the same environment after the cue cards were moved apart, together, and after one card was removed. Experimental and 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 place cells are shown, numbers above each plot indicate peak rate, dark colours show high firing, yellow indicates zero firing. **b**) Simulated firing.

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 1Hz, this rose to 56 after learning, place fields became tighter and secondary fields were lost. Rotation of the cue cards apart or together affected place cell firing in a similar way, like Fenton *et al.* we observed subtle changes in the position and firing rate of place fields (Fig 3.9b). Fields close to one of the cue cards tended to rotate with the card, while fields equidistant to both cards tended to be translocated. All active cells remained 'on' (peak firing >=1Hz) 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 5.0Hz to 3.9Hz, a reduction of 21%, moving the cards together reduced firing by 15%. As expected, and for reasons outlined above, simulated place cell 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, non-significant, 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 3.10a&b). 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 1000 place cells 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: 0.147cm, together: 0.387cm), that movement of the cards together produced a small horizontal shift in field centroids (1.80cm), and that movement of the cards apart produced a larger negative shift (-4.84cm). We found the same pattern, mean vertical movement was close to zero (apart: -0.15cm, together: 0.16cm), mean horizontal movement produced by moving the cards together was positive (3.56cm), and moving the cards apart induced a negative shift of slightly larger magnitude (-3.66cm). 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.

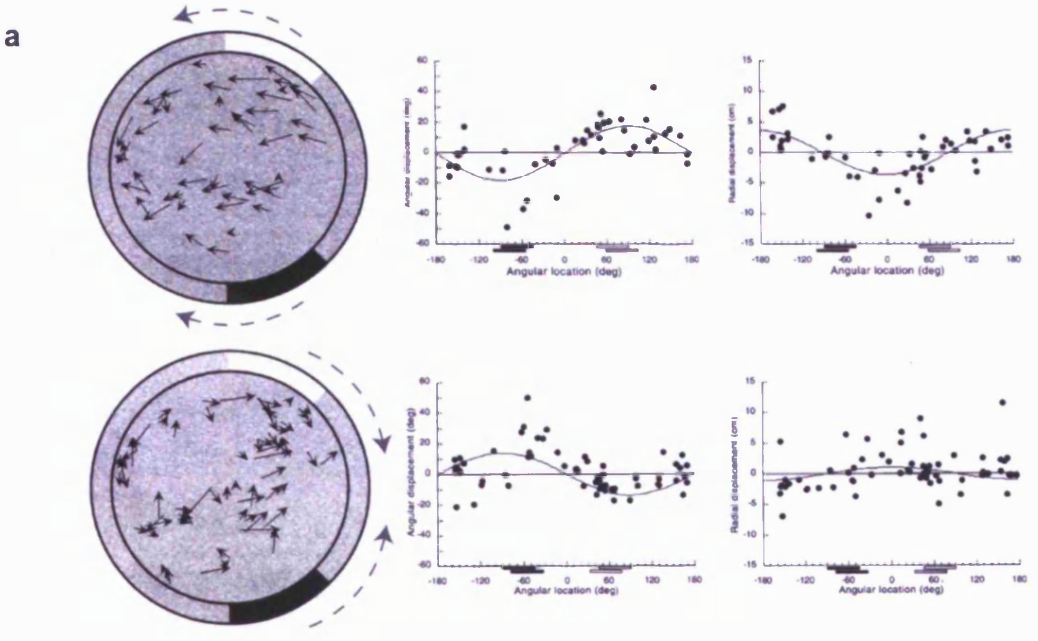
a
Angular location (deg)
Angular location (deg)
Angular location (deg)
Angular location (deg)

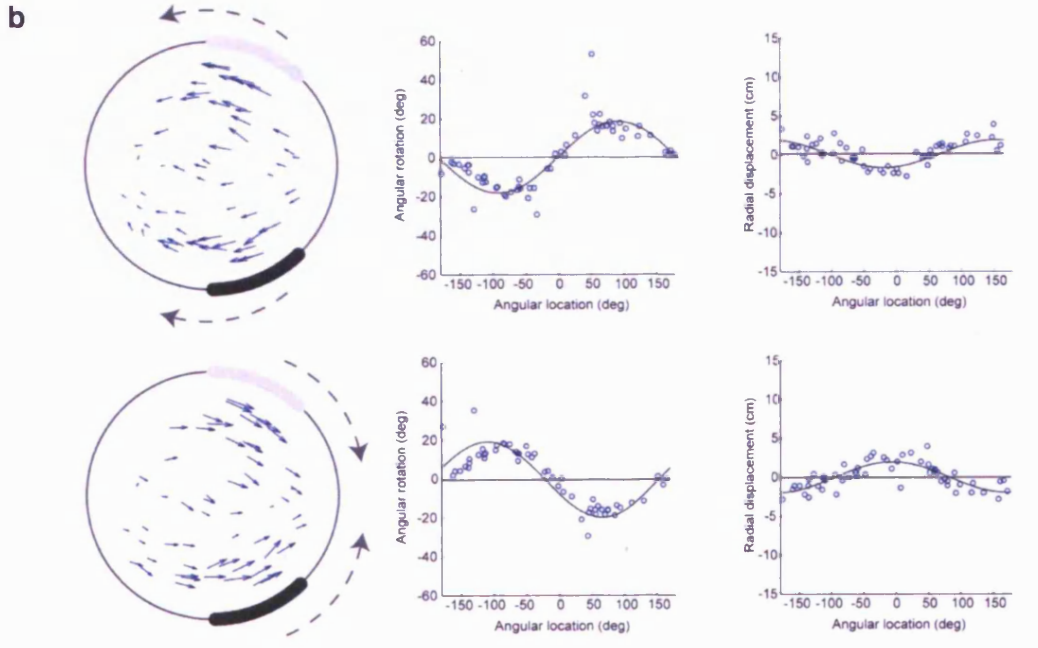
b
Angular rotation (deg)
Angular location (deg)
Radial displacement (cm)
Angular location (deg)
Angular rotation (deg)
Angular location (deg)
Radial displacement (cm)
Angular location (deg)

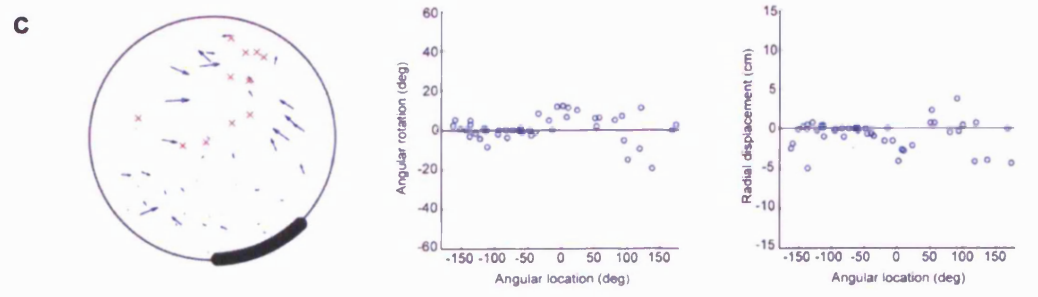
c
Angular rotation (deg)
Angular location (deg)
Radial displacement (cm)
Angular location (deg)

Fig 3.10 (previous page). 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 centre 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 centre of the environment such that movement directly away from the centre 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 minimise squared error. **c**) Simulation of removal of one cue card.

To characterise the contribution of the deflection to the head direction system, we recalculated the firing of the 1000 place cells described above, 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 principal difference induced by the head direction deflection is to accentuate the angular displacement of place fields produced by both manipulations (Fig 3.11b&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 Fig 3.10 & Fig 3.11). In the cards apart condition, taking account of the deflection in the head direction system increases mean horizontal field translocation by 39% (3.56cm vs. 2.54cm). Horizontal field movement in the cards together condition is similarly increased (-3.66cm vs -2.61cm, 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 place cell firing after removal of one of the cards. Using 100 place cells, we simulated removal of the white cue card. Simulated cards have arbitrary colour, 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 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 < 1Hz), whereas none of the experimental cells turned off. All the cells that became inactive had fields located close to the missing card (Fig 3.10c). Fenton *et al.* actually considered a cell to be active if it had a place field of at least nine contiguous pixels with above 0Hz firing. Even using this less stringent criterion for the simulation, nine of 54 cells became inactive (17%).

Observation of the inactive place cells 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 place cell 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 3.10c), 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 centre 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 2.0Hz to 1.7Hz (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%.

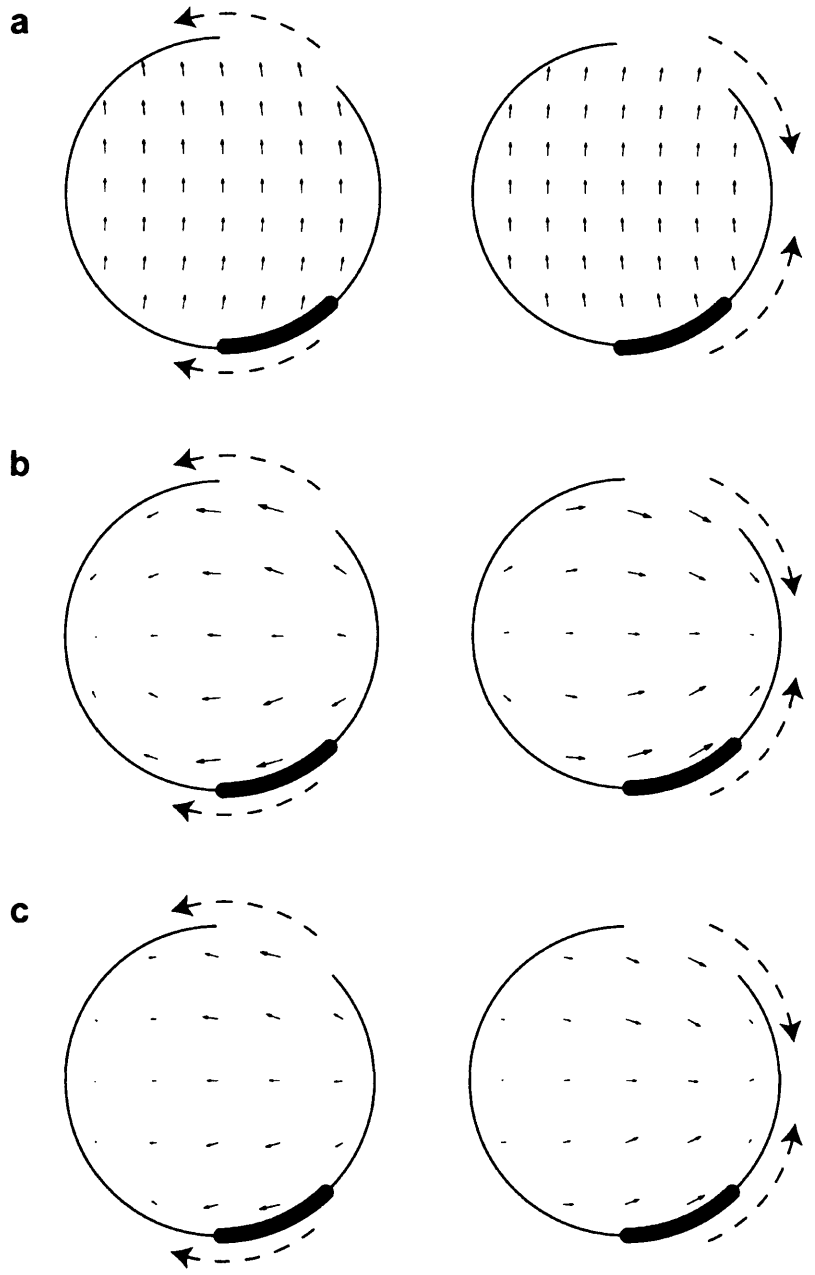

Fig 3.11. a) Distortion in the head direction system induced by card movement. Arrows show perceived North – in the standard condition all arrows point upwards. **b&c**) Average movement of place field centroids produced by increasing (left) and decreasing (right) cue card separation, simulated data. Average displacement was calculated using a population of 1000 place cells, 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.

3.7 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 place cells. Now, by incorporating synaptic plasticity mediated by the BCM rule, the model is able to describe changes that occur to place cells 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 place cells 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 et al., 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 place cells is reduced. *In vivo* three possible mechanisms for adaptation seem plausible: dynamic reduction of inhibition in response to reduced excitation (Marr, 1971; McNaughton & Nadel, 1990; Kali et al., 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 (Kali et al., 2000; Marr, 1971; Rolls, Stringer, & Trappenberg, 2002; Nakazawa et al., 2002; Wills et al., 2005). 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 utilised 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, Kjelstrup, Treves, Moser, & Moser, 2003; Leutgeb et al., 2004b), is best described as a wholesale reorganisation of place fields, such that individual fields undergo position and rate changes. It stands in contrast to the slow, cumulative changes in place cell activity discussed above (Barry *et al.*, 2006; Lever *et al.*, 2002b) and seems to be a distinct phenomena. 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 above, 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. The model, however, does not specify how a novel set of BVCs would be recruited to respond to the novel black cue card and white walls.

A final pair of question to be answered is, 'do BVCs really exist and how do they learn?' Burgess and O'Keefe (1996c) 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, in some situations where a place field has been induced to become bimodal after expansion of an environment, the firing of those fields appear to have two components: one reflecting sensory input from the boundary (be it visual, tactile

or acoustic) and the other reflecting additional path integrative information from the boundary most recently visited (O'Keefe and Burgess, 1996; Gothard *et al.*, 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 place cells (Sargolini et al., 2006; O'Keefe et al., 2005) 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, Hayman, Burgess, & Jeffery, 2007). 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, Barry, and O'Keefe (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 place cells 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 while learning about the configuration of those objects would occur within the hippocampus. Clearly the model suggests that, in the complete absence of BCM mediated plasticity, place fields should respond like those seen in the basic BVC model. A number of workers have recorded place fields after disabling the action of NMDAR mediated plasticity. For example, Kentros et al. (1998) showed that in the presence of CPP, a competitive antagonist of NMDARs, place fields fail to stabilise. Conversely, Nakazawa et al. (2002) deleted CA3 NMDARs in mice and found that place fields were largely unaffected, although subsequent removal of a subset of cues resulted in a reduction in field coherence. Obviously neither result accords perfectly with the model's

predictions. Never-the-less, the physiological locus of plasticity described by the model is unclear and so it is hard to know if the results of these and other studies are relevant.

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 place cells can encode aspects of episodic memory. In essence the firing of a set of place cells 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.

4 Environmental Rescaling of Entorhinal Grids[3]

4.1 Introduction

Grid cells in dorso-lateral medial entorhinal cortex (dlMEC) of freely moving rats show regular grid-like firing across the environment (Hafting et al., 2005; Sargolini et al., 2006). The uniformity of this pattern, rapidity with which it appears in a novel environment, and apparent independence of environmental scale suggests that grid cell firing provides an intrinsic metric for space (Hafting et al., 2005). Furthermore, a number of authors have already suggested that an entorhinal network, including grid cells, might be the site of a neural path integrator (Sargolini et al., 2006; O'Keefe et al., 2005; McNaughton et al., 2006). Indeed, several computational models encapsulate these ideas and provide a plausible account of the grid's origin as a product of self-motion information (O'Keefe et al., 2005; Fuhs et al., 2006). As described previously these models fall into two broad camps, oscillators-interference (O'Keefe et al., 2005; Blair et al., 2007) and recurrent networks (Fuhs et al., 2006; McNaughton et al., 2006), both of which currently receive adequate support from the experimental literature (Sargolini et al., 2006; Hafting, Fyhn, Moser, & Moser, 2006; Kohler, 1986; Klink & Alonso, 1997; Giocomo et al., 2007).

Despite their intrinsic nature, grid firing is reproducible across trials, indicating an association to environmental information (Hafting et al., 2005). Phrased differently, behavioural data (Etienne et al., 2004; Etienne et al., 1996) and computational considerations (Samsonovich et al., 1997; O'Keefe et al., 2005; Zhang, 1996; Burgess, Barry, Jeffery, & O'Keefe,) suggest that a pure path integrator would drift in the absence of reliable sensory input. Grid firing, however, has been shown to be stable both within, and between trials, suggesting that sensory information can be used to update the path integrative network (Fyhn et al., 2004; Hafting et al., 2005). In particular O'Keefe and Burgess (2005) (also Burgess *et al.* (in press)) have suggested that sensory input to grid cells might be mediated via places cells. Though contrary to some perceptions of entorhinal-hippocampal information flow, this view sees evidence from several sources, for example: Computational considerations suggest that it is easier to associate imprecise sensory information with a unitary place field, than with grid firing that covers an environment (O'Keefe et al., 2005). Anatomical

[3] The data and analysis in this chapter appear in Barry, Hayman, Burgess and Jeffery (2007).

projections from CA1 directly to the deep layers of the MEC and projections from these deep layers to superficial layers, suggest how spatial information might directly reach the grid cell network (Kloosterman, Van Haeften, Witter, & Lopes da Silva, 2003; Kohler, 1986; Kloosterman, Van Haeften, & Lopes da Silva, 2004). However, contrary to this proposal, while place cell firing responds parametrically to deformations of the environment (O'Keefe et al., 1996; Hartley et al., 2000; Barry et al., 2006), grid cell firing apparently does not (Hafting et al., 2005).

To clarify, O'Keefe and Burgess's (2005) suggestion that place cells anchor grid cells to the environment leads to a clear predication; geometric manipulation of an environment sufficient to produce parametric changes in place cell firing should induce complementary, though not necessarily commensurate changes in the 'grids'. By contrast, an environmental change that induced remapping in the place cells (e.g. changing the colour of the enclosure (Bostock et al., 1991)), would see the grids temporarily decoupled from external sensory input, causing a novel pattern of regular grid firing to be instantiated. This second aspect of the proposal, that place cell remapping should be accompanied by non-parametric changes in grids, is supported by Fyhn et al.'s (2007) demonstration that grid realignment is coincident with global remapping in CA3, though obviously this finding does not demonstrate causality.

However, although O'Keefe and Burgess saw the sensory input to grid cells mediated via place cells, the contradiction highlighted above holds true for any sensory input to the medial entorhinal network. Put baldly: subtle discrepancy between the putative path-integrator (the grid cell network) and sensory input should result in distortion of the regular grid metric. The actual magnitude of distortion created in the grids is likely to depend on the relative efficacy of sensory and path integrative systems. Here we test this proposal and show that grid cell firing patterns distort parametrically in response to deformations of a familiar enclosure.

4.2 Method

See General Method for details of animal housing, surgical procedure, basic recording strategy, spike sorting, and histology.

4.2.1 Animals

Six male Lister Hooded rats (250-400g in weight and three to six months of age at implantation) were used in this study[4].

4.2.2 Recording & behaviour training

Entorhinal activity was recorded while the animals foraged for honey-sweetened rice in a four-sided enclosure placed centrally on the floor of the experimental room. Four 120 x 50cm Perspex sheets backed with white plastic made the walls of the enclosure and a grey polyvinyl sheet (100 x 120cm) the floor. Walls were held in place with corner clamps allowing the Perspex sheets to be moved past one another to form rectangles with different length sides. Four configurations of the enclosure were used: large square (100 x 100cm); vertical rectangle (70 x 100cm); horizontal rectangle (100 x 70cm) and; small square (70 x 70cm). In all conditions the walls were 50cm high and afforded the animals a view of the experimental room containing rich orientation cues including the recording system, several posters and a desk. Before commencing data collection, rats were familiarised with the environment in a 'baseline' configuration, for three animals this was the large square (100 x 100cm) and for the other three it was the vertical rectangle (70 x 100cm). Animals were considered to be familiar with the enclosure after at least three, 20 minute exposures on separate days. All subsequent screening took place with the enclosure in the baseline configuration. Animals were not exposed to the other configurations (the probe enclosures) unless a cell recorded in the baseline configuration was visually confirmed to have a spatial firing field consistent with it being a grid cell.

Each experimental session consisted of five, 20 minute trials. The first and final trial always had the enclosure in the baseline configuration. The three interim trials saw the rat exposed once each to the three configurations which were not its baseline. These three probe trials were delivered in a random order. For example, a rat with the large square as its baseline condition would see this configuration first, then it might

[4] In total 17 rats were implanted for this study, 13 by the author (five of which contributed grid cells) and a further four by Dr. Robin Hayman (one of which contributed grid cells).

be run in the vertical rectangle, small square, and horizontal rectangle before being returned to the large square in the final trial of the session. The recording enclosure was disassembled between trials and reassembled with the walls in random positions, similarly the floor was cleaned with water and rotated by 0°, 90°, 180° or 270°. After each session electrodes were advanced at least 50μm. In subsequent sessions spike waveforms and the locations of peak firing were used, where possible, to confirm that cells were not submitted to the analysis on multiple occasions. Each rat experienced a maximum of one complete session per day and only complete sessions, consisting of five trials, were submitted to analysis (a single session was excluded as the recording system failed during a probe trial).

4.2.3 Data analysis

Inter-trial stability

The reliability of spatial firing between baseline trials was assessed by calculating the spatial correlation between the two smoothed ratemaps. The Pearson product moment correlation coefficient was calculated between equivalent bins in the two trials, unvisited bins were excluded from the measure. Cells with a spatial correlation of less than 0.5 between baseline trials were excluded from the analysis as being unstable. Twelve cells were excluded in this way and take no further part in the analysis. The positional and spike data from the remaining 54 cells was combined across baseline trials, yielding, for each cell, a single baseline ratemap constructed from 40 minutes worth of data.

Directionality

Watson's U^2-test (Sargolini et al., 2006; Zar, 1999; Watson, 1962; Fisher, 1996) was used to assess the extent to which each cell's firing was modulated by head direction. The U^2 statistic provides a non-parametric measure of departure between circularly distributed datasets. Accordingly the test distribution was constructed using the rat's orientation which was logged at 50Hz by the recording system. The sample distribution was taken as the animal's orientation at each occasion that a spike was emitted. To test the stability of directional firing two polar firing ratemaps were constructed from interleaved minutes of each baseline trial (e.g. ratemap 1 from minute 1, 3, 5 etc. and ratemap 2 from minute 2, 4, 6 etc.). Pearson's product moment correlation was calculated between the two polar ratemaps. Following Sargolini *et al.*

(2006) a cell was considered to be directional if both the U^2 statistic and the correlation were significant.

Gridness, grid scale and grid orientation

Spatial autocorrelograms of ratemaps (Hafting et al., 2005; Sargolini et al., 2006) were used to assess the periodicity, regularity and orientation of cells with multiple firing fields. Spatial autocorrelograms were estimated using unsmoothed ratemaps. Specifically the spatial autocorrelogram was defined as:

$$r(\tau_x,\tau_y)=\frac{n\sum\lambda(x,y)\lambda(x-\tau_x,y-\tau_y)-\sum\lambda(x,y)\sum\lambda(x-\tau_x,y-\tau_y)}{\sqrt{n\sum\lambda(x,y)^2-\left(\sum\lambda(x,y)\right)^2}\sqrt{n\sum\lambda(x-\tau_x,y-\tau_y)^2-\left(\sum\lambda(x-\tau_x,y-\tau_y)\right)^2}} \quad (1)$$

Where $r(\tau_x,\tau_y)$ is the autocorrelation between bins with spatial offset of τ_x and τ_y. $\lambda(x,y)$ is the firing rate in the bin at (x,y) and n is the number of bins over which the estimate was made. The autocorrelogram was then smoothed with a two dimensional Gaussian kernel of width 2.5 bins. The six peaks surrounding the central peak on the autocorrelogram were considered to be the local maxima closest to, but excluding, the central peak. The extent of each peak was defined as the contiguous set of bins around the peak with a value greater than half the value of the peak bin. In particular, the spatial autocorrelograms constructed from the combined unsmoothed baseline ratemaps were used to estimate the orientation, grid scale, gridness and regularity of each cell (Sargolini et al., 2006; Hafting et al., 2005). Orientation was the angle between a nominal horizontal reference line and an axis defined by the centre of the spatial autocorrelogram and the peak closest to the reference line in an anti-clockwise direction. Grid scale, or grid spacing, was the median distance from the central peak to the six surrounding peaks. Gridness, a measure of spatial periodicity, was calculated by defining a mask of the spatial autocorrelation centred on but excluding the central peak and bounded by a circle passing around the outside edge of the outermost of the six central peaks. This area was rotated in 30° increments up to 150°, and for each rotation the Pearson product moment correlation coefficient was calculated against the un-rotated mask. Gridness was then expressed as the lowest correlation obtained for rotations of 60° and 120° minus the highest correlation obtained at 30°, 90° or 150°. Cells with gridness of 0.3 or greater were classified as grid cells, a threshold that closely matched the experimenter's subjective judgment. A

further 14 cells that did not fulfil this criterion were removed from the analysis reducing the number to 40.

Measuring rescaling

The firing ratemaps of the stable grid cells were analysed to identify the presence of rescaling in the grid spacing between baseline and probe trials (see Fig 4.1 for an illustration of the process used to estimate rescaling, a description follows). Each smoothed ratemap from the three probe trials in a session were compared with the composite ratemap derived from the two baseline trials. Spatial correlations, described above, were used as a measure of similarity. The probe ratemaps were stretched independently in both the horizontal and vertical dimension by factors of 0.57 to 1.43, in steps of 2cm. Transformations were achieved by two dimensional linear interpolation using a nearest neighbour method. Hence, considering a probe trial run in the horizontal rectangle (100 x 70cm), the ratemap is 100cm long in the horizontal dimension, this would be 'squashed' and 'stretched' between 57cm and 143cm in 2cm steps during this analysis. At the same time the vertical dimension would be stretched between 39.9cm and 99.9cm. Forty three transformations were tested for each dimension, resulting in 1849 (43 x 43) combinations for each probe trial, each of which was compared with the baseline ratemap. In the event that a grid did rescale as a result of changes made to the geometry of the environment the focus of expansion or contraction would not necessarily be the centre of the enclosure. To mitigate against such effects of translational mis-alignment, correlations were calculated between the transformed probe ratemap and baseline ratemap for a complete range of spatial shifts at 2cm intervals along both the vertical and horizontal axis. This is equivalent to calculating the crosscorrelogram (Hafting et al., 2005) between the two ratemaps:

$$r(\tau_x,\tau_y)=\frac{n\sum\lambda_1(x,y)\lambda_2(x-\tau_x,y-\tau_y)-\sum\lambda_1(x,y)\sum\lambda_2(x-\tau_x,y-\tau_y)}{\sqrt{n\sum\lambda_1(x,y)^2-\left(\sum\lambda_1(x,y)\right)^2}\sqrt{n\sum\lambda_2(x-\tau_x,y-\tau_y)^2-\left(\sum\lambda_2(x-\tau_x,y-\tau_y)\right)^2}},$$

where $\lambda_1(x,y)$ is the firing rate in bin x,y of the baseline ratemap and $\lambda_2(x,y)$ the firing rate in the transformed probe ratemap. The value of the crosscorrelogram at bin (τ_x, τ_y) corresponds to the spatial correlation between ratemaps shifted by (τ_x, τ_y) relative to each other. The value for the spatial correlation between the baseline and

probe trial for each transformation was taken as the peak of the crosscorrelogram closest to the centre. The transformation that best matched the probe ratemap to the baseline was taken as the pair of horizontal and vertical stretches that gave the highest spatial correlation. The amount of rescaling exhibited was then the inverse of each of these transformations. For example, a baseline ratemap recorded in the large square (100 x 100cm) that 'squashed' in the vertical axis to fit in a horizontal rectangle (100 x 70cm) would show a rescaling of 1.00 in the horizontal axis and 0.70 in the vertical axis [1.00, 0.70]. Similarly a baseline ratemap from the vertical rectangle (70 x 100cm) that 'stretched' along the horizontal axis when recorded in large square (100 x 100cm) would have rescaled by [1.43, 1.00]. If the best correlation found between a probe trial and the baseline was below 0.5 that cell was removed from the analysis on the basis the estimated rescaling was not reliable; a further 2 cells were removed in this way reducing the total to 38.

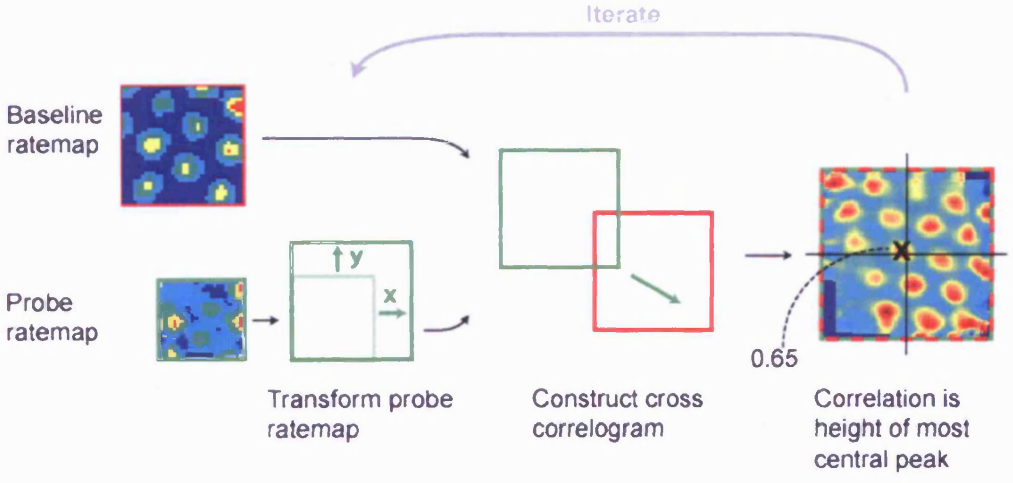

Fig 4.1. Process by which rescaling was estimated. Each probe ratemap was transformed independently in the X and Y axis to produce a series of 'stretched' and 'squashed' ratemaps. Each of these transformed ratemaps was compared against the baseline ratemap by calculating the crosscorrelogram; equivalent to finding the spatial correlation between the two for a range of offsets. Specifically the correlation, was taken to be the height of the most central peak in the crosscorrelogram. The correlations for all transformations of the probe ratemap were compared and the transformation that yielded the highest correlation was stored. Rescaling is assumed to be the inverse of the stored transformation.

Preliminary analyses were conducted using the raw rescaling values described above. For some secondary analyses (e.g. section 4.3.4) the raw values were collapsed across trials to provide an estimate of rescaling per cell; manipulated and un-manipulated axes were treated separately. To achieve this, 'stretches' and 'squashes' were first normalised to provide a comparable measure of rescaling. This was achieved by

converting each rescaling value into a proportion of the fractional change to the environment. Dimensions in which the enclosure was shortened in the probe relative to the baseline were treated according to equation 2:

$$S_{norm} = \frac{\left(\frac{1}{S_{raw}}\right) - 1}{0.43} \quad (2)$$

and, dimensions in which the probe enclosure was lengthened according to equation 3:

$$S_{norm} = \frac{S_{raw} - 1}{0.43} \quad (3)$$

where S_{raw} is the raw rescaling value and S_{norm} the normalised value. After this treatment a ratemap that did not rescale would have a normalised value of 0.0 and one which rescaled commensurately with the enclosure a value of 1.0. It follows that a negative value indicates that the ratemap rescaled in the opposite direction to the enclosure (e.g. squashed when the enclosure was lengthened). For probe trials in which one dimension remained unchanged relative to the baseline, rescaling along that dimension was converted in the same way as the manipulated dimension (being expressed as a proportion of the change to the other dimension). For example, the estimates of vertical and horizontal rescaling obtained from an animal trained in the large square (100 x 100cm) and probed in the vertical rectangle (70 x 100cm) would both be normalised using equation 2.

Grid regularity

Regularity was defined as the distance from the centre of the spatial autocorrelogram to the peak closest to the horizontal axis divided by the distance from the centre to the peak closest to the vertical axis (Fig 4.2). A regular grid would yield a value of 1.0, vertical compression of the grid would be expressed as a value above this and horizontal compression as a value less than 1.0. These values were used to compare the regularity of grids as a function of experience (e.g. grid regularity in the large square for rats familiar with this configuration as compared to those familiar with the vertical rectangle).

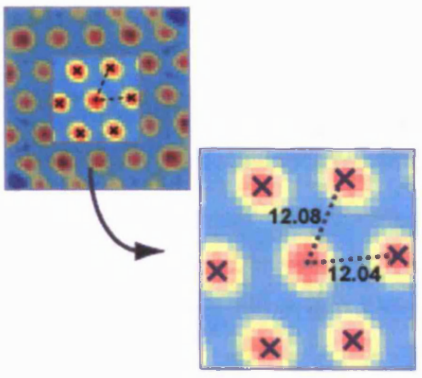

Regularity derived from central peaks of spatial autocorrelogram:

$$\text{Grid regularity} = \frac{\text{Distance to peak closest to X-axis}}{\text{Distance to peak closest to Y-axis}}$$

$$\text{Grid regularity} = 1.00$$

Fig 4.2. Example showing how grid regularity is calculated from the spatial autocorrelogram produced for each baseline ratemap.

Running speed analysis

Rescaling of the grid in response to a transformation of the environment might arise simply due to changes in running speed. To exclude this possibility spikes recorded in each baseline trial were divided into two groups, those emitted when the rat was running faster than the median speed in that trial and those emitted when it was running slower. These two sets were used to produce ratemaps and the rescaling between the maps was estimated.

4.2.4 Histology

Two out of 6 rats had successful recordings from multiple layers such that: one animal had recording locations in layer II, four in layer III, two in layer V and one in layer VI (Table 4.1 & Fig 4.3).

Table 4.1. Recording location and number of accepted grid cells by rat.

Rat	Layers recorded from	No. grid cells
214	III	6
216	III	10
217	II & III	12
228	VI	4
236	III & V	4
407	V	2

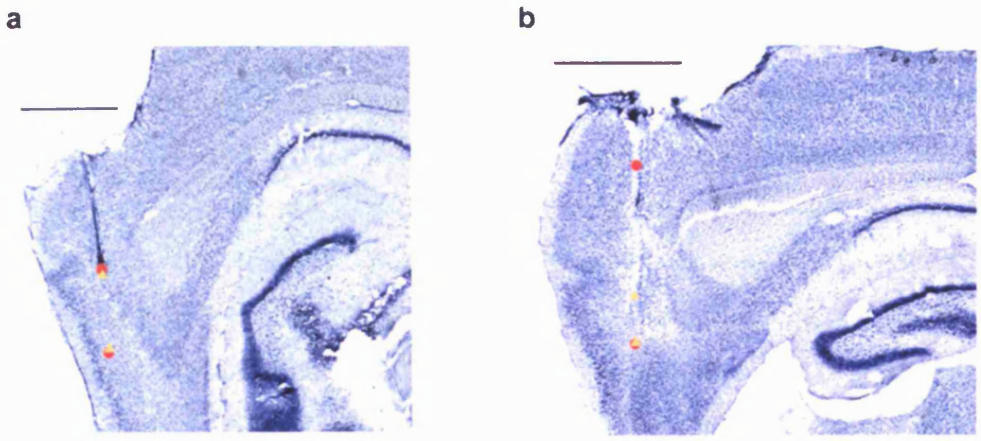

Fig 4.3. Cresyl violet stained sagittal sections showing typical recording locations. Red dots indicate the range of electrode locations (depth at implant to final depth) and yellow triangles the range over which grid cells were found. Scale bar indicates 2,000μm. **a**) r236, layer V moving to layer III. **b**) r216, layer III.

4.3 Results

4.3.1 Behavioural observations

During their first exposure to the baseline enclosure the rats were initially hesitant to explore, remaining close to the walls while repeatedly rearing and defecating. These signs of novelty detection habituated quickly and all animals foraged proficiently for rice by the end of the third, 20 minute familiarisation trial. The animals did not display the same novelty reaction when presented with the enclosure in its probe configurations.

4.3.2 Ratemap transformations

In total 66 putative grid cells were recorded during 41 recording sessions from 6 rats. Thirty eight cells from 3 rats familiarised with the large square (100x100cm) and 28 from 3 animals familiar with the vertical rectangle (70x100cm). Forty cells fulfilled the acceptance criteria set out in the method; 12 cells were excluded on the basis of low inter-trial stability and further 14 with low gridness scores (Hafting et al., 2005) were also removed. Each cell was recorded in five consecutive 20 minute trials. Recording sessions started with a trial conducted with the enclosure in its baseline configuration, this was followed by three probe trials, each with a different configuration, the final trial saw the enclosure returned to its baseline configuration. Data from the two baseline trials were pooled before comparison with the probes. Each cell is submitted to the analysis once.

Visual comparison between raw spikes recorded in the baseline trials and those recorded in probe trials show, for some cells at least, a change in spacing between peaks and also in the dimension of individual peaks (see Fig 4.4a&b also Fig 4.6 first column). The same effect is also evident in ratemaps and spatial autocorrelograms (Fig 4.4c-f also Fig 4.6 second and third column). To quantify this rescaling the vertical and horizontal transformation that best mapped each probe ratemap onto its baseline was estimated. Goodness of fit between the baseline ratemap and each transformed probe ratemap was measured using spatial correlations. Cells that yielded a correlation of less than 0.5 between the baseline and any of the three probe trials were removed from the analysis on the basis that the estimated transformations were unreliable. Two cells were removed in this way reducing the total to 38.

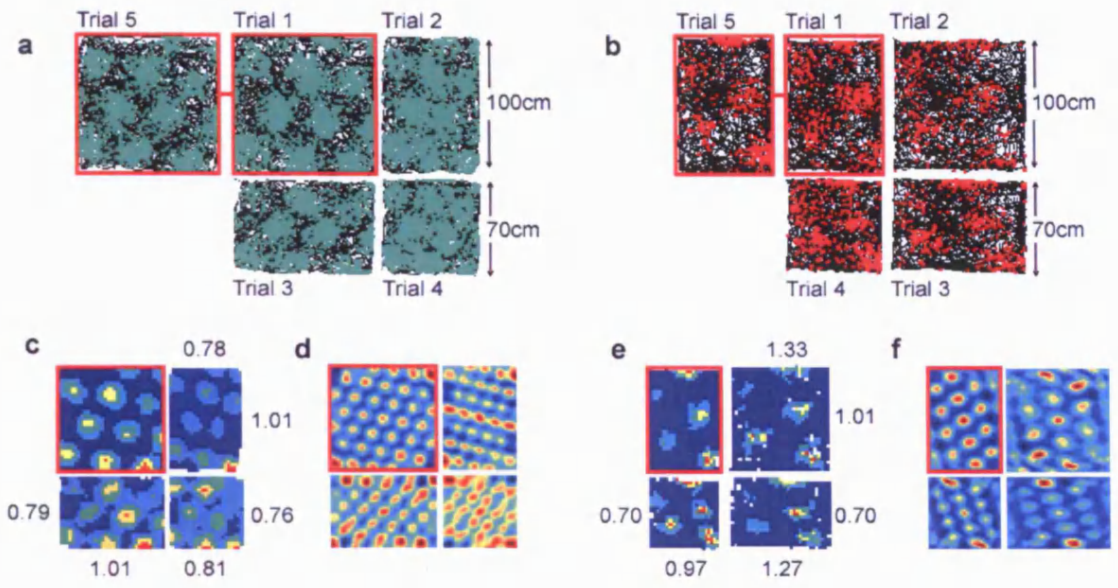

Fig 4.4. Rescaling of grid cell firing in response to geometric changes made to a rat's enclosure. Data from two cells recorded from different animals is shown. **a&b**) Raw spikes (green/red) superimposed on the animals' paths (black). Each cell was recorded in five trials, trials 1&5 being in the animals' baseline enclosure (outlined in red); large square for one animal (a), vertical rectangle for the other (b). **c&e**) Ratemaps constructed from a&b respectively (red indicates high firing, dark blue low). Baseline ratemaps, outlined in red, are composed of data from trials 1&5. Labels show, for each dimension, the rescaling of the ratemap relative to the baseline (>1 expansion, <1 contraction). **d&f**) Spatial autocorrelograms constructed from unsmoothed ratemaps.

The transformations obtained for each pairing of baseline and probe enclosure are summarised in Fig 4.5. A transformation of 1.0 indicates that the probe ratemap has not changed relative to the baseline, values less than this show contraction of the probe ratemap and values greater than 1.0 indicate expansion. Turning first to cells recorded from rats with the large square as a baseline (Fig 4.5a) (n=28). It is clear that reducing the size of the animals' enclosure induces, on average, a contraction of the ratemap in the same direction. This is most clearly seen in the vertical rectangle and horizontal rectangle. In the case of the vertical rectangle its horizontal dimension has been shortened from 100 to 70cm. Accordingly the ratemaps recorded in the vertical rectangle are, on average best fit by baseline ratemaps that have been reduced to 83% of their original width. Similarly the vertical extent of ratemaps recorded in the horizontal rectangle is best fit by the baseline ratemap after reduction to 81% of its original height. Ratemaps from the small square are contracted in both dimensions (88% in the horizontal and 83% in the vertical). One-tailed t-tests confirmed that each of these effects were highly significant; mean less than one (see Table 4.2). In the two

rectangular probes, rescaling of the ratemaps is also evident in the unchanged dimension. The horizontal dimension of the vertical rectangle which remained at 100cm shows a mean transformation of 1.01, a small expansion relative to the baseline. The vertical dimension of the horizontal rectangle shows a larger transformation of 1.09. Only the latter effect is significant (Table 4.2).

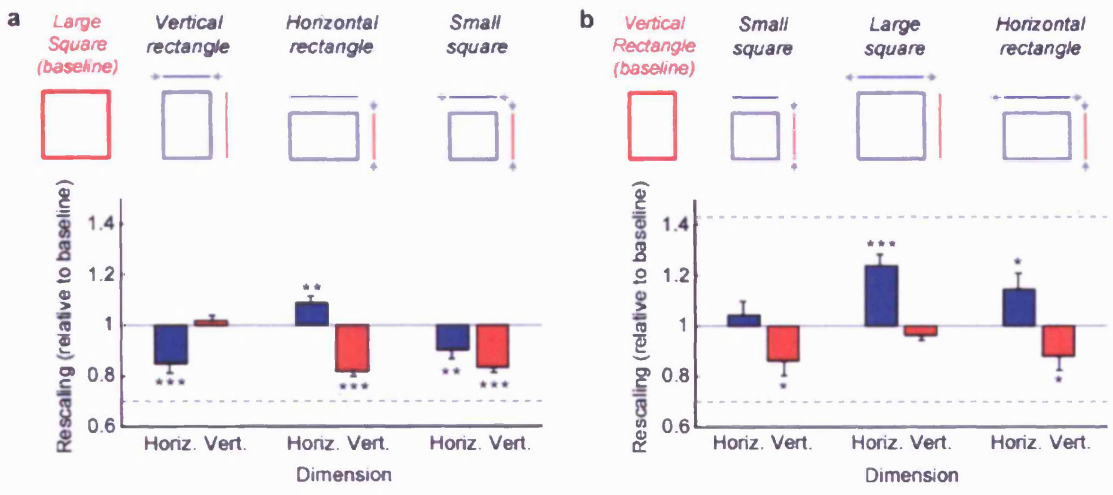

Fig 4.5. Mean vertical & horizontal transformations by enclosure. Values less than 1 indicate compression of the ratemap relative to baseline, values greater than 1 indicate expansion. **a**) Cells recorded from rats with the large square as baseline. **b**) Cells with the vertical rectangle as baseline. Error bars show SEM. In all cases: * p<0.05, ** p<0.01 and, *** p<0.001. T-tests against mean=1, one-tailed for transformed dimensions, two-tailed otherwise.

Cells recorded from rats with the vertical rectangle as a baseline show complementary results (Fig 4.5b) (n=10). Considering probes in the small square first, the mean transformation in the vertical dimension, reduced from 100 to 70cm, is 0.86. Probes in the large square show a transformation of 1.24 in the horizontal dimension which was extended from 70 to 100cm. Relative to the baseline shape the horizontal rectangle is stretched in the horizontal dimension and shortened in the vertical, ratemaps from this enclosure mimic that change showing a mean transformation of 1.14 and 0.88 respectively. Each of these effects is significant; again single tailed t-tests show a departure of the mean from 1.0 (see Table 4.2). Neither of the unchanged dimensions, the horizontal axis of the small square (70cm) and the vertical axis of the large square (100cm) showed a significant effect (Table 4.2), means were 1.04 and 0.96 respectively.

Table 4.2. Horizontal and vertical rescaling in each probe enclosure was assessed for significance. Single sample t-tests were made against a null hypothesis of no change (mean transformation=1).

Enclosure	Baseline	Horizontal Dimensions			Vertical Dimensions		
		P	T stat	DF	P	T stat	DF
Horizontal rectangle	Large square	0.003 (2 tail)	3.31	27	<0.0001 (1 tail)	-8.44	27
Vertical rectangle	Large square	<0.001 (1 tail)	-3.85	27	0.365 (2 tail)	0.92	27
Small square	Large square	0.008 (1 tail)	-2.60	27	<0.0001 (1 tail)	-8.25	27
Horizontal rectangle	Vertical rectangle	0.0238 (1 tail)	2.29	9	0.029 (1 tail)	-2.16	9
Large square	Vertical rectangle	<0.0001 (1 tail)	5.44	9	0.105 (1 tail)	-1.81	9
Small square	Vertical rectangle	0.448 (2 tail)	0.79	9	0.022 (1 tail)	-2.33	9

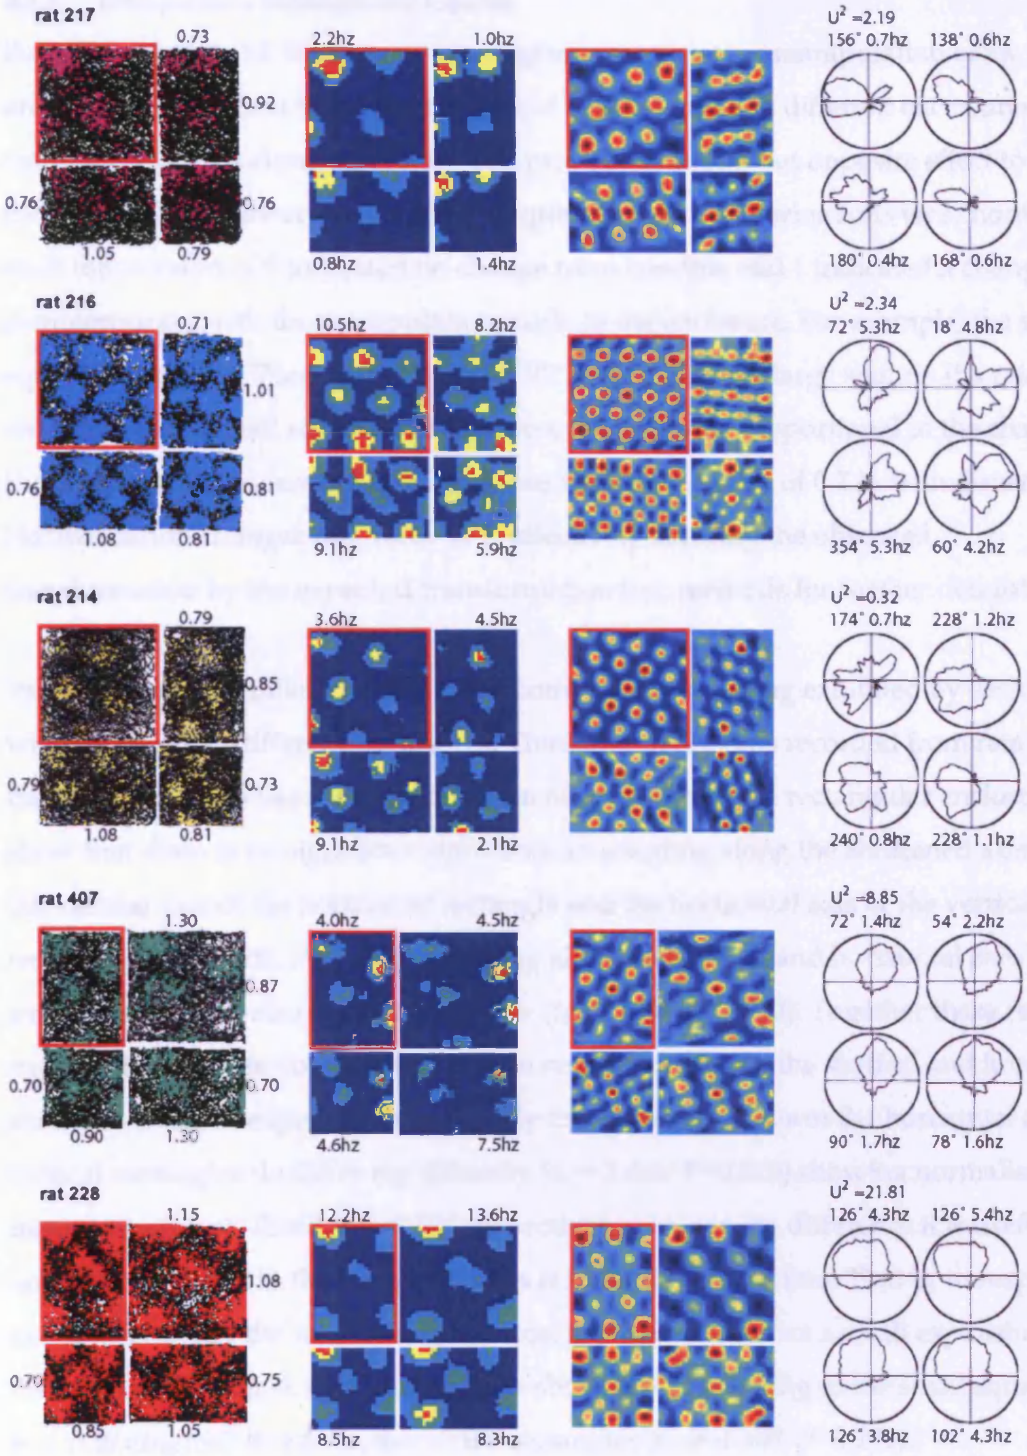

Fig 4.6. Examples of grid cell rescaling in response to geometric changes made to animals' enclosures. Each row shows a cell from a different rat, the first three rows from animals with the large square as baseline condition, the bottom two from animals with the vertical rectangle as a baseline. The first column shows raw spikes superimposed on the animals' paths. In each case the baseline condition is outlined in red. Numeric labels against the probe trials show, for each dimension, rescaling relative to the baseline (>1 expansion, <1 contraction). The second column shows ratemaps corresponding to the data in the first column, peak firing rate is indicated against each map. Autocorrelograms are shown in column three. The final column shows polar ratemaps, the preferred firing direction and peak rate is shown next to each map and the Watson U^2 statistic for the baseline condition is shown above the maps.

4.3.3 Comparison between enclosures

Rescaling of grid cell firing occurs in response to geometric manipulation of an animal's environment but is the amount of rescaling seen in different enclosures comparable? Does elongation of an axis provoke a similar but opposite effect to contraction? To answer these and other questions all transformations were normalised such that a value of 0 indicated no change from baseline and 1 indicated a change commensurate with the manipulation made to the enclosure. For example, the small square has sides of 70cm, a reduction of 70% relative to the large square. If a ratemap recorded in the small square exhibited rescaling directly proportional to the change in the enclosure then we would expect to see a transformation of 0.7 in both dimensions. Normalisation changes this value to 1, effectively dividing the observed transformation by the expected transformation (see methods for further details).

Paired t-tests (two tailed) were used to compare the rescaling exhibited by each cell when exposed to different enclosures. Thinking first of cells recorded from rats with the large square as baseline. Comparison of results from the rectangular enclosures show that there is no significant difference in rescaling along the shortened axes (i.e. the vertical axis of the horizontal rectangle and the horizontal axis of the vertical rectangle) (t_{27}=*-0.508, P=0.616*). Rescaling along the vertical and horizontal axes of the small square were also indistinguishable (t_{27}=*-1.937, P=0.063*). Together these results suggest that there is no particular bias in rescaling between the vertical and horizontal axes, this is not unexpected. Surprisingly the unchanged axes of the horizontal and vertical rectangles do differ significantly (t_{27}=*-2.496, P=0.019*) showing normalised mean rescaling of -0.160 and -0.020 respectively. Despite the difference it is useful to note that the effect in the two enclosures is in the same direction. That is, in response to a shortening of the walls along one axis, grid spacing shows a small expansion in the orthogonal axis. A two-sample t-test showed that rescaling in the small square was indistinguishable from that in the rectangles (t_{110}=*-1.107, P=0.271*).

Data from rats with the vertical rectangle as a baseline show no systematic differences in the magnitude of rescaling between manipulated axes. Specifically, rescaling along shortened axes (i.e. the vertical axes of the small square and horizontal rectangle) was not significantly different (t_9=*-0.441, P=0.669*). Similarly the lengthened axes (i.e. horizontal axes of the large square and horizontal rectangle) did not differ (t_9=*-1.391,*

P=0.198). Cells recorded in the horizontal enclosure were subject to simultaneous shortening of the vertical axis and extension of the horizontal axis, the magnitude of rescaling in these dimensions did not differ either (t_9*=0.319, P=0.757*). Comparison of rescaling along the horizontal axis of the large square and vertical axis of the small square also showed no difference despite the former expanding and the latter contracting (t_9*=0.417, P=0.687*). Finally rescaling along the unchanged vertical side of the large square was indistinguishable from that along the unchanged horizontal side of the small square (t_9*=-0.3431, P=0.739*). Again note that both unchanged axes show a non-significant trend for rescaling to occur in the opposite direction to the orthogonal, manipulated axes; normalised mean rescaling was -0.0849 and -0.0498 for the unchanged axes of the large and small square respectively. Hence in response to shortening of the vertical axis (small square) grids expand slightly in the horizontal direction and when the horizontal axis is expanded (large square) grid spacing in the vertical axis reduces.

4.3.4 Rescaling on a cell-by-cell basis

For rats with the same baseline enclosure, comparisons between changed axes showed no significant difference between normalised rescaling measures. Accordingly data were combined to provide a single mean estimate of rescaling for each cell in response to manipulation of the enclosure. A two sample t-test was used to compare data from animals with the square baseline and rectangular baseline, no difference was found and in subsequent analyses data is combined (t_{36}*=0.3153, P=0.754*); means were 0.480 and 0.445. Results from the unchanged axes were also combined and a comparison made between rats with different baselines. However, data from the unchanged axis of the horizontal rectangle was exclude as rescaling in this axis has been shown to be significantly different to rescaling in the unchanged axis of the vertical rectangle and also to be the only unchanged axis with rescaling significantly different to the baseline condition. Comparison of the other unchanged axes showed no difference between rats trained in the large square and vertical rectangle (t_{36}*=0.594, P=0.557*); means were -0.020 and -0.067 respectively.

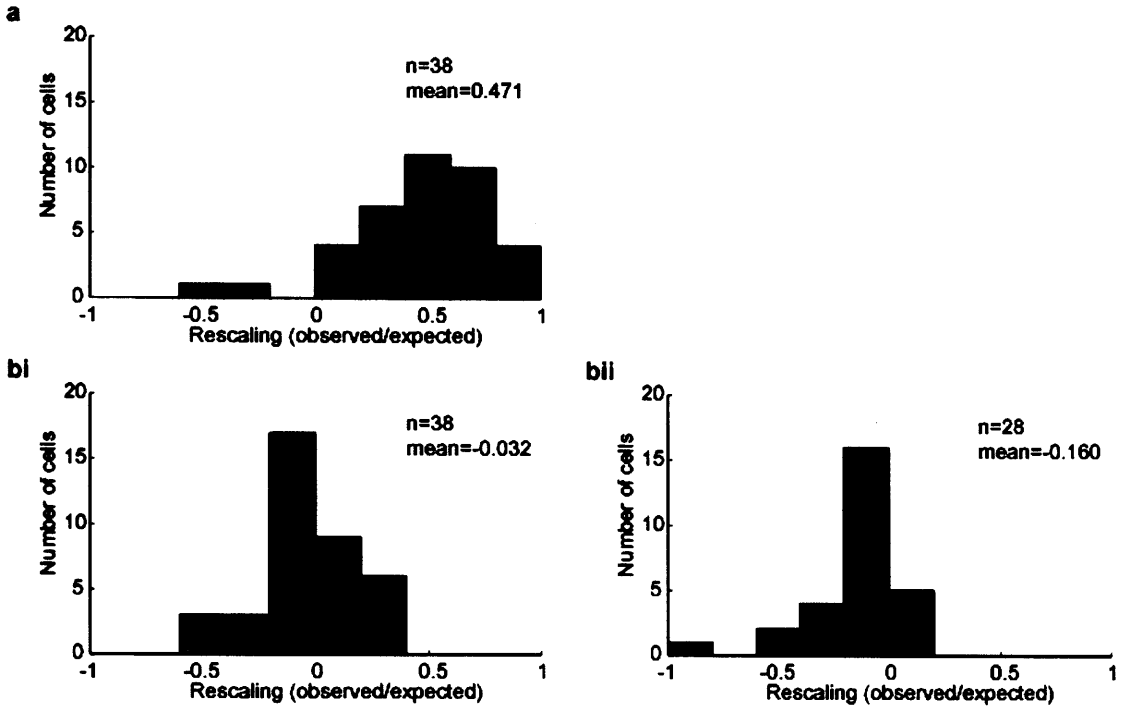

Fig 4.7. Normalised rescaling along manipulated and non-manipulated axes. Rescaling measures were normalised and combined to provide for each cell an estimate of rescaling along axes that had been changed relative to baseline and for those that remained the same. **a)** Frequency distribution of rescaling along manipulated axes **bi)** Frequency distribution of rescaling along unchanged axes – excluding data from the horizontal axis of the horizontal rectangle with the large square as baseline **bii)** Frequency distribution of rescaling along the unchanged horizontal axis of the horizontal rectangle with the large square as baseline.

The frequency histogram for rescaling along manipulated axes (Fig 4.7a) exhibits a unimodal distribution and precludes the possibility that some cells show rescaling and others do not. In the same figure the two outliers on the left hand side of the plot arise due to poor matches between the baseline ratemap and transformed ratemaps. In both cases the ratemaps in the probe trials have few areas of robust firing relative to the baseline and as a result a spurious match is found (see appendix 1 for examples of the two cells). Fig 4.7bi shows a frequency histogram for rescaling along unchanged axes. Data for the horizontal axis of cells recorded in horizontal rectangle with the large square as baseline are not included but instead are shown in Fig 4.7bii. Both distributions are similar, being unimodal with negative means close to 0.

One sample t-tests were used to assess whether the three distributions depicted in Fig 4.7 differ significantly from zero. Considering data from all cells, rescaling along manipulated axes is highly significant ($t_{37}=9.860$, $P<0.0001$). Whereas rescaling along un-manipulated axes, but excluding data from the horizontal rectangle, was not

distinguishable from zero (t_{37}=*-0.929, P=0.359*). Nevertheless, the negative mean (-0.032) suggests a tendency for un-manipulated axes to exhibit weak remapping opposite to that observed in the orthogonal axis. Data from the un-manipulated axis of the horizontal rectangle, as confirmed previously, departs significantly from zero (t_{37}=*-3.733, P<0.0001*) with a mean of -0.160. See Fig 4.8 for means and confidence intervals.

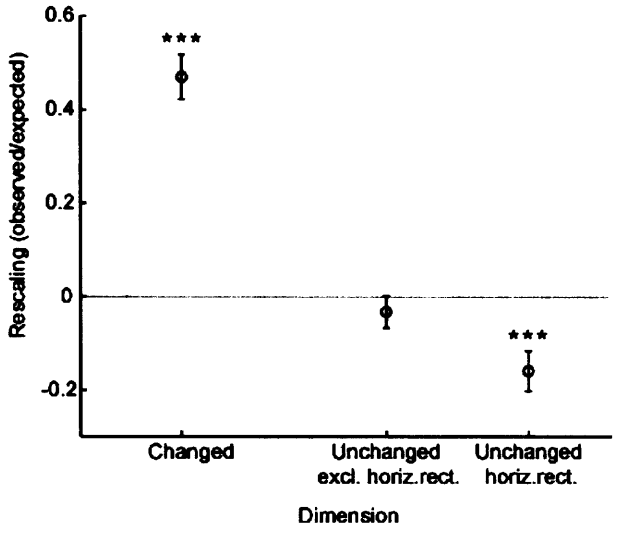

Fig 4.8. Mean rescaling along enclosure dimensions that were manipulated and along those that were unchanged (rescaling data for the horizontal axis of the horizontal rectangle with the large square as baseline is shown separately to other unchanged axes). Error bars show SEM. *** indicates significance at the P<0.001 level.

Finally as a conservative test of significance, data from the six rats were combined such that each animal contributed a single estimate of rescaling along manipulated and un-manipulated axes (Table 4.3). On this occasion all un-manipulated axes were combined to provide a single 'un-manipulated' measure of rescaling for each cell. Along the manipulated axes, grids rescaled by 47.9% of the change made to the enclosure (t_5=*4.949, P=0.004, mean=0.479*). Along un-manipulated axes grids rescaled by 7.9% in the opposite direction to the manipulated axis (t_5=*-2.920, P=0.033, mean=-0.079*).

Table 4.3. Mean normalised rescaling by rat for manipulated axes and unchanged axes. Rescaling was normalized to a proportion of the change to the enclosure along that axis, and the mean over all cells and manipulations calculated for each rat. A negative value for the unchanged axes indicates that rescaling occurred in the opposite direction to the change made to the manipulated axes.

Rat	No. grid cells	Manipulated axes		Unchanged axes	
		Mean rescaling	SEM	Mean rescaling	SEM
214	6	0.221	0.181	-0.186	0.106
216	10	0.691	0.030	-0.040	0.017
217	12	0.433	0.053	-0.083	0.051
228	4	0.292	0.104	-0.101	0.152
236	4	0.403	0.170	-0.073	0.032
407	2	0.835	0.029	0.011	0.032

4.3.5 Layer analysis

Grid cells were recorded from all layers of the MEC, though layer III was by far over represented and layer II was represented by a single cell (Table 4.4). Considering only cells from layers III, V and VI (layer II was neglected due to lack of data), there was no significant difference in rescaling between layers (*one-way ANOVA, $F(2,34)=2.24$, $P=0.070$*). A two-sample t-test between cells recorded from shallow (II & III) and deep layers (V & VI) also returned a null result ($t_{36}=-0.5696$, $P=0.573$).

Table 4.4. Number of grid cells accepted into the analysis and mean normalised rescaling by layer.

Layer	No. Grid Cells	Mean rescaling	SEM
II	1	0.196	-
III	29	0.465	0.055
V	4	0.756	0.067
VI	4	0.292	0.104

4.3.6 Directional cells

Previous work (Sargolini et al., 2006) has shown that the firing of some MEC grid cells is modulated by an animal's head direction. The amount of modulation was variable with some cells showing no modulation while others showed directional tuning similar to postsubicular head direction cells (Taube et al., 1990b). Additionally the

proportion of these 'conjunctive cells' was found to vary between layers; while layer II seemingly has none, grid cells in layer III seemed to be almost exclusively conjunctive (Sargolini et al., 2006). We follow Sargolini *et al.* (2006) and define directional cells as those with stable directional firing that departs significantly from circular (see method for details).

Applying these criteria we find that the majority of cells were directional (n=24) with about a third showing no directional modulation (n=14); proportions compatible with those reported by Sargolini *et al.* (2006). A two sample T-test found no significant difference in the amount of rescaling shown by directional and non-directional cells (t_{36}=*-0.94, P=0.356*).

4.3.7 Gridness

'Gridness', a measure of how well spatial firing conforms to an idealised hexagonal pattern, has been used previously to identify grid cells (Sargolini et al., 2006). Accordingly, gridness was calculated for all cells in this analysis and has already been used as an exclusion criterion; cells with gridness less than 0.3 being excluded from the analysis. Note that the formula used to calculate gridness departs slightly from that used by Sargolini et al (2006) (see methods section 4.2.3 for details). For the purpose of the following analysis the requirement that cells must have gridness greater than or equal to 0.3 is temporarily relaxed, all other criteria remain in place.

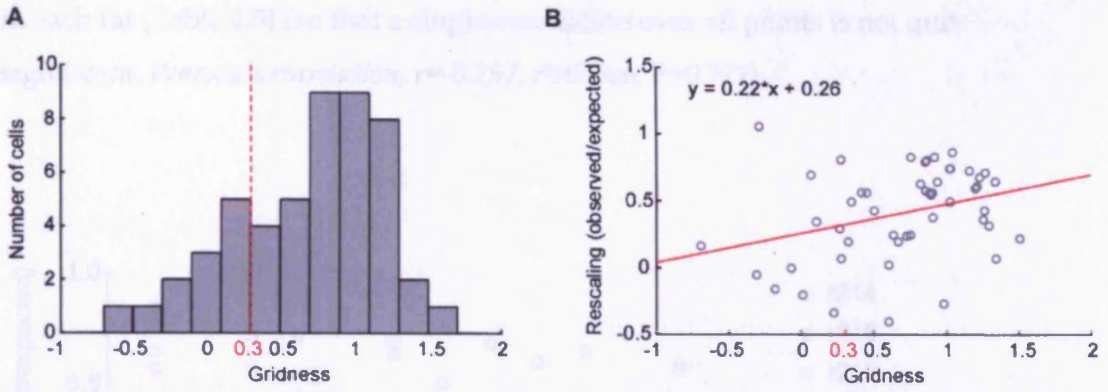

Fig 4.9. a) Frequency distribution for gridness (n=50). Twelve cells that were previously excluded as having low gridness (<0.3) are included. **b)** Scatterplot showing gridness vs normalised rescaling, again cells with low gridness are included (*Pearson's correlation, r=0.320, r^2=0.102, P=0.024*)

The gridness scores for analysed cells approximates a unimodal distribution from which the exclusion criterion removes the left-hand tail (Fig 4.9a). This may suggest that the cells currently identified as 'grid cells' are part of a continuum and simply represent a particularly regular subset of spatially periodic cells in the MEC. It is possible that grid cells with highly regular firing, this is high gridness, exhibit more rescaling than those which are not regular. A significant Pearson's correlation between gridness and rescaling (*r=0.320, r^2=0.102, P=0.024*) confirmed this to be the case. If cells with gridness less than 0.3 are excluded from the analysis the same trend is evident though the correlation is not significant (*r=0.178, r^2=0.032, P=0.285*). Mean rescaling for cells with gridness less than 0.3 was 0.224 and for cells with gridness greater than or equal to 0.3 it was 0.471 (t_{48}=2.247, *P=0.029*).

4.3.8 Rescaling as a function of exposure

Rescaling is possibly mediated by learnt associations between the relative positions of boundaries in an animal's environment. If this is the case, repeated exposure to the probe enclosures might cause new associations to be learnt and so reduce the amount of rescaling. To tackle this question the correlation between exposure to the protocol and rescaling was examined. Fig 4.10 shows a scatter-plot of exposure measured in experimental days, against rescaling; data were collapsed across simultaneously recorded cells such that a single value per rat per day is shown. In all five animals with cells recorded in multiple sessions, rescaling extent exhibited a negative relationship with session (*binomial, P = 0.031*), but with differential rates of reduction

in each rat (Table 4.5) (so that a single correlation over all points is not quite significant, *Pearson's correlation, r=-0.257, r^2=0.066, P=0.171*).

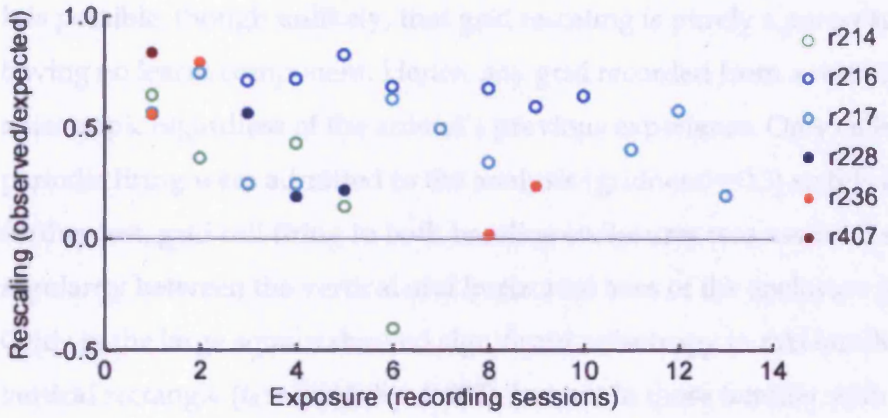

Fig 4.10. Effect of experience on grid rescaling. Mean normalised rescaling is shown per rat per day against the number of recording sessions experienced. Empty circles represent animals with the large square as baseline, filled circles represent those with the vertical rectangle as baseline.

Table 4.5. Pearson's correlation between exposure (number of recording sessions) and normalised rescaling for each rat. Concurrently recorded cells were pooled to provide a single mean estimate of rescaling per session. Individually, none of the correlations is significant but all animals with cells from multiple sessions had a negative relationship between exposure and rescaling such that rescaling becomes less pronounced with time (*binomial, P = 0.031*).

Rat	No. grid cells	No. sessions	r	r^2	P
214	6	5	-0.861	0.741	0.171
216	10	7	-0.669	0.447	0.061
217	12	10	-0.292	0.085	0.100
228	4	3	-0.830	0.689	0.376
236	4	4	-0.861	0.741	0.139
409	2	1	-	-	-

The possibility that rescaling reduces within each recording session (i.e. during the course of the five trials that were run each day) was also examined. To clarify, rescaling in the first probe in a session might be more pronounced than in the final probe, if this did occur rescaling would be expected to recover at least partially between sessions. Mean rescaling was calculated for each probe trial and the position

of the trial in the session was compared against rescaling. A Pearson's correlation between rescaling and position in the session found no significant relationship ($r=0.032$, $r^2=0.001$, $P=0.737$).

4.3.9 Grid regularity

It is possible, though unlikely, that grid rescaling is purely a perceptual process, having no learnt component. Hence, any grid recorded from a rectangle would be anisotropic regardless of the animal's previous experience. Only cells with regular periodic firing were admitted to the analysis (gridness>=0.3) so this is unlikely. As a further test, grid cell firing in both baseline enclosures was assessed specifically for regularity between the vertical and horizontal axes of the enclosure (see methods). Grids in the large square showed significant anisotropy in rats familiar with the vertical rectangle ($t_9 = 3.684$, $P = 0.005$), but not in those familiar with the large square ($t_{27} = -1.605$, $P = 0.120$). Conversely, grids were anisotropic in the vertical rectangle in the three rats familiar with the large square ($t_{27} = -5.682$, $P < 0.0001$) but symmetrical in the rectangle familiar to the other three rats ($t_9 = 0.035$, $P = 0.973$). In both cases, anisotropy was significantly greater in the unfamiliar than familiar configuration (*squares:* $t_{36} = -3.851$, $P < 0.001$; *rectangles:* $t_{36} = 3.280$, $P = 0.002$, see

Fig 4.11). Thus, grid structure in the deformed environments reflects experience and not simply the shape of the current environment.

The degree to which grid anisotropy is dependent upon experience is demonstrated nicely by a single grid cell that was recorded in two different rooms (Fig 4.12). Note, this example is provided purely for clarification and should be considered as a subjective observation. The cell in questions was recorded across six consecutive, 20 minute trials. The first two trials were run in the main experimental room using the usual white Perspex environment. Specifically, trial one saw the environment configured as the large square (the animal's baseline configuration), in trial two it was configured as the vertical rectangle. As expected the grid was regular in the square and anisotropic in the rectangle (see methods section 4.2.3 for details of how regularity was calculated). The next two trials took place in a separate recording room using a black, painted environment. The rat had only seen this room twice before and had largely been run with the environment configured as a vertical rectangle (70 x

100cm). Accordingly, trial three was run with the novel black environment configured as a vertical rectangle; in this case the grid was regular. The novel environment was reconfigured as a large square (100 x 100cm) for trial four, this induced the grid to become anisotropic, being 'stretched' in the horizontal dimension. For the final two trials the animal was returned to the first room and run in the vertical rectangle and finally in the large square; the original firing patterns were reinstated. Hence it seems that grid rescaling is both experience-dependent and, for this cell at least, independent between contexts. Interestingly there also appears to be a slight increase in absolute grid spacing in the novel environment; contrast trial three with trials one and six. Fyhn *et al.*'s (2007) noted a similar effect; a small number of grids expanded after grid realignment was induced by moving recording room or changing the animal's enclosure.

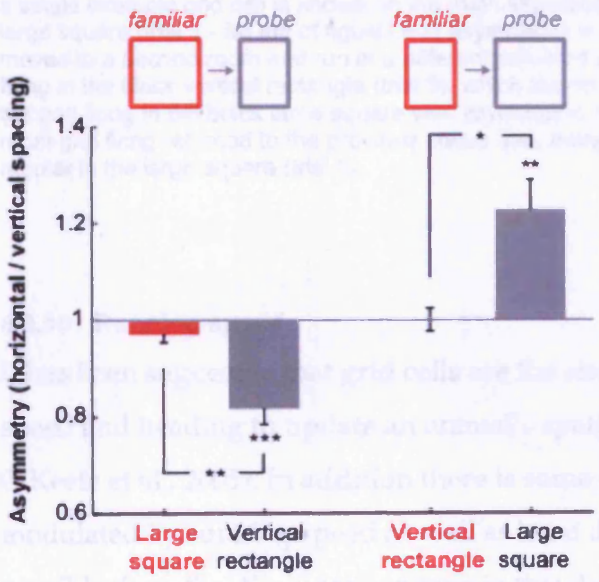

Fig 4.11. Induced grid asymmetry as a product of experience. Asymmetry was assessed for recordings made in the large square and vertical rectangle (mean ±SEM). Grids are symmetrical in the familiar enclosure regardless of its shape. Asymmetry was induced when grids were recorded in the probe enclosures. T-tests compare observed values to a mean of 1.0 (where 1.0 indicates no asymmetry; two-tailed, * indicates P<0.05, ** indicates P<0.01, *** indicates P<0.001).

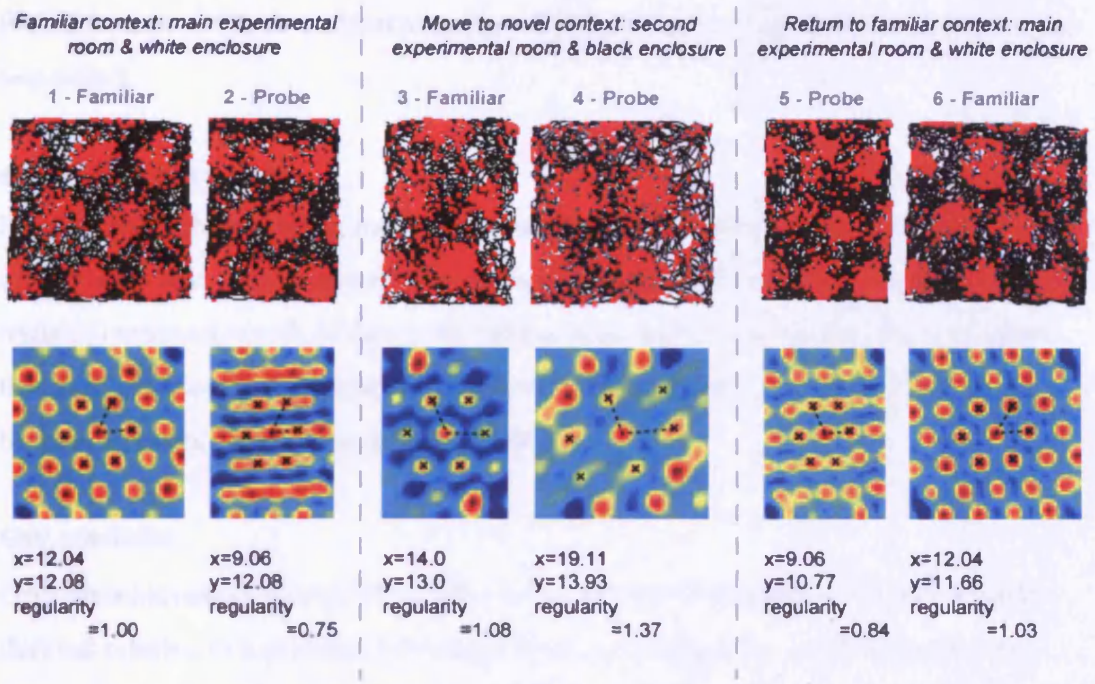

Fig 4.12. Grid rescaling is experience-dependent and apparently independent between contexts. Firing of a single example grid cell is shown. In the main experimental room, grid firing was regular in the familiar large square (trial 1 - far left of figure) and asymmetric in the vertical rectangle (trial 2). The rat was then moved to a second room and run in a different coloured enclosure (black). In this second context grid firing in the black vertical rectangle (trial 3), which the rat had experienced on previous days, was regular but grid firing in the black large square was asymmetric (trial 4). Once returned to the main experimental room grid firing returned to the previous status quo, being asymmetric in the vertical rectangle (trial 5) and regular in the large square (trial 6)

4.3.10 Running speed

It has been suggested that grid cells are the site of a neural path integrator that uses speed and heading to update an animal's spatial representation (Hafting et al., 2005; O'Keefe et al., 2005). In addition there is some evidence that grid cell firing is modulated by running speed as well as head direction (Sargolini et al., 2006). It is possible then, that the rescaling seen in this dataset is produced by changes in the animals' running speed, themselves resulting from changes in the enclosure dimensions. To exclude this possibility spikes recorded in each baseline trial were divided into two groups, those emitted when the rat was running faster than the median speed in that trial and those emitted when it was running slower. These sets were used to produce two ratemaps for each grid cell and the rescaling from the fast to the slow ratemap estimated. One sample t-tests confirmed that there was no rescaling between the fast and slow ratemaps in either the horizontal (t_{37}=0.247,

P=0.808, mean=1.00) or vertical axes (t_{37}*=-0.722, P=0.475, mean=0.99*), both tests were two tailed.

4.3.11 Grid cell metrics

In addition to the rescaling measures described above, the distribution of grid orientations was also assessed. Measures were extracted from the combined baseline trials (40 minutes worth of data). As before, cells were accepted into the analysis if they had gridness greater than or equal to 0.3 and a spatial correlation between the baseline conditions of 0.5 or greater (n=40).

Grid orientation

Grid orientations were identified from each cell's spatial autocorrelogram and are defined relative to a nominal horizontal line (see method section 4.2.3 for further information). A single cell was excluded from the analysis as its grid orientation was calculated as 87.8°, in excess of the plausible maximum of 60°. Inspection of the cell confirmed that its firing field was not a regular hexagon.

Inspection of the grid orientations suggests that grids within each rat are, to an extent, oriented in the same way (Fig 4.13). This is most evident in the first three plots (r214, r216 & r217). Rayleigh tests, treating each cell as an independent measure, confirmed that the distribution of grid orientations in rats 214, 216, 217 and 228 was significantly clustered (P values depicted in Fig 4.13). The result for r236 was almost significant (P=0.056), indicating a trend for clustering of orientations. In the case of r407 there were too few data points (n=2) to conduct an effective test.

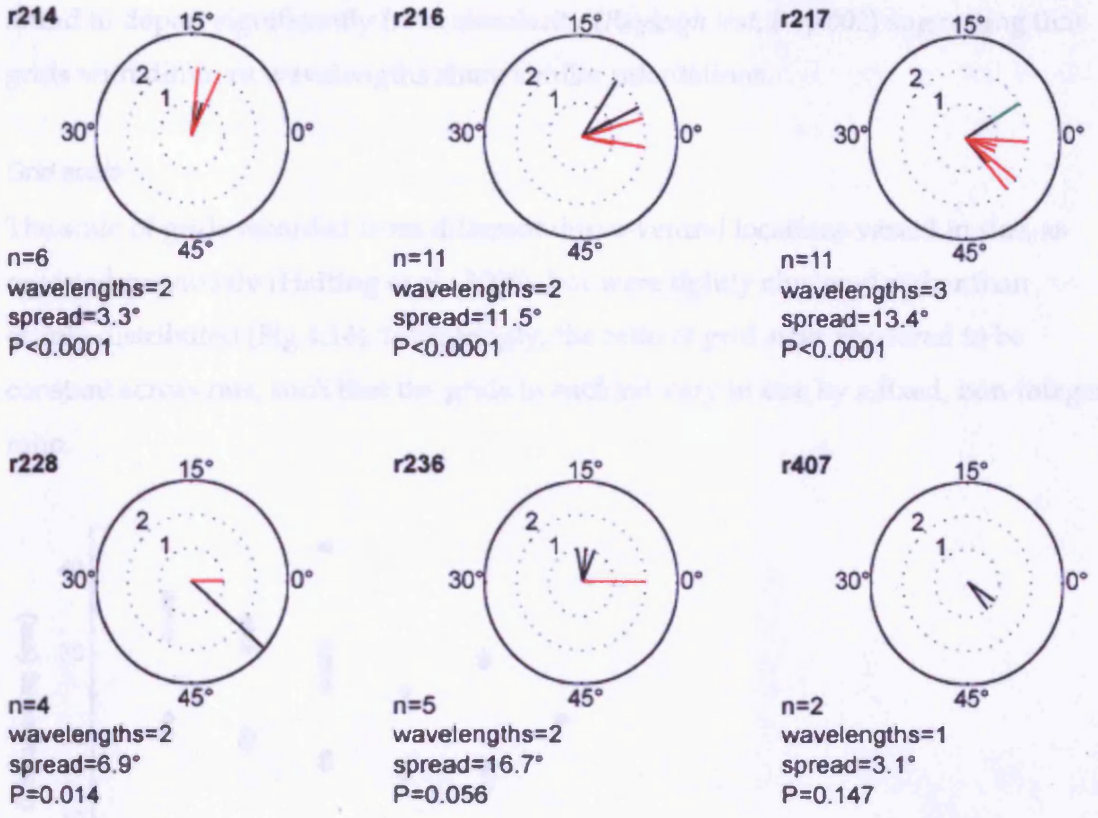

Fig 4.13. Orientation of grids in each rat. Orientation was identified from baseline spatial autocorrelograms and is defined relative to a nominal horizontal axis. Plots represent angles between 0° and 59° such that 0°=60°. Each cell is represented by a line of length 1, additional cells with the same orientation extend that line. Same coloured lines denote grid cells with similar spacing. The number of cells and number of distinct spacings (wavelengths) is shown below each plot, the angular distance (spread) between the most disparately oriented cells is also shown. P values for the Rayleigh test of circular non-uniformity are also shown; the first four rats show significant clustering of grid orientations.

Previous work has identified a topographical organisation of grid cells in the MEC such that adjacent cells tend to have similar orientations (Hafting et al., 2005). Furthermore work from O'Keefe and Burgess (2005) and others (Witter & Moser, 2006) suggest that proximal grid cells may form local circuits, the members of which have the same spacing and orientation. Hence a conservative view might be that the orientations of grids which have similar spacings should not be treated as independent measures. With this in mind the grid cells from each rat were sorted into sets with similar spacings; cells with spacings that differed by less than 20% were considered to be part of the same set (in Fig 4.13 sets are coloured differently). The mean orientation of each set was calculated and from this the angular distance between sets. In animals with more than two sets, the maximum distance between each set and the others was used. Angular distance was pooled across rats and was

found to depart significantly from circularity (*Rayleigh test, P=0.002*) suggesting that grids with different wavelengths share similar orientations.

Grid scale

The scale of grids recorded from different dorso-ventral locations varied in size, as reported previously (Hafting et al., 2005), but were tightly clustered rather than evenly distributed (Fig 4.14). Intriguingly, the ratio of grid sizes appeared to be constant across rats, such that the grids in each rat vary in size by a fixed, non-integer ratio.

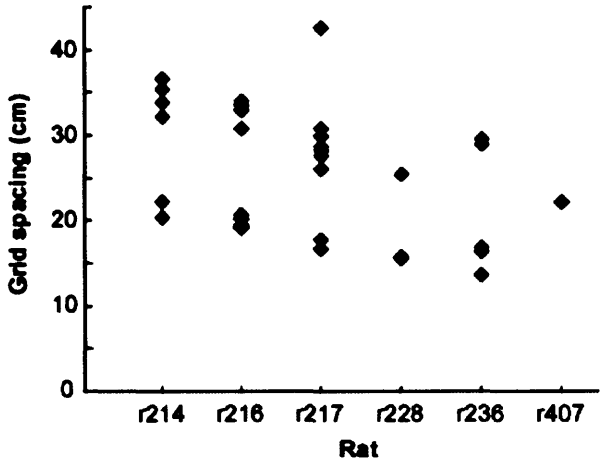

Fig 4.14. Grid scale, measured from familiar enclosures, by rat. The length of grids recorded from different dorso-ventral locations in each rat show a tendency to cluster.

To study relative grid scale in each rat, we grouped grids that had similar spacings (spacing differed by less than 20%) and calculated the median spacing for each group. The spacing of groups from each rat were then normalised by dividing by the spacing of the shortest group from that rat. Normalisation revealed a clustering of grid lengths not obvious from the un-normalised data; compare Fig 4.15a, a histogram of normalised scale, with Fig 4.15b, raw lengths. The second peak in Fig 4.15b indicates a tendency for the second-shortest group of grids from each rat to be approximately 1.7 times the length of the shortest. One rat had grids of three different scales and shows an additional peak close to 2.5 times the length of the shortest group. To determine if the clustering of the second peak was statistically significant we applied a non-parametric procedure as follows. First, the variance of the second peak was calculated. Second, for all five rats that had grids of at least two distinct lengths, the

length of the shortest group and second-shortest group were shuffled so that all possible pairings (i.e. 120 combinations) of shortest and second-shortest were created. Third, for each combination the normalisation process was reapplied and the variance of the second peak recalculated. Finally, variances were ranked and compared with the unshuffled condition. The unshuffled condition was found to have the lowest variance (*P=0.008*), indicating that grids within the same rat have a tendency to cluster at a fixed, non-integer ratio.

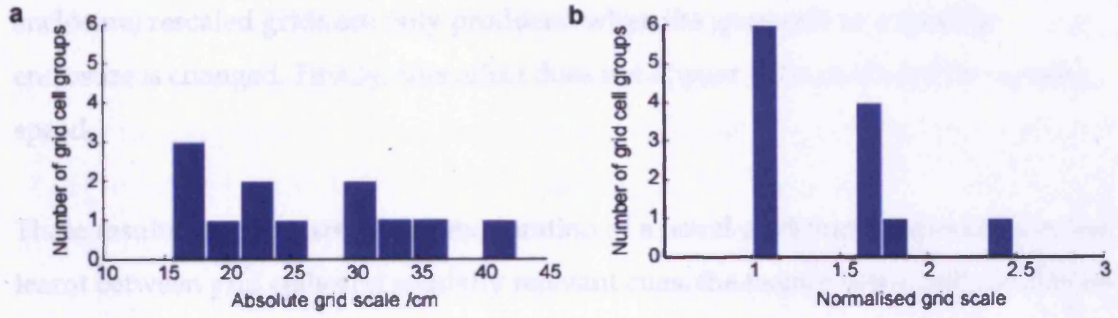

Fig 4.15. Grid scale clustering within rats. **a**) Histogram of absolute grid scale (cm) from all rats shows no obvious clustering. **b**) Histogram of normalised grid scale shows a tendency for grid scale to cluster at a fixed, non-integer ratio.

4.4 Discussion

The principal finding of this study is that the spatial firing of grid cells changes scale in response to geometric manipulations made to a familiar enclosure; the grids rescale. Rescaling occurs in the same dimension and direction as the change made to the enclosure and can be produced in both axes independently. There was no difference in the relative magnitude of rescaling exhibited in vertical and horizontal axis or between shortening and lengthening of the enclosure. The amount of rescaling exhibited by a cell was, on average, about half (47.9%) of that predicted by the change in geometry, though varied between cells and rats, and reduced with continued exposure to the protocol. Crucially, rescaling is a product of experience and not of the enclosure; rescaled grids are only produced when the geometry of a familiar enclosure is changed. Finally, this effect does not appear to be mediated by running speed.

These results suggest that during exploration of a novel environment associations are learnt between grid cells and spatially relevant cues, the former being initially driven purely by path integrative processes. Furthermore, that sensory input is sufficiently potent to continually update grid cell firing and does not simply function as a reset when an animal first enters a familiar enclosure. As expected, sensory input does not wholly control grid cell firing; the animal's position, as encoded by grid cells at least, appears to derive from a mixture of sensory and path integrative information. Hence grids do not rescale commensurately with the change made to the enclosure. A similar dichotomy has been noted in place fields recorded on a shortened linear track: Place fields were positioned according to a path integrative strategy at the start of the track but towards the end of the track were increasingly positioned according to a sensory strategy (Gothard et al., 2001; Gothard et al., 1996). Interestingly, O'Keefe and Burgess (1996) also noted a path integrative influence on place field positions in a geometrically manipulated enclosure. Although our results suggest that grids rescale uniformly (the measure of rescaling implicitly assumes this), it seems plausible that if animals were constrained to run in one-dimension an effect similar to the one observed by Gothard *et al.* (1996) might be apparent; rescaling might be more pronounced towards the end of a shortened track than at the start.

The slow, rat-specific reduction in rescaling is particularly interesting and might also be understood in terms of the relative efficacy of path integrative and sensory inputs. It seems possible that continued exposure to situations that see these two sources of information conflicted leads the system to prefer the internal path integrator. It follows that the ability of sensory cues to control grid cell firing is, for this reason, reduced, and so we observe less rescaling. However, the timescale of this reduction is notably similar to that of slow remapping seen by Lever *et al.* (2002b); they observed that over several days, continued exposure to geometrically dissimilar environments lead to a gradual change in the location and rates of place cell firing. Are these phenomena related? If, for example, place cells receive path integrative information from grid cells, then the gradual reversion of grids to their intrinsic dimensions might well drive steady changes in the firing of place cells. This view predicts that Lever's slow remapping can only be produced in enclosures that have different geometry and not by other contextual differences (Anderson et al., 2003). However, O'Keefe and Burgess' (2005) (also (Burgess, Barry, & O'Keefe,)) idea that place cells arbitrate between grids cells and the sensory environment suggests an alternative explanation. If place cells anchor grids to the environment, then the gradual reorganisation of place fields during slow remapping would allow grids to revert to their intrinsic shape; essentially a slow version of the process we anticipate during complete remapping (see below). Clearly both ideas are speculative; without hippocampal recordings it is impossible to know if the reduction in rescaling is paralleled by changes in place cell firing. However, the key ambiguity is the extent to which grid cells influence place cells and vice versa, it is to this point that we now turn.

So do these results prove O'Keefe and Burgess' (2005) proposal that place cells mediate sensory input to grid cells? The blunt answer is no. Although grid rescaling is compatible with place cell to grid cell connectivity, it does not preclude the possibility that grid cells receive direct sensory input from another source shared with place cells, possibly postrhinal cortex (Hargreaves, Roa, Lee, & Knierim, 2005), or indeed that grid cells drive place cells. However, the place cell-grid cell hypothesis is still attractive for several reasons. First, as mentioned previously, because grid cells fire in multiple locations, it is more difficult to see how sensory input specific to a single locale would be associated to a given cell (O'Keefe et al., 2005). On the other hand, place cells with their unitary fields can easily be described as a function of sensory

input (Hartley et al., 2000; Barry et al., 2006). Second, if place cells do 'pin' grid cells to the environment then any environmental change that induces complete remapping in the hippocampus would temporarily relax sensory control of grid cell firing, possibly causing them to realign (Fyhn et al., 2007). As such, rescaling may have been absent from Hafting *et al.*'s original study because their change from a large environment to a different small environment triggered remapping among place cells, resulting in the establishment of a new, symmetrical grid. Similarly, the slow, rat-specific, reduction of rescaling with repeated experience of the probe configurations implies an intrinsic tendency towards regular grids of a specific scale, and is consistent with the slow remapping of place cells between configurations of different shape (Lever et al., 2002b). Third, and perhaps most compelling, preliminary results from the Mosers' lab show that when the hippocampus is temporarily inactivated, the firing of layer II and III grid cells becomes unstable, to the extent that regular grid firing is no longer detectable (Bonnevie et al., 2006) (a similar but less pronounced effect was seen after NMDA induced hippocampal lesions (Fyhn et al., 2004)). One interpretation of this result is that, in the absence of sensory input from place cells, grids drift relative to the environment; visualisation of the data over short time periods might thus reveal a stable grid structure which would not be expected to rescale with geometric changes. However, other interpretations are possible; it might be that any interruption of the cortico-hippocampal processing loop leads to disrupted firing in the remaining structures, regardless of the site of sensory integration. A tentative conclusion then is that the activity of a population of place cells is necessary to maintain stable grid cell firing. What remains unclear, however, is if direct projections from the EC are sufficient to maintain reliable place cell firing in CA1 or if an intact cortico-hippocampal circuit is also required; dentate gyrus and CA3 are apparently not required (Brun et al., 2002). Furthermore, the relative importance of the MEC and LEC in generating a spatial signal in the hippocampus is also unclear though bilateral lesions of both structures are known to disrupt place cell firing (Miller & Best, 1980).

Can we infer anything about the timescale over which sensory cues are bound to the grid? In our study rats received at least three, 20 minute familiarisation trials on separate days before being exposed to manipulated enclosures. Not surprisingly this was sufficient time for sensory control over grids to become apparent. However, from other experiments we know that grids in a novel environment are almost immediately

stable (Hafting et al., 2005) and can be reinstated on a subsequent visit after 20 minutes of exploration (Fyhn et al., 2007). Most likely then, sensory control is established rapidly, possibly within the first few minutes of exploration. The establishment of stable place cell firing in a novel environment follows a similar time course (Wilson et al., 1993), though it is not possible to infer a causal relationship between these events.

We did not find any significant difference in the propensity of grid cells from different layers to rescale, although the result of an ANOVA between layers was close to significance ($P=0.070$). In particular the mean rescaling shown by layer V and VI cells (0.756 and 0.292, respectively) differ most strongly (but still not significantly). However, both these populations are represented by just four cells and are confounded with any rat specific effects that might exist. More instructive is the comparison between superficial layers, which project to the hippocampus, and deep layers, which receive input from the hippocampus: no significant difference was found. In a similar vein, we found no difference in the amount of rescaling displayed by directional and non-directional grid cells. Surprisingly we did see a significant, positive correlation between grid regularity, measured by gridness, and rescaling. A clear interpretation of this last finding is not obvious. The correlation is only significant when cells with gridness less than 0.3 are considered and, as such, these may simply not be grid cells (mean rescaling for cells with gridness greater than or equal to 0.3 was 0.471, and for cells with gridness less than 0.3 it was 0.224). Such a conclusion is problematic, although little is known about the mechanism by which grids form, our own results and results from other labs (Sargolini et al., 2006) suggest that gridness is a fairly continuous measure, by implication there is likely to be no distinct population of grid cells. Thus it is not valid to conclude that some cells are definitively grid cells, whereas others are not. Furthermore, it is difficult to explain why the cells in our sample with gridness less than 0.3 were not parametrically controlled by spatial cues; after all they did exhibit stable spatial firing. A more plausible account is that cells with low gridness scores are likely to have less well defined firing fields and fewer regular peaks. As such, the estimate of rescaling derived from these ratemaps is likely to be less precise and this, given the small number of these cells, may have led to an underestimate of rescaling. Hence, we conclude that rescaling is a phenomenon common to grid cells throughout the MEC.

In addition to rescaling in the same dimension as changes made to the environment we also observed a tendency for an opposing change in the orthogonal axis. Although this effect was only significant in one enclosure (horizontal rectangle), a trend in the same direction was evident in each of the enclosures in which a single axis was manipulated. With data from each rat pooled, the 'orthogonal reaction' was still significant and on average grids rescaled by 7.9% in the opposite direction to the manipulated axis. This reaction may indicate an intrinsic tendency to preserve overall grid scale, possibly resulting from recurrent dynamics (Fuhs et al., 2006; Kumar et al., 2007). Alternatively, an intriguing proposal is that it results from perturbation of the head direction system. For instance, assuming the head direction system employs the corners of an enclosure as orientation cues, then geometric manipulation of a single axis (or opposite manipulations of both axes) will cause some directions to be over represented relative to others. At the level of the path integrator, travel in the over represented directions will likely be over estimated and vice versa. As such this mechanism would contribute both to the orthogonal reaction and also the principal rescaling effect. One prediction then is that rescaling produced by a change made to a single axis should be greater than that induced by the same change made to both axes. Examination of data from rats familiar with the large square shows that our data does not show such an effect, rescaling in the small square is indistinguishable from that in the rectangles (t_{110}=*-1.107, P=0.271*). However, we note that the difference is expected to be the same magnitude as the orthogonal reaction (7.9%) and thus quite small. In fact mean normalised rescaling in the rectangles was 45.9% and in the small square it was 35.1%, a difference of 10.9%, a fair match to the prediction. Nevertheless, our first proposal, that the orthogonal reaction is a result of recurrent dynamics, also suggests an effect in the same direction and it is difficult to separate the models on the basis of this data alone.

The dataset collected here also provided an opportunity to assess several general features of the grid cell code for self-location. In particular, grids recorded from each rat were found to share a common orientation, including those of different sizes. This result stands in contrast to initial reports (Hafting et al., 2005) but accords with more recent observations (Fyhn et al., 2007). We also found that the scales of grids recorded from different dorso-ventral locations were clustered and that the ratio of scales appeared to be a fixed, non-integer across rats. These results chime with Fyhn *et al.*'s

(2007) observation that grids maintain coherent offset and orientation after realignment, and suggest that in addition to local recurrent circuits (Witter et al., 2006), grids of different scales are also intimately associated, a relationship that may form during development (McNaughton et al., 2006). Even so, it is not clear why grids should have a common orientation and specific, relative lengths. Indeed, the system suggested by the data runs contrary to preliminary expectations (McNaughton et al., 2006; Witter et al., 2006) poses problems for some models of place cell formation (Soldstad et al., 2006), and as pointed out by Fyhn *et al.*, seems to suffer from an intuitive flaw: if grids are bound together, then grid firing alone is not sufficient to disambiguate any two arbitrarily large environments (Fyhn et al., 2007). A novel theoretical approach from Brookings *et al.* (2006) suggests a possible resolution (see (Gorchetchnikov & Grossberg, 2007) for a similar approach). They characterised grid cells as a residue number system (a mathematical system predicated on modulo arithmetic), and showed that under this assumption location is optimally encoded when the lengths of similarly orientated grids are co-prime to one another (share no common factors). This theoretical framework also shows that Fyhn *et al.*'s fear regarding unique encoding of location is not founded, as a relatively small set of different grid lengths can accurately specify the location of a rat within a very large area (with 12 different lengths, 6cm precession can be reached across a $4km^2$ environment (Brookings, Burak, & Fiete, 2006)). Hence, the co-prime hypothesis is consistent with our demonstration that grid scales are quantised at fixed ratios and helps to explain why grids should share a common orientation.

In conclusion, our results suggest that the interaction between sensory and path integrative information that is needed for accurate self-localization may be mediated by entorhinal grid cells. Observed parallels in phenomenology and time course between grid cells and place cells suggest that this mediation may result from experience-dependent interactions between hippocampus and entorhinal cortex. Understanding the interactions between these regions will likely be critical to understanding spatial memory and, more generally, human episodic memory.

5 General Discussion

In general, the aim of this thesis was to explore how space is represented in the brain. More specifically, I wanted to understand how the spatial firing of place cells and grid cells is defined. How are firing fields established in the first place and what role does learning have to play in their maintenance? To tackle these questions I adopted two distinct approaches. In chapter three I described a computational model of place cell firing. The model portrays place fields as a function of environmental geometry and incorporates plasticity between place cells and their inputs. In chapter four I reported single unit recordings made from entorhinal grid cells while rodents explored a familiar enclosure with variable geometry. In the following sections I revisit each of those chapters and briefly restate their principal conclusions. I then examine the outstanding issues relating to each experiment. Finally, I describe how the experiments relate to each other, ask if the results are complementary, and look at the implications for spatial memory.

5.1 Learning and the BVC model

The boundary vector cell (BVC) model proposed by Hartley *et al.* (2000) describes place fields in terms of the barriers throughout an animal's environment. As such it explains parametric changes in place fields arising from changes made to environmental geometry (O'Keefe et al., 1996; Lever et al., 1999). The model also successfully captures the presence of crescent shape place fields in circular enclosures (Muller et al., 1987a), and predicts the duplication of place fields after barrier insertion (Barry et al., 2006), as well as the search location of human subjects in a variable geometry virtual reality environment (Hartley et al., 2004). However, the model employs a static feed-forwards architecture that does not allow learning, and so cannot capture experience-dependent changes in place cell firing (Lever et al., 2002b; Hayman et al., 2003). Neither does the model provide any basis, other than position, on which to distinguish different surfaces; Burgess and Hartley (2002) presented a preliminary solution to this latter point. To address these concerns I described an updated model incorporating BVCs which discriminate between environmental surfaces. The new model also allows the connection strengths between BVCs and place cells to be updated through iterative application of the BCM learning rule (Bienenstock et al., 1982). This formulation captures plastic changes in place cell activity arising during continued exposure to a static environment. Specifically: the initial refinement and stabilisation of place fields; as well as the production and subsequent reduction of duplicated fields in the presence of a centrally placed barrier. In addition, the model also shows that after learning place cell firing can be described in terms of the relative position of specific, spatially constrained cues. Hence, perturbation of those cues produce graded changes in the firing rate and position of place fields similar to those observed experimentally by Rivard *et al.* (2004) and Fenton *et al.* (2000a). As I highlighted in the discussion, the updated model performs less well in situations where cues are removed. Specifically, the model predicts that cells will become silent when experimental data shows that they continue to fire. I suggested three possible mechanisms that might account for the maintenance of activity *in vivo:* dynamic reduction of inhibition in response to reduced excitation; pattern completion mediate by CA3; and path integration.

What are the implications of these findings? At the very least they assert the validity of the standard BVC model. In particular, the graded translocation of place fields

reported by Rivard *et al.* and Fenton *et al.* accord with two core tenets of the model: BVCs respond in allocentric coordinates; and place fields are controlled more strongly by BVCs responding to adjacent barriers than those responding to distant barriers. More generally, the updated model provides an account of how place fields are defined relative to specific cues in a familiar environment. Even if some of the 'remembered' cues are removed or displaced, place cells continue to fire in positions that closely fit the expected configuration of cues. Put a different way: after learning, the model's synapses effectively encode the position of objects relative to the environment. This process happens as a result of ongoing plasticity at the level of individual synapses, using a biologically plausible learning rule that does not require an explicit learning signal. In many ways then, the model describes an allocentric memory system similar to that proposed by O'Keefe and Nadel (1978); it learns during exploration, is stable with time, and encodes position relative to a constellation of cues (locale navigation). Furthermore, given such a system it is possible to specify a network that can re-experience an egocentric view from a specific position within the environment. For example, Burgess, Becker, King and O'Keefe (2001) described a model in which an allocentric representation stored in the hippocampus was accessed as an egocentric view after undergoing coordinate transform in the parietal cortex.

Currently the model deliberately avoids the question of remapping (Muller et al., 1987a; Bostock et al., 1991). Indeed, the gradual evolution of simulated place fields seems very different to the rapid reorganisation of firing typical of global remapping. Are the slow, cumulative changes reported by Lever *et al.* (2002b) and others, any more approachable? Simulated cells do exhibit several characteristics seen during slow remapping: individual fields undergo independent, iterative changes in firing rate and to a lesser extent, position. As such, cumulative change in the efficacy of synapses between place cells and their inputs do seem to capture some aspects of slow remapping. That said, I did not present simulations of environments that are known to induce remapping experimentally. For example, Lever *et al.* described a slow divergence between place cell representations recorded in a square and circular environment. Can the updated model address this specific result? The answer would seem to be 'yes' but with one important modification. In the simulations presented in chapter three, BVCs learnt to respond differentially to extended cues with distinct sensory qualities. However, in Lever *et al.'s* experiment the same morph box was used

to create the circular and square environment. Hence, although the geometry of the environment changed, the specific qualities of the walls did not. A possible solution is that we allow the characteristics of an environment as a whole to influence the firing of BVCs. In this example then, each BVC set would consist of two BVCs, one which responded in the circular environment and one in the square. As such place fields would initially be similar in the two environments but would start to diverge as learning occurred. The idea that geometry can function as a contextual cue is not new; Jeffery and Anderson (2003) have already suggested that Lever's environments might be treated as distinct contexts. The same authors also speculated that BVC activity might be modulated by context but did not suggest a mechanism.

In effect then, it seems that the updated BVC model might be applicable to some forms of remapping, including those induced by contextual changes (e.g. odour, texture, light level). Still, without conducting specific simulations it is difficult to know how productive this approach would be. Furthermore, it seems unlikely that the current model can capture the global changes in place cell firing that seem to be associated with realignment of entorhinal grids (Fyhn et al., 2007).

5.2 Sensory control of grid cells

Grid cells in the dorso-lateral medial entorhinal cortex (dlMEC) of freely moving rats show regular grid-like patterns of firing across the environment (Hafting et al., 2005). These are thought to provide an absolute metric whereby an animal can update its own location using self-motion information (O'Keefe et al., 2005; Fuhs et al., 2006; McNaughton et al., 2006; Sargolini et al., 2006; Hargreaves et al., 2005). In contrast to previous reports (Hafting et al., 2005) I described a strong experience-dependent environmental influence: the grids' spatial scales vary parametrically with changes to the size and shape of a familiar environment. Thus it seems that grid cell firing reflects an interaction between intrinsic, path-integrative, calculation of location, and learned associations to the external environment. The 'balance' of this interaction was shown to vary with repeated experience of the probe configurations, suggesting a tendency for the system to revert to an intrinsic grid scale. Interestingly, grids from the same animal appear to be aligned and have fixed relative sizes; observations consistent with grids being optimised to encode position (Brookings et al., 2006; Gorchetchnikov et al., 2007). I also examined the evidence for O'Keefe and Burgess's (2005) proposal that the association of sensory information to grid cells is mediated via place cells. I concluded that the present experiment was consistent with an intimate association between the two cell types but that it did not demonstrate causality. However, I presented other evidence to support the notion, not least that inactivation of the hippocampus causes grids to become unstable (Bonnevie et al., 2006).

How do these findings relate to models of grid cell firing? Unsurprisingly all the published models of grid cell firing focus on the regular, hexagonal firing pattern typical of these cells (O'Keefe et al., 2005; Fuhs et al., 2006; McNaughton et al., 2006; Blair et al., 2007; Burgess, 2007). With the exception of Blair *et al.*, each of these models see grid firing as the product of a network performing path integration. As such they accept, or at least tacitly accept, the necessity of sensory input to maintain the grids' alignment with sensory cues (Hafting et al., 2005; Etienne et al., 1996). For example, O'Keefe and Burgess (2005) and Burgess, Barry and O'Keefe (2007) explicitly described a mechanism by which place cells could reset the dual oscillators on which their model is predicated. It can be seen then, that any manipulation that changes the position of place fields would be expected to induce a complementary change in grids. Thus it seems that given sufficient reset points, grids produced by the dual

oscillator model would exhibit rescaling broadly similar to that seen experimentally; simulations would be necessary to estimate the exact magnitude of rescaling. The other models however, did not describe precise mechanisms by which grid alignment could be maintained with the environment. Hence it is impossible to say whether they predict rescaling. Nevertheless based on the results of this experiment it is possible to specify two requirements for the 'reset' mechanism. First, because grids appear to rescale continuously across the environment, the reset mechanism should also function across the environment. For example, it is not sufficient for grids to be realigned with sensory cues when the animal returns to a single, specific position. Second, because grids did not rescale exactly proportionate to the changes made to the enclosure, the reset mechanism should not be absolute. That is, it should not replace entirely the animal's estimate of its position based on path integration.

Clearly an important question to be addressed in future research is the nature of the associations between grids and the environment. Although I have demonstrated that grid firing is partially controlled by environmental geometry the origin of that signal remains unclear. Connectivity and electrophysiological data (Fyhn et al., 2007) suggest an intimate association between grids and place cells, and one cell type might mediate sensory input to the other. Indeed both eventualities are currently plausible: sensory information, specifically a representation of the distance and direction to landmarks, might reach the hippocampus via grid cells; alternatively it might arrive via the LEC and only reach grids through hippocampal projection to the MEC.

5.3 Synergies between these findings

In this final section I will look at the results from both experiments in parallel. In particular I want to address two issues: First, are the two sets of results reported here complementary? Second, when considered together, what are the implications of these results, what can they tell us about the way space is represented in the brain?

In chapter three, and again in the preceding text, I suggested that the BVC model's failure to support place cell firing after removal of specific cues might be addressed by the inclusion of path integration into the model; grid cells present an obvious means by which this might be achieved. Models of grid cell firing already exist and several authors have shown that grids can interact to produce receptive fields similar to place fields (Soldstad et al., 2006; Blair et al., 2007). As such it is possible to imagine how these ideas might be combined with the BVC model to describe place cells as a function of sensory and path integrative information. Indeed one goal of future research should be to develop such a model and verify its performance in Rivard *et al's* (2004) and Fenton *et al's* (2000a) experiments. Clearly such work is far from trivial, in some situations, such as Rivard *et al's* translocation of a barrier, BVC firing would have to dominate grid cell firing. Conversely, when a cue card was removed in Fenton *et al's* experiment, place cell firing would apparently have to be driven more strongly by grid cells. This problem does not seem insurmountable and a possible solution was expressed in the discussion of chapter four. In that chapter I briefly touched on an expectation that grid rescaling on a linear track might take a form similar to that reported in place cells by Gothard *et al.* (1996); place fields were positioned according to path integration at the start of a run but with distance traveled were increasingly defined by external sensory cues. Hence, it seems that with respect to place cells, the relative potency of path integrative and sensory systems can vary dynamically depending on the animal's position.

What might be the mechanism underlying this effect? An intriguing possibility is that the hippocampus is able to combine information about an animal's location weighted according to its perceived accuracy, and as a result extract a 'best guess' as to the animal's true position. There is evidence for this view, for example, Deneve *et al.* (1999) showed that a recurrent network is able to extract a maximum likelihood estimate of the variables encoded in noisy signals received from other brain areas. In a

similar vein, Touretzky *et al.* (2005) implemented a recurrent CA3 network in their model of place cell firing; given conflicting information the network essentially constrained place cells to fire at the animal's most likely position. It seems that a similar framework might address the concerns expressed above: in the absence of sensory drive (e.g. after cue removal in Fenton *et al's* experiment) the model's estimate of position would be dictated by grid cells; when sensory information was still present though (e.g. after rotation of the barrier in Rivard *et al's* experiment), location would be defined more strongly in terms of the cues surrounding the rat.

This framework might also explain a particular result noted by O'Keefe and Burgess (1996). In their experiment, expansion of the recording enclosure caused a number of place fields to become bimodal, in at least one case a single peak of the bimodal field was favored on each run (Fig 5.1). Interestingly the peak that exhibited the highest firing rate was always the one the rat encountered first. Hence, when running from East to West the Eastern peak was most evident. The standard BVC model predicts that place fields can become bimodal after expansion of an enclosure but does not suggest a directional modulation. In terms of the framework described above, place cell activity would represent both external sensory input (presumably mediated via BVCs), as well as path integrative information received from grid cells. Assuming the path integrator is reset by contact with the wall. Then, as the rat starts its run from East to West, the grid cells will drive the place cell to fire at a distance from the wall equal to the place field's position relative to that wall in the smaller, familiar enclosure. In effect the BVCs specify that the place field should be bimodal but the grid cells specify that the place cell should only fire when the animal encounters the first peak. The outcome is that the cell fires most strongly at the position where both sources of information agree.

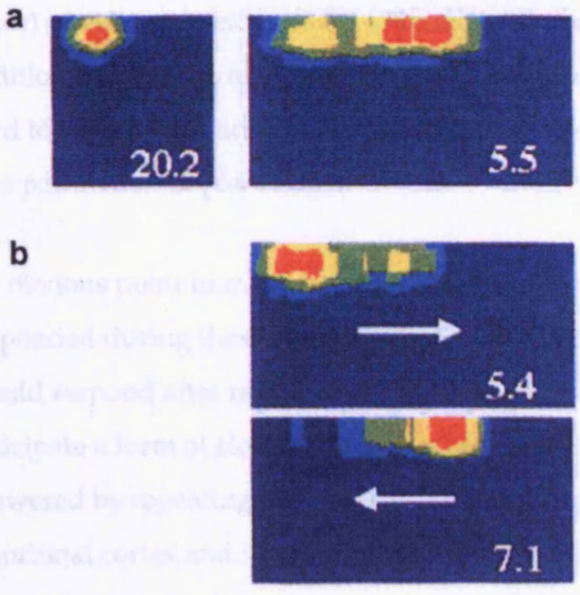

Fig 5.1. Bimodal place field produced by expansion of the recording environment, image adapted from O'Keefe and Burgess (1996). **a**) The same place field recorded in a small square enclosure (left) and horizontal rectangle (right). Elongation in the horizontal plane induces bimodality. Firing rate in Hz is indicated in white. **b**) Data from the horizontal rectangle. Top ratemap shows runs in an Easterly direction, bottom ratemap shows runs in a Westerly direction. In each case the peak that the animal encounters first is most pronounced.

Can a similar model address the results from chapter four; that grids rescale when the geometry of a familiar environment is altered? The standard BVC model accurately describes that movement of place fields under similar manipulations. Hence, following O'Keefe and Burgess's (2005) suggestion that place cells mediate sensory attachment to grids we can rephrase the question as: is grid rescaling consistent with the changes in place fields observed after geometric manipulation of an environment? Generally speaking the answer would seem to be 'yes'; place fields exhibit parametric changes in position and size consistent with that seen in grid cells. However, it is possible to identify specific examples that seem to be problematic. For example, as mentioned above place fields can become bimodal in an enlarged environment, grid cells do not appear to exhibit an analogous effect. In addition, place fields close to the edge of an enclosure are largely unaffected by environmental manipulations, they maintain their position relative to the adjacent walls (see O'Keefe and Burgess (1996b) and Barry *et al.* (2006) for examples). This last point is important as it would seem to imply that grid peaks close to the edge of an enclosure should also maintain their position relative to the walls. However, some grids that showed limited rescaling (e.g.

grids recorded after an animal had been subject to the experimental protocol several times) seem to gain extra peaks after expansion of the environment; effectively additional peaks are revealed when the enclosure is expanded. In essence it seems hard to understand how grids can exhibit less rescaling than that apparently exhibited by a population of place fields.

An obvious point to make is that we do not know specifically how place fields responded during these manipulations. In particular, it is unclear how place cells would respond after repeated exposure to the probe enclosures; plausibly we might anticipate a form of slow remapping (Lever et al., 2002b). These questions could be answered by repeating the experiment and making simultaneous recording from the entorhinal cortex and hippocampus. Nevertheless, the ideas expressed above also seem to be relevant here. Several models of grid cell firing incorporate the idea of a recurrent network connecting grids with similar spacing (Fuhs et al., 2006; McNaughton et al., 2006; Witter et al., 2006; Kumar et al., 2007). As such, the firing of grid cells likely represents the integrated contribution of a wide range of inputs, such as self motion cues and sensory cues. Put simply, grid firing is not dictated by a single place field. In other words, if after environmental enlargement, sufficient place fields retain a single firing field then grid firing is unlikely to be fragmented. However, in a situation where most place fields exhibit dual fields, for example in Skaggs and McNaughton's (1998) joined environments, then grid firing would also be expected to be duplicated in the two environments. Interestingly, it is not clear what effect insertion of a centrally placed barrier would have on grids. The recordings reported in chapter three suggest that when a barrier is first introduced about one third of place cells exhibit duplicate fields; this may be sufficient to cause grid firing to be replicated, at least initially, on either side of the barrier. What of the relationship between grids and place fields positioned at the edge of the enclosure? This is more difficult to resolve verbally and certainly seems to relate to the rat specific reduction in rescaling seen with continued exposure to the protocol. Superficially though, it appears that because grids are positioned according to information derived from path integrative and spatial cues. Hence, it is not problematic that grids do not always align exactly with place fields at the edge of the enclosure.

The previous paragraphs have effectively outlined two separate models: the first saw place cells as the product of BVCs and grid cells; the second described grid cells as the product of self-motion cues and, via place cells, BVCs. Are both arrangements mutually compatible? To a certain extent I have already dealt with this question in the discussion of chapter four. I concluded that grid cells and place cells might well be part of a circuit such that grid cells project to place cells and place cells ultimately project back to grid cells. Hence self-motion cues are integrated by grid cells and thus place cells have access to path integrative information. I was not able to conclude whether sensory information reached place cells directly, perhaps from the LEC, or if it was routed via grid cells. Now, in view of the previous discussion it is to this last point that I want to return and pose a final question: Are grid cells BVCs? Clearly the multi-peaked firing pattern of a single grid cell is unlike that of a putative BVC. However, the fact that grid firing is sensory bound and not purely a path integration metric leaves the possibility that grids are functionally equivalent to BVCs. Indeed, Burgess *et al.* (2007) suggested that directional grids in the deeper layers of the MEC might be analogous to BVCs if they were reset by contact with a boundary behind the animal as it moved in the direction of preferred firing. Nevertheless, as I concluded previously, it is currently not possible to identify the point at which sensory information about the distance and direction to landmarks enters the hippocampal processing loop.

5.4 Conclusions

In conclusion, the findings presented here suggest a plausible mechanism by which the hippocampus and entorhinal cortex can represent specific locales both in terms of their relative positions and of their positions relative to the surrounding cues. Understanding the precise role of these two regions, how they interact, and how that relationship is affected by experience will be critical to understanding spatial memory and human episodic memory.

References

Agnihotri, N. T., Hawkins, R. D., Kandel, E. R., & Kentros, C. (2004). The long-term stablility of new hippocampal place fields requires new protein synthesis. *Proceedings of the National Academy of Sciences of the United States of America* **101**[10], 3656-3661.

Alonso, A. & Linas, R. R. (1989). Subthreshold Na+ - dependent theta-like rhythmicity in stellate cells of entorhinal cortex layer II. *Nature,* **342**, 175-177.

Amaral, D. G. & Lavenex, P. (2007). Hippocampal neuroanatomy. In *The hippocampus book* (pp. 37-114). Oxford University Press.

Amaral, D. G. & Witter, M. P. (1989). The three-dimensional organization of the hippocampal formation: A review of anatomical data. *Neuroscience* **31**[3], 571-591.

Amaral, D. G. & Witter, M. P. (1995). Hippocampal Formation. In P.Andersen, R. G. M. Morris, D. G. Amaral, T. V. Bliss, & J. O'Keefe (Eds.), *The Rat Nervous System* (**2nd** ed., pp. 443-486). Academic Press.

Anderson, M. I. & Jeffery, K. J. (2003). Heterogeneous modulation of place cell firing by changes in context. *The Journal of Neuroscience* **23**[26], 8827-8835.

Barnes, C. A., Suster, M. S., Shen, J., & McNaughton, B. L. (1997). Multistability of cognitive maps in the hippocampus of old rats. *Nature,* **388**, 272-275.

Barry, C. & Burgess, N. (2007). Learning in a geometric model of place cell firing. *Hippocampus* **In press**.

Barry, C., Hayman, R., Burgess, N., & Jeffery, K. J. (2007). Experience-dependent rescaling of entorhinal grids. *Nature Neuroscience,* **10**, 682-684.

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. *Reviews in the Neurosciences* **17**[1-2], 71-97.

Bayer, S. A., Yackel, J. W., & Puri, P. S. (1982). Neurons in the rat dentate gyrus granular layer substantially increase during juvenile and adult life. *Science,* **216**, 890-892.

Best, P. J. & Thompson, L. T. (1984). Hippocampal cells which have place field activity also show changes in activity during classical conditioning. *Society for Neuroscience Abstract* **10**[125].

Best, P. J. & Thompson, L. T. (1989a). Persistence, reticence, and opportunism of place-field activity in hippocampal neurons. *Psychobiology* **17**, 230-235.

Best, P. J. & Thompson, L. T. (1989b). Persistence, reticence, and opportunism of place-field activity in hippocampal neurons. *Psychobiology,* **Vol 17**, 230-235.

Bienenstock, E. L., Cooper, L. N., & Munro, P. W. (1982). Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience* **2**[1], 32-48.

Bingman, V. P., Siegel, J. J., Gagliardo, A., & Erichsen, J. T. (2006). Representing the richness of avian spatial cognition: properties of a lateralized homing pigeon hippocampus. *Reviews in the Neurosciences* **17**[1-2], 17-28.

Blackstad, T. W. (1956). Commissural connections of the hippocampal region in the rat, with special reference to their mode of termination. *Journal of Comparative Neurology,* **105**, 417-537.

Blair, H. T., Welday, A. C., & Zhang, K. (2007). Scale-invariant memory representations emerge from Moire interference between grid fields that produce theta oscillations: a computational model. *Journal of Neuroscience,* **27**, 3211-3229.

Bliss, T. V. & Lomo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *The Journal of Physiology* **232**[2], 331-356.

Bonnevie, T., Fyhn, M., Hafting, T., Moser, E. I., & Moser, M. B. (2006). Misalignment of entorhinal grid fields after hippocampal inactivation. *SfN Abstract* **68.1/BB9**.

Bostock, E., Muller, R. U., & Kubie, J. L. (1991). Experience-dependent modifications of hippocampal place cell firing. *Hippocampus* **1**[2], 193-206.

Brookings, T., Burak, Y., & Fiete, I. R. (2006). Triangular lattice neurons (grid cells) may encode rat position using an advanced numeral system. */arXiv.org/q-bio.NC/0606005*.

Brun, V. H., Otnaess, M. K., Molden, S., Steffenach, H., Witter, M. P., Moser, M. et al. (2002). Place cells and place recognition maintained by direct entorhinal-hippocampal circuitry. *The Journal of Neuroscience* **296**, 2243-2246.

Brunel, N. & Trullier, O. (1998). Plasticity of directional place fields in a model of rodent CA3. *Hippocampus,* **8**, 651-665.

Burgess, N. (2007). Computational models of the spatial and mnemonic functions of the hippocampus. In P.Andersen, R. G. M. Morris, D. G. Amaral, T. V. Bliss, & J. O'Keefe (Eds.), *The hippocampus book* (pp. 715-749). Oxford: Oxford University Press.

Burgess, N., Barry, C., & O'Keefe, J. (2007). An oscillatory interference model of grid cell firing. *Hippocampus,* **In press**.

Burgess, N., Becker, S., King, J. A., & O'Keefe, J. (2001). Memory for events and their spatial context: models and experiments. *Philos.Trans.R Soc Lond B Biol Sci,* **356**, 1493-1503.

Burgess, N., Cacucci, F., Lever, C., & O'Keefe, J. (2005). Characterizing multiple independent behavioral correlates of cell firing in freely moving animals. *Hippocampus* **15**[2], 149-153.

Burgess, N. & Hartley, T. (2002). Orientational and geometric determinants of place and head-direction. In T.G.Dietterich, S. Becker, & Z. Ghahramani (Eds.), (14 ed., pp. 165-172). MIT Press.

Burgess, N., Jackson, A., Hartley, T., & O'Keefe, J. (2000). Predictions derived from modelling the hippocampal role in navigation. *Biological Cybernetics* **83**, 301-312.

Burgess, N. & O'Keefe, J. (2002). Hippocampus: Spatial models. In *The handbook of brain theory and neural networks* (2nd ed., pp. 539-543). MIT Press.

Burgess, N. & O'Keefe, J. (1996b). Cognitive graphs, resistive grids and the hippocampal representation of space. *The Journal of General Physiology* **107**, 659-662.

Burgess, N. & O'Keefe, J. (1996a). Neuronal computations underlying the firing of place cells and their role in navigation. *Hippocampus,* **6**, 749-762.

Burgess, N. & O'Keefe, J. (1996c). 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.

Burwell, R. D. & Amaral, D. G. (1998a). Perirhinal and postrhinal cortices of the rat: interconnectivity and connections with the entorhinal cortex. *Journal of Comparative Neurology*, **391**, 293-321.

Burwell, R. D. & Amaral, D. G. (1998b). Cortical afferents of the perirhinal, postrhinal, and entorhinal cortices of the rat. *Journal of Comparative Neurology*, **398**, 179-205.

Burwell, R. D., Witter, M. P., & Amaral, D. G. (1995). Perirhinal and postrhinal cortices of the rat: a review of the neuroanatomical literature and comparison with findings from the monkey brain. *Hippocampus* **6**[3], 390-408.

Cacucci, F., Lever, C., Wills, T. J., Burgess, N., & O'Keefe, J. (2004). Theta-modulated place-by-direction cells in the hippocampal formation of the rat. *The Journal of Neuroscience* **24**[38], 8265-8277.

Chicurel, M. E. & Harris, K. M. (1992). Three-dimensional analysis of the structure and composition of CA3 branched dendritic spines and their synaptic relationship with mossy fiber boutons in the rat hippocampus. *Journal of Comparative Neurology*, **325**, 169-182.

Cressant, A., Muller, R. U., & Poucet, B. (1997). Failure of centrally placed objects to control the firing fields of hippocampal place cells. *J.Neurosci.*, **17**, 2531-2542.

Deneve, S., Latham, P. E., & Pouget, A. (1999). Reading population codes: a neural implementation of ideal observers. *Nature Neuroscience*, **2**, 740-745.

Ekstrom, A. D., Kahana, M. J., Caplan, J. B., Fields, T. A., Isham, E. A., Newman, E. L. et al. (2003). Cellular networks underlying human spatial navigation. *Nature* **425**, 184-187.

Etienne, A. S. & Jeffery, K. J. (2004). Path integration in mammals. *Hippocampus* **14**, 180-192.

Etienne, A. S., Maurer, R., & Seguinot, V. (1996). Path integration in mammals and its interaction with visual landmarks. *Journal of Experimental Biology* **199**, 201-209.

Fenton, A. A., Csizmadia, G., & Muller, R. U. (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. *The Journal of General Physiology* **116**, 211-221.

Fenton, A. A., Csizmadia, G., & Muller, R. U. (2000a). Conjoint control of hippocampal place cell firing by two visual stimuli I. The effects of moving the stimuli on firing field positions. *Journal of General Physiology* **116**, 191-209.

Fisher, N. I. (1996). *Statistical analysis of circular data.* Cambridge University Press.

Frank, L. M., Stanley, G. B., & Brown, E. N. (2004). Hippocampal plasticity across multiple days of exposure to novel environments. *Journal of Neuroscience* **24**, 7681-7689.

Fuhs, M. C. & Touretzky, D. S. (1999). Synaptic Learning Models of Map Separation in the Hippocampus. *Preprint submitted to Elsevier Preprint* , 1-6.

Fuhs, M. C. & Touretzky, D. S. (2006). A spin glass model of path integration in rat medial entorhinal cortex. *Journal of Neuroscience* **26**[16], 4266-4276.

Fyhn, M., Hafting, T., Treves, A., Moser, M. B., & Moser, E. I. (2007). Hippocampal remapping and grid realignment in the entorhinal cortex. *Nature* **446**, 190-194.

Fyhn, M., Molden, S., Witter, M. P., Moser, E. I., & Moser, M. (2004). Spatial representation in the entorhinal cortex. *Science* **305**, 1258-1264.

Germroth, P., Schwerdtfeger, W. K., & Buhl, E. H. (1989). Morphology of identified entorhinal neurons projecting to the hippocampus: A light microscopical study combining retrograde tracing and intracellular injection. *Neuroscience,* **30**, 683-691.

Giocomo, L. M., Zilli, E. A., Fransen, E., & Hasselmo, M. E. (2007). Temporal frequency of subthreshold oscillations scales with entorhinal grid cell field spacing. *Science,* **315**, 1719-1722.

Gloor, P. (1997). The hippocampal system. In *The temporal lobe and limbic system* (pp. 325-589). Oxford: Oxford Universtiy Press.

Gorchetchnikov, A. & Grossberg, S. (2007). Space, time and learning in the hippocampus: How fine spatial and temporal scales are expanded into population codes for behavioural control. *Neural Networks,* **20**, 182-193.

Gothard, K. M., Hoffman, K. L., Battaglia, F. P., & McNaughton, B. L. (2001). Dentate gyrus and CA1 ensemble activity during spatial reference frame shifts in the presence and absence of visual input. *The Journal of Neuroscience* **21**[18], 7284-7292.

Gothard, K. M., Skaggs, W. E., & McNaughton, B. L. (1996). Dynamics of mismatch correction in the hippocampal ensemble code for space: Interaction between path integration and environmental cues. *Journal of Neuroscience* **16**[24], 8027-8040.

Guzowski, J. F., McNaughton, B. L., Barnes, C. A., & Worley, P. F. (1999). Environment-specific expression of the immediate-early gene Arc in hippocampal neuronal ensembles. *Nature Neuroscience,* **2**, 1120-1124.

Hafting, T., Fyhn, M., Moser, M., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature* **436**, 801-806.

Hafting, T., Fyhn, M., Moser, M. B., & Moser, E. I. (2006). Phase precession and phase locking in entorhinal grid cells. *SfN Abstract* **68.8/BB16**.

Hargreaves, E. L., Roa, G., Lee, I., & Knierim, J. J. (2005). Major dissociation between medial and lateral entorhinal input to dorsal hippocampus. *Science* **308**, 1792-1794.

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**[1], 39-75.

Hayman, R. A., Chakraborty, S., Anderson, M. I., & Jeffery, K. J. (2003). Context-specific acquisition of location discrimination by hippocampal place cells. *European Journal of Neuroscience* **18**, 2825-2834.

Hetherington, P. A. & Shapiro, M. L. (1997). Hippocampal place fields are altered by the removal of single visual cues in a distance-dependent manner. *Behavioral Neuroscience* **111**[1], 20-34.

Hill, A. J. (1978). First occurrence of hippocampal spatial firing in a new environment. *Experimental Brain Research* **62**, 282-297.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the United States of America* **79**, 2554-2558.

Hori, E., Tabuchi, E., Matsumura, N., Tamura, R., Eifuku, S., Endo, S. et al. (2003). Representation of place by monkey hippocampal neurons in real and virtual translocation. *Hippocampus,* **13**, 190-196.

Insausti, R., Herrero, M. T., & Witter, M. P. (1997). Entorhinal cortex of the rat: cytoarchitectonic subdivisions and the origin and distribution of cortical efferents. *Hippocampus,* **7**, 183.

Ishizuka, N., Weber, J., & Amaral, D. G. (1990). Organization of intrahippocampal projections originating from CA3 pyramidal cells in the rat. *Journal of Comparative Neurology,* **295**, 580-623.

Jeffery, K. J., Anand, R. L., & Anderson, M. I. A role for terrain slope in orienting hippocampal place fields. *Exp.Brain Res.,* (in press).

Jeffery, K. J. & Anderson, M. I. (2003). Dissociation of the Geometric and Contextual Influences on Place Cells. *Hippocampus* **13**, 868-872.

Jeffery, K. J., Gilbert, A., Burton, S., & Strudwick, A. (2003). Preserved Performance in a Hippocampal-Dependent Spatial Task Despite Complete Place Cell Remapping. *Hippocampus* **13**, 175-189.

Jeffery, K. J. & O'Keefe, J. M. (1999). Learned interaction of visual and idiothetic cues in the control of place field orientation. *Experimental Brain Research* **127**, 151-161.

Jung, M. W. & McNaughton, B. L. (1993). Spatial selectivity of unit activity in the hippocampal granular layer. *Hippocampus* **3**[2], 165-182.

Kali, S. & Dayan, P. (2000). The involvement of recurrent connections in area CA3 in establishing the properties of place fields: a model. *The Journal of Neuroscience* **20**[19], 7463-7477.

Kee, N., Teixeria, C. M., Wang, A. H., & Frankland, P. W. (2007). Preferential incorporation of adult-generated granule cells into spatial memory networks in the dentate gyrus. *Nature Neuroscience,* **10**, 355-362.

Kentros, C., Hargreaves, E., Hawkins, R. D., Kandel, E. R., Shapiro, M., & Muller, R. V. (1998). Abolition of Long-Term Stability of New Hippocampal Place Cell Maps by NMDA Receptor Blockade. *Science* **280**, 2121-2126.

Klink, R. & Alonso, A. (1997). Morphological characteristics of layer II projection neurons in the rat medial entorhinal cortex. *Hippocampus* **7**, 571-583.

Kloosterman, F., Van Haeften, T., & Lopes da Silva, F. H. (2004). Two reentrant pathways in the hippocampal-entorhinal system. *Hippocampus,* **14**, 1026-1039.

Kloosterman, F., Van Haeften, T., Witter, M. P., & Lopes da Silva, F. H. (2003). Electrophysiological characterization of interlaminar entorhinal connections: an essential link for re-entrance in the hippocampal-entorhinal system. *European Journal of Neuroscience* **18**[11], 3037-3052.

Knierim, J. J. (2003). Hippocampal remapping: implications for spatial learning and navigation. In K.J.Jeffery (Ed.), (1 ed., pp. 226-239). Oxford: Oxford University Press.

Knierim, J. J. (2002). Dynamic interactions between local surface cues, distal landmarks, and intrinsic circuitry in hippocampal place cells. *Journal of Neuroscience* **22**[14], 6254-6264.

Knierim, J. J., Kudrimoti, H. S., & McNaughton, B. L. (1998). Interactions between idiothetic cues and external landmarks in the control of place cells and head direction cells. *Journal of Neurophysiology* **80**[1], 425-446.

Kohler, C. (1985). Intrinsic projections of the retrohippocampal region in the rat brain. I. The subicular complex. *Journal of Comparative Neurology,* **236**, 504-522.

Kohler, C. (1988). Intrinsic connections of the retrohippocampal region in the rat brain. III. The lateral entorhinal area. *Journal of Comparative Neurology,* **271**, 208-228.

Kohler, C. (1986). Intrinsic connections of the retrohippocampal region in the rat brain. II. The medial entorhinal area. *Journal of Comparative Neurology* **246**, 149-169.

Kumar, S. S., Jin, X., Buckmaster, P. S., & Huguenard, J. R. (2007). Recurrent circuits in layer II of medial entorhinal cortex in a model of temporal lobe epilepsy. *Journal of Neuroscience* **27**[6], 1239-1246.

Lenck-Santini, P. P., Rivard, B., Muller, R. U., & 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, J. K., Leutgeb, S., Treves, A., Fyhn, M., Meyer, R., Barnes, C. A. et al. (2004a). Pattern completion and pattern separation in CA3 during morphing of two environments. *Poster* .

Leutgeb, S., Kjelstrup, K. G., Treves, A., Moser, M., & Moser, E. I. (2003). Differential representation of context in hippocampal areas of CA3 and CA1. *Abstract from Society for Neuroscience* .

Leutgeb, S., Leutgeb, J. K., Moser, E. I., & Moser, E. I. (2006). Fast rate coding in hippocampal CA3 cell ensembles. *Hippocampus* **16**, 765-774.

Leutgeb, S., Leutgeb, J. K., Treves, A., Moser, M. B., & Moser, E. I. (2004b). Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science* **305**[5688], 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? *Biological Cybernetics* **87**, 356-372.

Lever, C., Cacucci, F., Burgess, N., & O'Keefe, J. (1999). Squaring the circle: place cell firing patterns in environments which differ only geometrically are not unpredictable. *Society for Neuroscience Abstract* **24**, 556.

Lever, C., Wills, T. J., Cacucci, F., Burgess, N., & O'Keefe, J. (2002b). Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature* **416**, 90-94.

Li, X. G., Somogyi, P., Ylinen, A., & Buzsaki, G. (1994). The hippocampal CA3 network: An in vivo intracellular labeling study. *Journal of Comparative Neurology,* **339**, 181-208.

Lorente de No, R. (1933). Studies on the structure of the cerebral cortex. *J Psychol Neurol,* **45**, 381-438.

Ludvig, N., Tang, H. M., Gohil, B. C., & Botero, J. M. (2004). Detecting location-specific neuronal firing rate increases in the hippocampus of freely-moving monkeys. *Brain Research* **1014**[1-2], 97-109.

Markus, E. J., Qin, Y., Leonard, B. S. W. E., McNaughton, B. L., & Barnes, A. C. (1995). Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *The Journal of Neuroscience* **15**[11], 7079-7094.

Marr, D. (1971). Simple memory: A theory for archicortex. *Philos.Trans.R Soc Lond B Biol Sci* **262**[841], 23-81.

McNaughton, B. L. & Barnes, C. A. (1977). Pyysiological identification and analysis of dentate granule cell responses to stimulation of the medial and lateral perforant pathways in the rat. *Journal of Comparative Neurology,* **175**, 439-454.

McNaughton, B. L., Barnes, C. A., Gerrard, J. L., Gothard, K., Jung, M. W., Knierim, J. J. et al. (1996). Deciphering the hippocampal polyglot: The hipppocampus as a path integration system. *Journal of Experimental Biology* **199**, 173-185.

McNaughton, B. L., Barnes, C. A., & 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, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., & Moser, M. B. (2006). Path integration and the neural basis of the 'cognitive map'. *Nature Reviews Neuroscience* **7**, 663-678.

McNaughton, B. L., Knierim, J. J., & Wilson, M. A. (1994). Vector encoding and the vestibular foundations of spatial cognition: Neurophysiological and computational mechanisms. In *The cognitive neurosciences* (pp. 585-595). MIT Press, Boston.

McNaughton, B. L. & Nadel, L. (1990). Hebb-Marr networks and the neurobiological representation of action in space. In *Neuroscience and connectionist theory* (pp. 1-63). Hillsdale,N.J.

Mehta, M. R., Barnes, C. A., & McNaughton, B. L. (1997). Experience-dependent asymmetric expansion of hippocampal place fields. *Proceedings of the National Academy of Sciences of the United States of America* **94**, 8918-8921.

Miller, V. M. & Best, P. J. (1980). Spatial correlates of hippocampal unit activity are altered by lesions of the fornix and entorhinal cortex. *Brain Research* **194**[2], 311-323.

Morris, R. G. M., Garrud, P., Rawlins, J. N., & O'Keefe, J. (1982a). Place navigation impaired in rats with hippocampal lesions. *Nature,* **297**, 681-683.

Morris, R. G. M., Garrud, P., Rawlins, J. N. P., & O'Keefe, J. (1982b). Place navigation impaired in rats with hippocampal lesions. *Nature* **297**, 681-683.

Muller, R. U. (1996). A quarter of a century of place cells. *Neuron* **17**, 979-990.

Muller, R. U., Bostock, E., Taube, J. S., & Kubie, J. L. (1994). On the directional firing properties of hippocampal place cells. *Journal of Neuroscience* **14**, 7235-7251.

Muller, R. U. & Kubie, J. L. (1989). The firing of hippocampal place cells predicts the future position of freely moving rats. *The Journal of Neuroscience* **9**[12], 4101-4110.

Muller, R. U. & Kubie, J. L. (1987a). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *The Journal of Neuroscience* **7**[7], 1951-1968.

Muller, R. U., Kubie, J. L., & Ranck, J. B. (1987b). Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *The Journal of Neuroscience* **7**[7], 1935-1950.

Muller, R. U., Stead, M., & Pach, J. (1996). The hippocampus as a cognitive graph. *The Journal of General Physiology* **107**, 663-694.

Nakazawa, K., Quirk, M. C., Chitwood, R. A., Watanabe, M., Yeckel, M. F., Sun, L. D. et al. (2002). Requirement of Hippocampal CA3 NMDA Receptors in Associative Memory Recall. *Science* **297**, 211-218.

O'Keefe J & Conway, D. H. (1978). Hippocampal place units in the freely moving rat: why they fire where they fire. *Experimental Brain Research* **31**, 573-590.

O'Keefe, J. (1976). Place units in the hippocampus of the freely moving rat. *Experimental Neurology* **51**, 78-109.

O'Keefe, J. (1991). The hippocampal cognitive map and navigational strategies. In *Brain and space* (pp. 273-295). Oxford University Press.

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**, 853-866.

O'Keefe, J. & Conway, D. H. (1978). Hippocampal place units in the freely moving rat: why they fire where they fire. *Experimental Brain Research* **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 Research* **34**, 171-175.

O'Keefe, J. & Nadel, L. (1978). *The hippocampus as a cognitive map.* (1st ed.) Oxford: Oxford University Press.

O'Keefe, J. & Recce, M. L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* **3**[3], 317-330.

O'Keefe, J. & Speakman, A. (1987). Single unit activity in the rat hippocampus during a spatial memory task. *Experimental Brain Research* **68**[1], 1-27.

Ono, T., Nakamura, K., Fukuda, M., & Tamura, R. (1991). Place recognition responses of neurons in monkey hippocampus. *Neuroscience Letters* **121**[1-2], 194-198.

Quirk, G. J., Muller, R. U., & Kubie, J. L. (1990). The firing of hippocampal place cells in the dark depends on the rat's recent experience. *The Journal of Neuroscience* **10**[6], 2008-2017.

Quirk, G. J., Muller, R. U., Kubie, J. L., & Ranck, J. B. (1992). The positional firing properties of medial entorhinal neurons: Description and comparison with hippocampal place cells. *The Journal of Neuroscience* **12**[5], 1945-1963.

Recce, M. L. & O'Keefe, J. (1989). The tetrode: a new technique for multiunit extra-cellular recording. *Society for Neuroscience* **15, 1250**.

Redish, A. D. (1999). *Beyond the Cognitive Map: From Place Cells to Episodic Memory*. Cambridge MA: MIT Press.

Rivard, B., Li, Y., Lenck-Santini, P., Poucet, B., & Muller, R. U. (2004). Representation of objects in space by two classes of hippocampal pyramidal cells. *The Journal of General Physiology* **124**, 9-25.

Rolls, E. T. (1989). The representation and storage of information in neuronal networks in the primate cerebral cortex and hippocampus. In *The computing neuron* (pp. 125-159). Addison-Wesley, Workingham.

Rolls, E. T., Miyashita, Y., Cahusac, P. M., Kesner, R. P., Niki, H., Feigenbaum, J. D. et al. (1989). Hippocampal neurons in the monkey with activity related to the place in which a stimulus is shown. *The Journal of Neuroscience* **9**[6], 1835-1845.

Rolls, E. T., Stringer, S. M., & Trappenberg, T. P. (2002). A unified model of spatial and episodic memory. *Proceedings of the Royal Society of London Series B-Biological Sciences,* **269**, 1087-1093.

Samsonovich, A. & McNaughton, B. L. (1997). Path integration and cognitive mapping in a continuous attractor neural network model. *The Journal of Neuroscience,* **17**, 5900-5920.

Sargolini, F., Fyhn, M., Hafting, T., McNaughton, B. L., Witter, M. P., Moser, M. et al. (2006). Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science* **312**, 758-762.

Save, E., Cressant, A., Thinus-Blanc, C., & Poucet, B. (1998). Spatial firing of hippocampal place cells in blind rats. *The Journal of Neuroscience* **18**[5], 1818-1826.

Save, E., Nerad, L., & Poucet, B. (2000). Contribution of multiple sensory information to place field stability in hippocampal place cells. *Hippocampus* **10**, 64-76.

Scharfman, H. E., Witter, M. P., & Schwarcz, R. (2000). The parahippocampal region. Implications for neurological and psychiatrix diseases. Introduction. *Ann.N.Y.Acad.Sci.,* **911**, ix-xii.

Scoville, W. B. & Milner, B. (1957b). Loss of recent memory after bilateral hippocampal lesions. *J.Neurol.Neurosurg.Psychiatry,* **20**, 11-21.

Scoville, W. B. & Milner, B. (1957a). Loss of recent memory after bilateral hippocampal lesions. *J Neurochem* **20**[1], 11-21.

Sharp, P. E. (1991). Computer simulations of hippocampal place cells. *Psychobiology* **19**, 103-115.

Sharp, P. E. (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. *Behavioral Neuroscience* **113**[4], 643-662.

Sharp, P. E. & Green, C. (1994). Spatial correlates of firing patterns of single cells in the subiculum of the freely moving rat. *The Journal of Neuroscience* **14**[4], 2339-2358.

Skaggs, W. E., Knierim, J. J., Kudrimoti, H. S., & McNaughton, B. L. (1995). A model of the neural basis of the rat's sense of direction. *Adv Neural Inf Process Syst.*, **7**, 173-180.

Skaggs, W. E. & McNaughton, B. L. (1998). Spatial firing properties of hippocampal CA1 populations in an environment containing two visually identical regions. *The Journal of Neuroscience,* **18**, 8455-8466.

Soldstad, T., Moser, E. I., & Einevoll, G. T. (2006). From grid cells to place cells: A mathematical model. *Hippocampus,* **16**, 1026-1031.

Somogyi, P. & Klausberger, T. (2005). Defined types of cortical interneurone structure space and spike timing in the hippocampus. *The Journal of Physiology,* **562**, 9-26.

Swanson, L. W. (1983). *The hippocampus and the concept of the limbic system.* Academic Press, New York.

Swanson, L. W., Kohler, C., & Bjorklund, A. (1987). The limbic region. I. The septohippocampal system. In *Handbook of chemical neuroanatomy. Vol 5. Integrated systems of the CNS. Part 1.* (pp. 125-227). Elsevier, Amsterdam.

Swanson, L. W., Sawchenko, P. E., & Cowan, W. M. (1980). Evidence that the commissural, associational and septal projections of the regio inferior of the hippocampus arise from the same neuron. *Brain Research,* **197**, 207-212.

Taube, J. S. & Muller, R. U. (1998). Comparisons of head direction cell activity in the postsubiculum and anterior thalamus of freely moving rats. *Hippocampus,* **8**, 87-108.

Taube, J. S., Muller, R. U., & Ranck, J. B. (1990a). Head-direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations. *The Journal of Neuroscience* **10**[2], 436-447.

Taube, J. S., Muller, R. U., & Ranck, J. B. (1990b). Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *The Journal of Neuroscience* **10**[2], 420-435.

Thompson, L. T. & Best, P. J. (1989). Place cells and silent cells in the hippocampus of freely-behaving rats. *Journal of Neuroscience,* **9**, 2382-2390.

Tolman, E. C. (1948). Cognitive maps in rats and man. *Psychological Review,* **55**, 189-208.

Tonegawa, S., Tsien, J. Z., McHugh, T. J., Huerta, P., Blum, K. I., & Wilson, M. A. (1996). Hippocampal CA1-region-restricted knockout of NMDAR1 gene disrupts synaptic plasticity, place fields, and spatial learning. *Cold Spring Harbor symposia on quantitative biology* **61**, 225-238.

Touretzky, D. S., Weisman, W. E., Fuhs, M. C., Skaggs, W. E., Fenton, A. A., & Muller, R. U. (2005). Deforming the hippocampal map. *Hippocampus* **15**, 41-55.

Trullier, O., Wiener, S. I., Berthoz, A., & Meyer, J. A. (1997). Biologically based artificial navigation systems: review and prospects. *Progress in Neurobiology* **51**[5], 483-544.

Ulanovsky, N. & Moss, C. F. (2007). Hippocampal cellular network activity in freely moving echolocating bats. *Nature Neuroscience* **10**, 224-233.

Wan, H. S., Touretzky, D. S., & Redish, A. D. (1994). A rodent navigation model that combines place code, head direction, and path integration information. *Society for Neuroscience Abstract,* **20**.

Watson, G. S. (1962). Goodness-of-fit tests on a circle. II. *Biometrika* **49**, 57-63.

West, M. J., Slomianka, L., & Gundersen, H. J. G. (1991). Unbiased stereological estimation of the total number of neurons in the subdivisions of the rat hippocampus using the optical fractionator. *Anat Rec,* **231**, 482-497.

Wills, T. J., Lever, C., Cacucci, F., Burgess, N., & O'Keefe, J. (2005). Experience-dependent attractors in the hippocampal representation of the local environment. *Science* **308**, 873-876.

Wilson, M. A. & McNaughton, B. L. (1993). Dynamics of the hippocampal code for space. *Science* **261**, 1055-1058.

Witter, M. P., Groenewegen, H. J., Lopes da Silva, F. H., & Lohman, A. H. M. (1989). Functional organization of the extrinsic and intrinsic circuitry of the parahippocampal region. *Progress in Neurobiology,* **33**, 161-253.

Witter, M. P. & Moser, E. I. (2006). Spatial representation and the architecture of the entorhinal cortex. *Trends in Neurosciences,* **29**, 671-678.

Yoganarasimha, D., Yu, X., & Knierim, J. J. (2006). Head direction cell representations maintain internal coherence during conflicting proximal and distal cue rotations: Comparison with hippocampal place cells. *The Journal of Neuroscience* **26**[2], 622-631.

Zar, J. H. (1999). *Biostatistical Analysis.* (4th ed.) Upper Saddle River, NJ 07458: Prentice Hall.

Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. *The Journal of Neuroscience* **16**[6], 2112-2126.

Zipser, D. (1985). A computational model of hippocampal place fields. *Behavioral Neuroscience,* **99**, 1006-1018.

Zipser, D. (1986). Biologically plausible models of place recognition and goal location. In *Parallel distributed processing Volume 2* (1 ed., pp. 432-470). MIT Press.

Appendix 1

Screenshots from analysis program written by the author in Matlab 7.1 (The MathsWorks, Inc.). Data relating to two 'outlier' grid cells is shown, one cell per page (recorded on 04/04/2006 & 21/04/2006, both from r214). In each image the top row shows ratemaps in the order they were recorded. The sixth ratemap is a composite created from trials one and five; the baseline conditions. Both cells were assessed by the analysis software as showing rescaling in the opposite direction to that predicted by the changes made to the enclosure. We believe this assessment to erroneous; though rescaling in these examples is minimal it does not appear to be negative. In both cases the ratemaps in the probe trials (trials two to four) have few areas of robust firing relative to the baseline. As a result it is likely that the algorithm for assessing rescaling located a spurious match between the probe and baseline ratemaps.

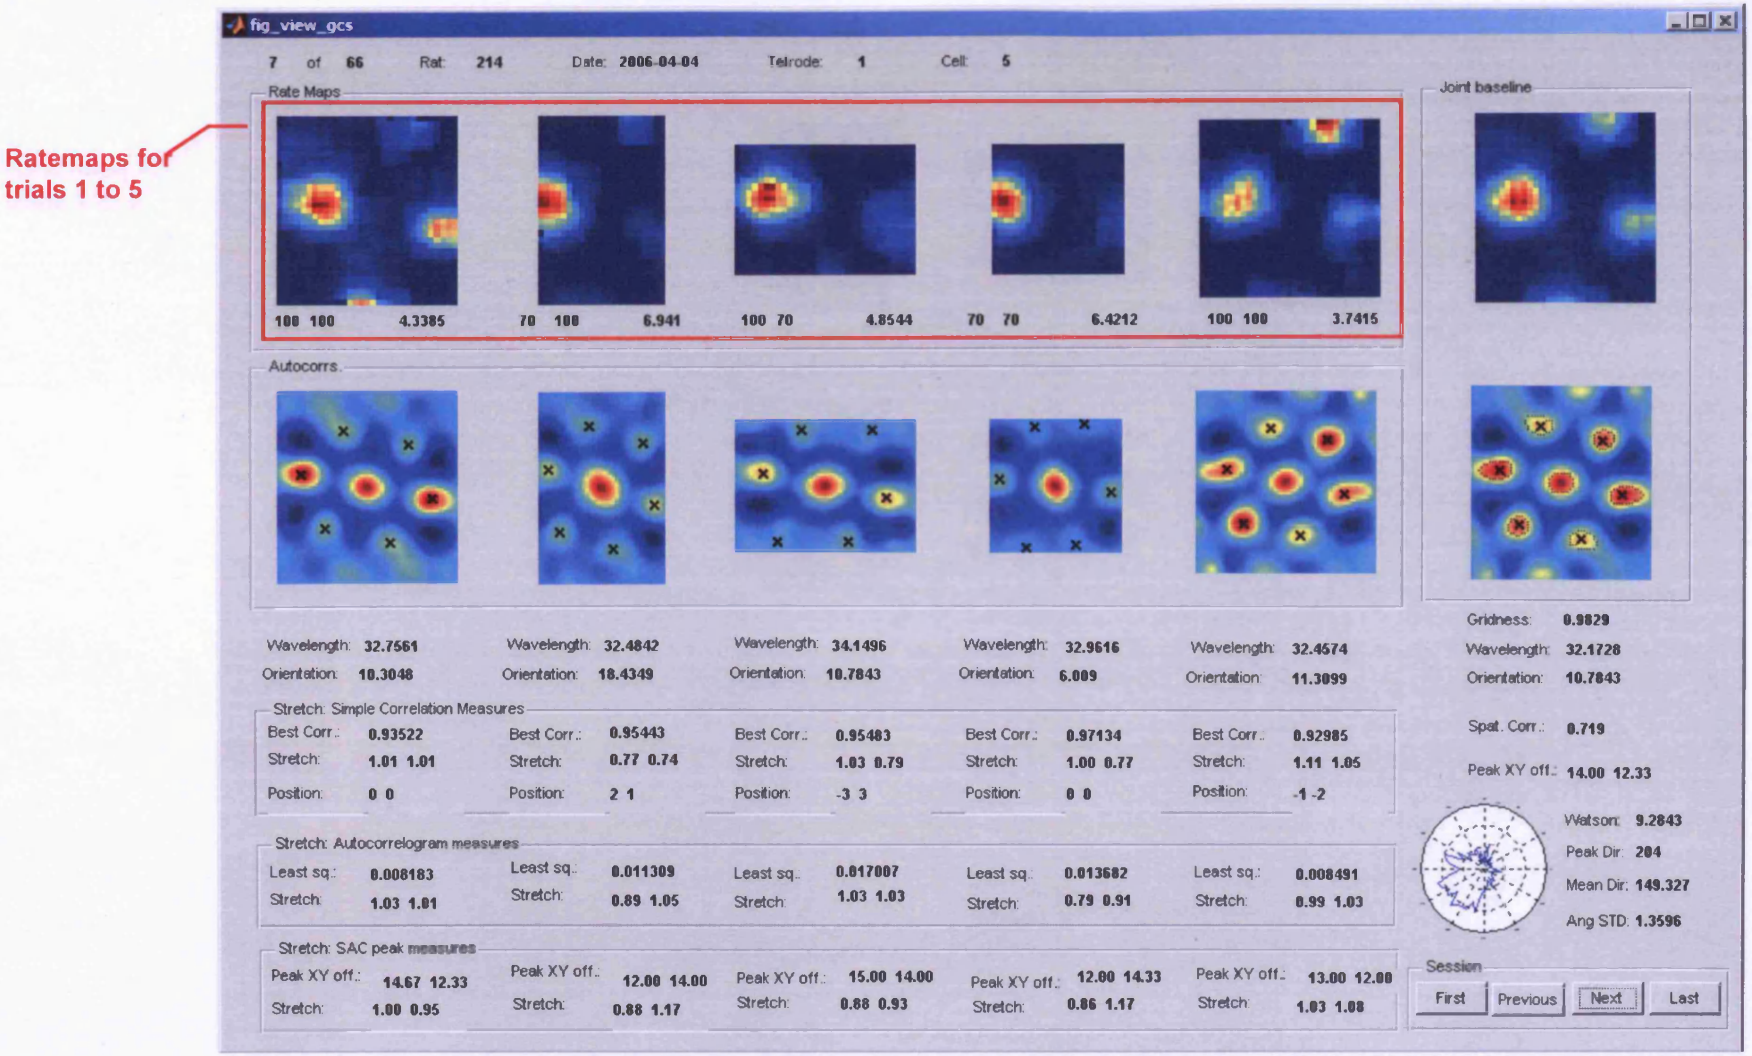
Ratemaps for trials 1 to 5
fig_view_gcs
7 of 66
Rat: 214
Date: 2006-04-04
Tetrode: 1
Cell: 5
Rate Maps
100 100 4.3385
70 100 6.941
100 70 4.8544
70 70 6.4212
100 100 3.7415
Joint baseline
Autocorrs
Wavelength: 32.7561
Orientation: 10.3048
Wavelength: 32.4842
Orientation: 18.4349
Wavelength: 34.1496
Orientation: 10.7843
Wavelength: 32.9616
Orientation: 6.089
Wavelength: 32.4574
Orientation: 11.3099
Gridness: 0.9829
Wavelength: 32.1728
Orientation: 10.7843
Stretch: Simple Correlation Measures
Best Corr.: 0.93522
Stretch: 1.01 1.01
Position: 0 0
Best Corr.: 0.95443
Stretch: 0.77 0.74
Position: 2 1
Best Corr.: 0.95483
Stretch: 1.03 0.79
Position: -3 3
Best Corr.: 0.97134
Stretch: 1.00 0.77
Position: 0 0
Best Corr.: 0.92985
Stretch: 1.11 1.05
Position: -1 -2
Spat. Corr.: 0.719
Peak XY off.: 14.00 12.33
Stretch: Autocorrelogram measures
Least sq.: 0.008183
Stretch: 1.03 1.01
Least sq.: 0.011309
Stretch: 0.89 1.05
Least sq.: 0.017007
Stretch: 1.03 1.03
Least sq.: 0.013682
Stretch: 0.79 0.91
Least sq.: 0.008491
Stretch: 0.99 1.03
Watson: 9.2843
Peak Dir: 204
Mean Dir: 149.327
Ang STD: 1.3596
Stretch: SAC peak measures
Peak XY off.: 14.67 12.33
Stretch: 1.00 0.95
Peak XY off.: 12.00 14.00
Stretch: 0.88 1.17
Peak XY off.: 15.00 14.00
Stretch: 0.88 0.93
Peak XY off.: 12.00 14.33
Stretch: 0.86 1.17
Peak XY off.: 13.00 12.00
Stretch: 1.03 1.08
Session
First
Previous
Next
Last

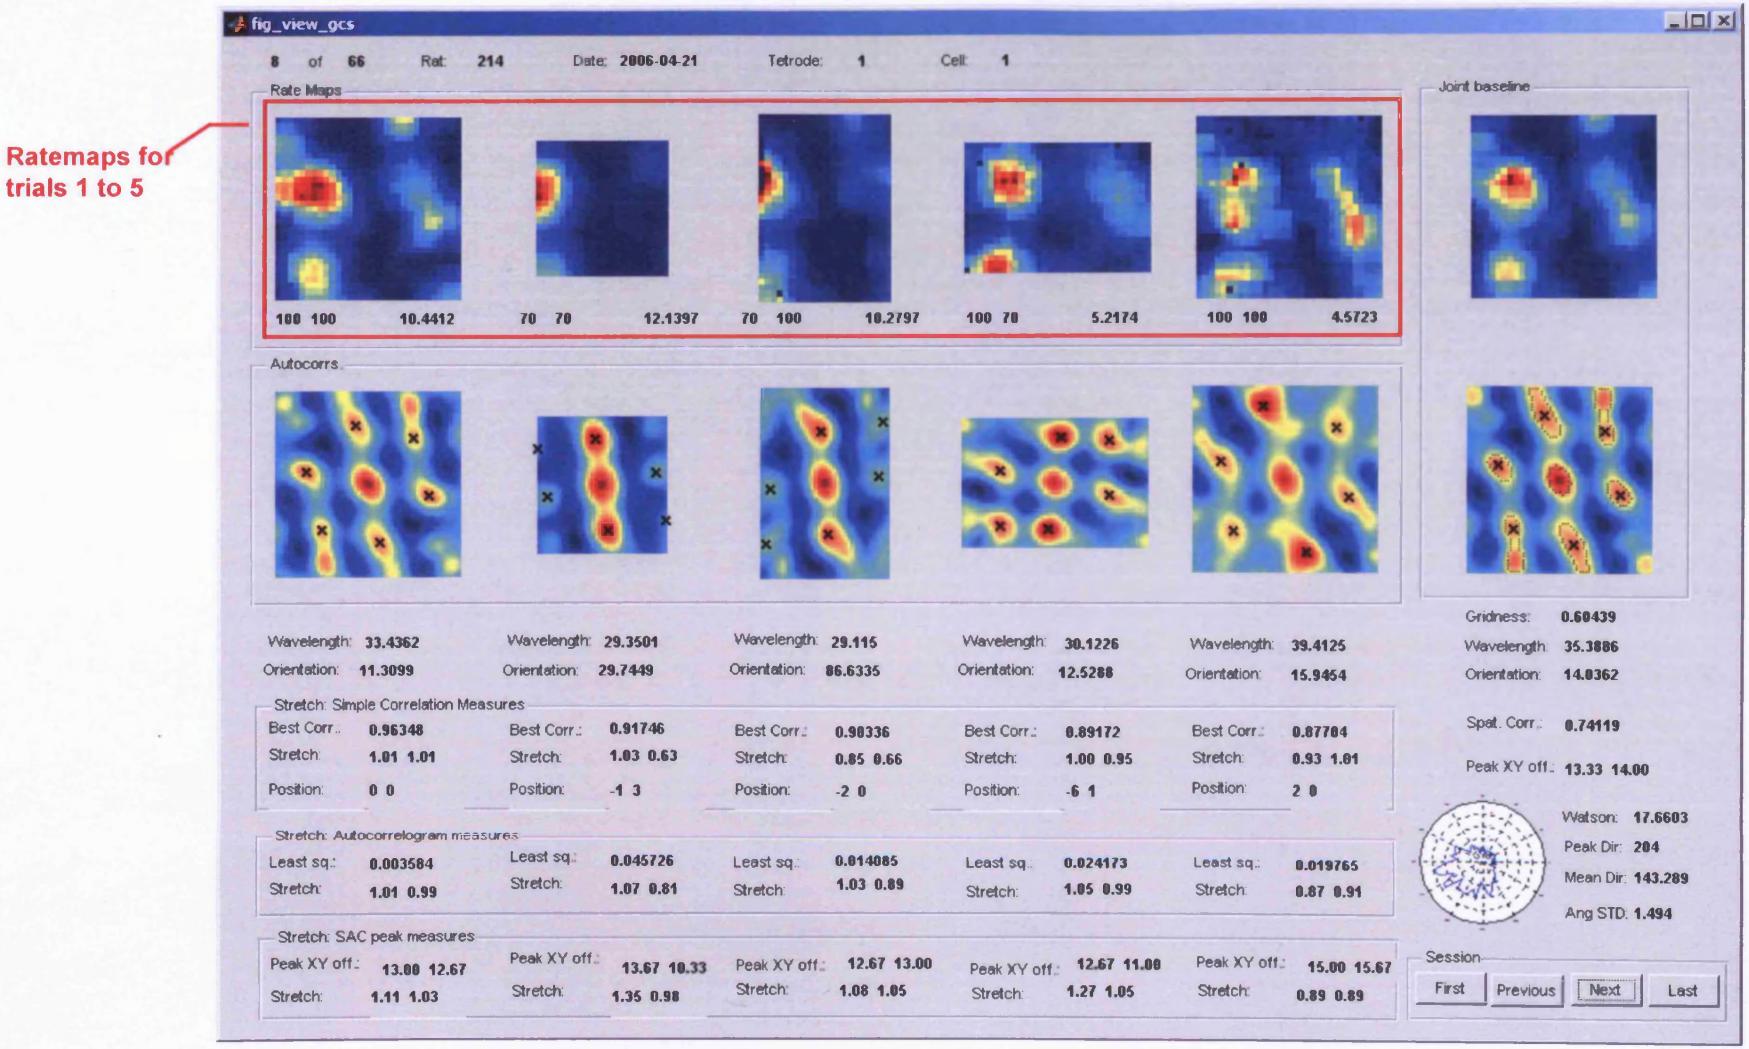
Ratemaps for trials 1 to 5
fig_view_gcs
8 of 66 Rat: 214 Date: 2006-04-21 Tetrode: 1 Cell: 1
Rate Maps
Joint baseline
100 100 10.4412
70 70 12.1397
70 100 10.2797
100 70 5.2174
100 100 4.5723
Autocorrs
Wavelength: 33.4362
Orientation: 11.3099
Wavelength: 29.3501
Orientation: 29.7449
Wavelength: 29.115
Orientation: 86.6335
Wavelength: 30.1226
Orientation: 12.5288
Wavelength: 39.4125
Orientation: 15.9454
Gridness: 0.60439
Wavelength: 35.3886
Orientation: 14.0362
Stretch: Simple Correlation Measures
Best Corr.: 0.96348
Stretch: 1.01 1.01
Position: 0 0
Best Corr.: 0.91746
Stretch: 1.03 0.63
Position: -1 3
Best Corr.: 0.98336
Stretch: 0.85 0.66
Position: -2 0
Best Corr.: 0.89172
Stretch: 1.00 0.95
Position: -6 1
Best Corr.: 0.87784
Stretch: 0.93 1.01
Position: 2 0
Spat. Corr.: 0.74119
Peak XY off.: 13.33 14.00
Stretch: Autocorrelogram measures
Least sq.: 0.003584
Stretch: 1.01 0.99
Least sq.: 0.045726
Stretch: 1.07 0.81
Least sq.: 0.014085
Stretch: 1.03 0.89
Least sq.: 0.024173
Stretch: 1.05 0.99
Least sq.: 0.019765
Stretch: 0.87 0.91
Watson: 17.6603
Peak Dir: 204
Mean Dir: 143.289
Ang STD: 1.494
Stretch: SAC peak measures
Peak XY off.: 13.00 12.67
Stretch: 1.11 1.03
Peak XY off.: 13.67 10.33
Stretch: 1.35 0.98
Peak XY off.: 12.67 13.00
Stretch: 1.08 1.05
Peak XY off.: 12.67 11.00
Stretch: 1.27 1.05
Peak XY off.: 15.00 15.67
Stretch: 0.89 0.89
Session
First
Previous
Next
Last