

Self-motion, environmental, and cholinergic influences on grid cell firing

Francis Carpenter

Supervised by Dr Caswell Barry and Professor Neil Burgess

Thesis submitted to University College London for the degree of Doctor of Philosophy in
Neuroscience

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

Acknowledgements

Publications Arising

The following publications were produced from work undertaken as part of this thesis:

Carpenter, F., D., Manson, K. J. Jeffery, N. Burgess, and C. Barry (2015). *Grid Cells Form a Global Representation of Connected Environments*. *Current Biology* 25, 1176-1182.

Carpenter, F. and C. Barry, (2016). *Distorted Grids as a Spatial Label and Metric*. *Trends in Cognitive Sciences* 20, 164-167.

Ólafsdóttir, H. F., F. Carpenter, and C. Barry (2016). *Coordinated Grid and Place Cell Replay During Rest*. *Nature Neuroscience* 19, 792-794.

Ólafsdóttir, H. F., F. Carpenter, and C. Barry. *Task Demands Predict a Dynamic Switch in the Content of Awake Hippocampal Replay* (Under Review).

Carpenter, F., N. Burgess and C. Barry. *Elevating medial septal cholinergic tone reduces medial entorhinal theta frequency without affecting speed or grid coding* (Under Review).

Abstract

As an animal navigates an environment, grid cells emit action potentials at the vertices of a tessellating triangular pattern. The distance between adjacent vertices varies between subsets of grid cells, with each subset forming a functional 'module'. Combined, this periodic firing and modular organisation is theoretically an extremely efficient means of encoding an animal's location within its environment, and may support goal-directed navigation in enabling the calculation of vectors connecting pairs of locations. The mechanisms by which grid cell firing patterns are generated and their functional contributions to cognition remain obscure however. In this thesis, I present the results of a number of experiments which shed light on these issues.

In the first experiment, I recorded from grid cells in rats exploring an environment containing two perceptually identical compartments connected by a corridor. This environment was designed to place the putative self-location and sensory inputs to grid cells in conflict, allowing inference of the relative contribution of each to the generation of grid cell firing patterns. During early exposures to the environment, firing patterns were replicated between the two compartments, demonstrating the significance of the identical sensory cues. However, firing patterns came to distinguish the compartments following prolonged exposure to the environment. Indeed, a single continuous pattern spanning both compartments eventually formed, indicating that self-motion cues gradually come to dominate sensory cues. Current models of navigation employing grid cells require regular and coherent firing patterns, and these results thus provide the first evidence that grid cells are capable of subserving such functions in complex environments.

Recent results suggest that the medial septum supports grid cell function through the encoding of one or more speed signals. Its exact contribution, and the roles of its constituent neuronal subtypes, remain undetermined however. In the second experiment, I therefore recorded from grid cells in mice while manipulating the activity of medial septal cholinergic neurones using Designer Receptors Exclusively Activated by Designer Drugs (DREADDs). Increasing septal cholinergic tone reduced the frequency of theta oscillations by shifting the theta frequency vs running speed relationship to lower frequencies, without affecting its depth of modulation. However, no change was observed in grid firing patterns in either familiar or novel environments, consistent with a lack of change in the putative speed signals used by grid cells. An absence of novelty-induced expansion of grid scale prevented testing of the theory that, in rats, this phenomenon depends on increases in cholinergic tone. However, increasing the excitability of medial septal cholinergic neurones led to a pattern of

behaviour normally seen on exposure to novel environments. That is, the effects on behaviour of increasing cholinergic tone were consistent with a role for acetylcholine in the signalling of novelty.

Finally, I present preliminary work in which recordings of grid cells from wild type mice demonstrate that the room in which a grid cell is recorded can affect the scale, temporal stability, and regularity of its firing pattern. As the local environment between the rooms was consistent, such differences likely indicate a role for distal sensory cues in anchoring and stabilising grid cell firing patterns.

Contents

1	Introduction.....	14
1.1	Anatomy of the hippocampal formation	14
1.1.1	<i>Historical perspectives; its evolution; its constituent components.....</i>	14
1.1.2	<i>The entorhinal cortex</i>	16
1.1.3	<i>The dentate gyrus.....</i>	20
1.1.4	<i>The hippocampus proper</i>	21
1.1.4.1	<i>CA3</i>	21
1.1.4.2	<i>CA2</i>	22
1.1.4.3	<i>CA1</i>	22
1.1.5	<i>The subiculum</i>	23
1.1.6	<i>The pre- and parasubiculum</i>	24
1.1.7	<i>Summary</i>	24
1.2	Representations of space in the hippocampal formation.....	25
1.2.1	<i>Historical perspective and overview</i>	25
1.2.2	<i>Navigation in mammals.....</i>	27
1.2.3	<i>Place cells, place by direction cells, boundary vector cells.....</i>	28
1.2.3.1	<i>Properties and determinants of place cell firing; remapping.</i>	28
1.2.3.2	<i>Computational models of place cells based on external sensory inputs ..</i>	31
1.2.3.3	<i>Computational models based on self-motion inputs</i>	32
1.2.3.4	<i>Using place cells for navigation.....</i>	34
1.2.3.5	<i>Summary</i>	35
1.2.4	<i>Head direction cells.....</i>	36
1.2.5	<i>Grid cells.....</i>	37
1.2.5.1	<i>Properties and determinants of grid cell firing patterns.....</i>	37
1.2.5.2	<i>Models of grid cells</i>	40
1.2.5.3	<i>Encoding across grid modules</i>	44
1.2.5.4	<i>Using grid cell representations for navigation.....</i>	47
1.2.5.5	<i>Grid cells as input to place cells.....</i>	47
1.2.6	<i>Summary: contributions of spatially-tuned neurones to the cognitive map</i>	48
1.3	Contributions of the medial septum to grid cell activity.....	49
1.3.1	<i>Anatomy, connectivity and neuronal subtypes of the medial septum.....</i>	49
1.3.2	<i>A role for the medial septum in spatial processing.....</i>	51
1.3.2.1	<i>...Through the encoding of running speed</i>	51
1.3.2.2	<i>...Through novelty signalling.....</i>	53
1.3.2.3	<i>Summary</i>	54

2	Grid cells form a global representation of connected environments	55
2.1	Introduction	55
2.2	Methods	56
2.2.1	<i>Animals.....</i>	56
2.2.2	<i>Microdrives and surgery.....</i>	56
2.2.3	<i>Electrophysiological recording and behavioural training.....</i>	57
2.2.4	<i>Experimental environment and protocol.....</i>	58
2.2.5	<i>Spike sorting and binning.....</i>	59
2.2.6	<i>Spatial autocorrelograms and grid cell inclusion criteria.....</i>	59
2.2.7	<i>Analyses.....</i>	61
2.2.7.1	<i>General.....</i>	61
2.2.7.2	<i>Correlations</i>	61
2.2.7.3	<i>Fitting of ideal grids.....</i>	62
2.2.7.4	<i>Analysis of head direction cells</i>	64
2.2.7.5	<i>Analysis of place cells.....</i>	65
2.2.7.6	<i>Phase offset analysis.....</i>	65
2.2.7.7	<i>Corridor grid cell classifier.....</i>	66
2.2.7.8	<i>Remapping analysis</i>	67
2.3	Results	67
2.3.1	<i>Grid cell firing disambiguates the compartments after prolonged experience .</i>	67
2.3.2	<i>Grid cell firing transitions from a local to a global representation with experience</i>	69
2.3.3	<i>The observed local to global transition cannot be explained by biases in sampling of grid modules across time</i>	73
2.3.4	<i>Transition to a global representation may depend on path integration</i>	74
2.3.5	<i>Grid cell firing patterns in the corridor</i>	75
2.3.6	<i>Head direction and boundary vector cells show consistent firing between compartments, regardless of experience</i>	76
2.3.7	<i>Place cell firing disambiguates the compartments after a similar timeframe to that seen in grid cells</i>	77
2.3.8	<i>Grid and place cells are anchored locally to the environment.....</i>	80
2.4	Discussion	82
2.4.1	<i>Summary of results and significance.....</i>	82
2.4.2	<i>Mechanisms underlying transition to a global representation</i>	82
2.4.3	<i>Grid firing is locally anchored to the environment.....</i>	83
2.4.4	<i>Impact of irregular firing on grid cells' capacity to act as a spatial metric.....</i>	84
2.4.4.1	<i>Grid cells robustly encode self-location despite distortions</i>	84
2.4.4.2	<i>Distorted grid firing introduces metric errors to vector navigation</i>	85

2.4.5	<i>Place cell representations in the multicompartment environment; their relationship with other spatially tuned cell types.....</i>	87
3	Modulating medial septal cholinergic tone alters the frequency of theta oscillations without affecting grid cell activity	88
3.1	Introduction	88
3.2	Methods	89
3.2.1	<i>Animals.....</i>	89
3.2.2	<i>Microdrives and surgery.....</i>	89
3.2.3	<i>Experimental protocol and electrophysiological recording.....</i>	90
3.2.4	<i>Analyses</i>	92
3.2.4.1	<i>Position and speed data</i>	92
3.2.4.2	<i>LFP.....</i>	93
3.2.4.3	<i>Theta modulation of single units</i>	94
3.2.4.4	<i>Behaviour</i>	96
3.2.4.5	<i>Grid cell analysis.....</i>	96
3.2.4.6	<i>Grid cell speed modulation.....</i>	97
3.2.4.7	<i>Grid cell phase precession.....</i>	97
3.2.4.8	<i>Hypothesis testing</i>	98
3.2.5	<i>Experimental protocol for control mice</i>	99
3.2.6	<i>Immunohistochemistry and cell counting.....</i>	100
3.3	Results	101
3.3.1	<i>Increasing septal cholinergic tone reduces the frequency of LFP theta oscillations</i>	103
3.3.2	<i>Increasing septal cholinergic tone also reduces the theta burst frequency of single units.....</i>	106
3.3.3	<i>Increasing septal cholinergic tone induces a pattern of behaviour consistent with the detection of novelty.....</i>	109
3.3.4	<i>Injection of CNO has no effect in control mice.....</i>	111
3.3.5	<i>Increasing septal cholinergic tone has no effect on grid firing patterns in familiar environments.....</i>	113
3.3.6	<i>Increasing septal cholinergic tone has no effect on grid cell speed signals ..</i>	115
3.3.7	<i>Increasing septal cholinergic tone has no effect on grid firing patterns in novel environments</i>	117
3.4	Discussion	120
3.4.1	<i>Summary of results and significance.....</i>	120
3.4.2	<i>Role of septal acetylcholine in theta oscillations.....</i>	120
3.4.3	<i>Non-linear relationship between theta frequency and running speed</i>	121
3.4.4	<i>No evidence for an involvement of septal acetylcholine in the generation of grid firing patterns</i>	121
3.4.5	<i>Absence of novelty-induced increases in grid scale</i>	122
3.4.6	<i>Possible role of septal acetylcholine in signalling novelty.....</i>	122

4	Distal sensory features of the environment determine the scale, stability, and regularity of grid firing patterns	124
4.1	Introduction	124
4.2	Methods	125
4.2.1	Animals.....	125
4.2.2	Microdrives and surgery.....	125
4.2.3	Experimental protocol and electrophysiological recording.....	126
4.2.4	Analyses	127
4.2.4.1	Clustering and ratemap generation.....	127
4.2.4.2	Grid cell analysis.....	127
4.2.4.3	Hypothesis testing	128
4.3	Results	128
4.3.1	Grid scale is determined by the distal sensory features of the environment .	128
4.3.2	Gridness is determined both by novelty and the distal sensory features of the environment.....	131
4.3.3	The temporal stability of grid firing patterns is determined by the distal sensory features of the environment	134
4.4	Discussion	136
4.4.1	Summary of results and significance.....	136
4.4.2	Factors giving rise to differences in grid properties between environments..	136
4.4.3	Reconciling results concerning the impact of novelty on grid firing.....	137
4.4.4	Significance of results for theories of grid-expansion in novelty	138
5	Discussion	139
5.1	Overview	139
5.2	Implications of the results described in this thesis	140
5.2.1	Implications for the mechanisms by which grid firing patterns arise	140
5.2.1.1	Importance of environmental cues	140
5.2.1.2	Importance of self-motion inputs	141
5.2.1.3	Absence of evidence for cholinergic influences on grid firing	142
5.2.2	Implications for the potential functions of grid firing patterns	143
5.2.2.1	Grid firing forms coherent global representations, even in complex environments.....	143
5.2.2.2	Novelty-induced grid expansion is unlikely to signal novelty/familiarity, or generate place cell remapping	144
5.3	Current understanding and remaining uncertainties regarding the encoding and decoding of grid cell activity	145
5.3.1	What do grid firing patterns encode	145
5.3.2	How is this information encoded?	145
5.3.2.1	...By individual cells and individual models	145

5.3.2.2	...Across groups of modules	146
5.3.3	How are grid patterns decoded?	147
5.3.4	Summary	148
6	References	149

List of Figures

<i>Figure 1.1 Position and orientation of the hippocampus and entorhinal cortex in rats and humans.</i>	14
<i>Figure 1.2 Relative positions of the sub regions of the hippocampal formation in rats.</i>	16
<i>Figure 1.3 Organisation of the entorhinal cortex's projections to the hippocampus.</i>	17
<i>Figure 1.4 Example firing of spatially modulated neurones of the hippocampal formation.</i>	29
<i>Figure 1.5 Theta phase precession examples in place and grid cells.</i>	34
<i>Figure 1.6 Properties of grid cell firing patterns; organisation of grid cells into modules.</i>	38
<i>Figure 1.7 Illustration of principles of continuous attractor network models of grid cells.</i>	41
<i>Figure 1.8 Position and shape of the medial septum in the mouse brain.</i>	50
<i>Figure 2.1 With increasing experience of the multicompartiment environment, grid cell firing patterns show reduced representation similarity between the compartments, and increased regularity.</i>	68
<i>Figure 2.2 Grid cell firing patterns transition from a local to a global representation with increasing experience.</i>	71
<i>Figure 2.3 Grid properties are consistent in the two compartments; grid phase analysis confirms local to global representation transition.</i>	72
<i>Figure 2.4 The transition from local to global representations cannot be explained by biases in the sampling of grid cells.</i>	74
<i>Figure 2.5 Firing patterns in the corridor are 'grid-like' and stable, but irregular, and co-occur with highly stereotyped behaviour.</i>	75
<i>Figure 2.6 Head direction and boundary vector cells firing replicates in the two compartments regardless of experience.</i>	77
<i>Figure 2.7 The range of place cell representation behaviours becomes more diverse with increasing experience of the multicompartiment environment.</i>	79
<i>Figure 2.8, Grid and place cells are locally anchored to the multicompartiment environment.</i>	81
<i>Figure 3.1 Confirmation hM3Dq is expressed almost exclusively in medial septal cholinergic neurones; pictorial representation of the experimental protocol.</i>	102
<i>Figure 3.2 Increasing medial septal cholinergic tone reduces LFP theta frequency by shifting its relationship with running speed to lower frequencies.</i>	105
<i>Figure 3.3 Examples of observed and artificial LFP theta frequency vs running speed plots demonstrate a clear non-linear relationship between the two.</i>	106
<i>Figure 3.4 Increasing medial septal cholinergic tone reduces the burst frequency of theta-modulated single units.</i>	108
<i>Figure 3.5 Conclusions of the effect of CNO on theta-modulated unit burst frequency are unchanged by performing analyses averaging at the level of cell.</i>	109
<i>Figure 3.6 Increasing medial septal cholinergic tone results in a pattern of behaviour normally seen during exposure to a novel environment.</i>	110
<i>Figure 3.7 Effects observed following injection of CNO in hM3DQ mice are not seen in DIO-mCherry control mice.</i>	112

<i>Figure 3.8 Effect of CNO on theta frequency and behaviour still seen when analysis is restricted to novelty days, in which only either CNO or saline were administered.....</i>	<i>113</i>
<i>Figure 3.9 Increasing medial septal cholinergic tone has no effect on grid cell firing patterns in familiar environments.</i>	<i>114</i>
<i>Figure 3.10 Increasing medial septal cholinergic tone has no effect on putative grid cell speed signals.</i>	<i>116</i>
<i>Figure 3.11 Increasing medial septal cholinergic tone has no effect on grid cell firing patterns in novel environments.</i>	<i>118</i>
<i>Figure 3.12 Conclusions of the effect of CNO on grid cell firing patterns are unchanged by repeating analyses averaging at the level of cell.</i>	<i>119</i>
<i>Figure 4.1 Illustration of experimental protocol for novelty recordings in wild type mice. ...</i>	<i>127</i>
<i>Figure 4.2 Grid scale is determined by the distal sensory features of the environment.....</i>	<i>130</i>
<i>Figure 4.3 Gridness is determined both by novelty and the distal sensory features of the environment.</i>	<i>133</i>
<i>Figure 4.4 Temporal stability of grid firing is determined by the distal sensory features of the environment.</i>	<i>135</i>

1 Introduction

1.1 Anatomy of the hippocampal formation

1.1.1 *Historical perspectives; its evolution; its constituent components*

The first known description of the hippocampus was published in 1587 by Giulio Cesare Aranzio, a student of Vesalius. From human dissections, he described a white growth extending anteriorly from the base of the lateral ventricle, remarking on its resemblance to both a seahorse (hippocampus) and a silkworm (vermis bombycinus) (Engelhardt 2016). While ‘The Vermis Bombycinus as a Cognitive Map’ sadly never came to pass, without a common name being settled upon it remained a possibility for centuries. For example, in reference to the Egyptian god Ammon-Ra, often depicted as having the head of a Ram, the hippocampus’ curved shape led Garengeot (1742) to propose the Latin term *cornu Ammonis* (Boccara et al. 2015). Now however, the term hippocampus is almost universally used, referring specifically to the allocortical region of the brain comprising sub regions CA3, CA2, and CA1; Garengeot’s proposal living on in the names of these subfields, as coined by Lorente de Nó (1934). Following Amaral & Lavenex (Andersen et al. 2006), the hippocampal formation in turn will here refer to a broader collection of structures which are densely interconnected with the hippocampus proper, including the dentate gyrus, subiculum, presubiculum, parasubiculum, and entorhinal cortex.

Figure 1.1 Position and orientation of the hippocampus and entorhinal cortex in rats and humans.
Schematic representation of the hippocampus (red) and entorhinal cortex (blue) in the rat (left) and human (right). Axes indicate the diagrams’ orientation. Figure adapted from Strange et al. 2014.

The hippocampus is held to be an evolutionarily ancient structure: The dorsomedial forebrain of birds and reptiles is thought homologous based on similarities in connectivity and function (Bingman 1992). If accurate, this homology implies that a precursor of the modern mammalian hippocampus existed at least 300 million years ago, in the last common ancestor of mammals and birds. Within mammals, the hippocampus is highly conserved, being both universally present and displaying little phylogenetic variation in basic architecture from monotremes to humans (Insausti 1993). The location of the hippocampal formation within the brain varies more obviously between species however. In the rodent, the hippocampus forms an extended 'C-shape', stretching from a rostro-dorsal 'septal pole' to a dorso-ventral 'temporal pole'. In the human and primate brain the hippocampus is more spatially restricted, residing entirely within the temporal lobe, with a long posterior-anterior axis equivalent to the septo-temporal axis in rodents (Figure 1.1). In discussing the anatomy of the hippocampal formation, I will predominantly describe that of the rodent, as this is both the experimental model in which hippocampal anatomy is best understood, and the model in which the experiments reported in this thesis were carried out. As such, I will refer to the long axis of the hippocampus as the septo-temporal axis, with a second axis, the transverse axis, normal to it. Within the hippocampus, a transverse section stained with the Golgi preparation reveals two C-shaped layers of cell bodies: the smaller one the dentate gyrus, interlocking with that of the hippocampus proper (Figure 1.2B). Each location in the transverse section is also described by a position on a proximo (close to the dentate gyrus) – distal (far from the dentate gyrus) axis. The subiculum, presubiculum, parasubiculum, and entorhinal cortex (together forming the parahippocampal cortex), sit posterior and in parts ventral to the hippocampus and dentate gyrus in broadly that order (Figure 1.2A). As such, at the septal pole of the hippocampus, a transverse section will contain only the dentate gyrus and hippocampus proper. In transverse sections moving towards the temporal pole, the subiculum appears first, followed by the pre- and parasubiculum, and eventually the entorhinal cortex.

Seminal early anatomical investigations into the hippocampal formation were carried out by Santiago Ramón y Cajal (1893, 1902) and his student Lorente de Nó (1933, 1934), using the Golgi stain. They argued that the hippocampal formation was defined by a unique 'loop' of largely unidirectional connections: the entorhinal cortex sending projections to the dentate gyrus, which in turn connects onto CA3; CA3 then projecting onto CA1, which sends the hippocampus' output via the subiculum and entorhinal cortex. While modern histological analysis has demonstrated that this view is somewhat simplistic, for example in not accounting for multiple parallel pathways or return projections (van Strien, Cappaert, and Witter 2009), the idea of the hippocampus as a loop in which sensory information is serially

processed remains a common analogy. In the following account of contemporary understanding of the circuitry of the hippocampal formation, I will proceed along the loop, discussing the anatomy and connectivity of each region in turn.

Figure 1.2 Relative positions of the sub regions of the hippocampal formation in rats. **A**, Two schematic representations of the constituent areas of the hippocampal formation. The left-hand image represents a predominantly sagittal view of the rat brain, while the right-hand image is predominantly coronal. Distinct sub regions are colour-coordinated according to the labels in **B**. The dark green region not labelled in **B** is the lateral entorhinal area. **B**, A horizontal slice through the rat brain is labelled to indicate the distinct substructures of the hippocampal formation. DG = dentate gyrus, Sub = subiculum, PrS = presubiculum, PaS = parasubiculum, MEA = medial entorhinal area, POR = postrhinal cortex. Figure adapted from Strien et al. 2009.

It should be noted, that while this description of the hippocampal formation's circuitry focusses largely on the interconnections amongst glutamatergic principal neurones, with only relatively brief consideration given to inhibitory interneurones, it is not because of an assumption of their greater significance. While interneurones are likely fundamental in hippocampal function through the regulation of excitatory connections and in controlling the temporal dynamics of firing, the huge diversity and lack of understanding of interneurone populations places a full treatment beyond the scope of this thesis. For extensive discussions of the diversity and functional contributions of interneurones in the hippocampal formation, I refer the reader to the existing comprehensive reviews of Freund & Buzsáki (1996) and Klausberger & Somogyi (2008).

1.1.2 The entorhinal cortex

Originally called the spheno-occipital ganglion (Ramón y Cajal 1902), the entorhinal cortex is named for being partially enclosed by the rhinal sulcus (Canto, Wouterlood, and Witter 2008). Ramón y Cajal identified and immediately understood the importance of the strong connectivity between the entorhinal cortex and hippocampus, stating that their interconnection meant the latter must follow the function of the former: if the entorhinal cortex processed visual information, so would the hippocampus (Ramón y Cajal 1902).

The entorhinal cortex was defined by Ramón y Cajal as comprising six layers (LI-LVI), though LV is here divided into LVa and LVb. The layers close to the pial surface (LI-III) are termed 'superficial', and the layers further away (LIV-VI) 'deep'. Brodmann further divided the entorhinal cortex into two distinct areas, which we now call the medial and lateral entorhinal cortex (Brodmann 1909). Each is roughly triangular in shape, with the lateral portion more rostro-lateral, and the medial portion defined by a ventral base that extends along the medial edge of the lateral entorhinal cortex to a dorsal border with the postrhinal cortex (Figure 1.2 and Figure 1.3) While there are some histological differences between the two entorhinal areas (LIV is clearer in the medial entorhinal cortex, and the LII/LIII boundary more pronounced in the lateral entorhinal cortex for example (Canto, Wouterlood, and Witter 2008)), the two are predominantly distinguished by their connectivity: the presubiculum only innervates the medial portion (Canto, Wouterlood, and Witter 2008), while only the lateral portion projects to distal CA1 and proximal subiculum, the medial entorhinal cortex projecting to proximal CA1 and distal subiculum (Naber, Lopes da Silva, and Witter 2001).

Figure 1.3 Organisation of the entorhinal cortex's projections to the hippocampus. The entorhinal cortex is labelled in shades of blue according to the position along the septo-temporal axis of the hippocampus with which each region predominantly sends and receives projections. Darker-shaded regions of the entorhinal cortex project primarily to septal portions of the hippocampus, lighter-shaded regions predominantly to temporal portions of the hippocampus. Black lines indicate borders between the perirhinal cortex (PR), postrhinal cortex (POR), lateral entorhinal cortex (LEC) and medial entorhinal cortex (MEC). Figure adapted from Burwell and Amaral (1998).

A striking topology exists in projections from the entorhinal cortex to the hippocampus: ventro-lateral entorhinal areas project to more temporal portions of the hippocampus, whereas dorso-medial areas tend to project to more septal areas of the hippocampus (Dolorfo and Amaral 1998). Return projections from CA1 and the subiculum are similarly organised: septal CA1 projects mainly to the dorso-medial entorhinal cortex, whereas

temporal CA1 mainly targets ventro-lateral areas. That is, while both medial and lateral entorhinal cortex project to each position in the septo-temporal axis of the hippocampus, there exists a clear banding joining specific portions of the dorso-medial to ventro-lateral entorhinal axis to specific portions of the septo-temporal hippocampal axis (Figure 1.3).

While differing in the specificity and strength of connections, tracing studies indicate that both medial and lateral entorhinal cortex receives input from extremely diverse cortical and subcortical areas (Canto, Wouterlood, and Witter 2008). Inputs to the entorhinal cortex have been extensively studied by Burwell and colleagues, whose investigations emphasise the importance of 'high-level' and multimodal sensory areas as inputs (Burwell and Amaral 1998). The lateral entorhinal area receives abundant projections from perirhinal, postrhinal, insular, and piriform cortex. Medial entorhinal cortex also receives dense input from the piriform and postrhinal cortex, but differs from lateral entorhinal areas in also receiving projections from retrosplenial, posterior parietal, and visual association areas (Burwell and Amaral 1998). Both regions also receive inputs from subcortical structures including the amygdala and medial septum, as well as neuromodulatory inputs from brainstem regions such as the locus coeruleus and raphe nucleus (Canto, Wouterlood, and Witter 2008).

The connectivity of the entorhinal cortex varies in a layer-specific manner. LI is predominantly acellular, mostly comprised of dendrites from neurones in deeper layers, and transverse-travelling axons both from within and without the entorhinal cortex (de Nó 1933). LI does however contain GABAergic interneurones, including the so-called horizontal and multipolar neurones (Canto, Wouterlood, and Witter 2008). LII is a densely-packed layer predominantly containing clusters of 'island' cells with pyramidal somata, which can be identified by staining either for calbindin or Wolfram syndrome 1 protein, as well as intermixed 'ocean' cells of a stellate shape, identified by staining for reelin (Varga, Lee, and Soltesz 2010; Kitamura et al. 2014; Ray et al. 2014). Island-pyramidal cells project to lacunosum-moleculare interneurones in CA1, which in turn inhibit CA1 principal neurones (Kitamura et al. 2014). In contrast, ocean-stellate cells project to the dentate gyrus and CA3, their axons forming the perforant path, so called due to its 'perforation' of the subiculum (Witter 2007; Varga, Lee, and Soltesz 2010; Kitamura et al. 2015). The two types of cell also differ in their intra-entorhinal connectivity: ocean-stellate cells project to pyramidal neurones in LVb of the entorhinal cortex, unlike LII island-pyramidal cells (Sürmeli et al. 2015). It should be noted however, that the island-pyramidal/ocean-stellate dichotomy has been challenged by Fuchs and colleagues, who have suggested that at least four types of principal neurone in LII exist (including island-pyramidal cells, ocean-stellate cells, and two intermediate types) (Fuchs et al. 2016). LII also contains a number of interneurones, including multipolar, bipolar, basket, and chandelier cells, which variably stain for

parvalbumin, cholecystokinin, and somatostatin (Canto, Wouterlood, and Witter 2008). Both excitatory and inhibitory neurones of LII send and receive reciprocal connections with the medial septum (Fuchs et al. 2016). Principal neurones of LIII are almost exclusively pyramidal in shape, sending axon collaterals to LI and LII as well as onto CA1 and the subiculum. LIV, also known as the lamina dissecans, is predominantly acellular, but occasional bipolar cells with horizontally arbourising dendrites, and axons diverging to both LIII and the deeper layers, are seen (Canto, Wouterlood, and Witter 2008). LV can be separated into LVa and LVb on the basis of the specific molecular staining and connectivity patterns of the pyramidal neurones found in each layer. LVa is specifically labelled by expression of *Etv1*, L5B by expression of *Ctip2* (Sürmeli et al. 2015). Further, while L5a sends major extra-hippocampal projections (including to the nucleus accumbens, retrosplenial cortex, amygdala, perirhinal cortex, and medial septum), LVb in contrast sends projections to no other telencephalic areas, apart from a possible minor projection to the thalamus. LVb is, however, preferentially targeted by hippocampal and LII ocean-stellate cell inputs (Sürmeli et al. 2015). LVI contains a largely heterogeneous group of principal neurones of varying cell sizes and shapes. LVI is notable, together with LV, for receiving input from CA1 however (Canto, Wouterlood, and Witter 2008).

In sum, the circuitry discussed above reflects a system in which superficial layers of the entorhinal cortex (LII and LIII) send input projections to various levels of the hippocampus including CA1, CA3, and the dentate gyrus. In contrast, CA1 and subiculum, as the output structures of the hippocampus proper, target deeper layers of the entorhinal cortex (LV and LVI). This organisation implies the entorhinal cortex acts as an 'input/output node' to the hippocampus, with information passing serially through the hippocampal formation loop, and remaining segregated in a laminar fashion in the entorhinal cortex. However, a range of anatomical data suggests that such an interpretation is at best simplistic (van Strien, Cappaert, and Witter 2009). For example, while LII is traditionally thought of as the exclusive origin of the perforant path, LIII, LV and LVI also contribute axons to a lesser extent (Deller et al. 1996). Similarly, the deep layers of the entorhinal cortex send projections to more superficial layers and vice versa, with principal neurones in each entorhinal layer also sending collaterals to its own layer (Köhler 1986; Sürmeli et al. 2015). The extent to which superficial or deep layers of the entorhinal cortex can be seen exclusively as 'input' and 'output' structures to the hippocampus respectively is unclear.

Finally, as well as variation between layers and media/lateral portions, more recent data have indicated the presence of histological variation along the dorso-ventral axis of the entorhinal cortex. For example, the density of parvalbumin expressing interneurones in LII (which strongly inhibit local ocean-stellate cells) varies in a continuum along the dorso-

ventral axis, with more parvalbumin expressing axon terminals in the dorsal portion of LII (Beed et al. 2013).

1.1.3 *The dentate gyrus*

The dentate gyrus is fundamental to hippocampal function in providing the primary input to the hippocampus proper, through its mossy fibre projections to CA3 (Amaral and Witter 1989). The dentate gyrus is trilaminar, comprising a predominantly cell-free molecular layer, the granule cell layer, which contains the majority of principal neurones, and the polymorphic cell layer (also known as the hilus). The granule cell layer appears as a 'U' or 'V' when a transverse section of the hippocampus is taken (depending on septo-temporal position), which 'interlocks' with the pyramidal cell layer of the hippocampus proper.

The principal neurone of the dentate gyrus is the granule cell, whose somata are restricted to the granule cell layer. Granule cells are unique within the hippocampal formation in that they continue to mature throughout adulthood, from a population of undifferentiated precursor cells in the subgranular zone (Kaplan and Hinds 1977). Easily identified by particularly strong staining for calbindin, granule cells are the only cells in the dentate gyrus that send projections to area CA3, doing so via the unmyelinated mossy fibre pathway (Bocacara et al. 2015). These glutamatergic projections terminate in the stratum lucidum, just superficial to the pyramidal cell layer of CA3. Mossy fibres culminate in large (up to 8µm) presynaptic terminals known as mossy fibre expansions (Amaral and Dent 1981). Each granule cell contacts relatively few CA3 pyramidal neurones. However, due to the large presynaptic terminals, the large number of terminals on each pyramidal neurone (up to ~37) (Chicurel and Harris 1992), and the proximity of terminals to the soma, the mossy fibre projection has long been thought to be a 'detonator synapse' which exerts strong control over CA3 firing (McNaughton and Morris 1987).

Granule cells receive their primary input from the perforant path, which predominantly originates in LII of the entorhinal cortex (Ramón y Cajal 1893). The perforant path terminates in the molecular layer, primarily onto the dendrites of granule cells, though interneurons may also be targeted (Freund and Buzsáki 1996). Tracing studies demonstrate that the dentate gyrus also receives inputs from the pre- and parasubiculum, which also project to the molecular layer (Köhler 1985).

The dentate gyrus also contains so-called mossy cells, glutamatergic neurones, whose somata are located in the polymorphic layer. Mossy cells receive strong input from ipsilateral granule cells, and exclusively send projections to contralateral granule cells (Hjorth-Simonsen and Laurberg 1977). Strikingly, mossy cell projections are not to the same

position along the septo-temporal axis as that containing the somata from which they originate. Mossy cells thus appear to link granule cells in distinct septo-temporal portions of the hippocampus through a disynaptic excitatory loop (Amaral and Witter 1989).

Both mossy cells and a number of interneurone types receive projections from medial septal cholinergic neurones (Lübke, Deller, and Frotscher 1997), which only sparsely innervate dendrites of granule cells in the molecular layer. GABAergic septal projections exclusively target inhibitory interneurons of the dentate gyrus (Amaral and Witter 1989). The dentate gyrus also receives broad neuromodulatory inputs including noradrenergic, serotonergic, and dopaminergic projections, from the locus coeruleus, raphe nucleus, and ventral tagmental area respectively (Amaral and Witter 1989).

1.1.4 *The hippocampus proper*

The hippocampus is divided into three subfields, CA3, CA2, and CA1, according to the distinct morphology and connectivity of principal neurones in each region, as determined by Lorente de Nó (1934). The area he termed CA4 is now thought to be an extension of the polymorphic layer of the dentate gyrus however (Amaral and Witter 1989). Each subfield shares a similar laminar structure. The deepest layer (furthest from the pial surface) is stratum oriens, comprising a number of inhibitory interneurone types and the basal dendrites of pyramidal cells. Stratum pyramidalae, or the pyramidal cell layer, is the seat of the pyramidal-shaped somata of the hippocampus' principal neurones, as well as that of a range of interneurone classes. Superficial to the pyramidal cell somata is stratum radiatum, containing their proximal apical dendrites, as well as the majority of excitatory axons from other principal neurones of the hippocampal formation (such as the schaffer collaterals and mossy fibres). Finally, the most superficial layer is stratum lacunosum moleculare, containing the distal apical dendrites of pyramidal cells, a number of interneurone classes, and axon terminals from LIII of the entorhinal cortex.

1.1.4.1 CA3

The large pyramidal neurones of CA3 receive projections from the perforant path onto their distal apical dendrites in stratum lacunosum moleculare. Apical dendrites closer to the somata, in a CA3-specific layer known as stratum lucidum, (and to a lesser extent basal dendrites in stratum oriens), are the site of inputs from mossy fibres originating in the dentate gyrus (Amaral and Witter 1989). Pyramidal neurones in CA3 project both ipsi- and contralaterally across a broad septo-temporal portion of the hippocampus, targeting other CA3 pyramidal neurones directly through the longitudinal associational bundle, and targeting

CA2 and CA1 pyramidal neurones through schaffer collaterals (de Nó 1934; Swanson, Wyss, and Cowan 1978). Indeed, a single pyramidal neurone may send both recurrent associational and CA2/1 projections (Swanson, Sawchenko, and Cowan 1980). Both associational and schaffer collateral projections predominantly target apical dendrites in stratum radiatum, but axon terminals can also be found in stratum oriens (Hjorth-Simonsen 1973). Strikingly, CA3 can also be observed sending 'return' projections to the polymorphic layer of the dentate gyrus, against the generally unidirectional pathways of the hippocampus (Laurberg 1979).

1.1.4.2 CA2

CA2 pyramidal neurones have similarly large somata to those in CA3, however CA2 has traditionally been distinguished from CA3 by an absence of stratum lucidum and mossy fibre afferents from the dentate gyrus (de Nó 1934). CA2 is also unique in receiving particularly strong input projections from the supramammillary nucleus (Haglund, Swanson, and Köhler 1984). More recently, results indicate that CA2 is also defined by its anatomical and electrophysiological connectivity within the hippocampal formation: CA2 receives particularly strong excitatory input from entorhinal cortex LII/III onto the distal apical dendrites of its pyramidal neurones, a connection which is highly plastic. In contrast, CA3 inputs onto CA2 proximal dendrites are dominated by feedforward inhibition, and show little sensitivity to plasticity protocols (Chevalleyre and Siegelbaum 2010). Research into CA2 is nascent however, and disagreements exist: other investigators report projections to CA2 from the dentate gyrus and an absence of projections from entorhinal LIII (Kohara et al. 2014). It may thus be that recently reported CA2-specific neurochemical expression patterns are more useful in segregating CA2 from other hippocampal subfields: reports suggest the markers RGS14, PCP4 and STEP uniquely define CA2 (Kohara et al. 2014). CA2 displays few of the recurrent associational projections which characterise CA3, with the majority of its pyramidal cells instead projecting to CA1 (Chevalleyre and Siegelbaum 2010), predominantly targeting basal dendrites in stratum oriens (Kohara et al. 2014).

1.1.4.3 CA1

As well as receiving inputs from CA3 and CA2 in stratum radiatum and stratum oriens, CA1 also receives direct projections from entorhinal cortex LIII onto the distal apical dendrites of pyramidal neurones in stratum lacunosum moleculare (Amaral and Witter 1989). More recently, LII island-pyramidal cells of the entorhinal cortex have been shown to target CA1 lacunosum-moleculare interneurones, which in turn inhibit CA1 pyramidal cells, creating a

feedforward inhibitory circuit (Kitamura et al. 2014). CA1 does exhibit recurrent associational connections (Amaral, Dolorfo, and Alvarez-Royo 1991), though these are sparser than those of CA3 (Amaral and Witter 1989). The majority of CA1 pyramidal cell axons instead travel through stratum oriens towards the subiculum and deep layers of the entorhinal cortex, particularly targeting L5b once there (Finch and Babb 1981; Sürmeli et al. 2015).

1.1.5 *The subiculum*

The subiculum (Latin for 'support'), was named by Santiago Ramón y Cajal (1909), who viewed it as a key auxiliary structure for the hippocampus proper. Ramón y Cajal's student, Lorente de Nó, gave the first strict definition of the subiculum, defining it as the trilaminar area adjacent to CA1, starting where CA1's pyramidal cell layer broadens and becomes less-dense, and stratum oriens disappears (de Nó 1934). Superficial to the principal cell layer are the deep and superficial molecular layers in turn. The superficial molecular layer is continuous with (but distinct from) CA1's lacunosum moleculare, with the deep molecular layer a continuation of CA1's stratum radiatum. The deep molecular layer predominantly receives its inputs from CA1, whereas the superficial portion is mainly targeted by axons from the entorhinal cortex (Amaral and Witter 1989). The principal cell layer contains at least two cell types: regular-spiking and intrinsically-bursting cells, differentiated from each other by their spontaneous firing properties (Greene and Totterdell 1997). While both neurone types are pyramidal in shape, and have dendritic trees that ramify throughout both molecular layers, only the intrinsically bursting-cells stain for somatostatin and project to the entorhinal cortex (Andersen et al. 2006). Most inputs to the subiculum arise from CA1 and entorhinal cortex LIII, though it also receives subcortical inputs from the medial septum and mammillary nuclei. The subiculum shows a large degree of associational connectivity, with principal neurones projecting onto other principal neurones in more temporal regions of the subiculum; giving rise to a system of unilateral associational connections along its septo-temporal axis (Swanson, Sawchenko, and Cowan 1981). The subiculum has long been seen as the output structure of the hippocampus proper, with its principal neurones projecting to the pre- and parasubiculum, and deep layers of the entorhinal cortex. (van Strien, Cappaert, and Witter 2009). The subiculum also sends axons to both cortical and non-cortical areas outside the hippocampal formation. Cortical regions targeted include the orbitofrontal, prelimbic, infralimbic, retrosplenial, and perirhinal cortex; while non-cortical projections are predominantly to the lateral septum, mammillary nuclei, and nucleus accumbens (van Strien, Cappaert, and Witter 2009).

1.1.6 *The pre- and parasubiculum*

Histologically, the presubiculum is easily distinguished from the subiculum by a densely packed 'external' cell layer of small pyramidal cells (Andersen et al. 2006). Although technically comprising six layers, LI is very thin, and LV and VI are difficult to distinguish from one another, and indeed from the deeper layers of the entorhinal cortex with which they share a border. LI and LIV are both acellular. The excitatory pyramidal neurone somata reside in LII/III (Boccarda et al. 2015).

As in the presubiculum, the pyramidal neurone somata of the parasubiculum also reside in Layers II and III, however the two regions can be histologically distinguished by a broader, less dense sheet of LII/III neurones in the parasubiculum (Boccarda et al. 2015). Further, the two regions differ in their expression of calbindin, which can be used to exclusively label the pyramidal neurones of the presubiculum (Boccarda et al. 2015). Like the presubiculum, the parasubiculum is multilaminar, with cell-sparse layers I and IV, and with a poorly defined LV/VI boundary (Boccarda et al. 2015).

Perhaps even more so than for the entorhinal cortex, it is hard to place the pre- and parasubiculum at a particular point in the hippocampal formation loop, or define these regions as either input or output structures to the hippocampus proper. Traditionally, the projections of both regions to the entorhinal cortex's superficial layers have been emphasised (presubiculum predominantly to LIII, almost exclusively of the medial entorhinal cortex; parasubiculum to LII (Köhler 1985; Caballero-Bleda and Witter 1993)), giving the impression more of input structures. Indeed, all layers of the pre- and parasubiculum also project to the molecular layer of the dentate gyrus (van Strien, Cappaert, and Witter 2009; Köhler 1985). However, both regions also send axons to areas classically considered to be 'output' structures, including the subiculum and deep layers of the entorhinal cortex (Canto, Wouterlood, and Witter 2008; Canto et al. 2012). The afferent connections to these regions fails to decide the issue: both are reciprocally connected with the peri- and postrhinal cortices (Furtak et al. 2007), and receive inputs from retrosplenial and visual cortical areas. Uniquely for the hippocampal formation, both also receive input from the anterior thalamic nuclei (van Strien, Cappaert, and Witter 2009). The majority of pre- and parasubicular inputs are from within the hippocampal formation however, with axons from the subiculum and entorhinal cortex terminating in LI of both areas (Andersen et al. 2006).

1.1.7 *Summary*

Presented here is a contemporary, though in parts simplified, description of the major excitatory connectivity of the hippocampal formation. The perception of the hippocampal

formation as comprising a predominantly unidirectional 'loop' continues to shape thinking regarding the potential functions of the hippocampus and the neural mechanisms by which they are achieved. In this scheme, highly processed sensory information from multimodal areas such as the peri- and postrhinal cortex is sent to the superficial entorhinal cortex. In turn, the entorhinal cortex routes this information to the hippocampus proper, through the dentate gyrus or direct projections to CA3. These regions then project to CA1, which sends the hippocampus' output via the subiculum and deep entorhinal cortex. However, such a simplistic model overlooks numerous complexities in the circuitry of the hippocampal formation, as discussed. For example, CA3 sends return projections to the polymorphic layer of the dentate gyrus, against the perceived direction of 'flow' of information (Laurberg 1979), while multiple parallel pathways are seen throughout the loop, such as in projections from entorhinal LII neurones to the dentate gyrus, CA3, CA1, and subiculum (van Strien, Cappaert, and Witter 2009). Similarly, while the recurrent excitatory connectivity of CA3 has led to hypotheses regarding its particular significance in the mnemonic functions of the hippocampus (Marr 1971), single and disynaptic recurrent excitatory connections are in fact seen at multiple levels of the hippocampal formation, including the dentate gyrus, CA1, subiculum, and entorhinal cortex (van Strien, Cappaert, and Witter 2009). Currently, it is unknown whether reductionist simplifications of neuroanatomy help or hinder our understanding of the brain. Through focusing on parsimonious neuroanatomical accounts, it is easier to 'see the wood from the trees', and so develop testable hypotheses regarding neural function. In contrast, it is clearly also conceivable that simplified descriptions of neural circuitry can hinder the advancement of understanding through obscuring salient connections. Certainly, however, the function and *modus operandi* of each brain region is intrinsically linked to its connectome. Discerning its circuitry, in full or in part, is thus an important first step in developing a full understanding of the hippocampal formation.

1.2 Representations of space in the hippocampal formation

1.2.1 *Historical perspective and overview*

Compelling early investigations into the functions of the hippocampal formation studied the effects on behaviour of damage to the region in monkeys and humans. In 1888, Brown and Schäfer described how bilateral lesions of the temporal lobe of a rhesus macaque gave rise to an apparent global amnesia:

"Every object with which he comes in contact, even those with which he was previously most familiar, appears strange and is investigated with curiosity. ... Even after having examined an object in this way with the utmost care and deliberation, he will, on again coming across

the same object accidentally even a few minutes afterwards, go through exactly the same process, as if he had entirely forgotten his previous experiments” (Brown and Schafer 1888).

Around a decade later, at a medical meeting in St Petersburg, von Bekhterev described a patient whose most striking clinical abnormality was a severe memory impairment. An autopsy subsequently revealed a bilateral ‘softening’ of the hippocampus and parahippocampal cortex (Bechterew 1900). Most famously, Brenda Milner reported the severe anterograde amnesia in patient H.M. that followed bilateral resection of the medial temporal lobes, a final attempt to treat his severe epilepsy (Scoville and Milner 1957). Subsequent investigations supporting these results have led to a broad consensus that the human hippocampus is central to episodic memory (Squire and Zola-Morgan 1991).

Later, John O’Keefe adapted microelectrode techniques developed to record from spinal and medullar neurones in freely moving animals, allowing him to record from CA1 as rats explored an enclosure. He observed that a subset of neurones, ‘place cells’, fired exclusively when the animal was in a specific area of the environment, remaining almost silent when the rat was elsewhere (O’Keefe and Dostrovsky 1971). In the decades since, research employing diverse experimental techniques, in a wide range of species, has added weight to the argument that the hippocampal formation is paramount to the mammalian brain’s representation of space, and animals’ ability to navigate. For example, rats with lesions to the hippocampus are impaired in learning to swim to a hidden platform, though perform normally if the platform is visible (Morris et al. 1982). Positron emission tomography scanning of cerebral blood flow in humans shows increased activity in the hippocampus when navigating to a goal location in a virtual reality environment, with hippocampal activity correlating with navigational accuracy across subjects (Maguire et al. 1998).

A wealth of evidence has thus made clear that the hippocampal formation is both central to episodic memory and the brain’s representation of space (Burgess, Maguire, and O’Keefe 2002). What follows is a review of the current understanding of the hippocampal formation’s spatial processing, and its role in the navigational abilities observed in animals. That is, to the extent to which it is possible to separate the spatial and mnemonic functions of the hippocampal formation, the focus here is on the former. While understanding how the brain represents space is interesting in and of itself, space also provides a uniquely tangible and interpretable domain in which to investigate the representations of a high-level and associational brain area. Understanding the neural mechanisms underlying the function of many such areas has been hampered by the abstract nature of the information they encode. For example, in the orbitofrontal cortex, single neurones non-linearly encode complex combinations of sensory, motor, motivational, and reward signals (Rigotti et al. 2013). That is, studying the neural processes underpinning high-level representations in the tangible

spatial domain may shed light on common principles which can help advance understanding in other, less approachable areas. Indeed, it may also be that the hippocampal formation itself uses common mechanisms to represent non-spatial variables, with the apparent spatial focus of the hippocampal formation an artefact of the domain in which experiments are performed. Neuroimaging suggests the human hippocampus may also encode social networks (Tavares et al. 2015), while place cells form ‘maps’ of non-spatial, task-relevant continuous variables (Aronov, Nevers, and Tank 2017); the modular, periodic system used by the entorhinal cortex to encode spatial variables may also be applied to abstract continuous variables (Constantinescu, O’Reilly, and Behrens 2016). Alternatively, rather than encode non-spatial variables through the same mechanisms by which it represents space, the hippocampal formation may aid non-spatial encoding through providing a spatial framework upon which other variables can be associated; as has been proposed to be the mechanism by which episodic memory arises (O’Keefe and Nadel 1978; Byrne, Becker, and Burgess 2007).

1.2.2 *Navigation in mammals*

The ability to navigate between remembered locations, such as between a nest and a food source, is of clear adaptive benefit to wild animals. The navigational capacity of rats was thoroughly investigated throughout the first half of the 20th century, as a model system for understanding the neural basis of learning. In 1929, Karl Lashley reported an anecdote in which, having learned an alley maze, two of his rats pushed back the cover of the start-box and ran directly across the top of the maze to the goal box, whereby they climbed down and ate the reward (Lashley 1929). The ability of rats to navigate directly to a goal, even along a path not previously taken, was quantitatively investigated by Edward Tolman. He taught rats to run to a goal box through a maze in which an open circular table was followed by a corridor with three right-angled turns. Following 12 training trials, an alternative maze was presented, with a number of straight corridors leading away from the open table, and the original corridor blocked off. In the majority of cases, rats chose a corridor which took them to a location within 4 inches of where the entrance to the goal box had been (Tolman, Ritchie, and Kalish 1946). That is, the rats were able to estimate the vector separating them from their goal. Importantly, as the rats had only previously run along the now-closed corridor, they could not have navigated to the goal through appetitive conditioning based on learned stimulus-response pairs. Tolman instead suggested that rat’s performance of the task depended on a ‘cognitive map’: an internal representation of the environment (Tolman 1948). Subsequent experiments showed that rats use both sensory cues and estimates of the distance and direction travelled (‘path integration’) to perform such navigation.

Maaswinkel and Whishaw investigated the relative contribution of visual, olfactory and self-motion cues in navigation. In their experiment, rats left a start-box in search of a food pellet in a circular arena, before returning to the familiar starting location to consume the pellet. In a probe trial in which rats started from a different location, and the arena was rotated relative to the start box to displace olfactory cues, sighted rats still returned to their familiar starting location, indicating the dominance of visual cues. In the probe trial, blindfolded rats in contrast showed a mixture of behaviours, either returning to the start box, or to positions offset from it by an angle that tracked the rotation of the olfactory cues. These results demonstrate that, in the absence of visual information, rats navigate both via olfactory cues and path integration (Maaswinkel and Whishaw 1999). In attempting to explain the neural underpinnings of mammalian navigation, we are therefore looking for a system which integrates environmental cues from multiple sensory modalities together with self-motion information, to form an internal map of the environment, such that the vectors relating current and goal locations can be inferred.

1.2.3 Place cells, place by direction cells, boundary vector cells

As discussed, place cells are hippocampal neurones with a receptive field tuned to a specific location in the environment (Figure 1.4A). That is, while firing at very low rates the majority of the time, they emit a burst of action potentials as an animal runs through a particular location, known as their 'place field' (O'Keefe and Dostrovsky 1971). Different place cells have place fields in different locations of the environment, such that within a population of place cells, the subset which are active together form a sparse identity code for the animal's location. As a result, place cells were suggested to encode a cognitive map of the environment, of the type envisioned by Tolman (O'Keefe and Nadel 1978). In the decades since their discovery, a large number of experiments have shed light on the determinants and properties of place cell activity. I will here review the major relevant findings, before subsequently turning to the question of the capacity of place cells to underpin the navigational behaviours described above.

1.2.3.1 Properties and determinants of place cell firing; remapping.

Place cells have been identified in a number of species, including bats (Ulanovsky and Moss 2007), monkeys (Ludvig et al. 2004), and humans (Ekstrom et al. 2003). Recordings from tens of place cells shows that their activity accurately encodes the animal's location within an environment, in that the vector of firing rates across the cells can be used to reconstruct the animal's position (Wilson and McNaughton 1993). The preferred firing locations of place

cells can be remarkably stable across time, with consistent spatial tuning seen across recordings separated by as much as 153 days (Thompson and Best 1990), though greater variance in tuning across time has also been observed (Ziv et al. 2013). The spatial modulation of place cells varies along the septo-temporal axis of the hippocampus. Moving towards the temporal pole, fewer cells are spatially modulated, their firing fields are larger, and there is more activity outside of the place field (Jung, Wiener, and McNaughton 1994; Kjelstrup et al. 2008). Not all place cells fire in each environment, with the majority tending to remain silent (Thompson and Best 1989). The number of place cells expressing place fields increases roughly with the logarithm of the size of the environment (Rich, Liaw, and Lee 2014). In larger environments, place cells may fire in multiple, irregularly spread locations, while the diameter of place fields tends to be larger (Fenton et al. 2008).




Figure 1.4 Example firing of spatially modulated neurones of the hippocampal formation. **A-C**, Left: black trace indicates the cumulative position of the animal as it navigates a 1m² screening environment. Overlaid are the spikes from single example neurones: a place cell (**A**), a grid cell (**B**), and a boundary vector cell (**C**). Right: binned firing rate maps of the same data. Hotter colours indicate higher firing rates. The number at the bottom of each ratemap indicates the peak firing rate. **D**, Blue line indicates the binned firing rate of a single head direction cell. Radial value indicates the firing rate in each directional bin.

Visual cues can be seen to exert a clear influence over place cell firing. For example, O'Keefe and Conway removed distal visual cues in various combinations within a curtained environment to infer the influence of each on place cell firing. The firing of some place cells depended on individual visual cues, while others were activated by more complex combinations of multiple cues (O'Keefe and Conway 1978). In addition, if the single visual cue outside of an environment is rotated to a different position, the location of place fields rotate so as to maintain a constant angle to the cue (Muller and Kubie 1987). Indeed, place

cell activity replicates between the compartments in an environment containing two visually identical compartments connected by a corridor (Skaggs and McNaughton 1998). While such results indicate that visual cues are used by the place cell system to self-locate, they are not necessary to maintain place cell activity: location specific firing continues even if recordings are made in total darkness (O'Keefe 1976). Indeed, congenitally blind rats have normal place fields (Save et al. 1998). These results, together with the fact that place cells fire when animals have run specific distances on a stationary treadmill (Pastalkova et al. 2008), suggests that their firing also depends on self-motion inputs. Indeed, a number of results indicate that although individual place cells may be modulated to a greater or lesser extent by one or the other, as a population, both external sensory and internal self-motion inputs determine place cell activity. For example, Gothard and colleagues recorded place cells in rats running on a linear track between a fixed reward site and a moving reward site, such that the length of the track could be varied. Place cells that fired at fixed distances from the moving reward site were observed. Place cells that did so on runs towards the moving goal must have been driven by visual cues, as the distance since leaving the fixed goal was variable between trials. In contrast, place cells that fired at a fixed distance on runs away from the moving site must have done so via path integration, as they could no longer see the moving goal. On trials in which the track was shortened significantly from its normal length, the population vector of activity across recorded place cells 'jumped' abruptly between representations of the two sites, indicating the conflicting influence of sensory and self-motion inputs (Gothard, Skaggs, and McNaughton 1996). Similarly, when visual and self-motion cues were placed in disagreement in a virtual reality environment, the relative influence of the two cue types varied widely across place cells. In around 25% of cells, visual information alone was sufficient to maintain localised firing, whereas in 75% of cells firing also depended on self-motion information (Chen et al. 2013).

Further evidence for the salience of external sensory cues is seen in the tendency of place cells to change the location of their firing fields, to 'remap', in novel environments (O'Keefe and Conway 1978), or following substantive changes to the sensory features of the environment (Muller and Kubie 1987). Smaller modifications of external sensory cues, such as to the odour or colour associated with the enclosure, generally gives rise to 'rate remapping', in which the in-field firing rate is altered without change in the cell's spatial tuning (Anderson and Jeffery 2003; Leutgeb et al. 2005). Other modifications, such as the colour of a distal cue card, or the shape of the enclosure, may evoke little initial remapping, with changes only emerging following prolonged experience (Bostock, Muller, and Kubie 1991; Lever et al. 2002). The variance between environments in the distinct set of place cells active, and their specific firing locations, implies that the place cell system forms a unique

spatial map for each environment encountered by the animal. Importantly, the location of a cell's firing in one environment is apparently unpredictable both from its own firing in a second environment, or from the firing of other place cells in either environment. That is, place fields are not easily predicted from the sensory features of environments, while place cells do not maintain consistent spatial relationships with one another across environments (Muller and Kubie 1987)

1.2.3.2 Computational models of place cells based on external sensory inputs

Investigators have employed computational models to assess the sufficiency of particular combinations of sensory information and network architecture to give rise to the location-specific firing of place cells. The earliest models described place fields as arising entirely from theoretical inputs encoding the distance and direction of the animal to distal sensory cues, either through hard-wired feedforward connections (Zipser 1985), or unsupervised competitive learning (Sharp 1991). While compelling in their simplicity, such models fail to account for the capacity of place cells to update their representation of self-location through self-motion cues, as is seen in the dark for example. While also failing to account for path integration, an alternative set of models can explain place cells' firing in darkness, as stemming from non-visual sensory inputs indicating the presence of borders. Such models followed the observation that, when recordings were made in a rectangular enclosure whose length could be varied on one or both sides, certain place fields stretched in the same dimension along which the enclosure was stretched (O'Keefe and Burgess 1996). That is, place cells appeared to fire at fixed distances to particular borders of the environment. O'Keefe and Burgess proposed a simple model in which place cell activity was described as being the thresholded sum of Gaussian curves tuned to specific distances of the rat from specific boundaries. This model was later updated, hypothesising the existence of 'boundary vector cells', which are tuned to borders at specific allocentric directions and distances. The thresholded linear sum of two or more such boundary vector cells is sufficient to account for the location specific firing of place cells (Hartley et al. 2000). In support of the model, the putative boundary vector cells which would give rise to place fields observed in one environment can be used to predict the place field locations of the same cells in novel environments that differ only in the shape of the enclosure (Hartley et al. 2000). Indeed, cells who fire along specific borders have subsequently been identified in the subiculum (Barry and Burgess 2007; Lever et al. 2009), medial entorhinal cortex (Solstad et al. 2008), and parasubiculum (Solstad et al. 2008) (Figure 1.4C), however cells which fire at a set distance away from specific allocentric borders are rarer (though see Supplementary Material of

(Koenig et al. 2011). The boundary vector cell model has been subsequently updated to include Hebbian learning, such that it can also account for the plasticity of place cell firing seen across time (Barry and Burgess 2007).

1.2.3.3 Computational models based on self-motion inputs

1.2.3.3.1 Continuous attractor network models

Alternative computational models have attempted to explain how place cell firing could emerge from self-motion inputs. McNaughton and colleagues suggested that place cells are connected so as to form a two-dimensional continuous attractor network, through excitatory connections between place cells in which the connection strength decreases as the distance between their preferred firing locations increases (McNaughton et al. 1996; Samsonovich and McNaughton 1997). Together with global inhibition, such connectivity gives rise to a stable ‘bump’ of activity in a hypothetical, topographically-organised sheet of place cells. The model envisions an additional set of neurones tuned to specific movement directions and speeds at specific locations in the environment. Through asymmetric excitatory projections to place cells offset in their preferred movement direction from those that are active at their preferred location, such cells can shift the bump of activity around the sheet in accordance with self-motion inputs specifying the animal’s velocity. That is, the model describes how place cells, together with the ‘shifter cells’, could perform path integration (McNaughton et al. 1996; Samsonovich and McNaughton 1997). Learned associations between distal sensory cues and place cells representing specific locations are used to reset noise accumulated in the system. Such a model is thus attractive in explaining how place cell firing could emerge from a combination of self-motion and sensory inputs. In accordance with the model, place-by-direction cells tuned to specific movement directions at specific locations have been observed in the pre- and parasubiculum (Cacucci et al. 2004). However, that place cells with overlapping fields can respond inhomogeneously to manipulation of the sensory features of the environment (O’Keefe and Conway 1978) suggests they are not organised in a continuous attractor: the excitatory connectivity between place cells in such a network would ensure that all cells are modulated in concert with one another. Further, the remapping of place fields to unpredictable locations between environments is problematic for the putative connectivity of such a network. For each environment, the connectivity between the ‘shifter cells’ and place cells would have to be relearned, and the system be able to switch between connectivity patterns according to which environment the animal was currently in.

1.2.3.3.2 *Oscillatory interference models*

Low-pass filtering of recorded voltage traces can be used to generate the 'Local Field Potential' (LFP) of the region of brain surrounding the electrodes. An epiphenomenon, the LFP likely represents the summation of individually minute electrical currents during spatiotemporally consistent flows of ions across the membranes of large populations of neurones (Buzsáki, Anastassiou, and Koch 2012). LFP recordings in the hippocampus indicate a significant oscillatory tendency, particularly in the theta (5-11Hz) range (Figure 1.5A) (Green and Arduini 1954), likely reflecting temporally coordinated network-level activity patterns (Buzsáki 2002). When animals run through a place field, place cells emit bursts of action potentials which span multiple cycles of the theta oscillation. Strikingly, spikes are emitted at progressively earlier theta phases as the animal runs through the field (Figure 1.5A-C). This phenomenon, termed 'theta phase precession', provides a temporal code for the distance travelled through the place field (O'Keefe and Recce 1993; Huxter et al. 2008). This temporal code is independent of the cell's firing rate code, with the latter better correlating with the animal's in-field running speed (Huxter, Burgess, and O'Keefe 2003). The frequency of theta oscillations in the LFP of the hippocampus increases as a function of the animal's running speed (Sławińska and Kasicki 1998; Jeewajee et al. 2008). As such, theta phase precession can be explained by computational models in which the place cell emits action potential bursts at frequencies which increase with a steeper slope than the LFP frequency-speed relationship: the cell's higher theta frequency causing it to fire at progressively earlier phases in successive theta cycles. As the difference in frequency between the LFP and cellular oscillation varies as a function of the animal's running speed, the amount of phase precession per unit time increases as the animal runs faster, such that distance travelled, rather than time elapsed since entering the field, is encoded (O'Keefe and Recce 1993). This fact entails that oscillatory interference models can also account for the capacity of place cells to show location specific firing based on self-motion inputs alone. For example, place cell firing can be modelled as being determined by a membrane potential that equates to the interference pattern of two oscillations, each coupled to the animal's running speed with differing slopes. When the two oscillations are in phase, their summation exceeds the firing threshold and spikes are emitted, defining the cell's place field. The different frequency-speed slopes of the oscillations means the change in phase-offset between them encodes the animal's speed, such that the cell's membrane potential and firing is updated across time by the distance travelled by the animal (Lengyel, Szatmáry, and Erdi 2003). While attractive in proposing a single-cell account of how place cells could perform path integration, the problem with such models is that, unlike recorded place cells,

modelled cells display periodic firing patterns due to the two oscillators moving cyclically in and out of phase as the animal moves.

Figure 1.5 Theta phase precession examples in place and grid cells. *A*, Light grey: raw LFP trace recorded in area CA1. Darker grey: the LFP trace low-pass filtered around the theta frequency. Black: phase of spikes from a single place cell as an animal runs through the cell's place field. On first entering the field, spikes are predominantly fired at later phases (left hand side). As the animal moves through the field, spikes are seen at earlier phases of the theta cycle (towards right hand side). **B**, **D**, Example ratemap and raw spike plots for the place cell (**B**) and grid cell (**D**) whose phase precession are shown. **C**, **E**, Each data point indicates the phase of the first spike in each theta cycle plotted as a function of the proportion of the distance along the run through the field. Red lines indicate the circular-linear regression line which best fits the data. Panels B-E are adapted from Jeewajee et al. 2014. Panel A is adapted from Pastalkova et al. 2008.

1.2.3.4 Using place cells for navigation

To act as the neural basis of the cognitive map envisioned by Tolman, a population of place cells must be capable of supporting the observed capacity of animals to navigate directly to a goal location. Several results implicate place cell firing in navigation. For example, CA1 place cell representations reorganise to track changes in the position of an un-cued goal

location (Dupret, O'Neill, and Csicsvari 2013). Further, place cells have also been observed to fire at specific distance and direction offsets to goal locations (Sarel et al. 2017). While such results indicate a role for place cells in representing navigational goals and vectors, they do not demonstrate that the vectors separating current and goal locations are computed by place cells. A number of models suggest how place cells could underlie vector navigation, with the majority involving comparison of the current pattern of place cell activity to that seen at the goal (Burgess and O'Keefe 1996; Touretzky and Redish 1996). For example, Burgess & O'Keefe's model uses one-shot Hebbian learning between a 'goal cell' representing the location of a target, and the place cells whose firing field overlap with the goal. As such, to navigate back to the goal, the animal can simply move in the direction which continually increase the activity of the goal cell, which will be driven to fire maximally when all the place cells with which it is wired are active, i.e. at the goal location. However, such models are unsatisfactory in a number of ways. First, as they suggest rats should navigate by progressively testing multiple movement directions, rather than by running in straight lines to the goal, as is reported by Lashley for example. Second, as the maximum navigational distance is limited by the diameter of the largest place field: if the goal is not within the place field of any currently active place cells, the goal cell will have a firing rate of zero, and will therefore be uninformative as to which direction the rat should move in. While place fields as large as 10m have been reported (Kjelstrup et al. 2008), this is still significantly smaller than the estimated foraging distances of rats (100m to 1km (Russell et al. 2005)).

1.2.3.5 Summary

To what extent do the data reviewed indicate that place cells alone underlie Tolman's cognitive map? In favour, even relatively small numbers of place cells can accurately encode locations within certain environments. Further, the activity of place cells depends on both external sensory cues and internal self-motion cues, in line with behavioural evidence that animals can navigate using both. However, it is not clear that place cells alone are capable of, or likely to, support either path integration or goal-directed navigation. The tendency of place cells to remap unpredictably across environments means that the connectivity supporting network-level path integration mechanisms would have to be relearned and adaptively applied across environments. Single-cell oscillatory interference models satisfactorily explain theta phase precession, but are unsatisfactory as models of path integration in suggesting that place cells should fire at multiple, periodically organised firing locations. In addition, models of goal-directed navigation employing place cells cannot easily account for navigation across distances greater than the diameter of the largest place field. Indeed, the identity code used by place cells (which cells are firing) is likely to be a fatally

inefficient scheme for specifying self-location in large environments: the number of distinct states (such as locations in space) which can be represented by an identity code increases only linearly with the number of neurones in the population. As such, the number of place cells required to specify an animal's self-location with constant precision increases linearly with the environment's size. The total number of place cells in the hippocampus of an animal therefore places an upper limit on the size of environment which could be effectively represented. The fact that place cells appear to only be recruited as a logarithmic function of the size of the environment (Rich, Liaw, and Lee 2014), instead suggests that place cells are fundamentally incapable of representing self-location in larger environments. For these reasons, we must look beyond place cells to explain animals' ability to path integrate, self-locate, and navigate, particularly in large environments. The question of what the function of place cells is, if not underpinning these capacities, is returned to in Section 1.2.6.

1.2.4 Head direction cells

Following results demonstrating that place cells can show location specific firing independent of the sensory features of the environment, O'Keefe hypothesised the existence of "*a navigational system... that could calculate subsequent positions in the environment on the basis of how far, and in what direction, the animal had moved in the interim*" (O'Keefe 1976). That is, he hypothesised the existence of a path integration system which incorporates distance and direction signals. Around a decade later, cells tuned to specific movement directions were reported in the presubiculum (Taube, Muller, and Ranck 1990a). These head direction cells fired when the animal moved in a specific direction, regardless of its location (Figure 1.4D). Such cells have since been identified in a number of regions, including the dorsal tegmental nucleus, lateral mammillary nucleus, anterior dorsal thalamus, entorhinal cortex, and retrosplenial cortex (Taube 2007). Lesion studies indicate the existence of a hierarchical circuit generating the head direction signal: lesions of the lateral mammillary nucleus and dorsal tegmental nucleus abolish directional tuning in the downstream anterior dorsal thalamus and presubiculum (Blair, Cho, and Sharp 1998, 1999). In turn, lesions to the anterior dorsal thalamus disrupt direction-selective firing in the presubiculum, but not vice versa (Goodridge and Taube 1997).

Like place cells, head direction cells are influenced by both sensory and self-motion cues: their preferred firing directions are anchored to distal visual cues in the environment, rotating together if the cue is moved (Taube, Muller, and Ranck 1990b). Head direction cells continue to show direction specific firing in the dark, but if the light is turned back on their firing will re-align if the cue has been rotated in the interim (Zugaro et al. 2003). Continuous attractor network models of head direction cells provide convincing accounts of how their

direction-selective firing arises (Skaggs et al. 1995; Zhang 1996; Redish, Elga, and Touretzky 1996). Equivalent to one-dimensional versions of the continuous attractor accounts of place cells described above, such models envision a hypothetical ring of head direction cells connected topographically, with excitatory connectivity between cells that falls off as the angular separation between their preferred firing directions increases. Together with global inhibitory connections, this connectivity profile gives rise to a stable bump of activity in the ring. A second layer of neurones are tuned to specific turning directions, modulated by the velocity of the animal's rotation, with inputs from head direction cells at a specific point on the ring. When activated by turning, these cells project asymmetrically back to the main ring, such that they shift the bump according to the direction and velocity of movement. Recordings in the anterior dorsal thalamus have identified such turning-velocity modulated cells (Blair and Sharp 1995; Taube 1995). Systems such as this, which integrate noisy signals across time, are sensitive to the accumulation of error. Models of head direction cells therefore use Hebbian plasticity between visual cues and head direction cells to learn direction-specific sensory input patterns which can reset the network periodically (Skaggs et al. 1995; Zhang 1996).

Continuous attractor network models well account for the direction-selective tuning of head direction cells, and their modulation by both internal self-motion and external sensory cues. That head direction cells are found throughout the hippocampal formation, and in areas with which it is highly interconnected, suggests that their encoding of the direction of current movement is central to the brain's representation of space, probably in supporting path integration. However, head direction cells are incapable of representing linear distances moved by an animal, such that, for a full account of path integration, we must look elsewhere.

1.2.5 *Grid cells*

1.2.5.1 *Properties and determinants of grid cell firing patterns*

Given its extensive connectivity with the hippocampus, in-vivo recordings were frequently made in the entorhinal cortex following the discovery of place cells. Early attempts observed cells whose firing was modulated by the animal's location in space, but with lower signal to noise than is seen in place cells (Quirk et al. 1992). Recordings throughout the entorhinal cortex showed a clear medio-lateral dissociation: medial entorhinal neurones were seen to display highly localised spatial firing, unlike those of the lateral entorhinal cortex (Hargreaves et al. 2005). Strikingly, recordings from more dorsal areas of the medial entorhinal cortex revealed cells which fire at multiple locations within the environment (Fyhn et al. 2004).

Indeed, when recorded in a larger enclosure, the individual firing fields of these cells form a repeating triangular pattern which tessellates to cover its extent (Figure 1.4B) (Hafting et al. 2005). That is, the periodic firing pattern forms a triangular grid, leading Hafting and colleagues to coin the name 'grid cells'. Subsequent investigations have identified grid cells in a range of species including mice (Fyhn et al. 2008), bats (Yartsev, Witter, and Ulanovsky 2011) and humans (Doeller, Barry, and Burgess 2010; Jacobs et al. 2013). Grid cells are seen in all layers of the entorhinal cortex (Sargolini et al. 2006), as well as in the pre- and parasubiculum (Boccarda et al. 2010).

Figure 1.6 Properties of grid cell firing patterns; organisation of grid cells into modules. **A**, Properties of grid cell firing patterns indicated on the ratemap of a single grid cell. Each cell's firing is defined by a phase (left: offset in each dimension of the peaks from some origin: ϕ_x, ϕ_y), scale (centre: distance between firing peaks: λ) and orientation (right: angle of the major grid axis relative to the horizontal: θ). **B**, Left: a Nissl-stained sagittal section through a mouse brain, with the medial entorhinal cortex encircled by a dashed line. Right: example firing rate maps from four grid cells, one pair from each of two example modules. As indicated, smaller scale grid cells are predominantly recorded in dorsal areas of the entorhinal cortex, larger scale grids are found in more ventral areas. The image in panel B is adapted from Paxinos & Franklin 2012.

The firing patterns of individual grid cells are parameterised by an orientation (angle of the major axes of the grid), scale (distance between the firing peaks), and phase (location of the firing peaks) (Figure 1.6A). When initially discovered, it was reported that co-recorded grid cells at particular positions in the entorhinal cortex's dorso-ventral axis tended to have identical orientations and scales, but different firing field locations (Hafting et al. 2005). Moving along the dorso-ventral axis, the scale of observed firing patterns increases (Hafting et al. 2005) (Figure 1.6B), matching the increase in scale of place cells seen along the

hippocampus' septo-temporal axis. Further investigations suggested that the scale of grid cells increases in discrete jumps along the dorso-ventral axis (Barry et al. 2007; H. Stensola et al. 2012). These observations of discontinuous jumps in grid scale stem from progressive recordings being made from distinct 'modules' of grid cells (Figure 1.6B). Each module shares a common orientation and scale, with its constituent cells displaying internal consistency in their firing patterns, but with distinct firing phases (H. Stensola et al. 2012). Recordings from the majority of the dorso-ventral axis, made while rats run on an 18m linear track, suggest the existence of 8-10 grid modules, with the most ventral module displaying a grid scale of 2-3m (Brun, Solstad, et al. 2008). When moved to a novel environment, a cell's firing pattern initially expands in scale (Barry, Ginzberg, et al. 2012), before settling back to the scale seen in the familiar environment, but with a distinct orientation and phase. Within a module, pairs of grid cells show consistency in their scale, orientation, and phase offset relationships between environments (Fyhn et al. 2007; Yoon et al. 2013).

As in place and head direction cells, both external sensory and internal self-motion cues modulate grid cell activity. For example, the orientation of grid cell firing is anchored to distal sensory cues: as the angle from a single cue card to the enclosure is altered, grid patterns rotate in concert (Hafting et al. 2005). Changes in grid firing patterns seen following translation or rotation of the enclosure's position within the recording room provides similar evidence for the anchoring of grid patterns to remote cues (Savelli, Luck, and Knierim 2017). In rats, the maintenance of the grid pattern in total darkness has been taken as evidence that path integration alone is sufficient to maintain normal grid activity, at least over short timeframes (Hafting et al. 2005). However, in mice, the contribution of visual cues may be greater, in that grid firing patterns quickly break down in total darkness (Chen et al. 2016; Pérez-Escobar et al. 2016). As in place cells, the boundaries of the enclosure appear to exert a particular influence over grid firing. Replicating the paradigm of O'Keefe & Burgess (1996) previously discussed, Barry et al. observed a similar anisotropic stretching of grid firing patterns that matched the axis of change in the enclosure length (Barry et al. 2007). Indeed, if a square environment is divided into multiple corridors connected by sequential hairpin turns, grid firing breaks into multiple fragments, often repeating across corridors in which the animal runs in the same direction (Derdikman et al. 2009). The shape of the enclosure can also determine the orientation of the firing pattern, both locally (T. Stensola et al. 2015) and globally (Krupic et al. 2015), depending on the size of the environment. Finally, the symmetry of the triangular pattern is disrupted in certain environments, such as the narrow end of a trapezoidal enclosure (Krupic et al. 2015). Grid cell firing patterns are thus determined both by distal sensory cues and the geometry of the enclosure, though may also be maintained by self-motion inputs alone, at least in some species.

1.2.5.2 Models of grid cells

From their initial discovery, the spatially periodic patterns expressed by grid cells – their firing every time the animal runs a specific distance in a specific direction – led to models which described their activity as arising from self-motion information. These models formulate how individual grid cells, or networks of grid cells, perform path integration based on inputs describing the animal's running speed and heading direction. Two major classes of grid cell model have been proposed, both extensions of equivalent models, discussed above, which describe how place cells could perform path integration.

1.2.5.2.1 Continuous attractor network models

Continuous attractor network models of grid cells are extensions of the one- and two-dimensional networks described, which were proposed to account for place and head direction cell updating via self-motion inputs. Though differing from one another in their details, continuous attractor models are unified in involving grid cells organised in a hypothetical two-dimensional topographic sheet, with the weight of excitatory connections between grid cells decreasing as the spatial offset of their firing locations increases (Figure 1.7A) (Fuhs and Touretzky 2006; McNaughton et al. 2006; Burak and Fiete 2009). In the model of McNaughton and colleagues for example, this pattern of local-excitation, together with longer-distance inhibitory connections, gives rise to a stable bump of activity on the sheet (Figure 1.7B) (McNaughton et al. 2006). Connections at the edges of the sheet 'wrap around', joining cells at opposite edges, such that each location on the sheet is equivalent to a position on the surface of a three-dimensional torus (Figure 1.7A, C). The bump of activity on the torus can be shifted according to the speed and direction of movement, by a second population of cells tuned to specific movement velocities at specific locations in the environment. Such cells receive inputs from grid cells at a specific point on the torus, and project back with asymmetric excitatory connections, to a nearby point offset in the direction of their velocity tuning (Figure 1.7D). As such, the heading direction and running speed information encoded by the firing rates of these 'shifter cells' can be used to move the bump of activity around the network of grid cells, updating their representation of self-location, and thus performing path integration (Fuhs and Touretzky 2006; McNaughton et al. 2006; Burak and Fiete 2009).

In favour of continuous attractor models, cells with triangular firing patterns which are also tuned to specific running directions and speeds, as is required of the 'shifter cells', have been described in the deep layers of medial entorhinal cortex (Sargolini et al. 2006). Further, grid cells within a module maintain fixed spatial offsets between their firing locations when

recorded in multiple environments (Fyhn et al. 2007; Yoon et al. 2013). The consistent spatial relationships between grid cells suggests that the connectivity between them is maintained across environments, implying the same continuous attractor underlies the activity in each. Finally, such models suggest that when moving through a firing field, the membrane potential of grid cells ramps up smoothly, as is observed in intracellular recordings from mice navigating in virtual reality (Schmidt-Hieber and Häusser 2013; Domnisoru, Kinkhabwala, and Tank 2013).

Figure 1.7 Illustration of principles of continuous attractor network models of grid cells. *A*, Illustration of the organisation of grid cells in a module into a two-dimensional sheet, with cells' position on the sheet defined by their phase of firing. Red arrows indicate the excitatory connectivity that exists between nearby cells of similar firing phase. Local excitatory connectivity is combined with more distant inhibitory connectivity in a Mexican-hat profile (*B*). At the edges of the sheet of cells, connections 'wrap around', such that positions on the two-dimensional sheet map to the surface of a three-dimensional torus, as is illustrated in *C*. *D*, Top layer represents the sheet of grid cells, who project to a separate population of 'shifter cells' in a location-specific manner, with the shifter cells also tuned to the animal's movement velocity. Once activated, the shifter cells shift the bump of activity in the sheet of grid cells through asymmetric excitatory projections onto grid cells offset in the direction of their velocity tuning from those they receive input from. Panel A, C and D are adapted from McNaughton et al. 2006. Panel B is adapted from Burak & Fiete 2009.

While criticisms have been made of continuous attractor networks, most relate to details of specific instantiations of particular models, rather than the concept of attractor networks themselves. For example, while ocean-stellate cells of LII of the entorhinal cortex do not show recurrent excitatory connections, grid cell firing patterns can also emerge from

networks with disynaptic inhibitory connections between grid cells (Couey et al. 2013). Similarly, while parvalbumin positive interneurons in medial entorhinal cortex do not show grid cell firing patterns, as is predicted by most continuous attractor network models (Buetfering, Allen, and Monyer 2014), other instantiations can produce normal grid cell activity with only weak spatial modulation amongst interneurons (Solanka, van Rossum, and Nolan 2015). Perhaps more problematically, the connectivity profiles required to support continuous attractor networks are highly complex, and the mechanisms by which such connections could develop are unclear. However, unsupervised learning using spike-time-dependent plasticity has been demonstrated to be capable of producing the required connectivity, though only if direction and location specific inputs are provided (Widloski and Fiete 2014).

1.2.5.2.2 Oscillatory interference models

As seen, oscillatory interference models of place cell firing are hampered by the tendency of the envelope of the interference pattern between two speed-modulated oscillators to be periodic. Of course, in explaining the cyclical activity of grid cells, this fact is a clear advantage of such models. In one-dimension, oscillatory interference models of grid cell firing simply replicate the proposals of previous place cell models (O'Keefe and Burgess 2005). For example, Burgess and colleagues envisioned two membrane potential oscillations, one somatic and one dendritic. The frequency of both oscillations increases linearly with running speed, but the dendritic oscillation with steeper slope. As such, the two oscillations move in and out of phase as the rat runs unidirectionally, such that their interference pattern causes the cell's membrane potential oscillation to periodically surpass its firing threshold, giving rise to the cell's firing fields. As the difference in frequency between the two oscillations increases with running speed, at higher speeds the two move in and out of phase at a higher rate, such that the cell's spatial firing tracks the distance travelled by the animal, rather than its speed (Burgess, Barry, and O'Keefe 2007). To generate two-dimensional firing, the oscillators' frequencies must also vary as a function of the animal's heading direction. That is, two-dimensional oscillatory interference models envision so-called velocity controlled oscillators, whose frequency varies as a linear function of running speed and as a cosine function of heading direction. If a single cell receives inputs from two or more velocity controlled oscillators, whose preferred firing direction are multiples of 60 degrees from one another, the interference pattern between them forms the characteristic triangular pattern seen in grid cells. Unsupervised learning between velocity controlled oscillators could give rise to such wiring, if the strength of their connections onto downstream grid cells were modulated so as to maximise post-synaptic activity (assuming

that pairs of oscillators with co-linear preferred firing directions are excluded) (Burgess, Barry, and O'Keefe 2007; Mhatre, Gorchetchnikov, and Grossberg 2012). In such models, the scale of the firing pattern expressed by a grid cell is determined by the difference in the slopes of the theta frequency vs running speed relationships of its component oscillators. A large difference in slope means the two oscillators move in and out of phase more quickly at each running speed, such that the scale of the cell is smaller than in cells with two oscillators of similar slope.

Oscillatory interference models therefore describe how path integration could be performed by individual grid cells based on self-motion inputs in which the speed and direction of movement is encoded in the frequency of theta oscillations (Blair, Gupta, and Zhang 2008; Hasselmo, Giocomo, and Zilli 2007). In favour of such accounts, the membrane potential of ocean-stellate cells of LII of the medial entorhinal cortex shows resonance at theta frequencies (Klink and Alonso 1997), while putative grid cells also display subthreshold theta oscillations in intracellular recordings made while mice navigate in virtual reality environments (Domnisoru, Kinkhabwala, and Tank 2013; Schmidt-Hieber and Häusser 2013). Indeed, a gradient in the resonant frequency of ocean-stellate cells along the dorso-ventral axis of the medial entorhinal cortex matches the gradient seen in grid cell scale along the same axis (Giocomo et al. 2007). The resonant frequency of these neurones is dependent on hyperpolarisation-activated cyclic nucleotide-gated (HCN1) channels (Giocomo and Hasselmo 2009), and mice with a genetic knockout of the HCN1 channel also show alterations in grid scale (Giocomo et al. 2011). Further, pairs of grid cells more similar in scale tend to fire bursts of action potentials at more similar theta frequencies (H. Stensola et al. 2012), while changes in the burst frequency of grid cells in novel environments covaries with the observed increase in grid scale (Barry, Ginzberg, et al. 2012). Perhaps most convincingly, inactivation of the medial septum, either with lidocaine or muscimol, abolishes theta oscillations, and simultaneously disrupts the periodic firing of grid cells (Brandon et al. 2011; Koenig et al. 2011). In addition, cells whose firing matches that of hypothesised velocity controlled oscillators have been observed in the medial septum and anterior dorsal thalamus (Welday et al. 2011). Cells with band-like firing patterns, potential 'building blocks' of grid firing patterns, and as would be observed in cells with input from a single velocity controlled oscillator, have also been recorded, in the entorhinal cortex and parasubiculum (Krupic, Burgess, and O'Keefe 2012). Oscillatory interference models can also account for the tendency of grid cells to show theta phase precession as animals run through their firing fields (Hafting et al. 2008; Jeewajee et al. 2014).

However, oscillatory interference models have been widely criticised, particularly for their reliance on precisely timed oscillations, especially given the noise in experimentally

observed oscillations (Giocomo and Hasselmo 2008; Dodson, Pastoll, and Nolan 2011; Fiete 2010). Further, intracellular recordings indicate that the location of firing fields of putative grid cells is better explained by depolarising ramps in their membrane potential, rather than the envelope of subthreshold theta oscillations (Schmidt-Hieber and Häusser 2013; Domnisoru, Kinkhabwala, and Tank 2013). In addition, grid cell firing patterns have been recorded in crawling bats, despite the absence of discernible theta oscillations (Yartsev, Witter, and Ulanovsky 2011), though these findings are disputed (Barry, Bush, et al. 2012).

While not discussed in detail here, both oscillatory interference and continuous attractor network models use learned conjunctions of location-specific sensory features to anchor firing patterns and minimise the accumulation of error stemming from the integration of noisy self-motion inputs. Apparently similar error-correction systems have been observed in grid cell recordings, with noise accumulating and then resetting between visits to the borders of the enclosure (Hardcastle, Ganguli, and Giocomo 2015).

1.2.5.2.3 Hybrid models

Though oscillatory interference and continuous attractor network models are often considered to be competing accounts of grid cell firing, many of their properties are not mutually exclusive, such that a number of hybrid models including features of both have been proposed (Navratilova et al. 2012; Hasselmo and Brandon 2012; Bush and Burgess 2014). For example, Bush & Burgess described a model in which velocity controlled oscillators are wired together in ring attractors, with each attractor sharing a preferred running direction, but with cells differing in the phase of their oscillations. Grid cells sharing inputs from velocity controlled oscillators with the same firing phases (and which therefore have overlapping firing fields) are in turn connected through recurrent inhibition in an attractor network. Such hybrid models are attractive in accounting for a wider range of experimental data than either type of model alone, such as both theta phase precession and the consistent phase offsets seen between pairs of cells across environments.

1.2.5.3 *Encoding across grid modules*

The models described above indicate that grid cells are well suited to perform path integration, suggesting that this may be their functional contribution to the hippocampal formation's spatial processing. However, to understand the genuine value of grid cells, it is necessary to consider what information downstream networks could decode from their representations. Alternative accounts of how populations of grid cells could encode the animal's self-location across modules demonstrate how the modular organisation and

periodic nature of grid cell firing can act as an extremely efficient encoding system, either to represent large environments (Fiete, Burak, and Brookings 2008) or small environments at high resolution (Mathis, Herz, and Stemmler 2012). That is, these theories indicate how populations of grid cells could form efficient cognitive maps of the environment, as envisioned by Tolman.

The orientation and scale of firing patterns within a module is consistent, but the firing locations of different grid cells are distributed evenly across space (Hafting et al. 2005). A grid module therefore encodes self-location as a pair of phases in a two-dimensional phase-space: which cells are active. Because grid cell firing is periodic, the phase of each module repeats once an animal traverses a distance larger than its grid scale. A repetition in phase mean that the module's encoding of the animal's location is ambiguous: as the module has the same phase in two locations the animal could be at either. However, the firing phase across modules of different scale can be combined to resolve this ambiguity. For example, in one dimension, a cell from a module with scale 20cm will fire every time the animal moves 20cm. The phase of a second module of 18cm will differ between adjacent firing fields of the cell in the first module however, enabling a decoding network to distinguish between them, even though the second module is itself ambiguous. Depending on the relative scale of the modules, the conjunction of phases across multiple modules (the 'population phase'), can be an extremely efficient code for self-location (Fiete, Burak, and Brookings 2008). Specifically, if grid scales across modules take integer values that are not a consistent multiple of one another, the number of unique population phases can be calculated as the lowest common multiple of their scales, according to the Chinese Remainder Theorem (Fiete, Burak, and Brookings 2008). Doing so demonstrates that the population phase could encode an animal's location with extreme efficiency, with the number of unique locations that can be encoded increasing exponentially with the number of grid modules.

Like the phases of single modules, the population phases of a set of grid modules representing a two-dimensional environment can be mapped onto the surface of a three-dimensional torus. Doing so illustrates an important feature of the 'combinatorial' system of encoding described above. Specifically, adjacent positions in population phase space map to non-adjacent positions in real space and vice versa (Sreenivasan and Fiete 2011). As such, small errors in the decoded population phase can give rise to catastrophically large errors in estimates of the animal's location (Mathis, Herz, and Stemmler 2012). As such, in environments similar in size to the encoding capacity of the grid network, the system will be intolerably sensitive to noise. However, if the navigable environment is considerably smaller, then an error-correction system can be envisioned. For example, if noise leads the population phase to encode the animal's self-location as being outside the navigable

environment, then the network may simply reset the population phase to the nearest plausible value (Sreenivasan and Fiete 2011).

Alternatively, 'nested' systems propose that grid cells encode an animal's self-location in environments smaller than the largest grid scale, but with high resolution (Mathis, Herz, and Stemmler 2012). Here, as the largest grid scale is larger than the size of the environment, its phase does not repeat, such that it provides an unambiguous code for the animal's location. A constant integer ratio between grid scales means that modules are iteratively 'nested' within the scale of the next largest module, specifying the animal's location at increasingly high resolutions. Such a system is equivalent to an analogue clock, in which the unambiguous but coarse unit of hours is subdivided sequentially into minutes and seconds. As animals are thought to forage over areas of 100m-1km (Russell et al. 2005), and the largest grid scale is 2-3m (Brun, Solstad, et al. 2008), the feasibility of such nested systems remains to be seen. Further, as the difference in grid scale across modules can be in the range of orders of magnitude in nested codes, it is not clear that a unitary velocity signal could provide the resolution required to accurately update both the smallest and largest grid scales concurrently (Fiete, Burak, and Brookings 2008).

More recently, an intermediate system in which the ratio of grid scales between modules is a fixed but non-integer value has been considered in detail (Wei, Prentice, and Balasubramanian 2015; Stemmler, Mathis, and Herz 2015). The non-integer ratio between scales means smaller modules do not perfectly subdivide larger ones, ensuring the population phase is unambiguous in environments larger than the largest scale. The number of unique population phases is smaller than the lowest common multiple of the constituent grid scales however (i.e. the Chinese Remainder Theorem is not obeyed), but this reduced capacity can be used as an alternative error-correction system. The fixed ratio of grid scales means that certain conjunctions of phase across module are invalid, a fact which could be used to correct errors stemming from the independent accumulation of noise across modules (Stemmler, Mathis, and Herz 2015).

In sum, the periodic nature of their firing means that by combining activity across multiple modules of different scale, grid cells can encode an animal's self-location with extreme efficiency. Depending on the relative scale of modules within the network, this efficiency could encode an animal's location in very large environments or in smaller environments with high resolution.

1.2.5.4 Using grid cell representations for navigation

As discussed, animals appear capable of navigating in straight lines to goal locations, even via routes not previously taken (Lashley 1929; Tolman, Ritchie, and Kalish 1946). A number of recent models have proposed how a population of grid cells could calculate the vector separating current and goal locations, each relying on the fact that grid cell's periodic firing means that particular distance and direction offsets in real-space equate to a particular offset in population phase space. Erdem & Hasselmo hypothesised that constant-speed look-ahead sweeps of activity could emanate from the grid cell representation of the animal's location. Once the population vector of activity in one of the sweeps matched that of the goal location, the direction of the sweep and the time taken for a match to occur could be used to infer the direction and distance to the goal respectively (Erdem and Hasselmo 2012, 2014). Bush and colleagues proposed three models of how biologically plausible decoding networks could directly calculate the vector to a goal from the grid cell population phase. In perhaps the most compelling, individual 'rate-coded vector cells' are each activated by a specific offset in population phase space (Bush et al. 2015). The rate-coded vector cells thus transform a pair of population phases specifying current and goal locations into an identity code for the real-space vector separating the two. Depending on the relationship between grid scales, such decoding could operate across grid modules (Bush et al. 2015), or independently in each (Edvardson 2016). Finally, Stemmler and colleagues note that the population phase encoding each point in space is effectively measured relative to an arbitrary origin (i.e. where the phase of each module is 0). By multiplying the phase of each module by an appropriate sinusoidal grating, they suggest that the arbitrary reference could be shifted between goal locations, such that the population phase continually encodes the animal's vector displacement from its goal (Stemmler, Mathis, and Herz 2015). While each of the proposed decoding networks suffers from its own particular drawbacks, they also demonstrate that a population of grid cells encodes sufficient information for biologically plausible networks to decode the vector separating an animal's location from its goal.

1.2.5.5 Grid cells as input to place cells

The highest density of grid cells is seen in the superficial layers of the medial entorhinal cortex (Sargolini et al. 2006). As the hippocampus receives dense input projections from these layers, it has been suggested that hippocampal place cells are the probable decoders of grid cell activity (Sreenivasan and Fiete 2011; O'Keefe and Burgess 2005; Fuhs and Touretzky 2006). A number of simple models propose how the location-specific activity of place cells could arise from the linear summation ('Fourier synthesis') of inputs from specific, spatially-overlapping grid cells of different scale. Specifically, such models propose that the

strength of input from each grid cell is proportional to its degree of overlap with the output place field. Different accounts propose various mechanisms by which the required connectivity could arise, including hard-wiring (Solstad, Moser, and Einevoll 2006), heterosynaptic Hebbian learning (Savelli and Knierim 2010), and competitive network interactions (Fuhs and Touretzky 2006).

Despite the simplicity of the connections which can generate place cell firing from a small number of grid cells, a range of experimental results suggest that place fields are not simply the product of the summation of multiple grid patterns (Bush, Barry, and Burgess 2014). First, the location-specific firing of place cells is seen at earlier stages of development than the triangular firing of grid cells (Wills et al. 2010; Langston et al. 2010). Further, Fourier synthesis models suggest that removal of inputs from smaller-scale grid cells in dorsal medial entorhinal cortex should increase the size of place fields, a result not observed experimentally (Miao et al. 2015). In addition, place cell firing continues in conditions in which inputs from grid cells are disrupted, including inactivation of the medial septum (Brandon et al. 2014; Wang et al. 2015) and lesions to the entorhinal cortex (Brun, Leutgeb, et al. 2008; Schlesiger et al. 2015). Finally, neurones in medial entorhinal cortex tend to show peak firing later in the theta cycle than principal neurones of the hippocampus, making a causal relationship unlikely (Mizuseki et al. 2009). Bush and colleagues therefore argued that place cell firing likely arises predominantly from boundary vector cell and direct sensory inputs, with grid cell connections helping to maintain place cell firing through path integration, as is required in contexts of novelty and sensory deprivation (Bush, Barry, and Burgess 2014).

1.2.6 Summary: contributions of spatially-tuned neurones to the cognitive map

As discussed, the identity code employed by place cells means that they are an inefficient means of encoding an animal's self-location. Further, their tendency to remap to unpredictable locations constrains their ability to perform path integration and vector navigation, in that the complex connectivity required would have to be relearned and adaptively applied across environments. In contrast, the periodic, modular activity of grid cells allows them to encode an animal's self-location with extreme efficiency, even in large environments. In addition, a consistent relationship between the population phase of grid cells and real space can support navigation, in enabling the decoding of vectors separating locations in the environment. Importantly, the consistent spatial offsets between pairs of cells across environments means the complex wiring patterns that support decoding need not be relearned.

If grid cells are capable of performing path integration, self-location, and navigation, to what end does the brain employ a place cell code? One possible explanation is that place cells are of value in providing a mechanism by which specific locations in space can be associated with non-spatial information, such as sensory inputs or events. While efficient, the grid cell code is highly distributed, such that the learning rules and connectivity patterns required to support bidirectional associations between locations and non-sensory features would be nontrivial. In contrast, place cells can easily form these associations through simple Hebbian learning. Having formed such associations, input from place cells could then reset accumulated noise in, and anchor grid cell firing patterns, as well as reactivate goal locations for navigation. In addition, location-specific associations formed by place cells could also form the basis of episodic memory, by encoding non-spatial features of mnemonic episodes within a spatial framework. In sum, the ‘distributed’ and ‘identity’ encoding schemes of grid and place cells may form complementary cognitive maps of the environment, with grid cells supporting self-location and navigation, and place cells enabling the association of specific locations with non-sensory features of the environment.

1.3 Contributions of the medial septum to grid cell activity

1.3.1 Anatomy, connectivity and neuronal subtypes of the medial septum

The basal forebrain is a loose collection of heterogeneous structures in the ventral and medial aspect of the telencephalon. The medial septum, one of the basal forebrain’s constituent nuclei, sits towards its anterior and ventral extremes. Horizontal and coronal sections reveal the medial septum to be chevron shaped: unitary about the midline in anterior and dorsal regions, but dividing bilaterally into a discrete ventro-lateral region in each hemisphere more posteriorly (Figure 1.8). In the coronal plane, the medial septum is generally divided into two regions, according to the distinct connectivity pattern of each. In this thesis, and as is common in the field (Swanson and Cowan 1979), the medial septum here refers both to the dorsal most region (the medial septum proper), as well as the more ventrally extending vertical and diagonal bands of Broca, due to their consistent connectivity. The most ventral and lateral regions seen more posteriorly, which exhibit distinct anatomical connections, are termed the horizontal bands of Broca. The medial septum sends dense projections to all subfields of the hippocampus and dentate gyrus (Swanson and Cowan 1979), and entorhinal cortex (Alonso and Köhler 1984), via the fornix. Further projections are seen to preoptic areas, the lateral hypothalamus, and the mammillary complex; as well as midbrain aminergic regions such as the ventral tegmental area and the raphe nucleus (Swanson and Cowan 1979). In contrast, the horizontal band of Broca sends projections predominantly to

the habenula and hypothalamus (Swanson and Cowan 1979). Both areas tend to show a large degree of reciprocal connectivity, with the inputs to each arising mainly from areas which they also send efferent projections to (Swanson and Cowan 1979; Alonso and Köhler 1984).

Early immunohistochemical analyses identified a population of cholinergic neurones within the medial septum (Kimura et al. 1980). These acetylcholine synthesising neurones were subsequently shown to project to the hippocampus (Mesulam et al. 1983) and entorhinal cortex (Alonso and Köhler 1984). Only a subset of all projections from the medial septum were identified as cholinergic however (Alonso and Köhler 1984), and later studies demonstrated the existence of an additional population of GABAergic neurones (Panula et al. 1984), which were subsequently shown to also project to the hippocampus and entorhinal cortex (Köhler, Chan-Palay, and Wu 1984). Only more recently has a third population of neurones been identified, expressing the neurotransmitter glutamate, and again projecting broadly throughout the hippocampal formation (Manns, Mainville, and Jones 2001). Each of the three types of neurone appear to be highly interconnected with each other within the medial septum. Electron microscopy together with immunohistochemical staining demonstrates the presence of bidirectional synaptic connections between cholinergic and GABAergic neurones (Leranth and Frotscher 1989). Both neurone types also show AMPA receptor mediated glutamatergic inputs from within the septum (Manseau, Danik, and Williams 2005).

Figure 1.8 Position and shape of the medial septum in the mouse brain. **A-B**, Coronal (**A**) and horizontal (**B**) mouse brain slices stained for acetylcholinesterase. Dashed black lines indicate the outline of the medial septum. D-V: dorso-ventral axis. M-L: medio-lateral axis. A-P: antero-posterior axis.

1.3.2 *A role for the medial septum in spatial processing...*

1.3.2.1 *...Through the encoding of running speed*

Unsurprisingly given its degree of reciprocal connectivity with the hippocampal formation, the medial septum appears salient to the capacity of animals to perform spatial tasks.

Excitotoxic lesions of the medial septum impair rats' ability to learn to navigate to a hidden platform in the Morris water maze (Hagan et al. 1988). Similarly, electrolytic lesions of the medial septum prevent rats from learning to identify a goal purely by its location in the room (Winson 1978). In the latter experiment, the deficit was suggested to be spatial rather than purely mnemonic, in that lesioned rats still recognised the goal when they came upon it by chance (Winson 1978). However, both tasks contain a significant mnemonic component (for example in encoding the goal location), such that neither rules out the possibility that the medial septum's contribution to such behaviours is predominantly through supporting learning and memory. More recently, temporary inactivation of the medial septum has been shown to impair rat's estimations of linear distances travelled. While the task contained a mnemonic component, in that rats had to stop at specific remembered distances along a linear track, that the deficit was observed particularly when rats had to estimate longer distances suggests the impairment was primarily spatial (Jacob et al. 2017). The effect of medial septum inactivation on the hippocampal formation's representations of space gives credence to the idea it supports the encoding of distance travelled. First, inactivation of the medial septum with lidocaine or muscimol causes the triangular firing pattern normally seen in grid cells to break down (Koenig et al. 2011; Brandon et al. 2011). Importantly, the directional tuning of grid cells is maintained during medial septum inactivation (Koenig et al. 2011; Brandon et al. 2011). Second, the distance-specific firing of place cells on a running wheel, but not their location-specific firing on the arms of a horseshoe maze, is disrupted by septal inactivation (Wang et al. 2015). That its inactivation also modifies the firing rate vs running speed profile of speed-modulated entorhinal neurones (Hinman et al. 2016), suggests that the medial septum contributes to distance encoding through provision of one or more signals specifying the animal's running speed (Hinman et al. 2016; Hayman and Burgess 2016). The exact roles of the medial septum and its constituent neuronal subtypes in supporting normal grid cell function remain unclear however.

One possibility is that the medial septum is central to grid cell function in its role as a 'pacemaker' of theta oscillations. As discussed, theta oscillations are paramount to the oscillatory interference class of models describing how grid cell firing could arise from multiple theta oscillations encoding an animal's movement velocity in their relative frequencies (Burgess, Barry, and O'Keefe 2007). Evidence of an involvement of the medial

septum in the rhythmogenesis of theta oscillations is longstanding. For example, lesions of the medial septum abolish theta oscillations (Green and Arduini 1954; Petsche and Stumpf 1960), electrical stimulation of the medial septum induces theta oscillations (Green and Arduini 1954; Petsche and Stumpf 1960; Gray and Ball 1970), and recordings from neurones in the medial septum demonstrates a tendency amongst a subset to emit bursts of action potentials at theta frequencies (Petsche, Stumpf, and Gogolak 1962; Ranck 1973; King, Recce, and O'Keefe 1998). Progress in elucidating the contributions of the distinct neuronal subgroups of the medial septum has been hampered by the existence of multiple types of theta oscillation however, which may or may not share common mechanisms. 'Type I' theta is observed during volitional movement, where it changes in frequency with running speed (Sławińska and Kasicki 1998), and is insensitive to the muscarinic receptor antagonist atropine (Kramis, Vanderwolf, and Bland 1975; Bland et al. 1984). Given its atropine sensitivity, Type II theta is thought to depend on cholinergic activity, but the involvement of acetylcholine in Type I theta is less clear. Lesions of cholinergic neurones in the medial septum reduce Type I theta power, but do not abolish it (Lee et al. 1994; Yoder and Pang 2005). Optogenetically stimulating cholinergic neurones in the medial septum in behaving mice spares theta frequencies while adjacent frequencies are attenuated (Vandecasteele et al. 2014), and muscarinic receptor blockade flattens the normal increase in theta frequency with running speed (Newman et al. 2013). That cholinergic neurones fire at low rates, have long after-spike hyperpolarisation potentials (400-700ms), and only small I_h currents (Markram and Segal 1990), suggests they are unlikely to be the ultimate 'pacemakers' of theta oscillations however. As is seemingly confirmed by in-vivo juxtacellular recordings showing that septal cholinergic neurones do not burst at theta frequencies (Simon et al. 2006). This pacemaker role is generally attributed to the GABAergic neurones, given their high firing rates, short hyperpolarisation potentials, and large I_h currents (Sotty et al. 2003), as well as their tendency to burst at theta frequencies in-vivo (Borhegyi et al. 2004; Simon et al. 2006). Further, GABAergic neurones lead hippocampal theta (Hangya et al. 2009), and are extensively connected with hippocampal interneurones (Freund and Antal 1988), which in turn pace theta oscillations (Amilhon et al. 2015). While the glutamatergic neurones can also pace theta oscillations (Fuhrmann et al. 2015; Robinson et al. 2016), this depends on local connectivity within the medial septum, probably through connections onto GABAergic neurones (Fuhrmann et al. 2015).

An alternative possibility is that the medial septum supports grid cell function through encoding the animal's running speed in the firing rates of neurones. A range of entorhinal neurones exhibit increases in firing rate as animals move faster (Sargolini et al. 2006; Sun et al. 2015; Kropff et al. 2015). As seen, in the continuous attractor class of grid cell models,

such activity is combined with directional information to shift the grid cell network's representation of self-location as animals move (Fuhs and Touretzky 2006; Burak and Fiete 2009). Firing rate coding of speed has been observed in glutamatergic neurones of the medial septum (Fuhrmann et al. 2015), and their entorhinal-projecting axons (Justus et al. 2017). Interestingly however, inactivation of the medial septum with muscimol strengthens the speed-rate coding of entorhinal neurones (Hinman et al. 2016). Alternatively, the medial septum may not itself encode the animal's running speed, but instead regulate the activity of cortical regions whose function contributes to the estimation of movement velocity. For example, Dannenberg and colleagues have proposed that septal cholinergic projections, which modulate activity in visual cortical areas, could influence the processing of optic flow, which in turn alters running speed estimation (Dannenberg, Hinman, and Hasselmo 2016).

1.3.2.2 ...Through novelty signalling

A possible involvement of medial septal cholinergic neurones in learning and memory has been intensely studied following the identification of reduced choline acetyltransferase and acetylcholinesterase activity in post-mortem examinations of patients with Alzheimer's disease (Davies and Maloney 1976). While pharmacological and lesion studies have been inconsistent in confirming such a role (Deiana, Platt, and Riedel 2011), numerous experiments indicate a cholinergic influence over hippocampal-dependent learning (Hasselmo 2006). For example, muscarinic receptor agonists and antagonists produce dose- and delay- dependent increases and decreases respectively in the performance of macaques in a delayed nonmatching-to-sample task (Aigner and Mishkin 1986). Similarly, systemic administration of the muscarinic receptor antagonist scopolamine produces delay-dependent impairments in a delayed matching-to-sample task in humans (Robbins et al. 1997).

A range of data suggests a specific involvement of acetylcholine in the encoding of novel information. For example, acetylcholine concentration in the hippocampus increases in novel environments (Aloisi et al. 1997), and muscarinic receptor antagonism specifically impairs the encoding of novel but not familiar odours in rats in a delayed nonmatching-to-sample task (McGaughy et al. 2005). As such, a particular role for acetylcholine in priming networks to encode novel information, rather than retrieve existing memories, has been proposed (Hasselmo 2006). Acetylcholine could regulate the tendency of the hippocampal formation to either encode or retrieve memories through: modulating the relative strength of Schaffer collateral and cortical inputs to CA1 (Hasselmo, Schnell, and Barkai 1995; Giocomo and Hasselmo 2005), modulating the dendritic and somatic inhibition seen in theta rhythms

(Hasselmo, Bodelón, and Wyble 2002), enhancing persistent spiking (Klink and Alonso 1997) and directly enhancing long term potentiation (Huerta and Lisman 1995).

Given evidence that acetylcholine promotes the encoding of novel information, an increase in cholinergic tone has also been proposed to underlie the increase in grid scale seen in rats exposed to a novel environment for the first time (Barry, Ginzberg, et al. 2012; Barry, Heys, and Hasselmo 2012). In continuous attractor network models, the scale of the grid cell firing pattern depends on the relative balance of excitatory and inhibitory connectivity between grid cells with similar spatial tuning, and the degree of excitation in the velocity signal from the 'shifter' cells (Fuhs and Touretzky 2006; Burak and Fiete 2009). Acetylcholine could therefore increase grid scale through its presynaptic-inhibition based suppression of glutamatergic inputs onto medial entorhinal LII neurones (Hamam et al. 2007). Alternatively, in oscillatory interference models, grid scale depends on the relative slope of the constituent oscillators' theta frequency vs running speed relationships (Burgess, Barry, and O'Keefe 2007). That a reduction in theta frequency is observed in novel environments (Jeewajee et al. 2008; Barry, Ginzberg, et al. 2012; Wells et al. 2013), and that systemic administration of scopolamine flattens the theta frequency vs running speed slope (Newman et al. 2013), has also led to suggestions that cholinergic modulation of theta oscillations could underlie grid scale increases in novelty (Barry, Ginzberg, et al. 2012).

1.3.2.3 Summary

Given its dense reciprocal interconnectivity with the hippocampus, evidence for involvement of the medial septum in both mnemonic and spatial processing is perhaps unsurprising. Even in non-spatial contexts, the medial septum, and particularly its cholinergic population, appears to play a significant part in supporting learning, possibly through promoting the encoding of novel information. However, the effects of inactivation of the medial septum also suggest a purely spatial role, possibly through the encoding of running-speed information. The relative contribution of the medial septum to mnemonic and spatial processing, and its precise role in each, remains poorly understood however. Though highly interconnected with each other, the distinct electrophysiological profiles and connectivity of the different genetic subgroups of septal neurones suggests some functional segregation between them. Further investigation is also required to better understand the involvement of each in the mnemonic and spatial functions of the hippocampal formation.

2 Grid cells form a global representation of connected environments

Sections of this chapter have been published previously, in Carpenter et al. 2015, and Carpenter et al. 2016.

2.1 Introduction

As discussed, grid cells appear well suited to encoding an animal's location within an environment, as well as the distances and directions separating the animal from specific goal locations. That is, their periodic, modular activity suggests that a population of grid cells provides a universal spatial metric. Models asserting the theoretical capacity of grid cells to act as a spatial metric have in general only considered what can be decoded from a population of cells with perfectly regular triangular firing patterns, as were observed in early recordings from grid cells (Hafting et al. 2005). However, more recent investigations have demonstrated that sensory cues, and particularly boundaries, can distort grid cell firing. For example, stretching the length of a rectangular enclosure along one axis results in a concomitant stretching of grid firing patterns in the same axis (Barry et al. 2007). Further, grid firing appears to be locally anchored to the boundaries of the enclosure, such that the orientation of the pattern can vary across the environment (T. Stensola et al. 2015). Indeed, specific enclosure geometries can disrupt the normally uniform triangular symmetry of grid firing: grid patterns can locally rotate and rescale in the narrow end of a trapezoidal environment (Krupic et al. 2015). Finally, when a square environment is subdivided by internal walls to form a sequence of corridors connected by hairpin turns, the grid pattern fragments into multiple sub-patterns, which tend to repeat across corridors in which the rat runs in the same direction (Derdikman et al. 2009).

Results demonstrating irregularities in grid cell firing patterns across an environment are potentially problematic for conceptions of how grid cells could encode self-location and support navigation. For example, a decoding network is unable to distinguish between locations where the conjunction of phases across modules (the population phase) repeats within an environment. In a hairpin maze, in which grid patterns have been observed to fragment and repeat across corridors (Derdikman et al. 2009), if all modules repeated, an animal would be prone to errors of ambiguity in using grid cells for self-location and navigation. Further, local distortions in grid firing patterns indicate that, across the environment, the population phase evolves by an inconsistent amount per unit distance travelled by the animal. As such, estimates of the distances separating pairs of locations

based on changes in the population phase will inevitably also be inconsistent across the environment. Though the impact of any distortion or discontinuity in grid firing patterns depends on how the irregularity manifests across modules, in general the population phase can only be fully accurate as a universal metric for space if the grid patterns are continuous and regular across the entire environment.

To establish whether grid firing is determined by local environmental cues, or provides a coherent global representation, we recorded from grid cells of the medial entorhinal cortex as rats foraged in an environment containing two perceptually identical compartments connected by a corridor. During initial exposures to the multicompartment environment, grid firing was dominated by local sensory cues, replicating between the two compartments. However, with prolonged experience, grid cell firing patterns formed a single, continuous representation that spanned both compartments. We therefore provide the first evidence that, even in a complex environment, grid firing can form the coherent global pattern necessary for them to act as a metric capable of supporting large-scale spatial navigation.

2.2 Methods

2.2.1 Animals

8 experimentally naïve, male Lister Hooded rats (275-400g at surgery) were housed communally under a 12:12 inversed light-dark cycle for at least one week prior to surgery. Following surgery, animals were held individually in Perspex cages and restricted to 90% of their free-feeding body weight. All work was carried out within the terms of the Animals (Scientific Procedures) Act 1986, and according to Home Office and institutional guidelines.

2.2.2 Microdrives and surgery

During surgery, rats received either one or two custom-made microdrives. Each microdrive consisted of 32 HML-coated 17- μ m platinum-iridium (90-10%) electrodes twisted into 8 separate tetrodes. The electrodes were mounted on a single drive-mechanism (tetrodes not independently adjustable) which allowed for their advancement through the turning of a trapped screw. Before surgery, electrode tips were electroplated in a Platinum solution to <150k Ω impedance.

Anaesthesia was maintained through an isoflurane-oxygen mix (1.5-4L/min). Viscotears Liquid Gel (carbomer) was used to cover the rats' eyes and so prevent corneal damage due to their drying out. Rats rested on a water-heated pad which maintained a consistent body

temperature throughout surgery. Carprive (carprofen – 5mg/kg) was injected subcutaneously at the start of surgery as an analgesic. Following surgery, Metacam (meloxicam - 1mg/kg) suspended in jelly, and Baytril (enrofloxacin – 4ml/100ml) dissolved in drinking water, were provided as analgesic and bactericidal agents respectively.

Following steady anaesthesia, as indicated by loss of toe-pinch and eyeblink responses, and slowed breathing, rats were head-fixed in a custom-made stereotaxic frame. An incision was made along the midline to expose the skull, which was then levelled in the antero-posterior and medio-lateral planes. Seven 1.2mm diameter screw holes were drilled, one for the ground-screw and six for screws aiding mechanical stability. In four rats, a trephine was used to drill a single craniotomy above dorso-lateral medial entorhinal cortex in one hemisphere. In two rats, craniotomies were drilled bilaterally above dorso-lateral medial entorhinal cortex in both hemispheres. In the final two rats, craniotomies were drilled above the entorhinal cortex in one hemisphere and the hippocampus in the other. Screws were then inserted into the previously drilled holes. For medial entorhinal implants, electrodes were angled anteriorly 8-10° and implanted 4.5mm lateral to bregma, 0.3mm anterior to the transverse sinus, and at a depth of 1.6mm from the surface of the brain. Following insertion, a protective metal sheath was lowered around the electrodes, with Vaseline applied to cover any tissue still exposed by the craniotomies. Dental cement was then applied to fix the microdrives to the skull and screws.

2.2.3 Electrophysiological recording and behavioural training

Electrophysiological and positional data were acquired using the DACQ USB (Axona Ltd., St Albans, UK) recording system. Rats were connected to the recording system via RC-coupled, unity-gain operational amplifiers and cables suspended on elastic thread. Each channel was amplified 9,000 to 20,000 times, bandpass filtered (360Hz – 7kHz) and recorded differentially against a reference channel on a separate tetrode. Spikes exceeding a trigger threshold (50-80μV) were sampled at 48kHz and time stamped with a 96kHz clock signal. One to two channels were also low-pass filtered (<500Hz), and recorded continually without a reference channel at 250Hz as LFPs. Two groups of infrared light-emitting diodes (LEDs) were mounted on a custom-made 'boom', which was attached to the implant during recording. The boom was positioned with four LEDs close to the animal's head, and two at the end of the boom closer to the tail. The animal's position and head direction was detected by a camera mounted to the ceiling above the enclosure, which monitored and recorded the position of the LEDs. The animal's position was inferred from the location of the larger group of LEDs, and the angle between the larger and smaller group of LEDs was used to infer the

animal's head direction. In this way, spike, LFP, and positional data were recorded as rats moved around enclosures in search of sweetened rice scattered by the experimenter.

2.2.4 Experimental environment and protocol

Each recording session began with a 20-minute baseline trial in a square 'screening' environment. Following identification of grid cells rats were placed in an enclosed white-plastic travel box while the experimental environment was prepared. Comprising two 90x90x50cm compartments connected by a 180x40x50cm corridor, the multicompartment environment was designed such that the two adjacent compartments would be as perceptually identical as possible (Figure 2.1). Painted matte black, the environment was placed on a transparent Perspex floor on top of a uniform black plastic groundsheet, and surrounded by four black curtains organised in a square. The experimental room beyond the curtains was kept dark during the experiment, with polarising auditory cues minimised. Each compartment contained a light at the centre of the south wall, 40cm from the floor. Each light was a battery-powered bike light mounted to the outside of the compartment and enclosed in aluminium foil. The light entered each compartment solely through a 3cm diameter window covered with translucent frosted Perspex. No other light source was present. Into the north face of each compartment was cut a trapezoidal doorway, measuring 10cm across at its base and 50cm across at its top. Each compartment was reinforced along its base with a 1cm strip of stainless steel, allowing them to be easily moved without changing shape.

At the beginning of each experimental session the floor of the environment was cleaned. The first experimental trial was then started, with the rat being placed in the corridor between the two compartments, facing the north wall. Throughout the trial the experimenter moved pseudo randomly around the circumference of the environment, distributing rice such that the rat explored the whole arena. The tether of the recording system was counterbalanced and suspended above the animal using 'runners', such that the centre of mass of the tether moved freely as the animal did, ensuring no directional cue was provided. After 40 minutes, or once coverage of the environment was deemed sufficient, the rat was returned to its travel box. The compartments were then removed and the Perspex floor was rotated 180° before being cleaned. Once dry, the compartments were replaced in the opposite positions to the first trial, such that the compartment in the west of the environment during the first trial was now in the east and vice versa. The animal was then again placed in the corridor facing the north wall and the second experimental trial began. Rats ran at most one session per day for a maximum of 20 sessions. Recordings in the multicompartment environment continued on subsequent days as long as at least one grid cell was still identifiable during screening. If no grid cell was present, the electrodes were again moved until further grid cells were found.

In one animal, following 16 sessions in the multicompartment environment, a ‘remapping’ probe session was recorded. Here, the first trial was recorded in the multicompartment environment as in all other sessions. In the second trial, the floor texture, wall colour, and brightness of the left-hand compartment was changed. In the third trial the environment was returned to its normal configuration. The middle trial of this probe session aimed to provoke remapping of recorded grid and place cells, so as to infer whether their firing fields were anchored locally or globally to sensory cues.

2.2.5 Spike sorting and binning

Spikes were assigned to putative clusters offline through fitting a mixture of Gaussians according to an expectation-maximisation algorithm using KlustaKwik (Kadir, Goodman, and Harris 2014). Putative clusters were then further analysed using the data analysis suite Tint (Axona Ltd., St Albans, UK). In particular, Tint was used to coordinate clusters across trials and to correct for over-clustering according to the amplitude, waveform, and temporal autocorrelation of spikes.

For two-dimensional firing rate maps, the animals’ recorded positions and spikes were assigned to 2x2cm bins covering the environment. Unsmoothed firing rate maps were calculated by dividing the number of spikes assigned to each bin by the cumulative dwell time in each bin. Smoothed ratemaps were constructed using a 5x5 bin boxcar filter, with the firing rate in bin i equal to the number of spikes in the kernel centred on i divided by the occupancy of the kernel.

For directional ratemaps, recorded heading directions and spikes were assigned to 4° directional bins between 0 and 360°. Smoothed circular ratemaps were generated using a 5 bin boxcar filter. The Kullback-Leibler Divergence (KL Divergence) was used to measure the difference between the resulting circular ratemap and a uniform circular distribution with the same mean firing rate. Cells were classified as head direction cells where the KL Divergence in the screening environment exceeded 0.15.

2.2.6 Spatial autocorrelograms and grid cell inclusion criteria

For inclusion in subsequent analysis, putative grid cells were assessed using a gridness measure, following Sargolini et al. 2006. For both the screening environment, and for each compartment of the multicompartment environment, a spatial autocorrelogram was constructed from the smoothed ratemap, according to:

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

Where $r(\tau_x, \tau_y)$ is the autocorrelation between bins with spatial offset τ_x and τ_y . $\lambda(x, y)$ is the firing rate in bin (x, y) , while n is the total number of bins. The six local maxima closest to but excluding the origin of the autocorrelogram were used to identify the orientation, scale, and gridness of the putative grid cell. Orientation was measured as the angle from a consistent arbitrary horizontal reference line to the first peak of the autocorrelogram in an anticlockwise direction. The median distance from the origin to the six peaks was used to estimate the scale of the putative cell. Finally, the gridness score was calculated by rotating the autocorrelogram in 30° steps for 150°, and taking the Pearson product-moment correlation coefficient between each rotated autocorrelogram and the un-rotated autocorrelogram, including only the region between the edge of the central peak and the envelope of the six closest peaks. The maximal correlation obtained at 30, 90, or 150° was subtracted from the minimum correlation found at 60 or 120° to produce the gridness score. When analysing the gridness of firing patterns in the multicompartment environment, gridness scores were calculated separately for each compartment and subsequently mean averaged.

For each putative grid cell, the gridness score obtained from the screening ratemap and the multicompartment ratemap were compared to the distribution of gridness scores observed following 1000 shuffles of the positional and spiking data of the same cell during the same trial. That is, for each shuffle, the relative timing of the position data relative to the spike data was circularly shifted by a random offset of at least 20 seconds, with the autocorrelogram and gridness then calculated using the ratemap resulting from each shift. To be considered a grid cell and included in subsequent analysis, a cell was required to have a gridness score in the 99th percentile of the distribution of shuffled gridness scores in either the screening or multicompartment environment.

To ensure that no single grid cell recorded across tetrodes was considered to be two separate cells, where any putative cells with similar firing fields were observed within or between tetrodes, a spatial cross-correlation was first performed. If the two putative grid cells had consistent spatial firing the temporal cross-correlation between the cells was subsequently analysed. Where pairs of putative grid cells had both high spatial and temporal cross-correlations they, were considered to be a single cell.

85 grid cells passed these inclusion criteria, with a mean average gridness of 0.89 in the screening environment.

2.2.7 Analyses

2.2.7.1 General

The 85 grid cells were recorded across 99 sessions. There were 59 instances of recording a cell in a session in which that cell had been recorded in at least one previous session. In total, there were 144 grid cell/session conjunctions therefore (85 unique cells + 59 instances of re-recording). In 115/144 of the cell/session conjunctions the grid cell was held across both multicompartment trials. Here, any calculated metric was mean-averaged across the two trials, yielding a single value for each cell in each session.

Unless specifically analysing the corridor, all other analyses discarded the corridor and analysed only ratemaps formed by the left and right compartments separated by the gap created by the central wall. As is discussed fully in Section 2.3.5, while activity in the corridor was grid-like, with peaks and troughs in firing (Figure 2.5A), and stable (Figure 2.5B), corridor firing patterns were also significantly less regular and hexagonal than in the compartments (Figure 2.5C). This irregularity is consistent with past recordings of grid cells in linear environments (Hafting et al. 2008; Domnisoru, Kinkhabwala, and Tank 2013), and was likely exacerbated by highly stereotyped behaviour: in the corridor animals tended to run in a ballistic fashion between the compartments, resulting in a highly uneven sampling of space and direction (Figure 2.5D-F). As any ideal grid could only be poorly fitted to the corridor, the comparison of alternative models through the fitting of ideal grids was uninformative.

2.2.7.2 Correlations

To measure the similarity of grid cell representations between the left and right compartments, a Pearson product-moment correlation coefficient was calculated comparing the firing rates in equivalent bins of the smoothed ratemaps of the two compartments. Bins were discarded if they were unvisited in either compartment, or if they had a firing rate of 0 in both compartments.

Further correlations were used to assess whether grid cell representations tracked the movement of the compartments between trials. For each cell recorded across both trials in a single day, Pearson product-moment correlation coefficients were calculated between firing rates in equivalent bins in the same compartment in successive trials ('compartment-wise' correlation), or between equivalent bins in the same position in global space in successive trials ('space-wise' correlation).

2.2.7.3 Fitting of ideal grids

Ideal grid patterns were fitted to ratemaps according to three models, to analyse whether the recorded grid representations could be best described as being determined by local or global features of the environment. The idealised grid patterns were fit by identifying that which achieved the highest spatial correlation with the recorded ratemap. The ‘independent’ model was fit first. Here, 900 ideal grids were generated, each the product of three cosine gratings with wave vectors k_1 , k_2 , and k_3 and phase offsets c_1 , c_2 , and c_3 . The wave vector $\vec{k} = \left(\frac{2\pi}{\lambda} \cos(\varphi), \frac{2\pi}{\lambda} \sin(\varphi)\right)$ where λ is the grating wavelength ($\lambda = G\sqrt{3}/2$, where G is the grid scale), and φ is the grating orientation. The three wave vectors form a regular triangular grid where orientations differ by $\frac{2\pi}{3}$ (i.e., 120°) and $c_1 + c_2 = c_3$. As such, firing rate as a function of location is given by:

$$f(\vec{x}) = A(1 + \cos(\vec{k}_1 \cdot \vec{x} + c_1))(1 + \cos(\vec{k}_2 \cdot \vec{x} + c_2))(1 + \cos(\vec{k}_3 \cdot \vec{x} + c_3)) \quad (2)$$

where A determines the peak firing rate, and does not affect spatial correlations with $f(\vec{x})$. Each of the 900 ideal grids was a unique combination of one of 30 increments of scale and 30 increments of orientation. Orientation varied in increments of 2° between 0 and 58° , while scale varied in equal increments between the scale of the cell as estimated from the screening autocorrelogram $\pm 20\%$. Spatial cross-correlograms between the ideal grid and the recorded ratemap were used to find the Pearson product-moment correlation at all spatial offsets, ignoring unvisited bins. From the spatial cross-correlogram the phase offsets which yielded the highest correlation for that orientation and scale combination were then identified. The independent fit value was then calculated as the highest correlation found between the independent model and the ratemap across all combinations of scale, orientation, and phase. The independent model was used to assess the scale and orientation of the recorded grid, with subsequently fitted ‘local’ and ‘global’ models constrained to the orientation and scale achieving the best fit in the independent model. Further, as the independent fit indicates the highest correlation that can be achieved between any ideal grid and the ratemap, grid cells with an independent fit < 0.45 were excluded from subsequent analysis to account for cells which could not be well fitted by any model. A criterion of 0.45 was chosen as it was near the median independent fit of 0.4702 . As such, just under half of grid cell/session conjunctions were excluded ($65/144$, 45.14%). Varying the required inclusion criterion between 0.4 and 0.5 had little effect on the overall conclusions drawn from the results.

To subsequently fit the local and global models, two additional ideal grids were generated, each with the same orientation and scale as was found to achieve the best independent fit. In the local model, the ideal grid had the same phase in the two compartments, whereas the global grid had a continuous phase across both compartments. The Pearson product-moment correlations between the firing rate map and the local and global ideal grids was then calculated at each phase offset as above. The maximum correlations achieved between the local and global grid and the recorded firing pattern were then identified. These maximum correlations were divided by the maximum correlation found under the independent fit, giving normalised local and global fit values, with which comparisons could be made across cells of differing regularity.

To assess the significance of the local and global fits, each recorded ratemap was also fitted with 1000 randomly phase offset grids. As in the local and global models, each of the randomly offset grids was constrained to the same orientation and scale as was found to achieve the best fit under the independent model. However, each had a random phase offset ($0-2\pi$ along the first two grid axes) between the two compartments. For each of the 1000 grids, a spatial cross-correlogram with the recorded ratemap was calculated to find the spatial offset which achieved the maximum correlation between that ideal grid and the ratemap. Again, as in the local and global model, the maximum correlation at any offset was divided by the independent fit to yield the fit value for each randomly offset grid. For each cell/session conjunction, the proportion of the 1000 randomly phase offset grids which achieved a fit greater than that of the local and global models was then calculated. That is, we identified where the local and global fits fell in the distribution of fits achieved by the 1000 randomly offset grids. These values were then collapsed (mean averaged) within animals to create a single value for each animal in every session in which grid cells were recorded. Under the null hypothesis of no consistent phase relationship between the grid firing in each compartment, a Wilcoxon signed rank test was used to examine whether the observed proportions of the 1000 grids with a better fit than the local or global model differed significantly from 0.5. This test was applied in the first and last five sessions.

To identify whether the apparent transition between local and global representations could be explained by biases in the sampling of different grid scales across time, we repeated the above analyses on particular subsets of the data (see Section 2.3.3). First, the analysis was repeated including only those cells with a scale of 45-55cm, as estimated from the spatial autocorrelogram calculated from the screening environment's ratemap. In addition, the analysis was repeated on only those grid cells from a single animal in which a number of cells with independent fits > 0.45 , and from a single module, were recorded over a number of sessions. Grid cells were assigned to putative modules based on the scale of the cell in

the screening environment. Cells were considered to be of the same grid module if the ratio of the larger to the smaller scale cell was less than 1.4 (H. Stensola et al. 2012).

Finally, we fitted recorded grids in the same way as described above, except the thirds of each compartment closest to and furthest from the corridor were fitted separately. Again, for each recording, we discarded the corridor to produce a ratemap comprising both compartments separated by the gap due to the dividing walls. Here, we then divided the two compartments into three sections, according to distance from the corridor, to produce three ratemaps of ~180x30cm. For each cell/trial conjunction, idealised grids were then fitted to the thirds closest to and furthest from the corridor, under the same principles as above. Again, fits were collapsed within animals such that each data point corresponded to the average of all cells recorded from one animal in one session. These third-compartment fits were then used to test the hypothesis that the thirds of each compartment closest to the corridor are more globally coherent/less locally coherent than the thirds furthest away. That is, as the thirds nearer the corridor are closer together, there is a reduced distance over which the path integrator may accumulate error, potentially resulting in a more globally coherent firing pattern. In the first five sessions, when the grids displayed a local representation, a paired, one-tailed t-test was used to assess whether the thirds furthest from the corridor were significantly more local than the thirds closest to the corridor. In contrast, in the last five sessions, when the grids displayed a global representation, a paired, one-tailed t-test asked whether the grid patterns were more globally coherent in the thirds closest to rather than furthest from the corridor.

2.2.7.4 Analysis of head direction cells

To analyse whether the activity of head direction cells disambiguated the two compartments, we calculated the angle of directional tuning and peak firing rate of head direction cells separately for the two compartments. Directional firing rate maps were calculated as described above, using positional and spiking data from the left or right compartments separately. We then produced an angular cross-correlogram between the two directional firing rate maps for each cell, by rotating one ratemap relative to the other and calculating the correlation between the two at each offset. To calculate the degree of angular offset between the directional tuning in each compartment, we calculated the distance to the closest positive peak of the cross-correlogram from the origin, such that an offset of 0 indicates the highest correlation in the cross-correlogram occurred with no rotation of either directional ratemap. Angular offsets were then plotted as a function of experience of the multicompartment environment. The correlation between angular offset and recording session was used to assess whether the representation similarity of head direction cells in

the two compartments changed with increasing experience. We next calculated the firing rate peak separately for the left and right compartments. A paired Wilcoxon signed-rank test was used to analyse whether head direction cell firing rates disambiguated the two compartments.

2.2.7.5 Analysis of place cells

In one animal (that in which the remapping experiment was performed), a number of CA1 place cells were recorded concurrently with contralateral grid cells. Place cells were identified by eye based on ratemaps recorded in the screening environment. Only those putative place cells with a firing field in one of the main compartments, and with a peak rate of at least 1Hz were analysed further. Firing fields were defined as areas of at least 20 contiguous bins whose firing rates were at least two times greater than the mean firing rate across the compartment.

For each valid place cell, a Pearson product-moment correlation coefficient was calculated between the firing rate maps of each compartment, in an identical fashion to the method described above for grid cells. A correlation value was calculated for each trial, and averaged together for each session. Correlation values for each session were plotted as a function of experience of the environment. The standard deviation of the correlation values of each session were also calculated, and plotted against experience in the same way.

In addition, the peak firing rates of each cell in each compartment were identified. The absolute difference in the compartments' peak firing rates was then calculated, and the standard deviation of the peak rate differences plotted as a function of experience of the environment.

2.2.7.6 Phase offset analysis

The conclusions drawn from the fitting of ideal local and global grids were verified through a separate analysis of observed grid phases. In particular, we identified the difference in grid phase between the observed firing patterns in the right-hand compartment, and that predicted from observed firing patterns in the left-hand compartment extended under a local or global representation. First, an ideal grid was fitted to the firing rate map recorded in the left-hand compartment of each grid cell, as per the fitting principles described above. The ideal fitted grid was used to analyse the scale, orientation, and phase of the recorded grid in the left compartment. The expected location of grid fields in the right-hand compartment were then predicted, based on the identified grid parameters observed in the left-hand compartment. That is, the local model predicts that firing fields replicate in the two

compartments, so the phase in the right-hand compartment should equal the left. In the global model a single grid spans both compartments, thus the phase in the left compartment was ‘projected’ out, to identify the phase in the right-hand compartment, such that placed adjacent to one another, a single grid spans both compartments. An ideal grid was then fitted to the firing pattern recorded in the right-hand compartment. Thus, the right-hand compartment’s recorded phase could be compared to the phase predicted under the local and global models. The reported ‘phase error’ between the data and the local and global models was then calculated as the magnitude of the vector connecting the observed and predicted phases in grid phase-space, and collapsed across all cells recorded in the same animal in the same session. Grid cells were again discarded where the previously calculated independent fit was <0.45 . Further, the 10 data points where the predictions made by the local and global model were most similar were also discarded, to remove cells in which the models could not be distinguished. This equated to (18.52% of the 54 collapsed data points with an independent fit >0.45).

2.2.7.7 Corridor grid cell classifier

The long and thin shape of the corridor meant analysis of the hexagonality of grid cell firing using a spatial autocorrelogram was unfeasible due to only sampling a small number of fields. To assess whether firing in the corridor was ‘grid-like’, we therefore used a grid cell classifier developed for analysing grid cells based on one-dimensional linear-track firing (Domnisoru, Kinkhabwala, and Tank 2013). A smoothed firing rate map of the corridor of the multicompartment environment was first calculated. For each 2x2cm bin, a shuffled distribution of firing rates was calculated by shifting the positional information relative to spike times by at least 20 seconds, and recalculating the firing rate in each bin for each of 1000 shuffles. The proportion of shuffles which had a firing rate greater than that of the original firing rate map was then calculated for each spatial bin. These proportions therefore allowed identification of bins in the original ratemap in which firing rates were higher or lower than would be expected by chance. Firing fields were defined as areas of at least 20 contiguous bins in which firing rates were greater than the 85th percentile of the shuffle distribution for those bins. Out of field regions were defined as areas with at least 20 contiguous bins where firing rates were less than the 5th percentile of the shuffled distribution. Cells were then classified as grid cells if: i. the corridor contained at least 3 firing fields as well as at least one out of field region. ii. 30% or more of the bins were assigned to either in-field or out-of-field areas. iii. The mean in-field to out-of-field firing rate ratio was greater than 2. The proportion of grid cells recorded in each day which satisfied these criteria was then plotted as a function of experience of the environment.

2.2.7.8 Remapping analysis

In one animal, as described, a remapping probe session was recorded. This session consisted of three trials, with trial 2 a probe in which the sensory features of the left-hand compartment were changed. Trials 1 and 3 were normal multicompartment trials, as recorded on all other days.

To quantitatively assess changes in firing patterns in trial 2, Pearson product-moment correlation coefficients were calculated, as described previously, separately for the left and right compartments, between trials 1 and 2, and trials 1 and 3.

2.3 Results

2.3.1 Grid cell firing disambiguates the compartments after prolonged experience

We investigated whether grid cell firing patterns are determined by local sensory cues, or if they provide a coherent global representation of space, by recording from 85 medial entorhinal grid cells in 8 rats, as they foraged within an environment containing two perceptually identical compartments connected via a corridor (Figure 2.1A). We hypothesized that if grid cell representations are dominated by sensory cues, their firing should replicate between the two compartments. Conversely, if grid cell activity is determined by the global spatial features of the environment, their firing patterns should distinguish the two compartments due to their different absolute positions in space.

During early sessions, periodic firing patterns typical of grid cells were present in the environment and were replicated between the two compartments (Figure 2.1B). However, with increasing experience, the similarity of the representations between the two compartments decreased (Figure 2.1C); apparent in a negative correlation between the session number and the spatial correlation of firing rates between the two compartments ($r = -0.67$, $r^2 = 0.45$, $p = 1.67 \times 10^{-12}$, Figure 2.1D). The decrease in representation similarity across sessions was accompanied by an increase in the triangular regularity of grid cell firing patterns within the compartments (Figure 2.1F). Quantitatively, in the first five, but not the last five sessions, gridness in the screening environment was greater than in the multicompartment environment (*one sample t-tests*, $t_{41} = 11.46$, $p = 2.33 \times 10^{-14}$ and $t_{25} = 0.852$, $p = 0.402$ respectively, Figure 2.1G); with the difference in gridness between the screening and multicompartment environments greater in the first than the last five sessions (*unpaired t-test*, $t_{66} = 5.28$, $p = 1.44 \times 10^{-6}$).




Figure 2.1 With increasing experience of the multicompartment environment, grid cell firing patterns show reduced representation similarity between the compartments, and increased regularity. **A**, Schematic representation of the multicompartment environment and protocol for each recording session. **B-C**, Example firing rate maps. The left and right ratemaps in each row are the same cell recorded in trial 1 and trial 2 respectively. Hotter colours indicate higher firing rates, unvisited bins are white. The correlation values above each plot are the spatial correlations of firing rates between the two compartments. **B**, Example grid cells recorded during early exposures to the multicompartment environment, where firing fields replicated between compartments. **C**, Example grid cells from late recording sessions, where firing patterns distinguished the compartments. **D**, Spatial correlations of grid cell firing rates between the compartments as a function of the animals' experience of the environment. Each data point represents the average correlation across all cells from one animal in one session, with different animals plotted in different colours. **E**, Spatial correlations between grid cell firing in equivalent absolute locations in successive trials ('space-wise correlation': e.g. compartment A trial 1 vs compartment B trial 2) or between equivalent locations within the same physical compartment in successive trials ('compartment-wise correlation': e.g. compartment A trial 1 vs compartment A trial 2), showing mean + SEM for cells in the last five sessions. **F**, The difference in gridness of firing patterns between the familiar square screening environment and the average of the gridness in each compartment (screening gridness – multicompartment gridness) as a function of experience. Plotted values are mean \pm SEM across all cells recorded in each session. **G**, Difference in gridness (screening gridness – multicompartment gridness) in the first and last five sessions, showing mean and SEM. *** = $p < 0.001$.

To eliminate the possibility that unidentified local sensory cues allowed disambiguation of the two compartments, we verified that firing was stable in global space, and did not track the physical compartments when their positions were switched between trials. Specifically, in the last five sessions, the inter-trial spatial correlation between compartments in the same location was greater than between the same physical compartments in their new positions (paired t -test, $t_{20} = 6.56$, $p = 2.16 \times 10^{-6}$ Figure 2.1E).

2.3.2 Grid cell firing transitions from a local to a global representation with experience

To determine whether grid firing in the two compartments predominantly reflected a local or global reference frame, we fitted ideal grid patterns to the recorded firing rate maps according to three models. Briefly, the grids were first fit by the 'independent' model, in which grid phase was allowed to vary freely between the two compartments, while orientation and scale were required to be consistent. The independent fit represented the best possible fit of an ideal grid pattern to the data, and was used to exclude grids too irregular to be well fit by any model; a necessary step given the reduced gridness seen in the multicompartment environment, particularly during early sessions (Figure 2.1F). The independent fit also determined the scale and orientation used in the 'local' and 'global' models. In the local model, the fitted grid had the same phase in each compartment, such that the firing fields replicated. The global model required the phase to be continuous across the compartments, such that a single grid spanned the two. Local and global fits were normalised by the independent fit to allow comparison across cells which differed in the regularity of their firing patterns.

During early exposures, grid cell firing in the two compartments was best described by the local model. (Figure 2.2A). However, with increasing experience, the local model's fit to the data decreased ($r = -0.591$, $r^2 = 0.349$, $p = 2.5 \times 10^{-6}$, Figure 2.2C). In contrast to the local model, the fit to the global model increased with experience, showing a positive correlation ($r = 0.419$, $r^2 = 0.175$, $p = 0.0016$, Figure 2.2B, D). A two-way ANOVA revealed an interaction between experience of the environment and the goodness of fit of the two models (*Session*Model*, $F_{(19,68)} = 1.89$, $p = 0.0293$, Figure 2.2E). We assessed whether, in the first and last five sessions, the local or global models fitted the data significantly better than would be expected under the null hypothesis of no consistent phase relationship between the grids in each compartment. Specifically, each recorded ratemap was fitted by 1000 ideal grids with random phase offsets between the two compartments. For each cell/session conjunction with an independent fit > 0.45 , we then calculated the proportion of the 1000 randomly phase offset grids which achieved a better fit than the local and global models. If no consistent phase relationship between the grids in each compartment existed, the local and global models would on average fall in the middle of the distribution of the randomly offset grids. In contrast, in the first five sessions, the proportion of the 1000 grids with a better fit than the local model was significantly lower than 0.5 (*Wilcoxon signed-rank test (WSRT)*, $z = -3.72$, $p = 1.96 \times 10^{-4}$, Figure 2.2F). However, in the last five sessions, the local model no longer fit the data better than would be expected by chance (*WSRT*, $p = 0.6377$, Figure 2.2F). Conversely, in the last five sessions, but not the first five, the global model's fit to the data was significantly better than expected from the null distribution (*WSRTs*, $p =$

0.0019 & $z = -0.109$, $p = 0.913$, Figure 2.2F). It is important to note that local and global representations are not mutually exclusive: grid patterns can be both identical in the two compartments and continuous across them both. As such, one would not necessarily expect consistently low local fits during late sessions nor consistently low global fits during early sessions.

An independent ‘phase offset’ analysis confirmed transition from a local to a global firing pattern (Figure 2.3G-H). Here, we calculated the difference in grid phase between that observed in the right-hand compartment, and that predicted from extending firing in the left-hand compartment to the right according to a local or global representation. The phase error was calculated as the magnitude of the vector in grid phase space which connected the observed and predicted phases. The phase error from the local model increased across sessions ($r = 0.627$, $r^2 = 0.394$, $p = 5.18 \times 10^{-7}$). Phase error from the global model instead fell as animals became more familiar with the environment ($r = -0.551$, $r^2 = 0.304$, $p = 1.05 \times 10^{-4}$).

Implicit in both analyses described above is the assumption that the scale and orientation of grid firing patterns are the same in the two compartments. Given that grid scale and orientation can vary across larger environments (T. Stensola et al. 2015; Krupic et al. 2015), we compared the differences in the autocorrelogram-based estimates of these values seen between the compartments within a trial, with the differences in estimates seen between the two trials in each session. In both cases, the differences in estimates of both grid scale and orientation were tightly clustered around 0, indicating both temporal and spatial consistency (Figure 2.3A-D). The differences in estimated values between compartments did not differ from those observed between trials in a session (*WSRTs*: Scale: $z = -0.739$, $p = 0.459$; Orientation: $z = 0.898$, $p = 0.369$). These results therefore indicate that the orientation and scale were no more different between the two compartments than they are between adjacent trials, suggesting that the assumption of equal scale and orientation is unlikely to be problematic.




Figure 2.2 Grid cell firing patterns transition from a local to a global representation with increasing experience. **A-B,** Fits of local and global models to grid cell firing patterns in the two compartments. The local model was an ideal grid constrained to replicate between the two compartments, whereas the global model was a single continuous grid spanning both compartments. Each row is one cell in one trial: the underlying ratemaps in the left and right columns are the same. The white rings overlaid indicate the best fitting local and global models in the left and right columns respectively. Fit values show the spatial correlations between the local or global models and the data, normalised by the independent model's fit. **A,** Examples of grid cells recorded during early sessions, where the local model best fit the data. **B,** Example grids recorded during late sessions, where the global model best fit the data. **C-D,** The fit between grid cell firing patterns and ideal local (**C**) and global (**D**) grids respectively, as a function of experience of the environment. **E,** The difference in the fit (global fit – local fit) between the global and local models across sessions. In **C, D,** and **E,** each data point represents the average fit for all cells with an independent fit > 0.45 recorded from one animal in one session. **F,** The proportion of 1000 ideal grids, with random phase offsets between compartments, with a better fit to the data than the local or global models. Values are mean + SEM across all cells with an independent fit > 0.45 in the first or last 5 sessions. Wilcoxon signed rank tests compare observed values to an expected median of 0.5. ** = $p < 0.01$, *** = $p < 0.001$.




Figure 2.3 Grid properties are consistent in the two compartments; grid phase analysis confirms local to global representation transition. **A-D**, Grid properties are consistent in the two compartments. For each grid cell recorded in both trials of a single session, the differences in grid scale (**A-B**), and orientation (**C-D**), between trial 1 and trial 2 (**A, C**), and between the left and right compartments (**B, D**), were calculated from each trial's spatial autocorrelogram. Only cells with a fit of the independent model of 0.45 or greater were included, to discard cells which were sufficiently irregular that estimates of grid properties were inaccurate. **E-F**, Both small and large scale grid cells transition from local to global representations. The difference in fits between the global and local model (global – local) is plotted as a function of experience of the environment separately for grid cells with scale greater than (**E**) or less than (**F**) the median scale. Significant positive correlations are observed in both smaller and larger scale grid cells (smaller: $r = 0.506$, $r^2 = 0.256$, $p = 0.0071$; larger: $r = 0.686$, $r^2 = 0.471$, $p = 0.0001$). **G-H**, The phase offset analysis confirms transition from a local to a global representation. Phase error equates to the magnitude of the vector in grid-phase space connecting the observed phase in the right-hand compartment to the phase predicted under local or global models based on grid parameters estimated from firing fields in the left-hand compartment. Each data point represents the average phase error for all cells recorded from one animal in one session.

2.3.3 *The observed local to global transition cannot be explained by biases in sampling of grid modules across time*

Grid cells are organized into functionally distinct modules (Barry et al. 2007; H. Stensola et al. 2012). These modules are distributed non-uniformly in the brain, with modules in dorsal areas of the medial entorhinal cortex often having smaller scale firing patterns than those found more ventrally (Hafting et al. 2005; Brun, Solstad, et al. 2008). Here, as is typical of recordings in the entorhinal cortex, electrodes were implanted dorsally and advanced ventrally to locate grid cells. As such, the transition from a local to a global representation could be due to biases in the sampling of modules across time. Specifically, if smaller scale grid modules formed local representations, and larger scale modules global representations, the dorsal to ventral progression of recordings could produce an artefactual local to global representation shift across time. To eliminate this possibility, we repeated the fitting analysis separately for cells with scale either above or below the median grid scale. Both groups exhibited a significant shift from local to global representations (Figure 2.3E, F). However, as we did not record any grid cells with a scale less than the median in the last 5 sessions, we further analysed the single grid scale (45-55cm) for which grid cells with an independent fit > 0.45 were recorded throughout the experiment. Again, the local model's fit decreased with experience ($r = -0.619$, $r^2 = 0.383$, $p = 0.0048$, Figure 2.4A), while the global model's fit increased ($r = 0.80$, $r^2 = 0.64$, $p = 3.90 \times 10^{-5}$, Figure 2.4B). Indeed, within this scale, grid firing patterns changed with experience from significantly more local to significantly more global than expected by chance between the first and last five sessions (Figure 2.4C).

The same transition was also evident when analysis was further restricted to grid cells from a single module within an individual animal. In the animal with the most sequential recordings of grid cells from a single grid module and with an independent fit > 0.45 (nine sessions), a two-way ANOVA revealed an interaction between experience of the environment and the goodness of fit of the two models (*Session*Model*, $F_{(8,25)} = 9.99$, $p = 0.0019$, Figure 2.4D, E), with the global model's fit increasing significantly with experience ($r = 0.719$, $r^2 = 0.517$, $p = 0.0056$, Figure 2.4E). That the transition from local to global representations is apparent within individual grid scales and modules demonstrates that it cannot be explained simply by biases in the sampling of grid scales and modules across time.




Figure 2.4 The transition from local to global representations cannot be explained by biases in the sampling of grid cells. **A-B,** The fit between recorded firing patterns of grid cells of a single scale, and ideal local and global grids respectively, as a function of experience of the environment. Only cells with a scale of 45-55cm in the screening environment are included. Each data point represents the average local and global fits across all 45-55cm cells with an independent fit > 0.45 recorded from one animal in one session. **C,** The proportion of 1000 ideal grids, with random phase offsets between the compartments, with a better fit to the cells in **A** and **B** than the local or global models. Values are mean + SEM across 45-55cm cells with an independent fit > 0.45 in the first or last 5 sessions. Wilcoxon signed rank tests compare observed values to an expected median of 0.5. **D-E,** The fit between recorded firing patterns of grid cells from a single module in a single animal and ideal local and global grids respectively, as a function of experience. Dashed lines extend the least-squares lines to predict local and global fits in unrecorded sessions. **F,** The best fit achieved by the local model in the first five sessions, and the global model in the last five sessions, to the grid patterns in the thirds of the compartments nearest to or furthest from the corridor. Values are mean + SEM of the collapsed average within animals of cells with an independent fit > 0.45. Paired, one-tailed *t*-tests assess whether the difference in observed means differs from an expected mean of 0. * = $p < 0.05$, ** = $p < 0.01$.

2.3.4 Transition to a global representation may depend on path integration

The firing of grid cells likely depends in part on their performing path integration based on inputs encoding the animal's direction and speed of movement (Fuhs and Touretzky 2006; Burgess, Barry, and O'Keefe 2007). We therefore asked whether there was any difference in the grid representations between the thirds of each compartment closest to and furthest from the corridor. We hypothesized that the reduced distance between the sections of the compartments nearest the corridor may result in the accumulation of less path integration error, and so produce more accurate global representations than in the sections furthest away. Confirming this hypothesis, in the first five sessions, grid patterns in the near third of each compartment were significantly less local than those in the furthest third (*paired one-*

tailed t -test, $t_{17} = -1.93$, $p = 0.0352$, Figure 2.4F). In contrast, in the last five sessions, grid patterns in the thirds closest to the corridor were significantly more global than those in the furthest thirds (paired one-tailed t -test, $t_{10} = 1.96$, $p = 0.0392$, Figure 2.4F).

2.3.5 Grid cell firing patterns in the corridor

The local vs global fitting analysis was performed on a composite ratemap encompassing only the main two compartments, with activity in the corridor excluded. While firing patterns in the corridor were stable and ‘grid-like’, consisting of peaks and troughs in firing, they were sufficiently irregular to make uninformative the comparison of the relative goodness of fits of different models. The irregularity of firing observed in the corridor was consistent with past recordings of grid cells in linear environments (Hafting et al. 2008; Domnisoru, Kinkhabwala, and Tank 2013).




Figure 2.5 Firing patterns in the corridor are ‘grid-like’ and stable, but irregular, and co-occur with highly stereotyped behaviour. **A**, The proportion of cells designated as grid cells by a grid cell classifier analysing firing in the corridor, as a function of experience of the environment. **B**, The spatial correlation of firing rates in trial 1 with the firing rates in equivalent positions in space in trial 2, calculated separately for the corridor and compartments. Data are presented as mean + SEM across all cells recorded in both trials of a single session. **C**, The best fit achieved between ideal grids and the recorded firing patterns in the corridor or compartments. Ideal grids fitted were required to have the same scale and orientation as those found to achieve the best independent fit. **D**, Top: cumulative position plot of the path taken by the animal in a typical 40-minute trial. Bottom: heat map of the unsmoothed dwell times for the same trial, hotter colours indicate longer dwell times. **E**, The coefficient of variation (σ / μ) of dwell time between bins, averaged across all recording trials. **F**, The mean proportion of time spent in each head direction bin, across all animals in the corridor (red) and compartments (blue).

The correlations of spatially-binned firing rates across the first and second trial were not significantly different between the corridor and compartments (*paired t-test*, $t_{114} = 1.28$, $p = 0.204$, Figure 2.5B), indicating similar stability across time. However, firing patterns in the corridor could only be poorly fit with an idealised grid pattern: only a very small number of trials had greater fits in the corridor than the compartments (Figure 2.5C), despite grid firing in the compartments being relatively irregular (Figure 2.1F). As in the main compartments, firing appeared to become more grid-like in the corridor with experience. A classification procedure designed for one-dimensional environments was used to assess how grid-like activity in the corridor was (Domnisoru, Kinkhabwala, and Tank 2013). A significant correlation was found between session number and the proportion of cells whose firing in the corridor was significantly grid-like according to the classifier ($r = 0.524$, $r^2 = 0.274$, $p = 0.0177$, Figure 2.5A). The irregular firing patterns observed in the corridor likely stemmed from the highly-stereotyped behaviour displayed there. Example cumulative position and binned dwell-time plots indicated that rats tended to run in a highly ballistic manner between the two compartments. Areas of the corridor not directly between the entrances to the main compartments were only sampled relatively poorly (Figure 2.5D). The coefficient of variation (standard deviation / mean) of binned dwell times in the corridor was significantly greater than in the compartments (*paired t-test*, $t_{85} = 3.73$, $p = 3.41 \times 10^{-4}$, Figure 2.5E), indicating less even positional sampling. Finally, a plot of the average time spent in each running-direction bin indicated that in the corridor rats spent much longer in bins corresponding to movements along its long axis; with the distribution of dwell times across directional bin being significantly less even than that seen in the main compartments (*Kolmogorov-Smirnov test*, $p = 4.35 \times 10^{-4}$, Figure 2.5F).

2.3.6 *Head direction and boundary vector cells show consistent firing between compartments, regardless of experience*

In contrast to grid cells, head direction cells continued to show the same directional tuning in the two compartments, regardless of experience: no significant correlation was seen between session number and the angular offset in directional tuning between the compartments ($r = -0.174$, $r^2 = 0.0304$, $p = 0.348$, Figure 2.6A). Neither was a difference observed between the peak firing rates in each compartment (*WSRT*, $z = 1.66$, $p = 0.0958$, Figure 2.6B). The firing of a single border cell recorded in a late session also replicated between the two compartments (Figure 2.6D-F).




Figure 2.6 Head direction and boundary vector cells firing replicates in the two compartments regardless of experience. **A**, The angular offset between the firing rate peaks of directional ratemaps from the left and right compartments. Each data point represents the average angular offset across the two trials recorded for each head direction cell in each session. A negative angular offset indicates the firing peak in the left compartment is anticlockwise of the firing peak in the right compartment. **B**, Mean + SEM peak directional firing rates across all head direction cells in the left and right compartments respectively. **C**, Example directional firing rate maps for two representative head direction cells from two rats. The directional firing rate map of the left compartment is plotted in blue and the right compartment in red. **D**, The firing rate map from the screening environment of a putative boundary vector cell recorded on the 11th day of exposure to the multicompartiment environment. **E**, The unbinned path of the animal as it explores the multicompartiment environment in black, overlaid in green with the location of the spikes of the boundary vector cell. **F**, Firing rate map of the same boundary vector cell in the same trial.

2.3.7 Place cell firing disambiguates the compartments after a similar timeframe to that seen in grid cells

In one animal, a number of place cells (range: 13-25) were recorded across 14 sessions, concurrent with grid cell recordings from the contralateral hemisphere. Only putative place cells with a firing field in one of the two compartments, and with a peak rate of at least 1Hz were included in further analysis. During early exposures to the environment, place cell firing replicated between the compartments (Figure 2.7D), as has been observed previously

(Skaggs and McNaughton 1998; Spiers et al. 2015). In contrast to previous reports however, experience of the environment caused place cell activity to increasingly disambiguate the two compartments; apparent in a significant negative correlation between session number and the spatial correlation of firing rates between the compartments ($r = -0.488$, $r^2 = 0.238$, $p = 5.34 \times 10^{-18}$, Figure 2.7A). The time course by which place cells disambiguated the compartments was similar to that observed in grid cells (Figure 2.1D). After 20 sessions, the correlation between compartments predicted from the regression line for grid cells and place cells were similar (grid cells: 0.0368, place cells: 0.0162). Place cells had higher initial correlation values however, with a greater reduction in correlation per session (grid cell regression line: $y = 0.533 - 0.0248x$, place cell regression line: $y = 0.792 - 0.0388x$). While the average place cell correlation between the two compartments fell across sessions, the range of correlation values seen in each session increased, with a significant relationship between the session number and the standard deviation of correlations across cells observed ($r = 0.818$, $r^2 = 0.669$, $p = 0.00035$, Figure 2.7B). No such increase in the standard deviation of correlation values was seen in grid cells ($r = 0.149$, $r^2 = 0.0223$, $p = 0.529$). Further, the standard deviation in the difference in absolute peak firing rates between the two compartments also increased across trials ($r = 0.668$, $r^2 = 0.446$, $p = 0.009$, Figure 2.7C). That is, the range of firing behaviours observed across place cells increased as animals became more familiar with the environment. In early sessions, place fields had similar firing rates and locations within the compartments (Figure 2.7D). In contrast, in later sessions, while some place cells continued to have similar firing fields in each compartment, others exhibited differences in firing rates and/or field locations (Figure 2.7E).




Figure 2.7 The range of place cell representation behaviours becomes more diverse with increasing experience of the multicompartment environment. **A**, The spatial correlation of place cell firing rate maps between the left and right hand compartments, as a function of experience of the environment. Each grey cross indicates the correlation for each place cell for each session, averaged across the two trials comprising each session. Blue lines indicate the mean \pm SEM correlation across all place cells recorded in each session. **B-C**, The standard deviation of correlation (**B**) and peak firing rate (**C**) values across all place cells recorded in each session, as a function of experience of the environment. **D**, Example firing rate maps from place cells recorded during the second session in the multicompartment environment. Note that place cell firing fields tend to be in similar positions with similar peak rates. **E**, Example firing rate maps from the 15th session in the multicompartment environment, demonstrating the range of representation similarities between the compartments seen in place cells in later sessions.

2.3.8 *Grid and place cells are anchored locally to the environment*

In the same animal in which place and grid cell recordings were made concurrently, a 'remapping' probe was performed after 16 sessions in the multicompartment environment. Here, the lighting, floor texture, and wall colour in the left-hand compartment were modified for one trial of the probe session (trial 2). Two normal multicompartment trials were recorded either side of the probe (trials 1 and 3). This probe aimed to infer whether grid and place cell firing was locally or globally anchored to the sensory features of the environment. If locally anchored, only firing in the left-hand compartment would be expected to change in the probe trial. In contrast, if firing across the environment was uniformly anchored to specific cues, then the whole ratemap would be expected to change in the probe. By eye, the firing of both grid and place cells appeared to change only in the left-hand compartment in the probe trial, indicating local anchoring (Figure 2.8B). To assess this quantitatively, correlation values were calculated between the ratemaps recorded in each compartment across the trials. Specifically, for the left and right compartments separately, correlations were performed between trials 1 and 2, and trials 1 and 3. Place cell correlation values in the left-hand compartment between trials 1 and 2 were significantly lower than between trials 1 and 3 (*paired t-test*, $t_{27} = -7.42$, $p = 5.52 \times 10^{-8}$, Figure 2.8A). Place cell correlations in the left-hand compartment between trials 1 and 2 were also significantly lower than in the right-hand compartment between the same trials (*paired t-test*, $t_{27} = -8.11$, $p = 1.027 \times 10^{-10}$, Figure 2.8A). In contrast, no differences were seen in the correlation values between trials 1 and 3 in the left- and right-hand compartments (*paired t-test*, $t_{27} = -0.288$, $p = 0.776$, Figure 2.8A). A similar pattern of results was observed in the grid cell correlation values (Figure 2.8C-D), however only two grid cells were recorded, making statistical testing difficult.




Figure 2.8, Grid and place cells are locally anchored to the multicompartiment environment. **A, C,** Spatial correlations of place cell (**A**) and grid cell (**C**) ratemaps separately for the left and right compartments between trials 1 and 2 and between trials 1 and 3 of the remapping day. Grey crosses indicate the correlation value for each place or grid cell. Blue bars indicate the mean \pm SEM across all cells for that correlation type. **B, D,** Firing rate maps (top) and raw spike plots (bottom) of an example place cell (**B**) and grid cell (**D**). The left and right hand plots are from trials 1 and 3, in which the multicompartiment environment was in its normal configuration. The middle plots represent recordings made during trial 2, in which the floor, wall, and lighting colour of the left-hand compartment only were modified.

2.4 Discussion

2.4.1 *Summary of results and significance*

The periodic, modular activity of grid cells theoretically enables a population of them to efficiently encode an animal's location within an environment, as well as the vector separating the animal from specific goal locations (Fiete, Burak, and Brookings 2008; Bush et al. 2015). Computational accounts of how grid cells could support such behaviour assume that their firing patterns are regular and continuous across the entire navigable environment. Grid firing failed to meet this assumption during initial exposures to the multicompartment environment, with grid patterns replicating between the two compartments, probably due to their being dominated by local sensory cues. However, with increasing experience, discontinuities in grid cell firing patterns between the compartments were incrementally reduced, such that a single representation spanning both compartments was eventually formed. This transition suggests grid cells adjusted their firing to produce the globally coherent representation required for them to act as an effective spatial metric. That is, we provide the first evidence that, even in complex environments, grid cell firing can form regular and continuous firing patterns capable of supporting self-location and navigation.

2.4.2 *Mechanisms underlying transition to a global representation*

For grid cells to form a coherent global representation, it is necessary for them to identify the relative position of all locations within the environment. Models of grid cell firing largely describe their activity as arising from inputs describing the animal's direction and speed of movement, and the integration of self-motion as the animal explores the environment (i.e. path integration) is one way in which the relative positions of points in space can be discerned. The more globally coherent firing patterns observed in the sections of the compartments closest together gives credence to this explanation: the reduced distance between these portions of the environment resulting in reduced accumulation of error in the cells' path integration, and thus more coherent firing. However, unidentified distal sensory cues on the side of the environment with the corridor could also explain the more global patterns seen towards that side. While distal cues were minimised as far as possible, rotation of the multicompartment environment within the curtained area could have provided further evidence that disambiguation of the compartments did not depend on extraneous sensory information. To fully account for all sensory cues, the experiment described here could be repeated in a virtual reality environment, in which one could more easily ensure that the only distinguishing feature of the compartments was their locations within the virtual environment. However, current rodent virtual reality systems involve the animal running in

place on low-friction treadmills (Harvey et al. 2009; Domnisoru, Kinkhabwala, and Tank 2013). As the animal's head is stationary in these setups, the otolith organs of the vestibular system remain inactive, unlike during normal movement, in which they signal linear acceleration (Angelaki and Dickman 2000). Vestibular inputs encoding linear and angular acceleration of the head likely form a subset of the self-motion inputs to grid cells, such that their absence could lead to a reduction in path integration accuracy in current virtual environments. The absence of non-visual sensory cues (such as somatosensory inputs to the body and whiskers) in virtual reality setups may further hinder accurate path integration. As such, the results of a virtual-reality version of the multicompartment experiment may need careful interpretation.

The highly-extended time frame over which grid cells transitioned from a local to a global representation made the continual recording from single cells across the whole period difficult. While the precise temporal dynamics remain unclear, that single grid modules (which are thought to form functional units (H. Stensola et al. 2012)) appear to adjust their representations gradually and continuously argues against an abrupt change. Further, the firing fields of one grid cell recorded across 15 consecutive sessions appeared to shift continuously, rather than undergo a sudden transformation (see Supplementary Video 1 of Carpenter et al. 2015). It is similarly unclear how a single firing pattern emerged from the two discontinuous patterns initially observed. For example, the ratemap of one compartment could have extended into the other, or the firing fields in both compartments could have similarly adjusted until continuous with one another. To fully understand the spatiotemporal dynamics of the transition from a local to a global representation, one would ideally follow the same population of cells throughout, perhaps possible through two-photon imaging of grid cells in a virtual-reality version of the experiment (Heys, Rangarajan, and Dombeck 2014).

2.4.3 Grid firing is locally anchored to the environment

That firing patterns were only locally modified following changes to the sensory features of just one compartment suggests that subsets of firing fields are independently anchored to local sensory cues. Sensory inputs are required for grid cell firing to be stable across multiple visits to the same environment (Hafting et al. 2005), and to prevent the accumulation of excessive noise in the path integration system (Hardcastle, Ganguli, and Giocomo 2015). Particularly given the latter, it is not altogether surprising that multiple sensory cues are used to locally anchor distinct portions of the grid pattern: if a single cue was used, noise would be continually accumulated away from that cue, likely resulting in deleteriously high levels of noise in large environments. The number of sensory cues used to

anchor firing may therefore depend on the size of the environment, perhaps explaining why the geometry of the enclosure appears to determine the orientation of grid firing globally in small environments (Krupic et al. 2015) and locally in larger environment (T. Stensola et al. 2015).

2.4.4 Impact of irregular firing on grid cells' capacity to act as a spatial metric

While sensory cues allow for the resetting of accumulated noise in path integration systems, they may also give rise to distortions of the triangular firing pattern normally expressed by grid cells (Krupic et al. 2015). In the current experiment, while grid firing eventually formed a continuous representation that spanned both compartments, grid patterns were highly irregular and discontinuous for long periods, even after many days of experience with the environment. The significance of distortions and discontinuities to the capacity of grid cells to act as a spatial metric is not well understood, and further experiments are required to relate the regularity of grid cell firing patterns to animals' navigational performance. Theoretically however, and as is discussed below, the impact of any irregularity depends on multiple factors, including the system by which information is encoded across modules, the consistency of any irregularities across modules, and the size of the environment relative to the ensemble's encoding capacity.

2.4.4.1 Grid cells robustly encode self-location despite distortions

A network downstream of grid cells, with access to each module's phase, can accurately decode self-location provided that the population phase is sufficiently dissimilar at each location in the environment. Distortions and discontinuities, which respectively equate to shifts in the rate of change and sudden jumps in the population phase, are immaterial unless they cause the population phase to take the same value at two or more points in the environment. Given the large estimated capacity of an ensemble of grid modules (Fiete, Burak, and Brookings 2008), random distortions and discontinuities are likely only to generate different but still unique population phases across the environment. As such, an ensemble of grid cells robustly encodes self-location, despite irregularities in their individual firing patterns. In the present experiment, as soon as grid firing across modules distinguished the two compartments, the population phase could be used to decode the animal's location. As adjacent positions in grid phase space map to divergent positions in real space (Sreenivasan and Fiete 2011), even apparently small differences in firing patterns between the compartments would unambiguously encode distinct locations.

Nevertheless, the present data confirm previous results indicating that grid patterns do replicate in environments comprised of multiple compartments of high perceptual similarity (Derdikman et al. 2009). Such environments are perhaps rare outside the laboratory, and the repetition of firing patterns observed here may therefore be of minimal ethological significance. Indeed, without concurrent recordings across all modules, it is unknown if the firing of each replicated in the same way, such that the population phase repeated. If this were the case, grid cell activity alone would be insufficient to disambiguate the compartments.

2.4.4.2 Distorted grid firing introduces metric errors to vector navigation

Conceptions of grid firing as a basis for navigation use the difference in population phase between two positions to calculate the vector that connects them in real space. This requires that all points separated by a consistent vector in real-space be separated by a consistent vector in phase-space. Thus, for accurate navigation, the population phase must change at a constant rate across the environment, as well as being unique at each location. Therefore, distortions and discontinuities in grid patterns will by definition introduce errors into navigation vectors that span those irregularities.

The nature of these errors depends primarily on whether any disruption of the grid-pattern is consistent across modules. If the distortion or discontinuity shift each module's phase in inverse proportion to its scale, the encoded location will remain in agreement across modules. In such cases, the population phase will shift by an amount proportional to the size of the disruption in real-space. Navigational vectors spanning such a region would be erroneous by an amount proportional to the magnitude of the disruption. In contrast, if the distortion or discontinuity affects modules inconsistently, the population phase can jump to any other value. Resulting navigational errors would be disproportionate to the magnitude of the disruption, and potentially catastrophic given the large capacity of the grid system.

As discussed, spare capacity in grid-networks has been proposed to provide a form of error correction (Sreenivasan and Fiete 2011). Such a scheme relies on the fact that, in enclosures smaller than the capacity of the grid population, a proportion of population phases are redundant, encoding locations outside the navigable enclosure. Disruptions resulting in invalid population phases could then in principle be identified as erroneous, with the phase returned to the most recent or closest plausible value (Sreenivasan and Fiete 2011). However, this requires a separate neural representation of which population phases are valid (Sreenivasan and Fiete 2011). Thus, errors stemming from misshapen grid-patterns may only be corrected if they arise after the set of valid phases has previously been

identified: distortions or discontinuities that appear on first exposure to an environment may be uncorrectable using this system.

Thus far, the impact of distortions on navigation has been considered based on the assumption that the difference in grid scale between adjacent modules is inconsistent, such that together they form a ‘combinatorial’ encoding system (Fiete, Burak, and Brookings 2008). As seen, in an alternative framework, the largest module alone coarsely encodes self-location, though only unambiguously in enclosures smaller than its scale (Mathis, Herz, and Stemmler 2012). A constant integer ratio between the scales of adjacent modules means that smaller-scale modules are ‘nested’ within the largest module, providing increased resolution to the self-location code. As the phase in larger-scale modules resolves ambiguity in smaller-scale modules, the impact of disruptions in grid-patterns again depends on how they are distributed across modules: distortions to a given module potentially render erroneous the contribution of all smaller-scale modules. Disruptions to small-scale modules would therefore result in small navigational errors, whereas disruptions to large-scale modules would potentially result in catastrophic errors.

In considering the implications of misshapen-grids for navigation, it has been assumed that the decoding network has limited capacity to account for such disruptions. While the possibility that all deformations are ‘mapped-out’ downstream of grid cells cannot be rejected, doing so would require accurate identification of the location, nature, and magnitude of all disruptions. If this were possible, it is unclear why the same information would not be used directly to correct grid-firing. Indeed, that grid-patterns are here seen to regularise with experience suggests that disruptions can be resolved at the level of grid cells in certain conditions. Conversely, if the decoding-network could account for all disruptions, it is unclear why grid-patterns would regularise.

In summary, the capacity of grid population codes, together with the requirement only for a unique population phase at each location in the environment, likely makes grid firing a robust code for self-location. In contrast, a requirement for a constant rate of change in population phase means that using distorted or discontinuous grid patterns for navigation is inherently prone to errors. In the multicompartiment environment, grid firing could unambiguously encode the animal’s location as soon as differences in firing between the compartments give rise to unique population phases in each. In contrast, navigation may be prone to significant errors until grid firing had formed a single, continuous firing pattern across the environment.

2.4.5 Place cell representations in the multicompartment environment; their relationship with other spatially tuned cell types

Here, both grid and place cells formed dissimilar representation of the two compartments following prolonged experience of the environment. In a past experiment, in which place cells were recorded in a similar environment, no change in the correlations between compartments was reported across sessions (Skaggs and McNaughton 1998). In the four rats from which data are presented, three had a wide range of correlation values in early sessions, with a similar range seen across all sessions. The fourth rat displayed a similar pattern of correlations to that observed here. Namely, consistently high correlation values were observed in early sessions, with an increasing range of correlations apparent in later sessions. That is, following prolonged experience, while some cells continued to show highly similar firing in the two compartments, others showed distinct firing rates or field locations. In the previous experiment, the greater range and lower average correlations seen in the first session may indicate that distal sensory cues were less well controlled. For example, the environment was not surrounded by curtains, such that rats could potentially distinguish the two compartments according to the angle from each to individual visual cues in the experimental room (Skaggs and McNaughton 1998).

The increase in the standard deviation of the correlation between compartments of place cell firing across sessions was not observed in grid cells, where a reduction in representation similarity was seen across all cells following extensive experience of the environment. That is, when familiar with the environment, whereas all grid cells disambiguated the two compartments, the similarity of firing patterns between the compartments was highly variable between place cells. This may reflect variance amongst place cells in the relative strength of inputs determined by sensory and self-motion cues. Specifically, place cells whose firing distinguished the two compartments may be predominantly driven by grid cell inputs. In contrast, place cells whose firing remained similar in the two compartments, even when the environment was highly familiar, may predominantly receive inputs from boundary vector cells (which also continue to fire consistently). Variance in the strength of afferent connectivity from grid and boundary vector cells, whose firing do and do not distinguish the compartments respectively, thus giving rise to the large range of representation similarities seen in place cells after prolonged experience. Following similar logic, the compartment-consistent firing seen in boundary vector and head direction cells implies that they do not receive location-specific inputs from grid cells.

3 Modulating medial septal cholinergic tone alters the frequency of theta oscillations without affecting grid cell activity

3.1 Introduction

The dense interconnectivity between the medial septum and the hippocampal formation has historically been linked with the generation of theta oscillations (Buzsáki 2002) and the formation of episodic memories (Hasselmo 2006). More recently however, evidence for an involvement of the medial septum in the encoding of animals' running speed has come to light. This latter potential function is of particular relevance to the spatially tuned firing of grid cells, whose tessellating triangular activity patterns are thought to arise from the integration of inputs conveying the animal's speed and heading direction (Fuhs and Touretzky 2006; Burgess, Barry, and O'Keefe 2007). The encoding of one or more speed signals by the medial septum would explain observations that its inactivation disrupts the regular firing patterns seen in grid cells, without affecting representations of the animal's heading direction (Koenig et al. 2011; Brandon et al. 2011); as well as the fact that medial septal inactivation impairs animals' estimates of linear distances travelled (Jacob et al. 2017), and abolishes the distance-specific firing of place cells on a running wheel (Wang et al. 2015).

While these data suggest a role for the medial septum in encoding running speed, its exact contribution to normal grid cell function, and indeed that of each of its constituent neuronal subtypes, remains unclear. We therefore recorded from medial entorhinal grid cells in familiar and novel environments while modulating the excitability of medial septal cholinergic neurones using Designer Receptors Exclusively Activated by Designer Drugs (DREADDs). Increasing septal cholinergic tone led to a reduction in the frequency of theta oscillations, without a change in the depth of their modulation by running speed. However, no change was observed in the firing patterns of grid cells. The absence of clear change in their spatial firing was consistent with a lack of change in putative speed signals: both theta phase precession and grid cell firing rate vs running speed relationships were unaffected. As such, the activity of medial septal cholinergic neurones does not appear to play a central role in signalling running speed nor modulating grid firing patterns in familiar contexts.

Medial septal cholinergic neurones are also implicated in the coordination of network dynamics which favour the encoding of new memories, rather than the retrieval of existing ones (Hasselmo 2006). Given the strong association between cholinergic tone and novelty, the increase in grid scale seen in rats during initial exposures to unfamiliar environments has been proposed to depend on an increase in acetylcholine concentration (Barry, Ginzberg, et

al. 2012; Barry, Heys, and Hasselmo 2012). We aimed to test this hypothesis by modulating the excitability of medial septal cholinergic neurones whilst recording from grid cells in mice exploring novel environments for the first time. In contrast to previous experiments in rats, no novelty-induced increase in grid scale was observed, making direct testing of the theory impossible. However, an increase in the proportion of time that mice spent not-moving and near the perimeter of familiar enclosures during elevation of cholinergic tone was reminiscent of patterns of behaviour seen in rodents on exposure to novel contexts. That is, the behavioural effects of increasing the excitability of medial septal cholinergic neurones suggest that doing so may induce a sense of novelty, consistent with a role for acetylcholine in biasing mnemonic systems towards encoding. However, such behaviours are also associated with anxiety, and while not mutually exclusive, the current data cannot conclusively indicate whether elevated cholinergic tone led to a perception of novelty or anxiogenesis.

3.2 Methods

3.2.1 Animals

15 experimentally naïve, male ChAT-IRES-Cre (B6;129S6-Chat^{tm2(cre)Low}/J) mice, bred from a homozygous pair obtained from The Jackson Laboratory (ME, USA. Stock no: 006410), gave rise to the data presented in this chapter. Mice were housed communally under a 12:12 inversed light-dark cycle with free access to food and water until surgery, at which point mice were 12-22 weeks old. Following surgery, mice were held individually in Perspex cages and food-restricted to a target weight of 85% of their pre-surgery weight. The target weight was then increased by 0.5g per week. All work was conducted within the terms of appropriate Home Office Project and Personal licences.

3.2.2 Microdrives and surgery

The custom-built microdrives used comprised 16 HM-L coated 17µm platinum-iridium (90-10%) electrodes (California Fine Wire, CA, USA) twisted into 4 tetrodes. Tetrodes were wired to an 18-pin Omnetics connector (Omnetics Connector Corp., MN, USA) and cemented to a drive-mechanism (Axona Ltd., St Albans, UK) which allowed for advancement of the tetrodes through turning of a trapped screw.

Anaesthesia was induced and maintained with an isoflurane-oxygen mix (1.5-3 L/min). During surgery, Viscotears Liquid Gel (carbomer) was applied to the mice's eyes to prevent corneal damage. Mice rested on a water-heated pad to maintain body temperature

throughout surgery. Analgesia was provided by Carprive (carprofen – 5mg/kg) injected subcutaneously at the start of surgery, and Metacam (meloxicam – 5mg/kg) suspended in jelly administered orally once a day for 4 days following surgery.

After induction of steady anaesthesia, the animal's head was fixed in a stereotaxic frame. An incision was made along the midline to expose the skull, which was subsequently levelled in the horizontal plane. Six 0.7mm diameter screw holes were drilled in the skull, one for the ground-screw and five for screws which provided mechanical stability. Craniotomies were drilled to expose the entorhinal cortex and transverse sinus in both hemispheres for insertion of electrodes. An additional craniotomy was drilled in line with but laterally offset from the medial septum, for injection of the virus. Viruses were obtained from UNC Vector Core (NC, USA). 10 mice were injected with the excitatory DREADD virus AAV2-hSyn-DIO-hM3D(Gq)-mCherry, two mice were injected with the inhibitory DREADD virus AAV2-hSyn-DIO-hM4D(Gi)-mCherry, and three mice were injected with the control virus AAV2-hSyn-DIO-mCherry. Viruses were injected at a rate of 100nl/min using an UMPIII Microsyringe Pump connected to a Nanofil Syringe with a 33g bevelled needle (World Precision Instruments, FL, USA). Four 300nl injections were made 0.7-0.9mm anterior to bregma, 0.8mm lateral to bregma, with an angle of 12° towards the midline, and at depths of 4.3, 3.9, 3.5, and 3.1mm ventral from the brain surface, starting at the most ventral site. The needle was held in place for five minutes after each injection, then moved either to the next site or removed from the brain. Ground and stabilising screws were then inserted and the electrodes moved into position over the entorhinal cortex. Four tetrodes were implanted in each hemisphere, mounted on separate drive mechanisms. Electrodes were inserted 0.3-0.5mm anterior to the anterior edge of the transverse sinus, 3.2mm lateral to lambda, and with an angle of 6° in the posterior direction. The dura was removed at the insertion site and the electrodes implanted 0.8-1mm below the brain surface. A protective metal sheath was lowered around the electrodes before dental cement was applied to fix the microdrives to the skull and screws. After surgery, mice were transferred to a heated chamber until recovered from anaesthesia.

3.2.3 Experimental protocol and electrophysiological recording

In the second week after surgery, mice were handled and exposed to the baseline environment for three to four sessions of 20 minutes, to familiarise them with the room and experimenter. The baseline or 'familiar screening' environment consisted of either a 90x90cm square or a 1m diameter circular enclosure, each with 50cm high walls. The enclosure was placed on a black metal ground-sheet, partially enclosed on one side by black curtains, with a single white cue card with a desk lamp shining on it, and open on the other side to the rest of the experimental room. Ultimately, recordings were also made in two

additional 'novel' environments. Each of the novel environments were at opposite ends of a different room to the baseline environment. The enclosures used in the novel environments were of a different shape to that in which baseline recordings were made, but the same shape as each other (e.g. baseline square → novel circles). The two novel environments were not surrounded by curtains, and each had a similar brightness and a similar number of distal cues. The order in which mice experienced the two novel environments as well as which environment the DREADD agonist was administered in was counterbalanced between animals.

At the start of the third week (i.e. at least 14 days after surgery), mice were connected via an RC-coupled unity-gain operational amplifier and a lightweight cable suspended on a pulley system to an Axona 'DACQ USB' recording system (Axona Ltd., St Albans, UK). Each channel was amplified 9,000 to 20,000 times, bandpass filtered (360Hz – 7kHz), and recorded differentially against a reference channel on a separate tetrode. Spikes exceeding a trigger threshold (50-80uV) were sampled at 48kHz and time stamped with a 96kHz clock signal. One to four channels were also low-pass filtered (0-500Hz), either with or without a notch filter around 50Hz, and recorded continually without a reference channel at 4.8kHz and 250Hz as LFPs. The mice's location was tracked by way of an infrared LED on the operational amplifier and a camera above the enclosure recording at 50Hz. In this way, spike, LFP, and positional data were collected while mice moved around the enclosure in search of droplets of sweetened formula baby milk scattered by the experimenter.

Each baseline or 'screening' trial lasted 15-30 minutes, terminated once the mouse was deemed to have covered the environment sufficiently to assess the spatial tuning of any recorded cells. Following recording, data was cut using an expectation-maximisation algorithm via the KlustaKwik software (Kadir, Goodman, and Harris 2014). Putative clusters were then investigated by eye using the Waveform software (d1manson.github.io/waveform), which was used to coordinate clusters across trials and to correct for over-clustering according to the amplitude, waveform, and temporal autocorrelation of putative cells.

In seven of the 10 hM3Dq mice, screening was continued until grid cells were recorded, with electrodes advanced 30-60µm between screening sessions, and with at least 4 hours between sessions. In three of the hM3Dq and the three control-mCherry mice, no grid cells were recorded, and the experiment was commenced once the electrodes had been advanced such that large amplitude theta oscillations were observed in the LFP. To allow for virus expression, experiments were only initiated after a minimum of three weeks following surgery.

See Figure 3.1 for a pictorial representation of the experiment protocols. In the hM3Dq mice, two types of experimental day were performed. On 'novel' days, following a baseline trial, the mouse was briefly (<30 seconds) anaesthetised with isoflurane and 3mg/kg of Clozapine-N-Oxide (CNO) or an equivalent volume of saline was injected intraperitoneally, before the mouse was returned to its home cage, with the order of the drugs across days counterbalanced between animals. One hour after the injection, a second trial in the familiar environment was recorded, termed a 'CNO probe' or 'saline probe' trial. Approximately 20 minutes after the end of the probe trial, the animal was transferred to one of the novel environments and three 'novel' trials were recorded, each 15-30 minutes in length, and separated by approximately 20 minute gaps. Finally, the animal was then returned to the familiar environment and a second familiar trial was recorded (not considered a probe trial nor used in plots/hypothesis testing), such that six trials were recorded on each 'novel' day. In seven of the 10 hM3Dq mice, recordings were made across four novel days, two with injection of CNO and two with injection of saline. The other three mice only had two novel days, one with CNO and one with saline.

In addition to the novel days, four mice also had 'familiar' day protocols to augment the data recorded in the familiar environment. Here, following the baseline trial, the mouse was first injected with saline, and after a one hour break, a saline probe trial was recorded in the familiar environment. Subsequently, the mouse was then injected with CNO, and after a second one hour break, a familiar CNO probe trial was recorded. Familiar days thus included three trials, with the saline probe trial always leading the CNO probe trial due to the long-lasting time course of CNO action (Alexander et al. 2009). Across both types of day, probe trials thus refer to the trial recorded in the familiar environment one hour after injection of either CNO or saline.

The CNO solution injected was made by dissolving 1mg of powdered CNO (Tocris Bioscience, Bristol, UK) in DMSO (Sigma-Aldrich, MO, USA) and then sterile 0.9% saline, to a final volume of 5ml with a concentration of 0.2mg/ml.

3.2.4 Analyses

3.2.4.1 Position and speed data

The recorded LED data was used to infer the position of the animal throughout the trial. Position data was smoothed separately in the x and y dimensions with a 400ms boxcar filter. The speed of the animal was calculated as the distance connecting the mouse's location in

adjacent position samples using Pythagoras' theorem, adjusting for the pixels per metre recorded by the camera, divided by the temporal offset between samples (20ms).

3.2.4.2 LFP

In mice where more than one LFP channel was recorded, an LFP channel from the same microdrive in which grid cells were recorded, and on which large amplitude theta could be observed, was used for all analyses. If no grid cells were recorded, then the LFP channel with maximal theta power was used. The same LFP channel was used for all trials in a single day.

The power in the LFP at different frequency bands was assessed by analysing the entire trial through constructing a power spectrum using Welch's method, with non-overlapping windows of 4096 samples, having down sampled the LFP from 4.8 to 1.2kHz. The power at each frequency was then log transformed. The peak theta frequency was defined as the frequency in the 6-10Hz band which had the highest power. The peak theta power was taken to be the log power of the peak in the 6-10Hz band. Theta signal to noise (S2N) was defined as the average logged power in a 2Hz band around the theta peak divided by the average logged power in the rest of the 3-25Hz frequency range.

Following Torrence & Compo 1998, Morlet wavelets were used to analyse changes in theta frequency as a function of running speed. 500 wavelet scales assessed the power at frequencies between 3 and 100Hz around each sample of the 250Hz LFP trace. The resulting spectrogram was then down sampled in time to match the 50Hz position sampling. The constituent power spectra were then binned into 120 bins of width 0.25cm/s, between 0 and 30cm/s, with the mean spectrum calculated for each speed bin. A maximum speed of 30cm/s was chosen as higher speeds were only poorly sampled. For each speed bin the frequency within a 3-11Hz window with maximum power was identified. The relationship between the frequency of maximum power and the average speed in each bin was then fit with both linear and three-term exponential functions using linear and non-linear least squares respectively. In the latter, parameter A, B, and C values were limited to the ranges 3-11, 0.5-4, and 0-1. The linear and non-linear functions fitted were defined as:

$$y = A + B * x \quad (3)$$

$$y = A - B * e^{(-C * x)} \quad (4)$$

Where x is the animal's running speed, y is the frequency of maximum power in the theta window, and A , B , and C are parameters to be fit. In one familiar trial and two novel trials (out of 221 trials overall), poor and uneven sampling of running speeds, together with the

narrow speed bins, meant that there were intermediate speed bins with zero samples, and these trials were not analysed further.

Each model's quality of fit was assessed using the F-test:

$$F = \frac{(\text{linRSS} - \text{expRSS})/(\text{linDF} - \text{expDF})}{(\text{expRSS}/\text{expDF})} \quad (5)$$

Where *linRSS* and *expRSS* are the residual sum of squared errors for the linear and exponential fits respectively, and *linDF* and *expDF* are the degrees of freedom of each model. Calculation of the F value enabled assessment of the significance of the difference in quality of fit between the two models through calculation of a p-value from the F distribution.

To assess whether the apparently non-linear relationship between theta frequency and running speed was an artefact of the analysis described, we repeated the procedure on synthetic LFP data. To generate the artificial LFP traces, a 50Hz-sampled vector of running speeds from a single baseline trial was linearly interpolated to 250Hz. A frequency value for each 250Hz sample was then calculated, varying as a linear, exponential, or sinusoidal function of the running speed values:

$$y = 4 + (0.15 * x) \quad (6)$$

$$y = 7.5 - 3 * e^{(-0.25 * x)} \quad (7)$$

$$y = 7.5 + 3 * \sin(x) \quad (8)$$

Where *x* is the animal's running speed, and *y* is the frequency of maximum power in the theta window. From the instantaneous frequency values, the phase-advance between each 250Hz LFP sample was calculated, and used to construct a continuous sinusoidal trace. This artificial LFP traces were then fit with the above-described procedure, to confirm that the extracted frequency vs run-speed relationship matched the artificial data, and that the exponential fit did not outperform the linear fit even when the underlying relationship was linear or sinusoidal.

3.2.4.3 *Theta modulation of single units*

The entire trial was divided into contiguous blocks of position samples where the speed remained above a minimum threshold of 2cm/s. For each contiguous block with at least two spikes, a temporal autocorrelogram of length 500ms was generated from the cell's spike train. The individual autocorrelograms were then weighted according to the length of the block, with a weighted average calculated across the trial. A power spectrum of the temporal

autocorrelogram was then calculated using the Fourier transform, and smoothed using a Gaussian kernel (width 187.5ms). The peak theta frequency was then defined as the frequency in the 6-10Hz band which had the highest power. Theta signal to noise (S2N) was defined as the average power in a 2Hz band around the theta peak, divided by the average power in the rest of the 3-25Hz frequency range.

The significance of a cell's theta modulation was assessed using a shuffle procedure. The above procedure was repeated 100 times on artificial data in which the cell's spike times had been randomly reassigned to another time within the trial. A cell was deemed to be significantly theta modulated if its theta S2N exceeded that seen in 99/100 of the shuffles. Because of the small number of grid cells which showed significant theta modulation (Figure 3.5), we instead focused on cells assessed by eye as having theta-modulated temporal autocorrelograms, but which were not spatially modulated (either by the animal's location or heading direction). Only the subset of these cells which had significant theta modulation in either a baseline or probe trial were used for hypothesis testing.

The change in the burst frequency of single units as a function of running speed was also assessed. To do so, the trial was first divided into two second windows, each with a one second overlap with the previous window. In each window with at least two spikes, the mean running speed and the power spectrum derived from the Fourier transform of the spike train's temporal autocorrelogram was calculated. Windows were then binned into 4cm/s speed bins between 0 and 24cm/s (24cm/s was chosen as the upper limit as it was the highest bin with reliable sampling of sufficient numbers of spikes), with the power spectra averaged together in each speed bin. Only bins up to the first bin which contained < 5% of the total number of valid windows were included. The frequency of maximum power in a 3-11Hz theta window was identified for each valid speed bin, and plotted as a function of the bin's average speed. The data were then fitted with linear and non-linear least squares as per the procedure described above for the LFP (see Section 3.2.4.2).

Again, as in the LFP, the difference in quality of the fits was assessed using the F test. Only cells which had significant theta modulation in either a baseline or probe trial, and only trials which had a significantly better ($p < 0.05$) exponential fit were included in plotting and hypothesis testing. That is, in assessing the null hypothesis of no-change in the parameter values which gave the best fits following CNO administration, only cells meeting these criteria were included. These requirements aimed to restrict analysis to only those trials in which the cell's exponential fit was sufficiently good for the best fitting parameter values to be meaningful.

3.2.4.4 Behaviour

To assess changes in the running behaviour of animals, the mean average speed in both the first five minutes and across the whole trial were calculated. A distribution of speeds was also generated, as the proportion of the whole trial spent in 2cm/s speed bins between 0 and 30cm/s. The amount of time the animal spent sitting still (running speed < 1 cm/s) was calculated for the first five minutes of the trial.

To quantitatively analyse the apparent increase in avoidance of the centre of the environment, the periphery of the environment was defined as the area within 10cm of the outermost extent of position samples recorded across the entire trial. 'Perimeter dwell', a measure of avoidance of the centre, was then calculated as the proportion of the first five minutes of the trial spent within this peripheral zone.

3.2.4.5 Grid cell analysis

Two-dimensional firing rate maps were calculated by assigning both the animal's recorded positions and the spikes of each cell to 2x2cm bins. The firing rate in bin i was then calculated by dividing the number of spikes in a 5x5bin boxcar kernel centred on bin i by the cumulative dwell in the kernel.

The smoothed ratemap of each putative grid cell was assessed through construction of a spatial autocorrelogram, following Sargolini et al. 2006, in which the ratemap was shifted relative to itself and the correlation calculated at each offset. Concretely:

$$r(\tau_x, \tau_y) = \frac{n \sum \lambda(x, y) \lambda(x - \tau_x, y - \tau_y) - \sum \lambda(x, y) \sum \lambda(x - \tau_x, y - \tau_y)}{\sqrt{n \sum \lambda(x, y)^2 - (\sum \lambda(x, y))^2} \cdot \sqrt{n \sum \lambda(x - \tau_x, y - \tau_y)^2 - (\sum \lambda(x - \tau_x, y - \tau_y))^2}} \quad (9)$$

Where $r(\tau_x, \tau_y)$ is the autocorrelation between bins with spatial offset τ_x and τ_y . $\lambda(x, y)$ is the firing rate in bin (x, y) , while n is the total number of bins. The rate map's degree of hexagonal regularity, or 'gridness', was calculated by rotating the autocorrelogram in 30° steps for 150° and taking the Pearson product-moment correlation coefficient between each rotated autocorrelogram and the un-rotated autocorrelogram, including only the area of the autocorrelogram extending from the edge of the central peak to the envelope of the six closest peaks. Gridness was defined as the maximum correlation at 30, 90, or 150° subtracted from the minimum at 60 or 120°. A cell was considered to be a grid cell if its gridness exceeded 0.3. The scale of each grid cell was also calculated, as the median distance (in cm) to the six peaks closest to the autocorrelogram's centre.

To assess the intra-trial stability of each grid cell's firing pattern, each trial was divided into two halves, with a ratemap generated separately for each, according to the above described procedure. A cross-correlogram was generated using the same correlation procedure discussed above, but instead using the ratemaps from the two halves of the trial, rather than two copies of a single ratemap. The value at the centre of the cross-correlogram was used as a measure of the stability of the cell's firing between the two halves of the trial.

Inter-trial stability was assessed between the baseline and familiar probe trials, again using the same procedure described above, but instead with the two ratemaps being those from the entirety of the first and second trials respectively. The degree of stability was again measured as the correlation value at the centre of the cross-correlogram between the two ratemaps.

For cells to be included in hypothesis testing of gridness and stability, the cell had to have gridness > 0.3 in either a baseline or probe trial. Only trials with gridness > 0.3 were used to assess grid scale, and therefore the cells included in grid scale analyses had gridness > 0.3 in both the baseline and probe trials.

3.2.4.6 Grid cell speed modulation

Following Kropff et al. 2015, to assess the firing rate vs speed modulation of grid cells, each cell's firing rate was calculated at the same frequency as the position sampling (50Hz), by binning spikes into 20ms bins and then smoothing with a Gaussian kernel of 60ms width and a sigma of 30ms. The regression line relating speed and firing rate was then fit for position samples above 2cm/s and below 30cm/s, with the correlation and slope identified. The correlation between speed and firing rate was termed the cell's 'speed score'.

To assess the significance of a cell's firing rate vs running speed relationship, each trial's speed score was compared to a distribution of 1000 shuffled speed scores calculated after circularly rotating the animal's speed relative to the cell's firing rate by a random temporal offset of at least 20 seconds. A grid cell was deemed to be significantly speed modulated if its speed score fell in the top 1% of the shuffled distribution of speed scores and had a gridness score above 0.3 in the same trial. Only grid cells with significant speed modulation in either a baseline or probe trial were included in hypothesis testing.

3.2.4.7 Grid cell phase precession

In order to analyse the tendency of grid cells to display phase precession, the procedure described in Jeewajee et al. 2014 was followed. Briefly, the 250Hz LFP was smoothed and

filtered with a Blackman filter between 6 and 10Hz. The Hilbert transform of the filtered LFP was then used to calculate the instantaneous phase of the theta oscillation for each LFP sample. The grid cell's firing rate map was produced as described above, but instead using 0.5x0.5cm bins and a Gaussian smoothing kernel of 100x100bins (sigma of 7.5cm). Grid firing fields were defined using a watershed function. Fields near the edges of the environment were only included if their area exceeded 75% of the mean area of fields in the centre of the environment. The time between entry and exit from a valid firing field was termed a 'run'. Individual runs through the field were normalised so as to fall within a unit circle, and rotated so that the average movement direction of each run was left to right. The 'proportion of way along run' was then calculated as the animal's signed distance to the field peak, accounting for running direction by projecting onto the vector joining current position in the field to the edge of the field, assuming no future change in running direction. Proportion of way along run varies from -1 as the animal enters to 1 as it exits the field, with a value of 0 when crossing the field's peak. A circular-linear regression was then performed, relating proportion of way along run to the phase of the first spike in each theta cycle in the run, with slope and correlation values identified.

Only grid cells that were not modulated by heading direction were included in the analysis. Directional ratemaps were generated by assigning heading directions and spikes to 6° bins between 0 and 360°, smoothing circularly with a Gaussian kernel of 5 bins, and dividing the number of spikes by the time spent in each bin. The Kullback-Leibler divergence (KL divergence) comparing the cell's directional ratemap to a uniform distribution of equal mean was then calculated. A cell was considered directionally modulated and not further analysed if its KL divergence exceeded 0.15.

The significance of the phase precession exhibited by a grid cell was determined using a shuffling procedure. A cell was considered to display significant phase precession if its correlation value was more negative than 95% of 1000 shuffles where distance-along-run and spike-phase were randomly repaired. To be included in hypothesis testing, a cell had to have KL divergence < 0.15, gridness > 0.3, and significant phase precession in either a baseline or probe trial.

3.2.4.8 Hypothesis testing

The analyses discussed above generated a number of measures for which we assessed whether or not there was sufficient evidence to reject the null hypothesis that there was no change in the measure following injection of CNO. Throughout, hypothesis testing here used Repeated Measures ANOVA (RM-ANOVA) at the level of animal. Specifically, for each

measure, we averaged together values from all appropriate cells and trials to generate four values for each animal: a CNO baseline value, a saline baseline value, a CNO probe value, and a saline probe value. CNO baseline and saline baseline values were distinguished by whether the baseline trial was followed by a CNO or saline probe trial. The Repeated Measures ANOVA test was then used to look for an interaction effect between the baseline vs probe groups and CNO vs saline groups. That is, to assess whether the change between baseline and probe averages was significantly different between the CNO and saline groups.

The only exception to the use of the RM-ANOVA test was in analysing the inter-trial stability of grid cell firing patterns, for which a paired t-test was used, as looking at the correlation between trials meant there was only a single value for each of the CNO and saline groups.

Analysis was performed at the level of animal as cells are often not independent of one another within animals, making statistical testing problematic (Fyhn et al. 2007; Yoon et al. 2013). For an animal to be included in the RM-ANOVA analysis it was required to have a data point for all four groups used in statistical testing (CNO baseline, saline baseline, CNO probe, and saline probe). Averaging at the level of animal therefore meant that for the LFP and behaviour analyses the statistical tests had an n of 10. Each of the cellular analyses had further criteria for inclusion as discussed above, which lowered the number of data points. As such, some of the analyses only had a small number of data points. When no change was seen in analyses averaging at the level of animal, we therefore repeated analyses averaging at the level of cell. As instability in recordings meant that the same cell was only rarely recorded across both saline and CNO days, cells were included as long as they had a value in both CNO baseline and CNO probe, or saline baseline and saline probe trials. As such, a mixed-design ANOVA was used for hypothesis testing, to allow inclusion of cells which did not have values for all four trial types.

3.2.5 Experimental protocol for control mice

Control mice injected with DIO-mCherry were used to assess whether any change seen in the hM3Dq mice was due to non-specific effects of CNO administration. Control mice were familiarised with the baseline environment and electrophysiological recordings made as described above for the hM3Dq mice.

At least three weeks after surgery, and once familiarised with the environment, three 'familiar' experiment days were carried out on consecutive days (that is, following the same familiar-day protocol used in the hM3Dq mice, see Figure 3.1). Briefly, a baseline trial was followed by intraperitoneal injection of saline. The animal was then returned to its home cage for one hour, with a saline probe trial subsequently run. Injection of CNO and a second one

hour pause then followed, before a CNO probe trial was performed; all in the familiar environment. On each day, an equivalent volume of saline to the volume of CNO solution was administered, with 1mg/kg of CNO on the first day, 3mg/kg of CNO on the second day, and 5mg/kg of CNO injected on the third day. As in the hM3Dq mice, injections were performed under brief isoflurane anaesthesia.

In assessing whether any effect was seen in control mice, the measures in which a significant difference was seen in the hM3Dq mice were examined by eye, averaging at the level of animal both across all drug concentrations, and separately for each drug concentration, to see if any dose-dependent effects could be observed.

3.2.6 Immunohistochemistry and cell counting

At the end of the experiment, mice were anaesthetised with isoflurane before intraperitoneal injection of 0.5ml of Euthatal (sodium pentobarbital). Following loss of toe-pinch and eye-blink reflexes, mice were trans-cardially perfused with 10ml of 0.9% saline followed by 20ml of 4% paraformaldehyde (PFA). Brains were extracted and left in PFA for 4-6 hours, and then transferred to a cryoprotectant solution of 30% sucrose in tris-buffered saline (TBS) until the brains had sunk.

Brains were then cut in the coronal plane into 50µm sections using a freezing microtome. Sections were collected in TBS before being washed three times for 10 minutes in a solution of TBS and 0.3% Triton X-100. Slices were subsequently incubated with DAKO Serum-Free Protein Block for one hour. Slices were then incubated on a shaker table for 24 hours at 4°C with primary antibodies for choline acetyltransferase (ChAT) (EMD Millipore, AB144P, 1:100) and mCherry (Abcam, AB167453, 1:1000). Slices were then again washed three times in TBS with 0.3% Triton X-100, before incubation at room temperature with secondary antibodies at 1:1000 for 18 hours (ChAT: ThermoFisher Scientific, A-10042 Alexa-488 conjugated. mCherry: ThermoFisher Scientific, A-11055 Alexa-568 conjugated). Slices were washed a final three times in TBS before being mounted in Vectashield and a coverslip applied.

Slices were viewed at 2.5x and 10x magnification using an epifluorescence microscope and at 40x magnification using a Leica SP8 confocal microscope. Confocal images of different planes of the slice were layered to produce z-stacks of the entire slice. Images of the medial septum were captured using filters appropriate for the 488 and 568nm dyes conjugated to the secondary antibodies. Selected captured images were exported as TIFF files using the Leica LAS X software.

In addition, cell counting was performed manually on composite images of microscope slides at 10x magnification, captured using a Zeiss Axioscan Slide Scanner. For each animal in which cell counting was performed (four hM3Dq mice and two hM4DI mice), the slide with the eight slices (400µm range) of the medial septum with maximal expression of mCherry was selected. Of these eight slices, cell counting was performed on one slice from each pair of adjacent slices (thus, in each animal, counting covered four slices over a minimum range of 300µm). The number of mCherry+, ChAT+, and mCherry+/ChAT+ cells were counted. This enabled calculation for each slice of the proportion of ChAT+ cells co-labelled for mCherry, and the proportion of mCherry+ cells co-labelled for ChAT. The values were then averaged together across the four slices counted to create average proportions for each animal.

3.3 Results

To investigate the function of medial septal cholinergic neurones, we enhanced their excitability using DREADDs while recording extracellularly from grid cells (Figure 3.2A). The Cre-dependent AAV2-hSyn-DIO-hM3D(Gq)-mCherry virus was injected into the medial septum of ChAT-IRES-Cre mice, so as to express hM3Dq exclusively in cholinergic neurones of the medial septum. Expression of hM3Dq (a modified Gq-coupled human M3 muscarinic receptor) leads subsequent intraperitoneal injection of its normally inert agonist, Clozapine-N-Oxide (CNO), to increase the excitability of hM3Dq expressing neurones (Alexander et al. 2009). Immunohistochemical staining for choline acetyltransferase (ChAT) and mCherry confirmed that expression of hM3Dq was limited to the medial septum (Figure 3.2B), while confocal imaging indicated a tight overlap in expression of ChAT and mCherry (Figure 3.2C). Cell counting revealed that mCherry was expressed in the majority of ChAT+ neurones of the medial septum (mean \pm SEM proportion of ChAT+ neurones co-labelled for mCherry = 0.57 ± 0.039 , Figure 3.1D), while almost all mCherry+ neurones were ChAT+ (mean \pm SEM proportion of mCherry+ neurones co-labelled for ChAT = 0.95 ± 0.0032 , Figure 3.1E), indicating that hM3Dq was expressed almost exclusively in cholinergic neurones. In contrast, in mice injected with the inhibitory DREADD hSyn-DIO-hM4D(Gi)-mCherry, a high proportion of mCherry expressing neurones did not co-stain for ChAT (mean \pm SEM proportion of mCherry+ neurones co-labelled for ChAT = 0.28 ± 0.14 , Figure 3.1B, E). As such, experiments were limited to the excitatory DREADD hM3Dq.




Figure 3.1 Confirmation hM3Dq is expressed almost exclusively in medial septal cholinergic neurones; pictorial representation of the experimental protocol. **A-B**, Example 10x epifluorescence image of coronal brain sections showing ChAT expression (green) and mCherry expression (red) following injection of hM3Dq (**A**) and hM4Di (**B**). Note the minimal overlap in ChAT and mCherry staining in **B**. **C**, Additional example 40x confocal z-stack images demonstrating tight overlap between ChAT and mCherry expression following injection of hM3Dq. **D-E**, Bars indicate mean \pm SEM across all animals, crosses indicate average value for each animal. **D**, Proportion of ChAT⁺ neurones co-labelled for mCherry, as indicated by manual cell counting. **E**, Proportion of mCherry⁺ neurones co-labelled for ChAT. **F**, Experimental protocol for recordings made on novel and familiar days. A full description of the protocol can be seen in Section 3.2.3.

3.3.1 *Increasing septal cholinergic tone reduces the frequency of LFP theta oscillations*

Electrophysiological recordings were made from the medial entorhinal cortex of 10 mice exploring a familiar environment during 'baseline' trials and one hour after injection of either CNO (3mg/kg) or an equivalent volume of saline ('CNO probe' and 'saline probe' trials respectively; see Figure 3.1E for full experimental protocol). Following injection of CNO, a clear reduction in the frequency of theta oscillations was observed (mean \pm SEM, peak theta frequencies, baseline = 8.13 ± 0.052 , CNO probe = 7.42 ± 0.064 , saline probe = 7.85 ± 0.067 , Figure 3.2D). Repeated Measures Analysis of Variance (RM-ANOVA) was used to assess the significance of differences between CNO and saline injections on the change in theta frequency between baseline and probe trials. CNO significantly reduced the frequency of theta oscillations (RM-ANOVA, trial*drug interaction; $F_{(1,9)} = 15.50$, $p = 0.003$, Figure 3.2E) but did not affect their absolute power, or the ratio of power in the theta band compared to adjacent frequency bands, termed signal to noise (S2N) (RM-ANOVA, trial*drug interaction; log power: $F_{(1,9)} = 0.39$, $p = 0.25$; S2N: $F_{(1,9)} = 0.23$, $p = 0.37$, Figure 3.2F&G). Hence, analysis of power spectra indicated that stimulating medial septal cholinergic neurones reduced the frequency of LFP theta oscillations without changing their power.

The frequency of theta oscillations vary as a function of the animal's running speed (Sławińska and Kasicki 1998; Jeewajee et al. 2008). We therefore analysed what impact elevating cholinergic tone had on the theta frequency vs running speed relationship. Wavelets were used to generate instantaneous estimates of the LFP power spectrum, which were combined to produce average power spectra in running speed bins of width 0.25cm/s from 0-30cm/s. For each speed bin, the theta-band (3-11Hz) frequency with maximum power was identified. The theta frequency vs running speed relationship consistently appeared to be non-linear (Figure 3.2H, Figure 3.3A). Analysis of artificial LFP traces confirmed this was not an artefact of the analysis procedure (Figure 3.3B). The data were therefore fit with both linear ($y = A + B * x$) and three-term exponential ($y = A - B * \exp[-C * x]$) functions using linear and non-linear least squares respectively. The effect of varying each parameter in the exponential function can be seen in Figure 3.3. The F-test was then used to assess whether there was a significant difference in the quality of the fits. In 100% of the 95 baseline, CNO, and saline probe trials performed, the exponential fit was significantly better (all $p < 0.001$). Following injection of CNO, the average theta frequency vs running speed relationship was shifted to lower frequencies (Figure 3.2H). This change was reflected in a significant reduction in the exponential fit parameter A value which best fitted the theta frequency vs running speed relationship (RM-ANOVA, trial*drug interaction; $F_{(1,9)} = 37.42$, $p = 0.00018$, Figure 3.2I). In contrast, there was no significant effect of CNO on parameter B or C (RM-ANOVA, trial*drug interaction; parameter B: $F_{(1,9)} = 1.15$, $p = 0.31$; parameter C: $F_{(1,9)} =$

0.0044, $p = 0.95$, Figure 3.2J&K). Increasing the excitability of medial septal cholinergic neurones therefore decreased the frequency of theta oscillations, by reducing the average frequency observed at each running speed, without attenuating the extent to which the animal's speed modulated theta frequency.




Figure 3.2 Increasing medial septal cholinergic tone reduces LFP theta frequency by shifting its relationship with running speed to lower frequencies. **A**, Schematic representation of the experiment design: the Cre-dependent excitatory DREADD hM3Dq was injected into the medial septum of ChAT-Cre mice. In the same animal, tetrodes were implanted in the medial entorhinal cortex. **B**, 2.5x epifluorescence images of coronal brain sections showing ChAT expression (green) and mCherry expression (red), with expression of the latter limited to the medial septum. **C**, 40x confocal z-stack images demonstrating tight overlap between ChAT and mCherry expression. **D**, Average LFP power spectra: mean \pm SEM log power at each frequency across all baseline, CNO, and saline probe trials respectively. **E-G**, Blue = CNO, black = saline. Bars indicate mean \pm SEM across all animals in each group. Crosses indicate average value for each animal in each group. Stars indicate RM-ANOVA baseline vs probe * CNO vs saline interaction significance. *** = $p < 0.001$, ** = $p < 0.01$, * = $p < 0.05$, ns = $p > 0.05$. **E**, Frequency of maximum power in the LFP power spectrum in the theta (6-10Hz) band. **F**, Maximum log power in the theta band of the LFP power spectrum. **G**, Theta S2N: the ratio of average power in a 2Hz window around the theta peak compared to average power in the rest of the 3-25Hz range of the LFP power spectrum. **H**, Mean \pm SEM frequency of maximum power in the 3-11Hz range at each running speed, averaged across all baseline, CNO, and saline probe trials respectively. **I-K**, Three-term exponential fit parameter A, B, and C values respectively that best fitted the LFP theta frequency vs running speed relationship.




Figure 3.3 Examples of observed and artificial LFP theta frequency vs running speed plots demonstrate a clear non-linear relationship between the two. **A-B**, Black crosses indicate theta frequency of maximum power in the 3-11Hz range in each running speed bin. Green line indicates the best three-term exponential fit fitted with non-linear least squares. Red line indicates the best linear fit fitted with linear least squares. Stars indicate whether the exponential model fits the data significantly better than the linear model, accounting for the increased number of free parameters, using the F-test. *** = $p < 0.001$, ** = $p < 0.01$, * = $p < 0.05$, ns = $p > 0.05$. **A**, Examples from three baseline trials of observed theta frequency vs running speed profiles **B**, Examples of artificial sinusoidal LFP plots whose frequency was set to vary as a linear (left), saturating-exponential (middle), and sinusoidal (right) function of running speed. **C**, Left to right: plots demonstrating the effect of systematically varying the parameter A, B, and C values respectively in the three-term exponential function used for fitting, while holding the other parameters constant. Darker colours indicate lower parameter values, lighter colours higher values.

3.3.2 Increasing septal cholinergic tone also reduces the theta burst frequency of single units

Theta oscillations can also be observed in the bursts of action potentials emitted by individual neurones of the entorhinal cortex (Jeffery, Donnett, and O'Keefe 1995). We therefore asked whether modulating the excitability of medial septal cholinergic neurones affected the theta bursting of single units. The theta band modulation of cells' firing was quantified from the power spectra of the temporal autocorrelogram of each neurone's spike

train. Theta modulated cells were defined as those whose signal to noise ratio (S2N) comparing theta to other frequency bands was greater than that seen in 99/100 shuffles where spike times were randomly set. As only a small number of grid cells were significantly theta modulated (27 out of 227 baseline and probe grid cell trials [11.95%: 11.86% of baseline trials, 10.34% of CNO trials, 14.00% of saline trials], Figure 3.5B), we instead focused on a population of theta bursting cells not modulated by the animal's position or heading direction (125 out of 145 baseline and probe trials with significant theta modulation [86.21%], Figure 3.5A). As in the LFP, after injection of CNO, the theta burst frequency of these cells was reduced, without change in the strength of theta modulation (RM-ANOVA, trial*drug interaction; frequency: $F_{(1,6)} = 30.93$, $p = 0.0014$; S2N: $F_{(1,6)} = 0.20$, $p = 0.67$, Figure 3.4A-C).

The relationship between theta burst frequency and running speed was also assessed. To do so, power spectra of the temporal autocorrelograms generated from spike trains seen in 2 second windows across the trial were averaged together into 4cm/s speed bins, with the theta frequency of maximum power identified in each bin. Administration of CNO shifted the speed-frequency relationship downwards (Figure 3.4D). Again, this change corresponded to a reduction in parameter A of the exponential fit, whereas no change was seen in parameters B and C (RM-ANOVA, trial*drug interaction; parameter A: $F_{(1,3)} = 40.40$, $p = 0.008$, parameter B: $F_{(1,3)} = 3.66$, $p = 0.15$; parameter C: $F_{(1,3)} = 5.97$, $p = 0.092$, Figure 3.4E-G). The same conclusions were reached when averaging at the level of cell rather than animal, to increase the number of data points (Figure 3.5E-G). Hence, as in the LFP, increasing cholinergic tone reduced the burst frequency of theta modulated neurones by shifting the theta frequency vs running speed relationship to lower frequencies, without affecting the depth of theta frequency modulation across different speeds. Repeating these analyses averaging at the level of cell, rather than animal, did not affect these conclusions (Figure 3.5C-G).




Figure 3.4 Increasing medial septal cholinergic tone reduces the burst frequency of theta-modulated single units. **A**, Example power spectra indicating the power at each frequency in the 3-25Hz range in temporal autocorrelograms of spike trains from two theta-modulated single units. Left: a unit in a baseline trial and the same unit in a CNO probe trial. Right: a unit in a baseline trial and the same unit in a saline probe trial. Inset: smoothed spike train temporal autocorrelograms from the same cells/trials from which the power spectra were derived. **B-C**, **E-G**, Blue = CNO, black = saline. Bars indicate mean \pm SEM across all animals in each group. Crosses indicate average value for each animal in each group. Stars indicate RM-ANOVA baseline vs probe * CNO vs saline interaction significance. *** = $p < 0.001$, ** = $p < 0.01$, * = $p < 0.05$, ns = $p > 0.05$. **B-C**, Only cells with significant theta modulation (assessed by a shuffle procedure) in either a baseline or probe trial are included. **B**, Theta burst frequency of single units, as inferred from the frequency of maximum power in the 6-10Hz theta window of each unit's power spectrum. **C**, Theta S2N: the ratio of average power in a 2Hz window around the theta peak compared to average power in the rest of the 3-25Hz range of each unit's power spectrum. **D**, Mean \pm SEM unit burst frequency of maximum power in the 3-11Hz range in each running speed bin, averaged across all baseline, CNO, and saline probe trials with significant theta modulation respectively. **E-G**, Parameter A, B, and C values respectively that best fit the single unit theta frequency vs running speed relationship. Only cells with significant theta modulation (assessed by a shuffle procedure) in either a baseline or probe trial, and only trials with a significantly better exponential than linear fit are included.




Figure 3.5 Conclusions of the effect of CNO on theta-modulated unit burst frequency are unchanged by performing analyses averaging at the level of cell A-B, Distribution of theta signal to noise (Theta S2N) as a function of location within a distribution of shuffled theta S2N values (shuffle percentile) for theta-modulated non-spatial cells (A) and grid cells (B). Only baseline and probe trials are included. Vertical dashed grey lines represent the significance threshold. Inset are example temporal autocorrelograms representative of specific S2N values. C-G, Blue = CNO, black = saline. Bars indicate mean \pm SEM across all cells in each group. Crosses indicate average value for each cell in each group. Only cells with significant theta modulation (assessed by the shuffle procedure) in either a baseline or probe trial are included. Stars indicate mixed-model ANOVA baseline vs probe * CNO vs saline interaction significance. *** = $p < 0.001$, ** = $p < 0.01$, * = $p < 0.05$, ns = $p > 0.05$. C, Theta burst frequency of single units. D, Theta S2N. E-G, Only cells with significant theta modulation (assessed by a shuffle procedure) in either a baseline or probe trial and only trials with a significantly better exponential than linear fit (according to the F-test) are included. E-G, Parameter A, B, and C values respectively that best fit the single unit theta frequency vs running speed relationship.

3.3.3 Increasing septal cholinergic tone induces a pattern of behaviour consistent with the detection of novelty

At the start of probe trials following injection of CNO, mice displayed a pattern of behaviour normally seen during exposure to novel environments. In the first five minutes of CNO probe,

but not saline probe trials, mice exhibited a clear tendency to run predominantly around the perimeter of the environment (Figure 3.6A). Indeed, injection of CNO led to a significant increase in the proportion of the first five minutes of the trial that mice spent within 10cm of the walls of the enclosure ('perimeter dwell', RM-ANOVA, trial*drug interaction; $F_{(1,9)} = 17.10$, $p = 0.0025$, Figure 3.6B), and not moving (RM-ANOVA, trial*drug interaction; $F_{(1,9)} = 7.25$, $p = 0.025$, Figure 3.6C). These changes were seen despite there being no effect of CNO on the average running speed, either in the first five minutes or across the whole trial (RM-ANOVA, trial*drug interaction, first five minutes; $F_{(1,9)} = 0.21$ $p = 0.66$; whole trial; $F_{(1,9)} = 0.11$, $p = 0.75$, Figure 3.6D). There was also no difference in the average distribution of running speeds across CNO and saline probe trials (Kolmogorov-Smirnov test, $p = 0.998$, Figure 3.6E). Therefore, in addition to modulating theta frequency, increasing the excitability of medial septal cholinergic neurones also led mice to spend a greater proportion of the beginning of subsequent trials near the perimeter of the enclosure and not moving.




Figure 3.6 Increasing medial septal cholinergic tone results in a pattern of behaviour normally seen during exposure to a novel environment. **A**, Cumulative position plots from two animals (one in each row), indicating coverage of the environment in the first five minutes of baseline, CNO, and saline probe trials. **B-D**, Blue = CNO, black = saline. Bars indicate mean \pm SEM across all animals in each group. Crosses indicate average value for each animal in each group. Stars indicate RM-ANOVA baseline vs probe * CNO vs saline interaction significance. *** = $p < 0.001$, ** = $p < 0.01$, * = $p < 0.05$, ns = $p > 0.05$. **B**, Proportion of the first five minutes of the trial spent within 10cm of the walls of the enclosure ('perimeter dwell'). **C**, Proportion of the first five minutes of the trial spent sitting still (running speed < 1 cm/s). **D**, Mean running speed across the whole trial. **E**, Mean \pm SEM proportion of the whole trial spent at each 2cm/s speed bin across all baseline, CNO, and saline probe trials respectively.

3.3.4 *Injection of CNO has no effect in control mice*

Control animals transfected with the fluorophore virus AAV2-hSyn-DIO-mCherry were exposed through intraperitoneal injections to 1, 3, and 5mg/kg doses of CNO, and equivalent volumes of saline. In contrast to mice injected with the hM3Dq virus, no effect of CNO was observed on peak LFP theta frequency (Figure 3.7A), LFP theta frequency vs running speed exponential fit parameter A (Figure 3.7B), perimeter dwell (Figure 3.7C), or the proportion of time spent sitting still (Figure 3.7D). That is, no difference was seen between baseline, CNO probe, or saline probe trials averaging across all concentrations at the level of animal.

Further, no dose-dependent effects of CNO were seen on any of these measures when values were separated according to the concentration of CNO administered (Figure 3.7E-H). Further, the changes observed in hM3Dq mice cannot be explained by the fact that CNO probe trials always followed saline probe trials on days in which a baseline, saline, and CNO probe trial were carried out: when only data from days in which only either a CNO or saline probe trial were performed, with the order of days counterbalanced between animals, all significant results were maintained (Figure 3.8A-D). These data thus give credence to the conclusion that the effects of CNO described here result from increases in the excitability of medial septal cholinergic neurones.




Figure 3.7 Effects observed following injection of CNO in hm3DQ mice are not seen in DIO-mCherry control mice. A-H, Grey = baseline, blue = CNO, black = saline. Bars indicate mean \pm SEM across all animals in each group. Crosses indicate average value for each animal in each group. A-D, Values are averaged across all three drug concentrations administered. E-H, Values are separated according to the drug concentration administered. A, E, Frequency of maximum power in the LFP power spectrum in the theta (6-10Hz) band. B, F, Three-term exponential fit parameter A values that best fitted the LFP theta frequency vs running speed relationship. C, G, Proportion of the first five minutes of the trial spent within 10cm of the walls of the enclosure ('perimeter dwell'). D, H, Proportion of the first five minutes of the trial spent sitting still (speed <1cm/s).




Figure 3.8 Effect of CNO on theta frequency and behaviour still seen when analysis is restricted to novelty days, in which only either CNO or saline were administered. A-C, Blue = CNO, black = saline. Bars indicate mean \pm SEM across all animals in each group. Crosses indicate average value for each animal in each group. Stars indicate RM-ANOVA baseline vs probe * CNO vs saline interaction significance. *** = $p < 0.001$, ** = $p < 0.01$, * = $p < 0.05$, ns = $p > 0.05$. **A**, Frequency of maximum power in the LFP power spectrum in the theta (6-10Hz) band. **B**, Three-term exponential fit parameter A values that best fitted the LFP theta frequency vs running speed relationship. **C**, Proportion of the first five minutes of the trial spent within 10cm of the walls of the enclosure ('perimeter dwell'). **D**, Proportion of the first five minutes of the trial spent sitting still (speed < 1 cm/s).

3.3.5 Increasing septal cholinergic tone has no effect on grid firing patterns in familiar environments

Grid cells were recorded in seven of the 10 mice. A cell was classified as a grid cell if it exhibited hexagonally periodic spatial firing (gridness score > 0.3) in either a baseline or probe trial. Following injection of CNO, little change was observed when comparing ratemaps of the same grid cells between the baseline and CNO probe trials (Figure 3.9A). Indeed, there was no significant change in gridness scores following injection of CNO (RM-ANOVA, trial*drug interaction; $F_{(1,6)} = 1.05$, $p = 0.34$, Figure 3.9B). Neither was there a significant effect on grid scale, assessed from the spatial autocorrelogram of the ratemap, when looking at cells with gridness > 0.3 in both a baseline and probe trial (RM-ANOVA, trial*drug interaction; $F_{(1,5)} = 0.021$, $p = 0.89$, Figure 3.9C).




Figure 3.9 Increasing medial septal cholinergic tone has no effect on grid cell firing patterns in familiar environments. **A**, Example ratemaps and spatial autocorrelograms from four grid cells (one in each row), in a baseline trial (left hand side) and a CNO probe trial (right hand side). **B-E**, Blue = CNO, black = saline. Bars indicate mean \pm SEM across all animals in each group. Crosses indicate average value for each animal in each group. Only cells with gridness > 0.3 in either a baseline or probe trial are included. **B-D**, Stars indicate RM-ANOVA baseline vs probe * CNO vs saline interaction significance. *** = $p < 0.001$, ** = $p < 0.01$, * = $p < 0.05$, ns = $p > 0.05$. **B**, Degree of hexagonal regularity of grid cell firing patterns, as measured by the gridness score. **C**, Grid cell scale in cm, measured as the median distance to the six peaks closest to the centre of the autocorrelogram. Only trials with gridness > 0.3 are included. **D**, Intra-trial stability of grid cell firing patterns, as measured by the value at the centre of a cross-correlogram constructed from two ratemaps, one from each half of the trial. **E**, Inter-trial stability of grid cell firing patterns, as measured by the value at the centre of a cross-correlogram constructed from two ratemaps, one from the baseline and one from the probe trial. No significant difference was found between the groups using a paired t-test.

In analysing whether CNO affected the intra-trial stability of grid-patterns, bin-wise Pearson correlations were calculated between ratemaps corresponding to the first and second half of each trial. No effect of CNO was observed (RM-ANOVA, trial*drug interaction; $F_{(1,6)} = 0.75$, $p = 0.42$, Figure 3.9D). Similarly, we assessed the inter-trial stability of grid cells using the same procedure to compare baseline ratemaps with those from saline and CNO probe trials. CNO was not found to affect the spatial stability of grid firing across trials (paired t-test, $t_6 = -0.38$, $p = 0.72$, Figure 3.9E). The same conclusions were reached when averaging at the level of cell, rather than animal, to increase the number of data points (Figure 3.12A-B). That is, while increasing cholinergic tone modulated theta frequency and behaviour, no effect was seen on grid cell firing patterns in familiar environments.

3.3.6 *Increasing septal cholinergic tone has no effect on grid cell speed signals*

Given the absence of change in grid cell firing patterns, we next asked whether stimulating medial septal cholinergic neurones affected putative speed signals seen in grid cells. First, we analysed changes in the firing rates of grid cells seen as a function of running speed. Specifically, following Kropff et al. 2015, we calculated the instantaneous firing rate of each grid cell, using linear regression to identify the slope and correlation ('speed score') of its relationship with running speed. A grid cell was identified as significantly speed modulated if, in either a baseline or probe trial, it had both gridness > 0.3 and a speed score exceeding 99% of a null distribution of 1000 speed scores generated by randomly permuting the cell's spike train relative to the animal's running speed. No effect of CNO could be seen in the binned average firing rate vs running speed plots of speed modulated grid cells (Figure 3.10A). Indeed, in animals with speed modulated grid cells, there was no effect of CNO on the speed score, nor the slope of the firing rate vs running speed regression line (RM-ANOVA, trial*drug interaction; speed score: $F_{(1,5)} = 0.00079$, $p = 0.979$; slope: $F_{(1,5)} = 0.41$, $p = 0.55$, Figure 3.10B&C). Next, we asked whether the putative encoding of speed in the difference in slope of intracellular and LFP theta frequency vs running speed relationships was affected by CNO injections. Such differences are thought to manifest in theta phase precession (O'Keefe and Burgess 2005), the tendency of grid cells to fire at progressively earlier phases of the theta cycle as the animal moves through grid fields (Hafting et al. 2008; Jeewajee et al. 2014). Following CNO injection, theta phase precession could still be observed in grid cells (Figure 3.10D), and there was no effect on the correlation or slope of phase precession in grid cells assessed as significantly phase precessing by a shuffle procedure (RM-ANOVA, trial*drug interaction; $F_{(1,3)} = 1.10$, $p = 0.38$; $F_{(1,3)} = 1.51$, $p = 0.31$, Figure 3.10E&F). Repeating the speed modulation and phase precession analyses averaging at the level of cell, rather than animal, to increase the number of data points, did

not change these conclusions (Figure 3.12C-F). The absence of change in grid cell firing patterns was therefore consistent with a lack of change in putative speed signals in grid cells following modulation of medial septal cholinergic tone.




Figure 3.10 Increasing medial septal cholinergic tone has no effect on putative grid cell speed signals. *A*, Mean \pm SEM firing rate of grid cells as a function of running speed, averaged across all baseline, CNO, and saline probe trials respectively. Only grid cell trials with significant speed modulation (determined by a shuffle procedure) are included. *B-C*, *E-F*, Blue = CNO, black = saline. Bars indicate mean \pm SEM across all animals in each group. Crosses indicate average value for each animal in each group. Stars indicate RM-ANOVA baseline vs probe * CNO vs saline interaction significance. *** = $p < 0.001$, ** = $p < 0.01$, * = $p < 0.05$, ns = $p > 0.05$. *B-C*, Only grid cells with significant speed modulation in either a baseline or probe trial are included. *B*, Grid cell speed scores: the Pearson correlation between a cell's instantaneous firing rate and the animal's running speed. *C*, Grid cell speed slope: the slope of the regression line which best fitted the relationship between instantaneous firing rate and running speed. *D*, Two example CNO probe trials in which significant grid cell phase precession was seen. One cell is plotted in each row. Left-hand side: smoothed firing rate map, with crosses indicating field peaks. Centre: watershed-isolated valid fields. Right: scatter of theta phase as a function of proportion of distance along run through the field. Red line indicates best fitting circular-linear regression slope. *E-F*, Only grid cells with significant phase precession in either a baseline or probe trial are included. *E*, Circular-linear correlation values relating theta phase to proportional distance through the field. *F*, The slope of the circular-linear phase precession regression line.

3.3.7 *Increasing septal cholinergic tone has no effect on grid firing patterns in novel environments*

Finally, to assess whether a lack of change in grid cell firing patterns following CNO injection was the result of their being stabilised by familiar sensory cues, we recorded from grid cells while mice explored novel environments after injection of either CNO or saline. In the first novel trial following both CNO and saline injections, some grid cells exhibited significantly hexagonal firing patterns (Figure 3.11A). While gridness was reduced in the first novel trial, there was not a significant interaction between CNO and saline in the change in gridness between baseline and novel trials (main effect trial (baseline vs novel); $F_{(1,9)} = 10.54$, $p = 0.023$; trial*drug interaction; $F_{(1,5)} = 1.07$, $p = 0.35$, Figure 3.11B). Neither was there a significant effect of CNO on grid scale, when looking at cells with gridness above 0.3 in both the familiar and novel environments (trial*drug interaction; $F_{(1,9)} = 0.92$, $p = 0.38$, Figure 3.11C). Again, analysis at the level of cell did not change these conclusions (Figure 3.12G-I). Thus, increasing the excitability of medial septal cholinergic neurones had similarly little effect on grid cell firing patterns in familiar and novel environments.




Figure 3.11 Increasing medial septal cholinergic tone has no effect on grid cell firing patterns in novel environments. **A**, Example ratemaps and spatial autocorrelograms from eight cells which had significant hexagonal firing patterns (gridness > 0.3) in the first trial in a novel environment following administration of CNO (left) or saline (right). **B-C**, Blue = CNO, black = saline. Bars indicate mean \pm SEM across all animals in each group. Crosses indicate average value for each animal in each group. Stars indicate RM-ANOVA baseline vs Novel * CNO vs saline interaction significance. *** = $p < 0.001$, ** = $p < 0.01$, * = $p < 0.05$, ns = $p > 0.05$. Only cells with gridness > 0.3 in either a baseline or novel trial are included. **B**, Degree of hexagonal regularity of grid cell firing patterns, as measured by the gridness score. **C**, Grid cell scale in cm, measured as the median distance to the six peaks closest to the centre of autocorrelogram. Only trials with gridness > 0.3 are included.




Figure 3.12 Conclusions of the effect of CNO on grid cell firing patterns are unchanged by repeating analyses averaging at the level of cell. **A-I**, Blue = CNO, black = saline. Bars indicate mean \pm SEM across all cells in each group. Crosses indicate average value for each cell in each group. Only cells with gridness > 0.3 in either a baseline or probe trial are included. Stars indicate mixed-model ANOVA baseline vs probe * CNO vs saline interaction significance. *** = $p < 0.001$, ** = $p < 0.01$, * = $p < 0.05$, ns = $p > 0.05$. **A-F**, Only baseline and probe trials in the familiar environment are included. **A**, Degree of hexagonal regularity of grid cell firing patterns, as measured by the gridness score. **B**, Grid cell scale in cm, measured as the median distance to the six peaks closest to the centre of the autocorrelogram. Only trials with gridness > 0.3 are included. **C-D**, Only grid cells with significant speed modulation in either a baseline or probe trial are included. **C**, Grid cell speed scores: the Pearson correlation between a cell's instantaneous firing rate and the animal's running speed. **D**, Grid cell speed slope: the slope of the regression line which best fitted the relationship between instantaneous firing rate and running speed. **E-F**, Only grid cells with significant phase precession in either a baseline or probe trial are included. **E**, Circular-linear correlation values relating theta phase to proportional distance through the field. **F**, The slope of the circular-linear phase precession regression line. **G-I**, Comparison of grid cell firing patterns between baseline and novel trials. **G**, Degree of hexagonal regularity of grid cell firing patterns, as measured by the gridness score, including only the first novel trial. **H**, Degree of hexagonal regularity of grid cell firing patterns, as measured by the gridness score, including all novel trials and days. **I**, Grid cell scale in cm, measured as the median distance to the six peaks closest to centre of autocorrelogram. Only trials with gridness > 0.3 are included, including all novel trials and days.

3.4 Discussion

3.4.1 *Summary of results and significance*

Aiming to illuminate the functions of the medial septum, and in particular its contribution to grid cell activity, we modulated septal cholinergic tone while recording from grid cells in mice exploring familiar and novel environments. Increasing the excitability of medial septal cholinergic neurones reduced theta frequency by shifting its saturating-exponential relationship with running speed to lower frequencies. However, no effect was seen on grid cell firing patterns in familiar environments, consistent with a lack of change in putative speed signals used by grid cells. While the regularity of grid firing decreased in novel environments, no novelty-induced increase in grid scale was observed, preventing testing of the hypothesis that, in rats, grid scale increases in novel contexts is a result of an increase in cholinergic tone. However, mice exhibited behaviours normally seen on exposure to novel environments during increases in the excitability of cholinergic neurones, consistent with a role for acetylcholine in the signalling of novelty.

3.4.2 *Role of septal acetylcholine in theta oscillations*

A reduction in theta frequency following an increase in cholinergic tone is concordant with evidence demonstrating that application of the muscarinic receptor agonist carbachol reduces the frequency of subthreshold membrane potential oscillations (Klink and Alonso 1997) and resonant frequency (Heys, Giocomo, and Hasselmo 2010) of entorhinal layer-II ocean-stellate cells in-vitro. Through application of the acetylcholinesterase inhibitor physostigmine, similar results have also been observed in-vivo (Tsuno, Schultheiss, and Hasselmo 2013). Blockade of muscarinic receptors with systemic administration of scopolamine has also been observed to reduce theta frequency in-vivo, though in contrast to the present results, through a flattening of the theta frequency vs running speed slope (Newman et al. 2013). Further, lesions specifically targeting septal cholinergic neurones have been found to reduce the power of Type I theta oscillations without affecting their frequency (Lee et al. 1994; Yoder and Pang 2005). The current results therefore complement a growing body of evidence demonstrating a role for medial septal cholinergic neurones in modulating Type I theta oscillations, though one in which their precise functional contribution remains unclear. Some variability in the described results likely stems from differences in the techniques employed both to modulate cholinergic tone and record theta. For example, recording in the hippocampus rather than the entorhinal cortex may lead to a change in power without a change in frequency of theta oscillations, because of an absence of the ocean-stellate cells whose subthreshold oscillatory dynamics are modulated by

acetylcholine. Similarly, non-linear responses to acetylcholine concentration, and adaptation over time, may mean that the impact of any intervention depends on whether it increases or decreases cholinergic tone, is acute or chronic. Further experiments are also required to identify whether the observed reduction in theta frequency is due to local connections of cholinergic neurones within the medial septum, or direct projections to the entorhinal cortex.

3.4.3 Non-linear relationship between theta frequency and running speed

The non-linear relationship between theta frequency and running speed observed here is unlikely to be an artefact of the analysis, as the same procedure was capable of correctly identifying linear and sinusoidal relationships in artificially-generated data. Further, though not remarked upon, a number of observations of non-linear increases in theta frequency have been made in the past (Jeewajee et al. 2008; Newman et al. 2013; Chen et al. 2016). The saturating-exponential relationship consistently observed here means that the fastest rate of change in theta frequency is seen at the lowest running speeds. It is possible that this rapid change in frequency represents a discontinuous transition between immobility-related Type II and mobility-related Type I theta. However, in the current data, analysis using narrow-width (0.25cm/s) speed bins consistently indicated a smooth increase in frequency, rather than a discontinuous jump between low- and high-frequency theta states. Oscillatory interference models, which use differences in the frequency vs running-speed slopes of multiple theta oscillations to encode an animal's velocity, are not in theory invalidated by non-linear slopes: to generate firing patterns that track the distance travelled by an animal, current models only require that the difference in the slopes of the constituent oscillators be linear across running speeds.

3.4.4 No evidence for an involvement of septal acetylcholine in the generation of grid firing patterns

In contrast to the observed effects on theta frequency, no change was seen in the firing patterns of grid cells recorded in familiar environments, consistent with an absence of change in their putative speed signals. The absence of change in firing rate vs speed profiles of grid cells seen here suggests such activity does not depend on cholinergic projections from the medial septum. Recordings from septal glutamatergic neurones indicate that their projections to the entorhinal cortex provide a firing rate signal which varies as a function of running speed (Fuhrmann et al. 2015; Justus et al. 2017), and it may be their activity which underlies a change in speed coding in entorhinal neurones following inactivation of the medial septum (Hinman et al. 2016). In the present experiment, increasing cholinergic tone had no effect on the depth of change in theta frequency seen as a function of running speed,

in contrast to the flatter profile seen following systemic administration of the muscarinic receptor antagonist scopolamine (Newman et al. 2013). This difference possibly reflects the fact that cholinergic tone was here increased, in contrast to the effective reduction in tone resulting from blockade of muscarinic receptors. The difference in effect on theta frequency-speed profiles may explain why little change was seen here in grid cell firing patterns, in contrast to the breakdown of periodic firing that has previously been reported following scopolamine administration (Newman, Climer, and Hasselmo 2014). The present results are therefore consistent with the idea that the functional contribution of the medial septum to grid cell activity is provision of one or more speed signals, though provides little additional evidence to support this hypothesis.

3.4.5 Absence of novelty-induced increases in grid scale

In agreement with previous results, a reduction in the regularity of grid firing patterns was seen during exposure to novel environments (Barry, Ginzberg, et al. 2012). In contrast, however, a temporary increase in the scale of grid firing patterns was not observed. The possible reasons for this difference are discussed fully in Section 4.4.3, following replication of this result in wildtype mice. An absence of novelty-induced expansion of grid scale prevented testing of the hypothesis that an increase in cholinergic tone underlies this phenomenon in rats (Barry, Ginzberg, et al. 2012; Barry, Heys, and Hasselmo 2012). While modulating the excitability of cholinergic neurones here had no effect on grid firing in novel environments, it is possible that this was due to acetylcholine concentration reaching saturating levels. That is, an endogenous increase in cholinergic tone due to novelty may have masked any effect of the artificial elevation of the excitability of septal cholinergic neurones. However, the observation of lower theta frequency in novel environments following injection of CNO (data not shown) suggests that this is not the case.

3.4.6 Possible role of septal acetylcholine in signalling novelty

The tendency of rodents to remain still and avoid the centre of a novel open enclosure has long been used as a test to identify anxiolytic effects of drugs (Prut and Belzung 2003), while electric shocks have been shown to increase rats' tendency to 'freeze', and remain near the perimeter of novel environments (Grossen and Kelley 1972). As such, these freezing and 'thigmotactic' behaviours are closely associated with both novelty and anxiety. In the current experiment, the increased proportion of time spent in the periphery of the enclosure and not-moving cannot be accounted for by an anxiogenic effect of the intraperitoneal injection alone, as the increase in both measures seen following injection of CNO was significantly

greater than that following injection of saline. Interestingly, the reduction in theta frequency seen here is similar to that observed following injection of anxiolytics in rats (Wells et al. 2013), while the behavioural changes are redolent of anxiogenesis or novelty. Acetylcholine is closely related to learning and memory (Deiana, Platt, and Riedel 2011), and observed increases in cholinergic tone in novel contexts (Aloisi et al. 1997) are thought to drive the encoding of new memories in the hippocampal formation (Hasselmo 2006). The novelty-like behaviours seen here after stimulation of septal cholinergic neurones could therefore be interpreted as the result of inducing a sense of novelty, consistent with a role for acetylcholine in orchestrating network dynamics that favour the encoding of new memories; although the current data cannot distinguish this interpretation from an increase in anxiety.

4 Distal sensory features of the environment determine the scale, stability, and regularity of grid firing patterns

4.1 Introduction

When rats are transferred to a novel environment, the spacing between the peaks of the grid pattern, the grid scale, increases (Barry, Ginzberg, et al. 2012). The expansion in grid scale is greatest in the first trial in the novel environment, with the pattern returning to its 'baseline' scale after several days of experience (Barry, Ginzberg, et al. 2012). An apparently robust result, this phenomenon has been replicated in multiple labs (Hafting et al. 2005; T. Stensola et al. 2015). A number of theoretical accounts of the utility of an increase in grid scale in novel contexts have been proposed. First, grid expansion may represent a compensatory mechanism which minimises the impact of increased uncertainty in novel environments (Towse et al. 2014). That is, on entering a previously unexplored space, the unfamiliarity of the environment means an animal is less able to self-locate using sensory cues alone. As such, the capacity of grid cells to use sensory cues to 'reset' noise accrued through path integration is reduced. If path integration were performed separately by distinct modules, each would accumulate noise independently, potentially leading to catastrophic errors in the animal's encoded location due to inconsistency across modules. In increasing their scale, grid cells may therefore sacrifice spatial resolution to increase their signal to noise ratio, and so reduce the chance of making egregious errors (Towse et al. 2014). Second, grid expansion may help generate place cell remapping (Barry, Ginzberg, et al. 2012). Temporary inconsistency between inputs to place cells from grid cells which expand, and boundary vector cell which do not (Solstad et al. 2008; Lever et al. 2009), could create a mismatch which causes place fields to shift. However, independent changes in the orientations and phases of distinct grid modules are probably alone sufficient to generate complete place cell remapping (Monaco and Abbott 2011). Finally, the expansion and subsequent contraction of the grid pattern could act as novelty and familiarity signals respectively to downstream regions (Barry, Ginzberg, et al. 2012).

As discussed, an increase in cholinergic tone has been suggested to underlie grid expansion in novelty (Barry, Heys, and Hasselmo 2012). To test this hypothesis, we recorded from grid cells in familiar and novel enclosures whilst modulating the activity of medial septal cholinergic neurones (Section 3.3). In the first recording made in a novel environment during increased cholinergic tone, the absence of an increase in grid scale led to the excited belief that this represented evidence for a role of acetylcholine in grid expansion. However,

following recordings in novel environments after the injection of saline, it became clear that a lack of grid expansion was not the result of altered cholinergic tone. Indeed, no evidence for increases in grid scale in novel contexts was observed in recordings across seven mice, regardless of the modulation of septal cholinergic tone (Section 3.3.7).

To assess whether the absence of novelty-based grid expansion could be attributed to the ChAT-Cre mice or medial septal injections used in the described experiments, we also recorded grid cells from C57/B6 wildtype mice in familiar and novel environments across multiple days of experience. Three mice were screened in Room 412, and tested in novel environments in Room 413. Two mice were screened in Room 413, and tested in a novel environment in Room 412. The room in which recordings were made had a profound influence on the properties of grid cell firing patterns, determining their scale, gridness, and intra-trial stability. Changes in these properties between rooms were seen despite a high degree of similarity in the immediate enclosure, and even when the same enclosure was used between rooms. Relatively little evidence for novelty-induced changes in grid scale was observed, however the regularity of firing patterns in the novel environment increased across days. These preliminary data therefore indicate that increases in grid scale in novel environments may be species specific, and that the distal sensory features of an environment can modulate salient properties of grid cell firing.

4.2 Methods

4.2.1 Animals

All work was conducted within the terms of appropriate Home Office Project and Personal licences.

Five C57/B6 ‘wildtype’ mice were used in this experiment. Prior to surgery, mice were housed communally under a 12:12 inversed light-dark cycle, with ad libitum access to food and water. After surgery, mice were transferred to individual Perspex cages, and their food was restricted such that their weight fell to 85% of that seen prior to surgery. Food was then provided such that the weight of each mouse increased by 0.5g per week.

4.2.2 Microdrives and surgery

Custom-built microdrives were produced and implanted as described in section 3.2.2, except no injections were made into the medial septum. All mice received two microdrives, each with four tetrodes, one targeting the medial entorhinal cortex of each hemisphere.

4.2.3 *Experimental protocol and electrophysiological recording*

Electrophysiological and position data were recorded as described in section 3.2.3.

The five mice were divided into two groups. The first group comprised three mice screened for grid cells in Room 412, the same room and configuration used for screening in the DREADD experiment. The other two mice were instead screened in Room 413, such that the screening environments were counterbalanced between mice.

On identification of stable grid cell recordings, mice were transferred to the alternative room for recordings to be made in a novel environment. On each experimental day, a baseline trial in the familiar screening environment was followed by three novel probe trials in the novel environment. A fifth trial was then performed with the mouse returned to the familiar screening environment. Trials lasted 15-30 minutes, terminated once the mouse was deemed to have covered the enclosure sufficiently to assess the spatial tuning of any recorded cells. Adjacent trials were separated by a 20-30-minute break in which the animal was returned to its home cage and the enclosure was cleaned with water. Recordings in each novel environment were continued for five days.

Screening of mice in Room 412 used a 1m diameter circular enclosure. The enclosure had 50cm high walls, was surrounded by black curtains over the western 180°, and open to the dimly-lit experimental room over the eastern 180°. A single white cue-card lit by a desk-lamp hung at the western pole. These mice were then recorded in two novel environments, one at either end of Room 413. The first novel environment contained a 90x90cm square enclosure. The same 1m diameter circle used for screening was the enclosure of the second novel environment. Both novel environments were open to the experimental room, but differed from one another in the colour of the floor and the distal cues present. All five days of recording were performed in the first novel environment before moving to the second.

Mice were screened in Room 413 using the 90x90cm enclosure and environment that acted as the novel square environment for mice screened in Room 412. In turn, upon identification of stable grid cell recordings, mice screened in Room 413 were transferred to Room 412. Recordings were there made in a second 90x90cm enclosure. Apart from being a 90x90cm square rather than a 1m diameter circle, the environment was otherwise identical to that used for screening the 412-first mice. Mice screened in Room 413 were only recorded in one, rather than two novel environments.

90x90cm square and 1m diameter circular enclosure were chosen so that the surface area of recordings in square and circular environments were consistent (8100cm² and 7854cm² respectively).

A summary of the experimental protocol and environments used can be seen in Figure 4.1 below.




Figure 4.1 Illustration of experimental protocol for novelty recordings in wild type mice. **A**, Schematic illustration of the environments used for screening and novel recordings in the two groups of mice. Mice screened in Room 412 each were recorded across 10 novel recording days, five in each novel environment in Room 413. Mice screened in Room 412 had just five novel recording days, all in the same environment in Room 413. **B**, Each novel recording day consisted of five trials: a baseline trial followed by three novel probe trials, and finally a second familiar trial (not used in analysis).

4.2.4 Analyses

4.2.4.1 Clustering and ratemap generation

Recorded spikes were organised into putative clusters as described in Section 3.2.3.

Position data and the spikes from each putative cluster were spatially binned to generate firing rate maps, as described in 3.2.4.5.

4.2.4.2 Grid cell analysis

From recorded ratemaps, the gridness, scale, and intra-trial stability of each putative cluster was calculated as in Section 3.2.4.5.

Any cell with a gridness of 0.3 in either a baseline or novel trial was considered a grid cell. For each grid cell, gridness and intra-trial stability values were mean-averaged across all

baseline and all novel trials. Across all unique grid cells, plots were then produced comparing gridness and stability in the baseline and novel environments, separately for the mice screened in Room 412 and Room 413.

In comparing grid scale changes between environments, scale values for each grid cell were mean-averaged across all baseline and novel trials in which gridness exceeded 0.3. This was done separately for mice screened in Room 412 and mice screened in Room 413, with the mean values from each recording room plotted against one another for each group.

For all three measures, the same analysis was also performed separately for each recording day. The mean and standard error of each value was then plotted as a function of experimental day, to assess changes in grid cell firing patterns that coincided with animals becoming more familiar with the novel environment.

4.2.4.3 Hypothesis testing

Given the small number of animals in this preliminary dataset, hypothesis testing was performed using values averaged at the level of grid cell, rather than animal. For mice screened in Room 412, data were pooled from both novel environments, as inspection of the data by eye indicated no differences between the two in any of the measures analysed. Paired t-tests were used to assess the null hypothesis of no difference in each measure between the baseline and novel environments, with tests performed separately for mice screened in Room 412 and Room 413. Further paired t-tests were used to assess the same null hypothesis, but with data only from specific recording days included. Finally, unpaired t-tests were used to assess, separately for the two groups of mice, the null hypothesis that the difference in particular grid measures between baseline and novel environments were not different between the first and fifth days in the novel environment.

4.3 Results

4.3.1 Grid scale is determined by the distal sensory features of the environment

To assess the influence of novel recording contexts on grid cell firing properties, we recorded grid cells from wildtype mice exploring familiar and novel environments in two experimental rooms, counterbalancing which room was used for screening between mice. In mice screened in Room 412, the grid scale of firing patterns in novel environments in Room 413 were consistently lower than those seen during screening (Figure 4.2A). The reduced grid scale was seen even when the exact same enclosure was used (Figure 4.2A). Grid cells

of different baseline scale appeared to show a similar absolute reduction in grid scale in the novel environment (mean reduction of 4.94cm, a mean decrease of 10.69% from baseline scale, Figure 4.2C). The decrease in grid scale was significantly greater than would be expected by chance (*paired t-test*: $t_{54} = 9.48$, $p = 4.39 \times 10^{-13}$). No consistent change in the difference in scale between the environments was seen across multiple days of recording, during which time the animal became increasingly familiar with the novel environments. There was not a significant correlation between recording day and the difference in grid scale between environments, nor a significant change in the difference in grid scales between the first and fifth recording days (*Spearman-rank correlation*: $\rho = 0.14$, $p = 0.216$. *Unpaired t-test, day 1 vs day 5*, $t_{27} = -0.0029$, $p = 0.998$, Figure 4.2D, left-hand panel).

In contrast, mice screened in Room 413 showed an increase in grid scale during recordings in a novel environment in Room 412 (Figure 4.2B). Again, the change in grid scale was seen despite the immediate enclosure being the same shape and size as that used for screening. Grid modules of different scale appeared to show a similar absolute increase in grid scale (a mean increase of 1.64cm, a 5.30% increase in grid scale compared to baseline, Figure 4.2C). Grid scales observed in the novel environment of Room 413 were significantly greater than those seen during screening in Room 412 (*paired t-test*, $t_{31} = -4.13$, $p = 2.55 \times 10^{-4}$). Unlike mice screened in Room 412, a fairly consistent reduction in the difference in grid scale between the baseline and novel environments was observed as the novel environment became more familiar (Figure 4.2D, right-hand panel). A significant correlation between the difference in grid scales between the environments across days was observed, and a significant difference seen in the change in scale between the first and fifth recording days (*Spearman-rank correlation*: $\rho = 0.51$, $p = 0.0083$. *Unpaired t-test*, $t_7 = -2.41$, $p = 0.0467$).

The grid scale differences observed here cannot be explained by differences in grid regularity between environments, as only trials in which the cell's gridness exceeded 0.3 were included in the analysis. These data thus demonstrate that the scale of grid firing patterns are determined by the distal sensory features of the recording environment. In contrast, only relatively weaker evidence was found for an effect of novelty on grid scale.




Figure 4.2 Grid scale is determined by the distal sensory features of the environment. **A-B**, Firing rate maps and spatial autocorrelograms of grid cells with gridness > 0.3 in both the baseline and ‘novel’ environments. Values indicate the grid scales observed. **A**, Grid cells from mice screened in Room 412 appeared to contract during recordings in novel environments in Room 413. **B**, Grid cells from mice screened in Room 413 instead expanded during recordings in Room 412. **C-D**, Blue/left: mice screened in Room 412. Black/right: mice screened in Room 413. **C**, Grid scale in the novel environment as a function of scale in the familiar environment. Stars indicate significance level of paired t-tests comparing baseline and novel grid scales. *** = $p < 0.001$. **D**, Mean \pm SEM difference in grid scale (baseline – novel) as a function of experience with the novel environment. Stars indicate significance level of unpaired t-tests comparing values on the first and fifth days of experience. * = $p < 0.05$; ns = $p > 0.05$.

4.3.2 Gridness is determined both by novelty and the distal sensory features of the environment

In mice screened in Room 412, grid cells showed both increases and decreases in gridness in novel environments in Room 413 (Figure 4.3A). As such, no significant difference was found in gridness values between baseline and novel environments (*paired t-test*, $t_{73} = 0.343$, $p = 0.733$, Figure 4.3C). However, a consistent pattern in the change in gridness values could be seen as the novel environment became more familiar (Figure 4.3D, left-hand panel). On the first novel day, gridness values were significantly lower in the novel environment (*paired t-test*, $t_{31} = 4.55$, $p = 7.70 \times 10^{-5}$, Figure 4.3D). However, following several days of experience, gridness values were on average higher in the novel environment, with significantly higher gridness observed on the fourth recording day (*paired t-test*, $t_{28} = -3.06$, $p = 0.005$, Figure 4.3D). Clear modulation of gridness by environmental familiarity could be seen in a significant correlation between the change in gridness across rooms and the recording day, and a significant difference in the change in gridness between the first and fifth days of experience (*Spearman-rank correlation*: $\rho = -0.35$, $p = 2.28 \times 10^{-5}$. *Unpaired t-test*, $t_{55} = 3.11$, $p = 0.0029$, Figure 4.3D, left hand panel).

In mice screened in Room 413, grid cells showed a reduction in gridness in the novel environment in Room 412 (*paired t-test*, $t_{43} = 8.05$, $p = 4.02 \times 10^{-10}$, Figure 4.3B-C). Indeed, not a single cell with significant gridness in the novel environment (gridness > 0.3) was observed which did not also have significant gridness in the baseline environment. A gradient of change in the difference in gridness values between the environments was observed as the novel environment became familiar, similar to that seen in mice screened in Room 412 (Figure 4.3D, right-hand panel). However, the plot of mean difference in gridness across recording days had a higher intercept for mice screened in Room 413: a significant reduction in gridness was seen on the first novel day (*paired t-test*, $t_{11} = 7.26$, $p = 1.61 \times 10^{-5}$, Figure 4.3D) which was larger than the reduction observed on the first novel day in mice screened in Room 412 (*unpaired t-test*, $t_{42} = -3.88$, $p = 3.62 \times 10^{-4}$, Figure 4.3D). After five days of recording, no difference in gridness was seen between the baseline and novel environments (*paired t-test*, $t_5 = 0.302$, $p = 0.77$, Figure 4.3D). As in mice screened in Room 412, the influence of environment familiarity on grid regularity could be seen in both a significant correlation between recording day and the change in gridness across rooms, and a difference between the first and fifth recording days in the change in gridness across rooms (*Spearman-rank correlation*: $\rho = -0.56$, $p = 2.76 \times 10^{-4}$. *Unpaired t-test*, $t_{16} = 4.04$, $p = 9.56 \times 10^{-4}$, Figure 4.3D, left hand panel).

These results indicate that both the distal sensory features of the environment, and its degree of novelty, modulate the regularity of grid firing patterns. Initial reductions in gridness

were seen in novel environments, regardless of the counterbalancing of which was room was used for screening. However, the magnitude of the initial reduction in gridness, and the gridness values observed once the animals had become familiar with the novel environment, were dependent on the order in which animals experienced the two recording rooms.




Figure 4.3 Gridness is determined both by novelty and the distal sensory features of the environment. A-B, Firing rate maps and spatial autocorrelograms indicating differences in regularity of grid firing patterns for early and late exposures to the novel environment. Values indicate gridness scores observed. **A,** For mice screened in Room 412, during the first day in the novel environment, grid firing tended to be less regular. By the fifth day, the majority of cells were more regular in the novel environment. **B,** For mice screened in Room 413, on the first day in the novel environment, a large reduction in gridness was seen. By the fifth day, cells had similar gridness values in the two environments. **C-D,** Blue/left: mice screened in Room 412. Black/right: mice screened in Room 413. **C,** Gridness in the novel environment as a function of gridness in the familiar environment. Stars indicate significance level of paired *t*-tests comparing baseline and novel gridness values. *** = $p < 0.001$; ns = $p > 0.05$. **D,** Mean \pm SEM difference in gridness (baseline – novel) as a function of experience with the novel environment. Stars indicate significance level of paired *t*-tests comparing difference in gridness values between environments to 0, and unpaired *t*-tests comparing baseline – novel gridness values on the first and fifth days of experience. *** = $p < 0.001$; ** = $p < 0.01$; ns = $p > 0.05$.

4.3.3 *The temporal stability of grid firing patterns is determined by the distal sensory features of the environment*

To assess the temporal stability of grid cell firing patterns, each trial was divided into two halves, with a ratemap generated for each. A spatial cross-correlogram between the ratemaps from each half was produced, with the value at the centre of the correlogram used as a measure of the stability of the firing pattern. For mice screened in Room 412, the temporal stability of grid firing was on average higher during recordings in novel environments in Room 413 (*paired t-test*, $t_{73} = -2.84$, $p = 0.006$, Figure 4.4A, C). As the novel environment became more familiar, a weak tendency for the difference in stability between environments to increase was observed (Figure 4.4D). However, the weak impact of environment familiarity on firing pattern stability was evidenced by the lack of a significant correlation between recording day and the difference in stability across environments, as well as the absence of a significant difference in change in stability between the first and fifth recording days (*Spearman-rank correlation*: $\rho = -0.14$, $p = 0.098$. *Unpaired t-test*, $t_{55} = 1.39$, $p = 0.168$, Figure 4.4D).

In contrast, mice screened in Room 413 showed a consistent reduction in the stability of firing patterns during recordings in the novel environment in Room 412 (*paired t-test*, $t_{43} = 6.41$, $p = 9.30 \times 10^{-8}$, Figure 4.4B-C). A weak tendency for the difference in stability between the recording rooms to decrease was observed as the novel environment became more familiar (Figure 4.4D). While significant correlation between recording day and the change in stability across environments was seen, as in mice screened in Room 412, the change between the first and fifth recording days was not significant (*Spearman-rank correlation*: $\rho = -0.39$, $p = 0.0167$. *Unpaired t-test*, $t_{16} = 1.42$, $p = 0.176$, Figure 4.4D).

That is, the temporal stability of grid firing patterns appears to be clearly determined by the sensory features of the recording environment: depending on which environment was used as the baseline, stability both increased and decreased during recordings in the alternate room. Only relatively little evidence indicating that stability was modulated by the degree of familiarity with the environment was observed.




Figure 4.4 Temporal stability of grid firing is determined by the distal sensory features of the environment. **A-B**, Example firing rate maps from the first and second half of trials, and spatial cross-correlograms based on these ratemap pairs. Values indicate stability observed, as measured by the correlation at the centre of the cross-correlogram. **A**, The firing patterns of grid cells from mice screened in Room 412 were more stable during recordings in Room 413. **B**, Grid cells from mice screened in Room 413 were instead less stable during recordings in Room 412. **C-D**, Blue/left: mice screened in Room 412. Black/right: mice screened in Room 413. **C**, Temporal stability of grid firing in the novel environment as a function of stability in the familiar environment. Stars indicate significance level of paired t-tests comparing baseline and novel grid stability values. *** = $p < 0.001$; ** = $p < 0.01$. **D**, Mean \pm SEM difference in stability (baseline – novel) as a function of experience with the novel environment. Unpaired t-tests compared difference in stability values between the first and fifth days of experience. ns = $p > 0.05$.

4.4 Discussion

4.4.1 *Summary of results and significance*

The results presented here indicate a propensity of the environment in which grid cells are recorded to powerfully modulate the scale, regularity, and temporal stability of their firing patterns. That changes in measures of these properties were seen despite the same shape and size enclosure being used in both environments (and indeed when exactly the same enclosure was used), suggests that the observed changes were likely due to differences in the distal sensory features of the environments. These results are therefore complementary to demonstrations that the geometry and size of the recording enclosure can similarly modulate grid firing properties (T. Stensola et al. 2015; Krupic et al. 2015).

Recording grid cells over five days of increasing familiarity with the novel environments allowed assessment of any effect of novelty on grid firing patterns. As reported previously, the regularity of grid firing was reduced during initial exposures to novel environments, a change that attenuated across days of experience (Barry, Ginzberg, et al. 2012). The effects of novelty and the distal features of the recording environments appeared to act independently on observed gridness scores. In mice moving from Room 412 to Room 413, gridness values were initially reduced in novelty, but after multiple days of experience were significantly greater than those seen in the baseline environment. In contrast, mice moving from Room 413 to Room 412 showed a larger initial reduction in gridness, with no difference between the environments after five days of experience. Evidence for an effect of novelty on the stability and scale of grid firing patterns was relatively weaker, with greater inconsistency in the change in these values seen across recording days, and between mice experiencing the environments in different orders.

4.4.2 *Factors giving rise to differences in grid properties between environments*

As noted, the shape and size of the proximal enclosure cannot account for the observed differences in grid firing patterns seen between environments. As such, it is probable that changes in grid properties stemmed from differences in the sensory features of the environment beyond the immediate boundaries of the enclosure. That is, differences in the availability of visual, auditory, and possibly even olfactory cues surrounding each enclosure may mean that animals' capacity to self-locate using external sensory cues differs in each, in turn leading to differences in the regularity, stability, and scale of grid firing. Room 412 was smaller than Room 413 (3.2x3m and 5.6x3.2m respectively), with half of the enclosure surrounded by black curtains. While the other side of the enclosure was open to the

experimental room, light levels were relatively low. In contrast, both enclosures in Room 413 were open to the room on all sides, with relatively higher brightness levels. This may have meant that more distal sensory cues, particularly visual ones, were available in Room 413. Indeed, recent experiments indicate a profound stabilising effect of visual input on grid firing, in that grid firing patterns in mice were seen to deteriorate in the total absence of visual cues (Chen et al. 2016; Pérez-Escobar et al. 2016). While the present results indicate that distal sensory features of the environment are an important determinant of grid firing properties, more carefully controlled experiments are necessary to confirm and specify the exact influence of distal cues. For example, grid cells could be recorded in curtained environments in which the availability of distal cues is quasi-parametrically varied. Such an experiment could also be performed in a virtual reality environment, such that all sensory cues could be tightly controlled and specified.

4.4.3 Reconciling results concerning the impact of novelty on grid firing

The present results replicates observed reductions in the regularity of grid patterns in novel environments (Barry, Ginzberg, et al. 2012). However, only weaker evidence was found for an influence of novelty on the stability and scale of grid firing. It is not clear why an increase in grid scale in novel environments was not observed here, despite replications by multiple labs previously (Barry, Ginzberg, et al. 2012; T. Stensola et al. 2015). One clear difference is the animal model used: the present data were collected from mice, whereas previous demonstrations of novelty-induced expansion have been based on experiments in rats. It could simply be that different neural mechanisms exist between the two species. Despite our hope that rats and mice are an effective model organism for humans, striking differences have been observed even between more closely related species. For example, bats do not show clear theta oscillations when locomoting (Yartsev, Witter, and Ulanovsky 2011). Further, normal activity is maintained by grid cells recorded in rats in total darkness (Hafting et al. 2005), in contrast to the deterioration of grid patterns observed in mice under similar conditions (Chen et al. 2016; Pérez-Escobar et al. 2016). Alternatively, it could be that grid cells in rats and mice do share common neural mechanisms, and that the lack of expansion seen here is a result of an interaction between animals' past experience and the size of the recording enclosure. For example, grid expansion in rats could be evidence of a Bayesian prior-estimate of the size of the enclosure, with the prior adjusted through learning as the animal explores its new surroundings, resulting in subsequent contraction of grid patterns. In a natural setting, rats and mice experience similar size environments. However, in the lab, mice are held in much smaller cages than rats. Lab-reared rats and mice may therefore have distinct learned prior-estimates of the size of enclosures, such that lab rats expect a novel

context to be larger than mice. This theory makes several testable predictions. For example, that grid scale is modulated by the size of the enclosure, and that the initial grid scale seen in novel contexts should be modifiable by exposing an animal to many novel enclosures of a particular size, such that its prior-estimate changes accordingly. A more parsimonious explanation for the lack of grid expansion seen here is that the effect of novelty is simply masked by the relatively stronger and less variable effects of distal sensory cues on grid scale.

4.4.4 Significance of results for theories of grid-expansion in novelty

A number of accounts of the utility of an increase in grid scale in novel contexts have been proposed, including to give rise to place cell remapping, to act as a novelty/familiarity signal to downstream areas, and as a mechanism to minimise navigational errors in the face of increased uncertainty (Towse et al. 2014; Barry, Ginzberg, et al. 2012). If it were the case that grid scale genuinely did not change in novel contexts in mice, it would argue against these proposed explanations of grid expansion. Place cells remap in novel environments in mice (McHugh et al. 1996), suggesting changes in grid scale are not a prerequisite; while place cells have been observed to remap normally during inactivation of the medial septum, which also abolishes periodic grid cell activity (Brandon et al. 2014). Indeed, areas downstream of the entorhinal cortex presumably have a similar need for novelty/familiarity signals in rats and mice, suggesting that the function of grid expansion in rats is not to provide such a signal. Further, both rats and mice likely have a comparable reduction in ability to reset accrued error in path-integration systems in novel contexts, due to unfamiliarity with their sensory features. A lack of grid expansion in mice would therefore suggest that, in rats, changes in grid scale are not in compensation for increased uncertainty. Noticeably however, in the data presented here, the regularity and stability of grid firing varied inversely with the scale of the grid pattern: increases in stability and regularity coincided with reductions in grid scale between environments. This inverse relationship has been previously reported in rats exploring novel environments (Barry, Ginzberg, et al. 2012), and is consistent with theories suggesting that changes in grid scale may relate to the amount of uncertainty-induced noise in grid firing (Towse et al. 2014). That is, differences in grid properties between environments seen here may relate to the degree of certainty with which mice could self-locate according to distal sensory cues. However, such an explanation suggests that grid firing properties are modulated by uncertainty, which presumably also increases in novel contexts, in contrast with the relatively little evidence found in support of an effect of environmental novelty on grid scale and stability.

5 Discussion

5.1 Overview

Throughout the 17th, 18th, and 19th centuries, the Watchmaker Analogy was a widely cited theological argument for the existence of a divine creator. It describes how, coming upon an intricate pocket-watch for the first time, in its complexity one can infer the existence of a designer. By analogy therefore, the complexity of biological organisms argued for the existence of God:

“I am like a man who sees the works of a watch for the first time; he is never weary of admiring the mechanism, though he does not know the use of the instrument and has never seen its face. I do not know what this is for, says he, but I see that each part of it is fitted to the rest, I admire the workman in the details of his work, and I am quite certain that all these wheels only work together in this fashion for some common end which I cannot perceive. Let us compare the special ends, the means, the ordered relations of every kind, then let us listen to the inner voice of feeling; what healthy mind can reject its evidence? Unless the eyes are blinded by prejudices, can they fail to see that the visible order of the universe proclaims a supreme intelligence?”.(Rousseau 1762)

By the 20th century, the theory of natural selection and the genetic revolution meant that, asked how biological complexity came to be, a scientist would likely answer with confidence “through evolution”. Biology and neuroscience instead turned to the questions for which Rousseau’s hypothetical man is less sure of the answer: though convinced the watch has some purpose, he cannot fathom what it might be. Such enquiries were the first made of the hippocampal formation: treating it as a black box, early neuroscientific investigations studied the effects on behaviour of damage to the area, in neuropsychological patients and lesioned animals. As the purpose and functions of a system come to be better understood, we begin to peer inside the black box, and ask how it works. As was in turn the case with the hippocampal formation, with electrophysiological recordings in animals and functional imaging in humans in particular employed to probe the mechanisms underlying its functions. The course of understanding an information processing system through these questions was formalised by David Marr (1975), who envisioned three levels of analysis: *computational*: what problem does the system solve? *Algorithmic*: what algorithms does the system employ to solve the problem? And *implementational*: how are the algorithms realised physically? Analysis at the computational level thus asks *what* a system does, while analyses at the algorithmic and implementational levels are concerned with *how* its functions are achieved.

While only loosely connected with one another, the experiments described in this thesis are unified in continuing to ask these ‘what-for’ and ‘how’ questions of the hippocampal formation, focusing in particular on the functions of medial entorhinal grid cells and the

mechanisms by which they are achieved. Tackling these questions in the reverse order, the present experiments specifically aimed to shed light on what factors influence grid cell firing patterns (how do the patterns arise?), and thus what potential functions their unique activity patterns sub serve (what are they for?). In this chapter, I will first review what insights into these questions can be gained from the results presented in this thesis. Second, I will use Marr's levels of analysis as a framework to consider current understanding and remaining uncertainties regarding both the input-output transformation by which grid cells encode spatial information, and the processes by which downstream networks decode grid firing.

5.2 Implications of the results described in this thesis

5.2.1 Implications for the mechanisms by which grid firing patterns arise

5.2.1.1 Importance of environmental cues

In Section 2.3, I described results from an experiment in which grid cells were recorded from rats exploring a multicompartment environment comprising two perceptually identical compartments connected by a corridor. The environment was designed to place self-motion and environmental cues in conflict with one another: if self-motion cues dominated grid firing, their activity patterns should distinguish the compartments, reflecting their differing absolute positions in space. In contrast, if sensory cues determine grid patterns, the perceptual consistence of the two compartments should result in replication of the firing patterns in each. During early exposures, it was unequivocally the case that grid patterns were identical in the two compartments, indicating the dominance of sensory cues. Grid patterns are stable across multiple visits to the same environment (Hafting et al. 2005), indicating that they are anchored to sensory features of the environment. Theoretically, variability both in velocity inputs to grid cells, and their path integration mechanism, results in the accrual of noise in their firing (Fuhs and Touretzky 2006; Burgess, Barry, and O'Keefe 2007). Sensory cues which allow the animal to unambiguously self-locate are one potential mechanism by which accumulated noise could be reset, as appears to happen at environmental boundaries (Hardcastle, Ganguli, and Giocomo 2015). The replication of grid firing during initial exposures to the multicompartment environment may therefore reflect the importance of the anchoring and denoising of grid cell firing patterns by sensory cues. The observation that place cell firing also replicated between the compartments during initial exposures to the multicompartment environment suggests that they could provide the location-specific sensory inputs which led grid cell firing to replicate. Consistent with this, grid patterns deteriorate during inactivation of the hippocampus (Bonnevie et al. 2013).

The multicompartment environment was designed such that proximal sensory features of the environment were identical in the two compartments, while the lighting and curtains used aimed to minimise the presence of all distal sensory cues. As such, replication of grid patterns was implicitly assumed to result from inputs reflecting the consistency of proximal sensory features, especially given contemporaneous results indicating the influence of boundaries over grid firing (Barry et al. 2007; Derdikman et al. 2009; Krupic et al. 2015). In Section 4.3, I described data demonstrating that, in mice, grid cell firing patterns are also influenced by the distal sensory cues that surround the enclosure. Even when the same enclosure was used, the regularity, stability, and scale of grid firing patterns varied between the two rooms in which recordings were made. While further experiments are necessary to infer the precise cause of the observed differences in firing, the two rooms noticeably differed in the availability of distal visual cues. These results may therefore demonstrate that location-specific distal cues are also used to anchor and denoise grid firing. That is, an increased number of useful distal cues would mean that inputs to grid cells would specify the animal's location more consistently and with greater accuracy, reducing the accumulation of noise and thus resulting in more stable and regular firing patterns. Interestingly, the scale of observed firing patterns varied inversely with the stability and regularity of grid firing, consistent with theories suggesting that grid scale is adjusted in proportion to the uncertainty with which an animal can use sensory cues to self-locate; the grid system sacrificing resolution in its code for the animal's location to reduce the chance of egregious errors (Towse et al. 2014).

5.2.1.2 Importance of self-motion inputs

While grid firing patterns initially replicated in the multicompartment environment, their phase of firing came to distinguish the two compartments after prolonged experience. Indeed, a single continuous pattern that spanned both compartments eventually formed, though only after many days of exposure to the environment (see Section 2.3.2). The present results cannot definitively rule out the possibility that animals slowly came to identify subtle sensory features of the environment that could be used to distinguish the two compartments, despite their overwhelming perceptual similarity. However, the observation that grid patterns were more coherently continuous in the thirds of the compartments closest to the corridor instead suggests that animals used self-motion inputs and path integration to distinguish the compartments, according to their absolute positions in space. That is, the reduced distance between the thirds of the compartments closest together may result in the accumulation of less path integration noise, and thus a more coherent single representation. Assuming that grid cells indeed used path integration to distinguish the compartments, the extent to which

other spatial cell types disambiguated the two can be taken to indicate the influence of self-motion inputs over their firing. Regardless of experience of the environment, boundary vector cells and head direction cells continued to display highly similar firing in each compartment, suggesting that their activity is not influenced by translational self-motion inputs, or indeed grid cells. After prolonged exposure to the multicompartment environment, place cells displayed a range of firing patterns, with some cells replicating between the compartments, while the firing of others distinguished the two. As such, place cells may vary on a continuum according to the relative influence of translational self-motion inputs (from grid cells) and sensory inputs (predominantly from boundary vector cells). Therefore, that all grid cells came to distinguish the compartments suggests that they are unique amongst the spatially modulated cell types of the hippocampal formation in the extent to which translational path integration determines their firing.

In sum, the results presented in this thesis indicate that grid patterns are strongly influenced by both proximal and distal location-specific sensory features of the environment, which are likely used to anchor and denoise grid firing. However, the multicompartment experiment also demonstrates that grid cells are unique in the hippocampal formation in the extent to which translational self-motion inputs govern their firing, likely indicating their role in path integration.

5.2.1.3 Absence of evidence for cholinergic influences on grid firing

Section 3.3 described results indicating that increasing the excitability of medial septal cholinergic neurones had no discernible effect on grid cell firing patterns in either familiar or novel environments. The data presented here therefore fail to confirm past results indicating a direct role of cholinergic function in the generation of regular grid firing (Newman et al. 2014). In that experiment, systemic administration of the muscarinic receptor antagonist scopolamine resulted in a reduction of the periodic firing of grid cells on an annular track and in an open-field enclosure. The difference between the two sets of results may indicate that grid cells require only a base level of cholinergic function, with increases of cholinergic tone above this level having little effect on grid firing. It is unlikely to be the case that the cholinergic system in general normally functions at saturating levels, as increasing the excitability of septal cholinergic neurones had marked effects on theta frequency and the animal's behaviour. Alternatively, the studies may differ due to the specificity with which cholinergic function was modulated. In the present experiment, only the excitability of cholinergic neurones in the medial septum was increased. In contrast, systemic administration of scopolamine perturbs the function of all neurones expressing muscarinic receptors, regardless of their location in the brain. As such, while the present results do not

rule out the possibility that grid firing is influenced by cholinergic tone, the data here provide no evidence for this being the case, and suggests that any relationship between grid function and acetylcholine concentration is non-linear or not specific to the septo-hippocampal system.

5.2.2 Implications for the potential functions of grid firing patterns

5.2.2.1 Grid firing forms coherent global representations, even in complex environments

The conjunction of phases across a set of grid modules, the ‘population phase’, can together encode an animal’s location, even if the individual modules are ambiguous due to repetitions in their phases across the environment (Fiete, Burak, and Brookings 2008). Further, the periodic activity of grid cells means that, if their firing patterns are regular, the population phase changes at a constant rate across the environment. Therefore, a specific change in the population phase equates to the animal having moved a specific distance. As such, population phase differences could be used to decode the vector separating pairs of locations in the environment, potentially the basis of goal-directed navigation (Bush et al. 2015). However, distortions and discontinuities, which respectively equate to instantaneous jumps in the population phase, and shifts in its rate of change, by definition mean that the relationship between change in population phase and distance traversed is inconsistent across the environment. Using irregular grid patterns to calculate navigational vectors will likely be prone to errors therefore. A number of recent results indicate that the requirement for regular firing patterns is not met in certain contexts (Derdikman et al. 2009; Krupic et al. 2015), leading to suggestions that grid cells do not provide a universal spatial metric capable of supporting navigation (Krupic et al. 2015). In contrast, the global firing patterns observed in the multicompartiment environment in Section 2.3.2 provides the first evidence that, even in complex environments, grid firing patterns can form the regular and continuous representations required for them to underlie accurate navigation.

Recordings in the multicompartiment environment also demonstrated, however, that grid patterns were discontinuous and distorted for prolonged periods. During this time, spatial inconsistencies in grid firing would likely lead to errors in navigation, assuming that the system which decodes navigational vectors has been trained on regular grid patterns. The nature and magnitude of navigational errors stemming from irregular grid cell activity will depend on how information is encoded across grid modules, as well as how the size of the environment compares to the scale of the constituent grid modules. Given present uncertainty regarding these issues, the behavioural impact of irregular grid patterns is

currently unclear. Potentially however, relating the nature and magnitude of navigational errors to concomitant observations of the irregularities in grid firing could provide novel insights into the mechanisms by which grid cells support navigation. For example, in nested grid codes, larger scale modules resolve ambiguity in smaller ones (Mathis, Herz, and Stemmler 2012; Stemmler, Mathis, and Herz 2015), such that navigational errors should vary in magnitude with the scale of the largest module whose firing is distorted. In contrast, combinatorial grid codes distribute location encoding evenly across modules (Fiete, Burak, and Brookings 2008), such that navigational errors should instead be disproportionately large where distortions or discontinuities are inconsistent across modules.

In sum, the continuous and regular firing patterns eventually formed by grid cells in the multicompartiment environment indicate that grid firing can form the globally coherent representations theoretically required for them to support accurate navigation. However, further experiments are necessary to infer how grid cell population activity is encoded across modules, and thus what impact initial irregularities in grid firing have on animals' navigational abilities. Conversely, experiments relating distortions in grid patterns to navigational errors could be used to shed light on these issues.

5.2.2.2 Novelty-induced grid expansion is unlikely to signal novelty/familiarity, or generate place cell remapping

In rats, increases in grid scale are seen when an animal is first exposed to a novel environment, with grid scale returning to its 'baseline' value after several days of experience (Barry, Ginzberg, et al. 2012). Various accounts of this phenomenon have been proposed, including that it provides a novelty/familiarity signal to downstream regions, and that it augments place cell remapping (Barry, Ginzberg, et al. 2012; Barry, Heys, and Hasselmo 2012). In Sections 3.3.7 and 4.3.1, results demonstrating the absence of novelty-induced increases in grid scale in mice were presented. It is possible that grid expansion does occur in mice during exploration of unfamiliar contexts, and that its relatively smaller and more variable influence on grid firing was simply masked by differences in the availability of distal cues between recording rooms. However, a genuine absence of increases in grid scale in mice exploring novel environments would argue against theories proposing that grid expansion acts as a novelty/familiarity signal, or sub serves place cell remapping. In that, it is not apparent why regions downstream of grid cells would be more in need of a novelty/familiarity signal in rats than mice. Similarly, remapping of place cells in novel environments in mice (McHugh et al. 1996), and during interruption of normal grid firing (Brandon et al. 2014), suggest that changes in grid scale are not a prerequisite.

5.3 Current understanding and remaining uncertainties regarding the encoding and decoding of grid cell activity

5.3.1 *What do grid firing patterns encode*

Marr's computational level of analysis aims to interrogate a system by asking what problem it solves. For grid cells, this level of analysis can thus be considered to ask 'what is the value added by the output they provide'? Several features of the unique firing pattern of grid cells is highly suggestive of their acting as a spatial metric. First, their periodic firing and modular organisation means that a set of modules form a computationally efficient system for encoding continuous variables such as location (Sreenivasan and Fiete 2011). When regular, the periodicity in their firing patterns also results in their encoding, in the phase offsets across modules, the directions and distances separating pairs of locations. Further, that pairs of grid cells in the same module maintain constant phase offsets across environments potentially means that the wiring patterns and weights required for decoding grid firing may not need be relearned (Yoon et al. 2013). These features of grid firing patterns have led to the belief that they act as a spatial metric being widely held (Hafting et al. 2005; Fiete, Burak, and Brookings 2008; Stemmler, Mathis, and Herz 2015). However, direct evidence that grid cells indeed act as a spatial metric is lacking. In large part, this reflects the current absence of an experimental intervention which allows grid cells alone to be targeted. Genetic manipulations which use genomic promoters to exclusively express exogenous proteins only in specific cell-types may be inappropriate, given that grid firing is observed in both ocean-stellate and island-pyramidal cells (Domnisoru, Kinkhabwala, and Tank 2013; Sun et al. 2015). Interventions based on calcium-imaging of activity patterns may thus be more suitable, however current 'read-write' technologies are likely unable to modulate the activity of sufficient numbers of neurones for an effective intervention (Packer et al. 2015).

5.3.2 *How is this information encoded?*

5.3.2.1 *...By individual cells and individual models*

Assuming that grid cells do indeed act as a spatial metric, Marr's algorithmic and implementational levels of analysis respectively consider what input-output function generates grid firing patterns, and how this transformation is implemented biologically. Recent evidence, including in this thesis, demonstrates the influence over grid patterns of both self-motion inputs conveying the animal's movement velocity, and sensory inputs conveying information about perceived features of the environment. The periodicity of grid

firing is suggestive of their integrating self-motion inputs such that their representation of self-location is updated by velocity signals as the animal moves. Pure path integration in this way would be prone to the accumulation of excessive noise, and sensory inputs are therefore likely useful in denoising grid firing, and in providing consistency in firing phases across multiple visits to an environment. Computational modelling indicates that the transformation of self-motion and sensory inputs to grid firing patterns could be implemented in various ways, including through interference between multiple velocity-dependent oscillations in individual cells (O'Keefe and Burgess 2005; Burgess, Barry, and O'Keefe 2007; Hasselmo, Giocomo, and Zilli 2007), or through a separate group of velocity-modulated neurones updating the population-representation of networks of grid cells connected through attractor dynamics (Fuhs and Touretzky 2006; McNaughton et al. 2006; Burak and Fiete 2009). While these biologically plausible models of path integration are consistent with large swathes of experimental data, further experiments are required to confirm the genuine mechanisms by which path integration is realised by grid cells. Evidence in favour of both classes of model has been observed, and hybrid systems incorporating elements of each have thus also been proposed (Navratilova et al. 2012; Bush et al. 2015). Given the mutual compatibility of these models, complete understanding of the neural implementation of path integration likely means moving beyond the tendency of past experiments to simply attempt to disprove one type or the other.

5.3.2.2 ...Across groups of modules

Their computationally tractable firing patterns have led to plausible models of how the activity of individual grid cells and grid modules could arise. However, the periodic firing of a single grid module is ambiguous in environments larger than its scale, as multiple locations in space share the same firing phase. A full account of the input-output transformation of grid cells thus requires understanding of how this ambiguity is resolved. One solution would be to use the conjunction of phases across multiple modules of distinct spatial scale. Two main types of model using this method have been proposed, differing in whether the ratio of scales between successive modules is constant or not. The consistency of this ratio determines the number of unique locations which can be encoded, and the mechanisms by which erroneous conjunctions of phase could be identified (Fiete, Burak, and Brookings 2008; Sreenivasan and Fiete 2011; Stemmler, Mathis, and Herz 2015). Models involving a fixed scale ratio are only unambiguous in environments bigger than the largest scale if the ratio is not an integer. Alternative models envisioning a fixed integer ratio propose that instead of resolving ambiguity in larger scales, each smaller module encodes location with increasing resolution (Mathis, Herz, and Stemmler 2012). Such proposals thus assume

either that the scale of the largest module adapts according to the size of the environment, or that grid firing is inherently ambiguous in large environments, such that sensory cues are required to distinguish between equivocal navigational vectors. Recordings from multiple grid modules suggest that the ratio between adjacent scales is between 1.4 and 1.7, though too few modules were recorded simultaneously to identify whether a single, precise ratio exists between all grid scales (Barry et al. 2007; H. Stensola et al. 2012; Krupic et al. 2015). As has traditionally been the case, these recordings were made with groups of tetrodes organised in a bundle, such that cells recorded concurrently tend to be close together, and are thus predominantly from a single or a small number of modules. High-density silicon probes currently in development have large numbers of electrical contacts organised linearly, potentially allowing recordings to be made simultaneously from the entire dorso-ventral extent of the medial entorhinal cortex, and thus from multiple modules. Together with environments of varying sizes, such recordings may help shed light on the nature of the scale ratio between modules, and thus the exact mechanisms by which the phase of individual grid modules are combined to encode spatial information.

5.3.3 *How are grid patterns decoded?*

Given the absence of causal evidence demonstrating an involvement of grid cells in spatial behaviour, we can only make assumptions about what information is decoded from entorhinal grid cells. That animals are capable of navigating in straight line vectors to unseen goals (Tolman, Ritchie, and Kalish 1946), suggests that they have a neural representations both of their own location and that of the goal, which can be used to decode the distance and direction separating the two. Given that individual locations and vector offsets can theoretically be decoded from grid firing patterns, it is probably safe to assume that the navigational behaviours described depend on neurones downstream of grid cells doing so. Decoding an animal's self-location from the conjunction of phases across a set of grid modules is trivial: a single layer of unsupervised Hebbian learning between grid cells and a layer of output neurones results in the latter demonstrating location-specific firing, which is updated by path-integration in the grid cells as the animal navigates. The firing properties of the output layer of this simple network looks remarkably like those of place cells, which, given the anatomical connectivity between the entorhinal cortex and hippocampus, has led to the belief that place cells decode self-location from grid cell firing being widely held (Bush et al. 2015). In contrast, networks capable of decoding the vector offsets between pairs of locations are more complicated. While models of how grid cell activity could be decoded to sub serve vector navigation have been proposed, most suffer drawbacks. For example, Erdem & Hasselmo's model is slow, requiring multiple potential directions to be iteratively

tested (Erdem and Hasselmo 2012, 2014). Bush and colleagues' 'distance cell' model requires an unviable number of neurones to decode grid firing, while their 'rate-coded vector cell' model depends on complex wiring that only enables decoding in single dimensions, such that band cells are a sufficient input, making grid patterns unnecessary (Bush et al. 2015). While there is a relative dearth of biologically plausible models by which navigational vectors could be decoded from grid firing, recent results suggest that CA1 place cells are capable of doing so: recordings in bats demonstrate that certain cells fire when a specific direction or distance separates the animal from a goal (Sarel et al. 2017).

5.3.4 *Summary*

Their unique firing patterns have been taken to imply much about the spatial processing of grid cells. Namely, their periodic activity suggests grid patterns arise through path integration, and are used for self-location and navigation. That is, at Marr's computational level it is widely held that grid cells act as a spatial metric, though there is little direct evidence to support this belief. Numerous plausible models demonstrating how the firing patterns of individual cells or modules could be derived from self-motion inputs have been proposed. While various algorithmic-level descriptions of the input-output transformation of grid cells therefore exist, further experimentation is required to infer the specifics of the neural implementation. Similarly, it remains unclear how the ambiguity inherent in the periodic firing of individual cells or modules is resolved. Presumably, this depends on combining the phases of a set of modules of distinct spatial scale, though the scheme employed remains to be determined. The decoding of self-location from a population of grid cells is conceptually easy, and is likely performed by place cells given their location-specific firing. Decoding vectors separating current and goal locations from grid firing is computationally much harder, and while recent evidence indicates that place cells appear capable of doing so, the underlying neural mechanism remains opaque. The current understanding of the encoding and decoding of grid firing described here is based on theories that assume grid patterns to be idealised in their regularity. As discussed, grid firing is instead often deformed by sensory cues. The significance of such irregularities in grid patterns is currently unclear: the implications depend on how spatial information is encoded and decoded across modules, which, as discussed, is one of the least well understood aspects of grid processing. As noted however, investigation of the relationship between navigational errors and grid disruptions could be one method by which current uncertainties regarding grid cell processing are resolved.

6 References

- Aigner, T. G., and M. Mishkin.** 1986. 'The Effects of Physostigmine and Scopolamine on Recognition Memory in Monkeys'. *Behavioral and Neural Biology* 45 (1): 81–87.
- Alexander, G. M., S. C. Rogan, A. I. Abbas, B. N. Armbruster, Y. Pei, J. A. Allen, R. J. Nonneman, et al.** 2009. 'Remote Control of Neuronal Activity in Transgenic Mice Expressing Evolved G Protein-Coupled Receptors'. *Neuron* 63 (1): 27–39.
- Aloisi, A. M., F. Casamenti, C. Scali, G. Pepeu, and G. Carli.** 1997. 'Effects of Novelty, Pain and Stress on Hippocampal Extracellular Acetylcholine Levels in Male Rats'. *Brain Research* 748 (1–2): 219–26.
- Alonso, A., and C. Köhler.** 1984. 'A Study of the Reciprocal Connections between the Septum and the Entorhinal Area Using Anterograde and Retrograde Axonal Transport Methods in the Rat Brain'. *The Journal of Comparative Neurology* 225 (3): 327–43.
- Amaral, D. G., and J. A. Dent.** 1981. 'Development of the Mossy Fibers of the Dentate Gyrus: I. A Light and Electron Microscopic Study of the Mossy Fibers and Their Expansions'. *The Journal of Comparative Neurology* 195 (1): 51–86.
- Amaral, D. G., C. Dolorfo, and P. Alvarez-Royo.** 1991. 'Organization of CA1 Projections to the Subiculum: A PHA-L Analysis in the Rat'. *Hippocampus* 1 (4): 415–35.
- Amaral, D. G., and M. P. Witter.** 1989. 'The Three-Dimensional Organization of the Hippocampal Formation: A Review of Anatomical Data'. *Neuroscience* 31 (3): 571–91.
- Amilhon, B., C. Y. L. Huh, F. Manseau, G. Ducharme, H. Nichol, A. Adamantidis, and S. Williams.** 2015. 'Parvalbumin Interneurons of Hippocampus Tune Population Activity at Theta Frequency'. *Neuron* 86 (5): 1277–89.
- Andersen, P., R. Morris, D. Amaral, T. Bliss, and J. O'Keefe, eds.** 2006. *The Hippocampus Book*. Oxford Neuroscience Series. Oxford, New York: Oxford University Press.
- Anderson, M. I., and K. J. Jeffery.** 2003. 'Heterogeneous Modulation of Place Cell Firing by Changes in Context'. *The Journal of Neuroscience* 23 (26): 8827–35.
- Angelaki, D. E., and J. D. Dickman.** 2000. 'Spatiotemporal Processing of Linear Acceleration: Primary Afferent and Central Vestibular Neuron Responses'. *Journal of Neurophysiology* 84 (4): 2113–32.
- Aronov, D., R. Nevers, and D. W. Tank.** 2017. 'Mapping of a Non-Spatial Dimension by the Hippocampal–entorhinal Circuit'. *Nature* 543 (7647): 719–22.
- Barry, C., and N. Burgess.** 2007. 'Learning in a Geometric Model of Place Cell Firing'. *Hippocampus* 17 (9): 786–800.
- Barry, C., D. Bush, J. O'Keefe, and N. Burgess.** 2012. 'Models of Grid Cells and Theta Oscillations'. *Nature* 488 (7409): E1–E1.
- Barry, C., L. L. Ginzberg, J. O'Keefe, and N. Burgess.** 2012. 'Grid Cell Firing Patterns Signal Environmental Novelty by Expansion'. *Proceedings of the National Academy of Sciences of the United States of America* 109 (43): 17687–92.
- Barry, C., R. Hayman, N. Burgess, and K. J. Jeffery.** 2007. 'Experience-Dependent Rescaling of Entorhinal Grids'. *Nature Neuroscience* 10 (6): 682–84.

- Barry, C., J. G. Heys, and M. E. Hasselmo.** 2012. 'Possible Role of Acetylcholine in Regulating Spatial Novelty Effects on Theta Rhythm and Grid Cells'. *Frontiers in Neural Circuits* 6: 5.
- Bechterew, W.** 1900. 'Demonstration Eines Gehirns Mit Zerstörung Der Vorderen Und Inneren Teile Der Hirnrinde Beider Schläfenlappen'. *Neurologisches Zentralblatt*, no. 19: 990–91.
- Beed, P., A. Gundlfinger, S. Schneiderbauer, J. Song, C. Böhm, A. Burgalossi, M. Brecht, I. Vida, and D. Schmitz.** 2013. 'Inhibitory Gradient along the Dorsoventral Axis in the Medial Entorhinal Cortex'. *Neuron* 79 (6): 1197–1207.
- Bingman, V. P.** 1992. 'The Importance of Comparative Studies and Ecological Validity for Understanding Hippocampal Structure and Cognitive Function'. *Hippocampus* 2 (3): 213–19.
- Blair, H. T., J. Cho, and P. E. Sharp.** 1998. 'Role of the Lateral Mammillary Nucleus in the Rat Head Direction Circuit: A Combined Single Unit Recording and Lesion Study'. *Neuron* 21 (6): 1387–97.
- . 1999. 'The Anterior Thalamic Head-Direction Signal Is Abolished by Bilateral but Not Unilateral Lesions of the Lateral Mammillary Nucleus'. *The Journal of Neuroscience* 19 (15): 6673–83.
- Blair, H. T., K. Gupta, and K. Zhang.** 2008. 'Conversion of a Phase- to a Rate-Coded Position Signal by a Three-Stage Model of Theta Cells, Grid Cells, and Place Cells'. *Hippocampus* 18 (12): 1239–55.
- Blair, H. T., and P. E. Sharp.** 1995. 'Anticipatory Head Direction Signals in Anterior Thalamus: Evidence for a Thalamocortical Circuit That Integrates Angular Head Motion to Compute Head Direction'. *The Journal of Neuroscience* 15 (9): 6260–70.
- Bland, B. H., M. G. Seto, B. R. Sinclair, and S. M. Fraser.** 1984. 'The Pharmacology of Hippocampal Theta Cells: Evidence That the Sensory Processing Correlate Is Cholinergic'. *Brain Research* 299 (1): 121–31.
- Boccaro, C. N., L. J. Kjonigsen, I. M. Hammer, J. G. Bjaalie, T. B. Leergaard, and M. P. Witter.** 2015. 'A Three-Plane Architectonic Atlas of the Rat Hippocampal Region'. *Hippocampus* 25 (7): 838–57.
- Boccaro, C. N., F. Sargolini, V. H. Thoresen, T. Solstad, M. P. Witter, E. I. Moser, and M-B. Moser.** 2010. 'Grid Cells in Pre- and Parasubiculum'. *Nature Neuroscience* 13 (8): 987–94.
- Bonnevie, T., B. Dunn, M. Fyhn, T. Hafting, D. Derdikman, J. L. Kubie, Y. Roudi, E. I. Moser, and M-B. Moser.** 2013. 'Grid Cells Require Excitatory Drive from the Hippocampus'. *Nature Neuroscience* 16 (3): 309–17.
- Borhegyi, Z., V. Varga, N. Szilágyi, D. Fabo, and T. F. Freund.** 2004. 'Phase Segregation of Medial Septal GABAergic Neurons during Hippocampal Theta Activity'. *The Journal of Neuroscience* 24 (39): 8470–79.
- Bostock, E., R. U. Muller, and J. L. Kubie.** 1991. 'Experience-Dependent Modifications of Hippocampal Place Cell Firing'. *Hippocampus* 1 (2): 193–205.
- Brandon, M. P., A. R. Bogaard, C. P. Libby, M. A. Connerney, K. Gupta, and M. E. Hasselmo.** 2011. 'Reduction of Theta Rhythm Dissociates Grid Cell Spatial Periodicity from Directional Tuning'. *Science* 332 (6029): 595–99.
- Brandon, M. P., J. Koenig, J. K. Leutgeb, and S. Leutgeb.** 2014. 'New and Distinct Hippocampal Place Codes Are Generated in a New Environment during Septal Inactivation'. *Neuron* 82 (4): 789–96.

- Brodmann, von K.** 1909. *Vergleichende Lokalisationslehre Der Großhirnrinde*. Leipzig: Barth.
- Brown, S., and E. A. Schafer.** 1888. 'An Investigation into the Functions of the Occipital and Temporal Lobes of the Monkey's Brain'. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 179 (January): 303–27.
- Brun, V. H., S. Leutgeb, H-Q. Wu, R. Schwarcz, M. P. Witter, E. I. Moser, and M-B. Moser.** 2008. 'Impaired Spatial Representation in CA1 after Lesion of Direct Input from Entorhinal Cortex'. *Neuron* 57 (2): 290–302.
- Brun, V. H., T. Solstad, K. B. Kjelstrup, M. Fyhn, M. P. Witter, E. I. Moser, and M-B. Moser.** 2008. 'Progressive Increase in Grid Scale from Dorsal to Ventral Medial Entorhinal Cortex'. *Hippocampus* 18 (12): 1200–1212.
- Buetfering, C., K. Allen, and H. Monyer.** 2014. 'Parvalbumin Interneurons Provide Grid Cell-Driven Recurrent Inhibition in the Medial Entorhinal Cortex'. *Nature Neuroscience* 17 (5): 710–18.
- Burak, Y., and I. R. Fiete.** 2009. 'Accurate Path Integration in Continuous Attractor Network Models of Grid Cells'. *PLoS Computational Biology* 5 (2): e1000291.
- Burgess, N., C. Barry, and J. O'Keefe.** 2007. 'An Oscillatory Interference Model of Grid Cell Firing'. *Hippocampus* 17 (9): 801–12.
- Burgess, N., E. A. Maguire, and J. O'Keefe.** 2002. 'The Human Hippocampus and Spatial and Episodic Memory'. *Neuron* 35 (4): 625–41.
- Burgess, N., and J. O'Keefe.** 1996. 'Neuronal Computations Underlying the Firing of Place Cells and Their Role in Navigation'. *Hippocampus* 6 (6): 749–62.
- Burwell, R. D., and D. G. Amaral.** 1998. 'Cortical Afferents of the Perirhinal, Postrhinal, and Entorhinal Cortices of the Rat'. *The Journal of Comparative Neurology* 398 (2): 179–205.
- Bush, D., C. Barry, and N. Burgess.** 2014. 'What Do Grid Cells Contribute to Place Cell Firing?' *Trends in Neurosciences* 37 (3): 136–45.
- Bush, D., C. Barry, D. Manson, and N. Burgess.** 2015. 'Using Grid Cells for Navigation'. *Neuron* 87 (3): 507–20.
- Bush, D., and N. Burgess.** 2014. 'A Hybrid Oscillatory Interference/Continuous Attractor Network Model of Grid Cell Firing'. *The Journal of Neuroscience* 34 (14): 5065–79.
- Buzsáki, G.** 2002. 'Theta Oscillations in the Hippocampus'. *Neuron* 33 (3): 325–40.
- Buzsáki, G., C. A. Anastassiou, and C. Koch.** 2012. 'The Origin of Extracellular Fields and Currents--EEG, ECoG, LFP and Spikes'. *Nature Reviews. Neuroscience* 13 (6): 407–20.
- Byrne, P., S. Becker, and N. Burgess.** 2007. 'Remembering the Past and Imagining the Future: A Neural Model of Spatial Memory and Imagery'. *Psychological Review* 114 (2): 340–75.
- Caballero-Bleda, M., and M. P. Witter.** 1993. 'Regional and Laminar Organization of Projections from the Presubiculum and Parasubiculum to the Entorhinal Cortex: An Anterograde Tracing Study in the Rat'. *The Journal of Comparative Neurology* 328 (1): 115–29.
- Cacucci, F., C. Lever, T. J. Wills, N. Burgess, and J. O'Keefe.** 2004. 'Theta-Modulated Place-by-Direction Cells in the Hippocampal Formation in the Rat'. *The Journal of Neuroscience* 24 (38): 8265–77.

- Canto, C. B., N. Koganezawa, P. Beed, E. I. Moser, and M. P. Witter.** 2012. 'All Layers of Medial Entorhinal Cortex Receive Presubicular and Parasubicular Inputs'. *The Journal of Neuroscience* 32 (49): 17620–31.
- Canto, C. B., F. G. Wouterlood, and M. P. Witter.** 2008. 'What Does the Anatomical Organization of the Entorhinal Cortex Tell Us?' *Neural Plasticity* 2008 (August): e381243.
- Carpenter, F., and C. Barry.** 2016. 'Distorted Grids as a Spatial Label and Metric'. *Trends in Cognitive Sciences* 20 (3): 164–67.
- Carpenter, F., D. Manson, K. J. Jeffery, N. Burgess, and C. Barry.** 2015. 'Grid Cells Form a Global Representation of Connected Environments'. *Current Biology* 25 (9): 1176–82.
- Chen, G., J. A. King, N. Burgess, and J. O'Keefe.** 2013. 'How Vision and Movement Combine in the Hippocampal Place Code'. *Proceedings of the National Academy of Sciences of the United States of America* 110 (1): 378–83.
- Chen, G., D. Manson, F. Cacucci, and T. J. Wills.** 2016. 'Absence of Visual Input Results in the Disruption of Grid Cell Firing in the Mouse'. *Current Biology* 26 (17): 2335–42.
- Chevalleyre, V., and S. A. Siegelbaum.** 2010. 'Strong CA2 Pyramidal Neuron Synapses Define a Powerful Disynaptic Cortico-Hippocampal Loop'. *Neuron* 66 (4): 560–72.
- Chicurel, M. E., and K. M. Harris.** 1992. 'Three-Dimensional Analysis of the Structure and Composition of CA3 Branched Dendritic Spines and Their Synaptic Relationships with Mossy Fiber Boutons in the Rat Hippocampus'. *The Journal of Comparative Neurology* 325 (2): 169–82.
- Constantinescu, A. O., J. X. O'Reilly, and T. E. J. Behrens.** 2016. 'Organizing Conceptual Knowledge in Humans with a Gridlike Code'. *Science* 352 (6292): 1464–68.
- Couey, J. J., A. Witoelar, S.-J. Zhang, K. Zheng, J. Ye, B. Dunn, R. Czajkowski, et al.** 2013. 'Recurrent Inhibitory Circuitry as a Mechanism for Grid Formation'. *Nature Neuroscience* 16 (3): 318–24.
- Dannenberg, H., J. R. Hinman, and M. E. Hasselmo.** 2016. 'Potential Roles of Cholinergic Modulation in the Neural Coding of Location and Movement Speed'. *Journal of Physiology, Paris* 110 (1–2): 52–64.
- Davies, P., and A. J. Maloney.** 1976. 'Selective Loss of Central Cholinergic Neurons in Alzheimer's Disease'. *Lancet* 2 (8000): 1403.
- Deiana, S., B. Platt, and G. Riedel.** 2011. 'The Cholinergic System and Spatial Learning'. *Behavioural Brain Research* 221 (2): 389–411.
- Deller, T., A. Martinez, R. Nitsch, and M. Frotscher.** 1996. 'A Novel Entorhinal Projection to the Rat Dentate Gyrus: Direct Innervation of Proximal Dendrites and Cell Bodies of Granule Cells and GABAergic Neurons'. *The Journal of Neuroscience* 16 (10): 3322–33.
- Derdikman, D., J. R. Whitlock, A. Tsao, M. Fyhn, T. Hafting, M.-B. Moser, and E. I. Moser.** 2009. 'Fragmentation of Grid Cell Maps in a Multicompartment Environment'. *Nature Neuroscience* 12 (10): 1325–32.
- Dodson, P. D., H. Pastoll, and M. F. Nolan.** 2011. 'Dorsal-Ventral Organization of Theta-like Activity Intrinsic to Entorhinal Stellate Neurons Is Mediated by Differences in Stochastic Current Fluctuations'. *The Journal of Physiology* 589 (Pt 12): 2993–3008.
- Doeller, C. F., C. Barry, and N. Burgess.** 2010. 'Evidence for Grid Cells in a Human Memory Network'. *Nature* 463 (7281): 657–61.

- Dolorfo, C. L., and D. G. Amaral.** 1998. 'Entorhinal Cortex of the Rat: Topographic Organization of the Cells of Origin of the Perforant Path Projection to the Dentate Gyrus'. *The Journal of Comparative Neurology* 398 (1): 25–48.
- Domnisoru, C., A. A. Kinkhabwala, and D. W. Tank.** 2013. 'Membrane Potential Dynamics of Grid Cells'. *Nature* 495 (7440): 199–204.
- Dupret, D., J. O'Neill, and J. Csicsvari.** 2013. 'Dynamic Reconfiguration of Hippocampal Interneuron Circuits during Spatial Learning'. *Neuron* 78 (1): 166–80.
- Edvardsen, V.** 2016. 'A Passive Mechanism for Goal-Directed Navigation Using Grid Cells'. In , 13:191–98.
- Ekstrom, A. D., M. J. Kahana, J. B. Caplan, T. A. Fields, E. A. Isham, E. L. Newman, and I. Fried.** 2003. 'Cellular Networks Underlying Human Spatial Navigation'. *Nature* 425 (6954): 184–88.
- Engelhardt, E.** 2016. 'Hippocampus Discovery: First Steps'. *Dementia Neuropsychologia* 10 (1): 58–62.
- Erdem, U. M., and M. E. Hasselmo.** 2012. 'A Goal-Directed Spatial Navigation Model Using Forward Trajectory Planning Based on Grid Cells'. *The European Journal of Neuroscience* 35 (6): 916–31.
- . 2014. 'A Biologically Inspired Hierarchical Goal Directed Navigation Model'. *Journal of Physiology, Paris* 108 (1): 28–37.
- Fenton, A. A., H-Y. Kao, S. A. Neymotin, A. Olypher, Y. Vayntrub, W. W. Lytton, and N. Ludvig.** 2008. 'Unmasking the CA1 Ensemble Place Code by Exposures to Small and Large Environments: More Place Cells and Multiple, Irregularly Arranged, and Expanded Place Fields in the Larger Space'. *The Journal of Neuroscience* 28 (44): 11250–62.
- Fiete, I. R.** 2010. 'Losing Phase'. *Neuron* 66 (3): 331–34.
- Fiete, I. R., Y. Burak, and T. Brookings.** 2008. 'What Grid Cells Convey about Rat Location'. *The Journal of Neuroscience* 28 (27): 6858–71.
- Finch, D. M., and T. L. Babb.** 1981. 'Demonstration of Caudally Directed Hippocampal Efferents in the Rat by Intracellular Injection of Horseradish Peroxidase'. *Brain Research* 214 (2): 405–10.
- Freund, T. F., and M. Antal.** 1988. 'GABA-Containing Neurons in the Septum Control Inhibitory Interneurons in the Hippocampus'. *Nature* 336 (6195): 170–73.
- Freund, T. F., and G. Buzsáki.** 1996. 'Interneurons of the Hippocampus'. *Hippocampus* 6 (4): 347–470.
- Fuchs, E. C., A. Neitz, R. Pinna, S. Melzer, A. Caputi, and H. Monyer.** 2016. 'Local and Distant Input Controlling Excitation in Layer II of the Medial Entorhinal Cortex'. *Neuron* 89 (1): 194–208.
- Fuhrmann, F., D. Justus, L. Sosulina, H. Kaneko, T. Beutel, D. Friedrichs, S. Schoch, M. K. Schwarz, M. Fuhrmann, and S. Remy.** 2015. 'Locomotion, Theta Oscillations, and the Speed-Related Firing of Hippocampal Neurons Are Controlled by a Medial Septal Glutamatergic Circuit'. *Neuron* 86 (5): 1253–64.
- Fuhs, M. C., and D. S. Touretzky.** 2006. 'A Spin Glass Model of Path Integration in Rat Medial Entorhinal Cortex'. *The Journal of Neuroscience* 26 (16): 4266–76.
- Furtak, S. C., S-M. Wei, K. L. Agster, and R. D. Burwell.** 2007. 'Functional Neuroanatomy of the Parahippocampal Region in the Rat: The Perirhinal and Postrhinal Cortices'. *Hippocampus* 17 (9): 709–22.

- Fyhn, M., T. Hafting, A. Treves, M-B. Moser, and E. I. Moser.** 2007. 'Hippocampal Remapping and Grid Realignment in Entorhinal Cortex'. *Nature* 446 (7132): 190–94.
- Fyhn, M., T. Hafting, M. P. Witter, E. I. Moser, and M-B. Moser.** 2008. 'Grid Cells in Mice'. *Hippocampus* 18 (12): 1230–38.
- Fyhn, M., S. Molden, M. P. Witter, E. I. Moser, and M-B. Moser.** 2004. 'Spatial Representation in the Entorhinal Cortex'. *Science* 305 (5688): 1258–64.
- Giocomo, L. M., and M. E. Hasselmo.** 2005. 'Nicotinic Modulation of Glutamatergic Synaptic Transmission in Region CA3 of the Hippocampus'. *The European Journal of Neuroscience* 22 (6): 1349–56.
- . 2008. 'Computation by Oscillations: Implications of Experimental Data for Theoretical Models of Grid Cells'. *Hippocampus* 18 (12): 1186–99.
- . 2009. 'Knock-out of HCN1 Subunit Flattens Dorsal-Ventral Frequency Gradient of Medial Entorhinal Neurons in Adult Mice'. *The Journal of Neuroscience* 29 (23): 7625–30.
- Giocomo, L. M., S. A. Hussaini, F. Zheng, E. R. Kandel, M-B. Moser, and E. I. Moser.** 2011. 'Grid Cells Use HCN1 Channels for Spatial Scaling'. *Cell* 147 (5): 1159–70.
- Giocomo, L. M., E. A. Zilli, E. Fransén, and M. E. Hasselmo.** 2007. 'Temporal Frequency of Subthreshold Oscillations Scales with Entorhinal Grid Cell Field Spacing'. *Science* 315 (5819): 1719–22.
- Goodridge, J. P., and J. S. Taube.** 1997. 'Interaction between the Postsubiculum and Anterior Thalamus in the Generation of Head Direction Cell Activity'. *The Journal of Neuroscience* 17 (23): 9315–30.
- Gothard, K. M., W. E. Skaggs, and B. L. McNaughton.** 1996. 'Dynamics of Mismatch Correction in the Hippocampal Ensemble Code for Space: Interaction between Path Integration and Environmental Cues'. *The Journal of Neuroscience* 16 (24): 8027–40.
- Gray, J. A., and G. G. Ball.** 1970. 'Frequency-Specific Relation between Hippocampal Theta Rhythm, Behavior, and Amobarbital Action'. *Science* 168 (3936): 1246–48.
- Green, J. D., and A. A. Arduini.** 1954. 'Hippocampal Electrical Activity in Arousal'. *Journal of Neurophysiology* 17 (6): 533–57.
- Greene, J. R., and S. Totterdell.** 1997. 'Morphology and Distribution of Electrophysiologically Defined Classes of Pyramidal and Nonpyramidal Neurons in Rat Ventral Subiculum in Vitro'. *The Journal of Comparative Neurology* 380 (3): 395–408.
- Grossen, N. E., and M. J. Kelley.** 1972. 'Species-Specific Behavior and Acquisition of Avoidance Behavior in Rats'. *Journal of Comparative and Physiological Psychology* 81 (2): 307–10.
- Hafting, T., M. Fyhn, T. Bonnevie, M-B. Moser, and E. I. Moser.** 2008. 'Hippocampus-Independent Phase Precession in Entorhinal Grid Cells'. *Nature* 453 (7199): 1248–52.
- Hafting, T., M. Fyhn, S. Molden, M-B. Moser, and E. I. Moser.** 2005. 'Microstructure of a Spatial Map in the Entorhinal Cortex'. *Nature* 436 (7052): 801–6.
- Hagan, J. J., J. D. Salamone, J. Simpson, S. D. Iversen, and R. G. Morris.** 1988. 'Place Navigation in Rats Is Impaired by Lesions of Medial Septum and Diagonal Band but Not Nucleus Basalis Magnocellularis'. *Behavioural Brain Research* 27 (1): 9–20.
- Haglund, L., L. W. Swanson, and C. Köhler.** 1984. 'The Projection of the Supramammillary Nucleus to the Hippocampal Formation: An Immunohistochemical and Anterograde

- Transport Study with the Lectin PHA-L in the Rat'. *The Journal of Comparative Neurology* 229 (2): 171–85.
- Hamam, B. N., M. Sinai, G. Poirier, and C. A. Chapman.** 2007. 'Cholinergic Suppression of Excitatory Synaptic Responses in Layer II of the Medial Entorhinal Cortex'. *Hippocampus* 17 (2): 103–13.
- Hangya, B., Z. Borhegyi, N. Szilágyi, T. F. Freund, and V. Varga.** 2009. 'GABAergic Neurons of the Medial Septum Lead the Hippocampal Network during Theta Activity'. *The Journal of Neuroscience* 29 (25): 8094–8102.
- Hardcastle, K., S. Ganguli, and L. M. Giocomo.** 2015. 'Environmental Boundaries as an Error Correction Mechanism for Grid Cells'. *Neuron* 86 (3): 827–39.
- Hargreaves, E. L., G. Rao, I. Lee, and J. J. Knierim.** 2005. 'Major Dissociation between Medial and Lateral Entorhinal Input to Dorsal Hippocampus'. *Science* 308 (5729): 1792–94.
- Hartley, T., N. Burgess, C. Lever, F. Cacucci, and J. O'Keefe.** 2000. 'Modeling Place Fields in Terms of the Cortical Inputs to the Hippocampus'. *Hippocampus* 10 (4): 369–79.
- Harvey, C. D., F. Collman, D. A. Dombeck, and D. W. Tank.** 2009. 'Intracellular Dynamics of Hippocampal Place Cells during Virtual Navigation'. *Nature* 461 (7266): 941–46.
- Hasselmo, M. E.** 2006. 'The Role of Acetylcholine in Learning and Memory'. *Current Opinion in Neurobiology* 16 (6): 710–15.
- Hasselmo, M. E., C. Bodelón, and B. P. Wyble.** 2002. 'A Proposed Function for Hippocampal Theta Rhythm: Separate Phases of Encoding and Retrieval Enhance Reversal of Prior Learning'. *Neural Computation* 14 (4): 793–817.
- Hasselmo, M. E., and M. P. Brandon.** 2012. 'A Model Combining Oscillations and Attractor Dynamics for Generation of Grid Cell Firing'. *Frontiers in Neural Circuits* 6: 30.
- Hasselmo, M. E., L. M. Giocomo, and E. A. Zilli.** 2007. 'Grid Cell Firing May Arise from Interference of Theta Frequency Membrane Potential Oscillations in Single Neurons'. *Hippocampus* 17 (12): 1252–71.
- Hasselmo, M. E., E. Schnell, and E. Barkai.** 1995. 'Dynamics of Learning and Recall at Excitatory Recurrent Synapses and Cholinergic Modulation in Rat Hippocampal Region CA3'. *The Journal of Neuroscience* 15 (7 Pt 2): 5249–62.
- Hayman, R., and N. Burgess.** 2016. 'Disrupting the Grid Cells' Need for Speed'. *Neuron* 91 (3): 502–3.
- Heys, J. G., L. M. Giocomo, and M. E. Hasselmo.** 2010. 'Cholinergic Modulation of the Resonance Properties of Stellate Cells in Layer II of Medial Entorhinal Cortex'. *Journal of Neurophysiology* 104 (1): 258–70.
- Heys, J. G., K. V. Rangarajan, and D. A. Dombeck.** 2014. 'The Functional Micro-Organization of Grid Cells Revealed by Cellular-Resolution Imaging'. *Neuron* 84 (5): 1079–90.
- Hinman, J. R., M. P. Brandon, J. R. Climer, G. W. Chapman, and M. E. Hasselmo.** 2016. 'Multiple Running Speed Signals in Medial Entorhinal Cortex'. *Neuron* 91 (3): 666–79.
- Hjorth-Simonsen, A.** 1973. 'Some Intrinsic Connections of the Hippocampus in the Rat: An Experimental Analysis'. *The Journal of Comparative Neurology* 147 (2): 145–61.
- Hjorth-Simonsen, A., and S. Laurberg.** 1977. 'Commissural Connections of the Dentate Area in the Rat'. *The Journal of Comparative Neurology* 174 (4): 591–606.

- Huerta, P. T., and J. E. Lisman.** 1995. 'Bidirectional Synaptic Plasticity Induced by a Single Burst during Cholinergic Theta Oscillation in CA1 in Vitro'. *Neuron* 15 (5): 1053–63.
- Huxter, J., N. Burgess, and J. O'Keefe.** 2003. 'Independent Rate and Temporal Coding in Hippocampal Pyramidal Cells'. *Nature* 425 (6960): 828–32.
- Huxter, J., T. J. Senior, K. Allen, and J. Csicsvari.** 2008. 'Theta Phase-specific Codes for Two-Dimensional Position, Trajectory and Heading in the Hippocampus'. *Nature Neuroscience* 11 (5): 587–94.
- Insausti, R.** 1993. 'Comparative Anatomy of the Entorhinal Cortex and Hippocampus in Mammals'. *Hippocampus* 3 Spec No: 19–26.
- Jacob, P-Y., M. Gordillo-Salas, J. Facchini, B. Poucet, E. Save, and F. Sargolini.** 2017. 'Medial Entorhinal Cortex and Medial Septum Contribute to Self-Motion-Based Linear Distance Estimation'. *Brain Structure & Function*, February.
- Jacobs, J., C. T. Weidemann, J. F. Miller, A. Solway, J. F. Burke, X-X. Wei, N. Suthana, et al.** 2013. 'Direct Recordings of Grid-like Neuronal Activity in Human Spatial Navigation'. *Nature Neuroscience* 16 (9): 1188–90.
- Jeewajee, A., C. Barry, V. Douchamps, D. Manson, C. Lever, and N. Burgess.** 2014. 'Theta Phase Precession of Grid and Place Cell Firing in Open Environments'. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* 369 (1635): 20120532.
- Jeewajee, A., C. Barry, J. O'Keefe, and N. Burgess.** 2008. 'Grid Cells and Theta as Oscillatory Interference: Electrophysiological Data from Freely Moving Rats'. *Hippocampus* 18 (12): 1175–85.
- Jeffery, K. J., J. G. Donnett, and J. O'Keefe.** 1995. 'Medial Septal Control of Theta-Related Unit Firing in the Entorhinal Cortex of Awake Rats'. *Neuroreport* 6 (16): 2166–70.
- Jung, M. W., S. I. Wiener, and B. L. McNaughton.** 1994. 'Comparison of Spatial Firing Characteristics of Units in Dorsal and Ventral Hippocampus of the Rat'. *The Journal of Neuroscience* 14 (12): 7347–56.
- Justus, D., D. Dalügge, S. Bothe, F. Fuhrmann, C. Hannes, H. Kaneko, D. Friedrichs, et al.** 2017. 'Glutamatergic Synaptic Integration of Locomotion Speed via Septoentorhinal Projections'. *Nature Neuroscience* 20 (1): 16–19.
- Kadir, S. N., D. F. M. Goodman, and K. D. Harris.** 2014. 'High-Dimensional Cluster Analysis with the Masked EM Algorithm'. *Neural Computation* 26 (11): 2379–94.
- Kaplan, M. S., and J. W. Hinds.** 1977. 'Neurogenesis in the Adult Rat: Electron Microscopic Analysis of Light Radioautographs'. *Science* 197 (4308): 1092–94.
- Kimura, H., P. L. McGeer, F. Peng, and E. G. McGeer.** 1980. 'Choline Acetyltransferase-Containing Neurons in Rodent Brain Demonstrated by Immunohistochemistry'. *Science* 208 (4447): 1057–59.
- King, C., M. Recce, and J. O'Keefe.** 1998. 'The Rhythmicity of Cells of the Medial Septum/Diagonal Band of Broca in the Awake Freely Moving Rat: Relationships with Behaviour and Hippocampal Theta'. *The European Journal of Neuroscience* 10 (2): 464–77.
- Kitamura, T., M. Pignatelli, J. Suh, K. Kohara, A. Yoshiki, K. Abe, and S. Tonegawa.** 2014. 'Island Cells Control Temporal Association Memory'. *Science* 343 (6173): 896–901.

- Kitamura, T., C. Sun, J. Martin, L. J. Kitch, M. J. Schnitzer, and S. Tonegawa.** 2015. 'Entorhinal Cortical Ocean Cells Encode Specific Contexts and Drive Context-Specific Fear Memory'. *Neuron* 87 (6): 1317–31.
- Kjelstrup, K. B., T. Solstad, V. H. Brun, T. Hafting, S. Leutgeb, M. P. Witter, E. I. Moser, and M-B. Moser.** 2008. 'Finite Scale of Spatial Representation in the Hippocampus'. *Science* 321 (5885): 140–43.
- Klink, R., and A. Alonso.** 1997. 'Ionic Mechanisms of Muscarinic Depolarization in Entorhinal Cortex Layer II Neurons'. *Journal of Neurophysiology* 77 (4): 1829–43.
- Koenig, J., A. N. Linder, J. K. Leutgeb, and S. Leutgeb.** 2011. 'The Spatial Periodicity of Grid Cells Is Not Sustained during Reduced Theta Oscillations'. *Science* 332 (6029): 592–95.
- Kohara, K., M. Pignatelli, A. J. Rivest, H-Y. Jung, T. Kitamura, J. Suh, D. Frank, et al.** 2014. 'Cell Type-Specific Genetic and Optogenetic Tools Reveal Hippocampal CA2 Circuits'. *Nature Neuroscience* 17 (2): 269–79.
- Köhler, C.** 1985. 'Intrinsic Projections of the Retrohippocampal Region in the Rat Brain. I. The Subicular Complex'. *The Journal of Comparative Neurology* 236 (4): 504–22.
- . 1986. 'Intrinsic Connections of the Retrohippocampal Region in the Rat Brain. II. The Medial Entorhinal Area'. *The Journal of Comparative Neurology* 246 (2): 149–69.
- Köhler, C., V. Chan-Palay, and J. Y. Wu.** 1984. 'Septal Neurons Containing Glutamic Acid Decarboxylase Immunoreactivity Project to the Hippocampal Region in the Rat Brain'. *Anatomy and Embryology* 169 (1): 41–44.
- Kramis, R., C. H. Vanderwolf, and B. H. Bland.** 1975. 'Two Types of Hippocampal Rhythmical Slow Activity in Both the Rabbit and the Rat: Relations to Behavior and Effects of Atropine, Diethyl Ether, Urethane, and Pentobarbital'. *Experimental Neurology* 49 (1 Pt 1): 58–85.
- Kropff, E., J. E. Carmichael, M-B. Moser, and E. I. Moser.** 2015. 'Speed Cells in the Medial Entorhinal Cortex'. *Nature* 523 (7561): 419–24.
- Krupic, J., M. Bauza, S. Burton, C. Barry, and J. O'Keefe.** 2015. 'Grid Cell Symmetry Is Shaped by Environmental Geometry'. *Nature* 518 (7538): 232–35.
- Krupic, J., N. Burgess, and J. O'Keefe.** 2012. 'Neural Representations of Location Composed of Spatially Periodic Bands'. *Science* 337 (6096): 853–57.
- Langston, R. F., J. A. Ainge, J. J. Couey, C. B. Canto, T. L. Bjerknes, M. P. Witter, E. I. Moser, and M-B. Moser.** 2010. 'Development of the Spatial Representation System in the Rat'. *Science* 328 (5985): 1576–80.
- Lashley, K. S.** 1929. *Brain Mechanisms and Intelligence: A Quantitative Study of Injuries to the Brain*. University of Chicago Press.
- Laurberg, S.** 1979. 'Commissural and Intrinsic Connections of the Rat Hippocampus'. *The Journal of Comparative Neurology* 184 (4): 685–708.
- Lee, M. G., J. J. Chrobak, A. Sik, R. G. Wiley, and G. Buzsáki.** 1994. 'Hippocampal Theta Activity Following Selective Lesion of the Septal Cholinergic System'. *Neuroscience* 62 (4): 1033–47.
- Lengyel, M., Z. Szatmáry, and P. Erdi.** 2003. 'Dynamically Detuned Oscillations Account for the Coupled Rate and Temporal Code of Place Cell Firing'. *Hippocampus* 13 (6): 700–714.
- Leranth, C., and M. Frotscher.** 1989. 'Organization of the Septal Region in the Rat Brain: Cholinergic-GABAergic Interconnections and the Termination of Hippocampo-Septal Fibers'. *The Journal of Comparative Neurology* 289 (2): 304–14.

- Leutgeb, S., J. K. Leutgeb, C. A. Barnes, E. I. Moser, B. L. McNaughton, and M-B. Moser.** 2005. 'Independent Codes for Spatial and Episodic Memory in Hippocampal Neuronal Ensembles'. *Science (New York, N.Y.)* 309 (5734): 619–23.
- Lever, C., S. Burton, A. Jeewajee, J. O'Keefe, and N. Burgess.** 2009. 'Boundary Vector Cells in the Subiculum of the Hippocampal Formation'. *The Journal of Neuroscience* 29 (31): 9771–77.
- Lever, C., T. Wills, F. Cacucci, N. Burgess, and J. O'Keefe.** 2002. 'Long-Term Plasticity in Hippocampal Place-Cell Representation of Environmental Geometry'. *Nature* 416 (6876): 90–94.
- Lübke, J., T. Deller, and M. Frotscher.** 1997. 'Septal Innervation of Mossy Cells in the Hilus of the Rat Dentate Gyrus: An Anterograde Tracing and Intracellular Labeling Study'. *Experimental Brain Research* 114 (3): 423–32.
- Ludvig, N., H. M. Tang, B. C. Gohil, and J. M. Botero.** 2004. 'Detecting Location-Specific Neuronal Firing Rate Increases in the Hippocampus of Freely-Moving Monkeys'. *Brain Research* 1014 (1–2): 97–109.
- Maaswinkel, H., and I. Q. Whishaw.** 1999. 'Homing with Locale, Taxon, and Dead Reckoning Strategies by Foraging Rats: Sensory Hierarchy in Spatial Navigation'. *Behavioural Brain Research* 99 (2): 143–52.
- Maguire, E. A., N. Burgess, J. G. Donnett, R. S. Frackowiak, C. D. Frith, and J. O'Keefe.** 1998. 'Knowing Where and Getting There: A Human Navigation Network'. *Science* 280 (5365): 921–24.
- Manns, I. D., L. Mainville, and B. E. Jones.** 2001. 'Evidence for Glutamate, in Addition to Acetylcholine and GABA, Neurotransmitter Synthesis in Basal Forebrain Neurons Projecting to the Entorhinal Cortex'. *Neuroscience* 107 (2): 249–63.
- Manseau, F., M. Danik, and S. Williams.** 2005. 'A Functional Glutamatergic Neuron Network in the Medial Septum and Diagonal Band Area'. *The Journal of Physiology* 566 (Pt 3): 865–84.
- Markram, H., and M. Segal.** 1990. 'Electrophysiological Characteristics of Cholinergic and Non-Cholinergic Neurons in the Rat Medial Septum-Diagonal Band Complex'. *Brain Research* 513 (1): 171–74.
- Marr, D.** 1971. 'Simple Memory: A Theory for Archicortex'. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* 262 (841): 23–81.
- Marr, D., and T. Poggio.** 1976. 'From Understanding Computation to Understanding Neural Circuitry'. *Science*, May.
- Mathis, A., A. V. M. Herz, and M. Stemmler.** 2012. 'Optimal Population Codes for Space: Grid Cells Outperform Place Cells'. *Neural Computation* 24 (9): 2280–2317.
- McGaughy, J., R. A. Koene, H. Eichenbaum, and M. E. Hasselmo.** 2005. 'Cholinergic Deafferentation of the Entorhinal Cortex in Rats Impairs Encoding of Novel but Not Familiar Stimuli in a Delayed Nonmatch-to-Sample Task'. *The Journal of Neuroscience* 25 (44): 10273–81.
- McHugh, T. J., K. I. Blum, J. Z. Tsien, S. Tonegawa, and M. A. Wilson.** 1996. 'Impaired Hippocampal Representation of Space in CA1-Specific NMDAR1 Knockout Mice'. *Cell* 87 (7): 1339–49.
- McNaughton, B. L., C. A. Barnes, J. L. Gerrard, K. Gothard, M. W. Jung, J. J. Knierim, H. Kudrimoti, et al.** 1996. 'Deciphering the Hippocampal Polyglot: The Hippocampus as a Path Integration System'. *The Journal of Experimental Biology* 199 (Pt 1): 173–85.

- McNaughton, B. L., F. P. Battaglia, O. Jensen, E. I. Moser, and M-B. Moser.** 2006. 'Path Integration and the Neural Basis of the "Cognitive Map"'. *Nature Reviews. Neuroscience* 7 (8): 663–78.
- McNaughton, B. L., and R. G. M. Morris.** 1987. 'Hippocampal Synaptic Enhancement and Information Storage within a Distributed Memory System'. *Trends in Neurosciences* 10 (10): 408–15.
- Mesulam, M. M., E. J. Mufson, B. H. Wainer, and A. I. Levey.** 1983. 'Central Cholinergic Pathways in the Rat: An Overview Based on an Alternative Nomenclature (Ch1-Ch6)'. *Neuroscience* 10 (4): 1185–1201.
- Mhatre, H., A. Gorchetchnikov, and S. Grossberg.** 2012. 'Grid Cell Hexagonal Patterns Formed by Fast Self-Organized Learning within Entorhinal Cortex'. *Hippocampus* 22 (2): 320–34.
- Miao, C., Q. Cao, H. T. Ito, H. Yamahachi, M. P. Witter, M-B. Moser, and E. I. Moser.** 2015. 'Hippocampal Remapping after Partial Inactivation of the Medial Entorhinal Cortex'. *Neuron* 88 (3): 590–603.
- Mizuseki, K., A. Sirota, E. Pastalkova, and G. Buzsáki.** 2009. 'Theta Oscillations Provide Temporal Windows for Local Circuit Computation in the Entorhinal-Hippocampal Loop'. *Neuron* 64 (2): 267–80.
- Monaco, J. D., and L. F. Abbott.** 2011. 'Modular Realignment of Entorhinal Grid Cell Activity as a Basis for Hippocampal Remapping'. *The Journal of Neuroscience* 31 (25): 9414–25.
- Morris, R. G., P. Garrud, J. N. Rawlins, and J. O'Keefe.** 1982. 'Place Navigation Impaired in Rats with Hippocampal Lesions'. *Nature* 297 (5868): 681–83.
- Muller, R. U., and J. L. Kubie.** 1987. 'The Effects of Changes in the Environment on the Spatial Firing of Hippocampal Complex-Spike Cells'. *The Journal of Neuroscience* 7 (7): 1951–68.
- Naber, P. A., F. H. Lopes da Silva, and M. P. Witter.** 2001. 'Reciprocal Connections between the Entorhinal Cortex and Hippocampal Fields CA1 and the Subiculum Are in Register with the Projections from CA1 to the Subiculum'. *Hippocampus* 11 (2): 99–104.
- Navratilova, Z., L. M. Giocomo, J-M. Fellous, M. E. Hasselmo, and B. L. McNaughton.** 2012. 'Phase Precession and Variable Spatial Scaling in a Periodic Attractor Map Model of Medial Entorhinal Grid Cells with Realistic after-Spike Dynamics'. *Hippocampus* 22 (4): 772–89. doi:10.1002/hipo.20939.
- Newman, E. L., J. R. Climer, and M. E. Hasselmo.** 2014. 'Grid Cell Spatial Tuning Reduced Following Systemic Muscarinic Receptor Blockade'. *Hippocampus* 24 (6): 643–55.
- Newman, E. L., S. N. Gillet, J. R. Climer, and M. E. Hasselmo.** 2013. 'Cholinergic Blockade Reduces Theta-Gamma Phase Amplitude Coupling and Speed Modulation of Theta Frequency Consistent with Behavioral Effects on Encoding'. *The Journal of Neuroscience* 33 (50): 19635–46.
- Nó, L. de.** 1933. 'Studies on the Structure of the Cerebral Cortex. I. Areal Entorhinalis.' *Journal of Psychology and Neurology* 45: 381–438.
- . 1934. 'Studies on the Structure of the Cerebral Cortex. II. Continuation of the Study of the Ammonic System.' *Journal of Psychology and Neurology* 46: 113–77.
- O'Keefe, J.** 1976. 'Place Units in the Hippocampus of the Freely Moving Rat'. *Experimental Neurology* 51 (1): 78–109.

- O'Keefe, J., and N. Burgess.** 1996. 'Geometric Determinants of the Place Fields of Hippocampal Neurons'. *Nature* 381 (6581): 425–28.
- . 2005. 'Dual Phase and Rate Coding in Hippocampal Place Cells: Theoretical Significance and Relationship to Entorhinal Grid Cells'. *Hippocampus* 15 (7): 853–66.
- O'Keefe, J., and D. H. Conway.** 1978. 'Hippocampal Place Units in the Freely Moving Rat: Why They Fire Where They Fire'. *Experimental Brain Research* 31 (4): 573–90.
- O'Keefe, J., and J. Dostrovsky.** 1971. 'The Hippocampus as a Spatial Map. Preliminary Evidence from Unit Activity in the Freely-Moving Rat'. *Brain Research* 34 (1): 171–75.
- O'Keefe, J., and L. Nadel.** 1978. *The Hippocampus as a Cognitive Map*. Oxford: Clarendon Press.
- O'Keefe, J., and M. L. Recce.** 1993. 'Phase Relationship between Hippocampal Place Units and the EEG Theta Rhythm'. *Hippocampus* 3 (3): 317–30.
- Ólafsdóttir, H. F., F. Carpenter, and C. Barry.** 2016. 'Coordinated Grid and Place Cell Replay during Rest'. *Nature Neuroscience* 19 (6): 792–94.
- Packer, A. M., L. E. Russell, H. W. P. Dalgleish, and M. Häusser.** 2015. 'Simultaneous All-Optical Manipulation and Recording of Neural Circuit Activity with Cellular Resolution in Vivo'. *Nature Methods* 12 (2): 140–46.
- Panula, P., A. V. Revuelta, D. L. Cheney, J. Y. Wu, and E. Costa.** 1984. 'An Immunohistochemical Study on the Location of GABAergic Neurons in Rat Septum'. *The Journal of Comparative Neurology* 222 (1): 69–80.
- Pastalkova, E., V. Itskov, A. Amarasingham, and G. Buzsáki.** 2008. 'Internally Generated Cell Assembly Sequences in the Rat Hippocampus'. *Science* 321 (5894): 1322–27.
- Paxinos, G., and K. B. J. Franklin.** 2013. *The Mouse Brain in Stereotaxic Coordinates*. Fourth. London: Academic Press.
- Pérez-Escobar, J. A., O. Kornienko, P. Latuske, L. Kohler, and K. Allen.** 2016. 'Visual Landmarks Sharpen Grid Cell Metric and Confer Context Specificity to Neurons of the Medial Entorhinal Cortex'. *eLife* 5 (July).
- Petsche, H., and C. Stumpf.** 1960. 'Topographic and Toposcopic Study of Origin and Spread of the Regular Synchronized Arousal Pattern in the Rabbit'. *Electroencephalography and Clinical Neurophysiology* 12 (August): 589–600.
- Petsche, H., C. Stumpf, and G. Gogolak.** 1962. 'The significance of the rabbit's septum as a relay station between the midbrain and the hippocampus. I. The control of hippocampus arousal activity by the septum cells'. *Electroencephalography and Clinical Neurophysiology* 14 (April): 202–11.
- Prut, L., and C. Belzung.** 2003. 'The Open Field as a Paradigm to Measure the Effects of Drugs on Anxiety-like Behaviors: A Review'. *European Journal of Pharmacology* 463 (1–3): 3–33.
- Quirk, G. J., R. U. Muller, J. L. Kubie, and J. B. Ranck.** 1992. 'The Positional Firing Properties of Medial Entorhinal Neurons: Description and Comparison with Hippocampal Place Cells'. *The Journal of Neuroscience* 12 (5): 1945–63.
- Ramón y Cajal, S.** 1893. 'La Rétine Des Vertébrés'. *Cellule* 9: 121–255.
- . 1902. 'Sobre Un Ganglio Especial de La Corteza Esfeno-Occipital.' *Trab Del Lab de Invest Biol Univ Madrid* 1: 189–206.

- . 1903. 'Un Sencillo Método de Coloración Del Retículo Protoplásmico Y Sus Efectos En Los Diversos Centros Nerviosos de Vertebrados E Invertebrados.' *Trab Del Lab de Invest Biol Univ Madrid* 2: 129–221.
- . 1909. *Histologie Du Système Nerveux de L'homme & Des Vertébrés*. Paris: Maloine.
- Ranck, J. B.** 1973. 'Studies on Single Neurons in Dorsal Hippocampal Formation and Septum in Unrestrained Rats. I. Behavioral Correlates and Firing Repertoires'. *Experimental Neurology* 41 (2): 461–531.
- Ray, S., R. Naumann, A. Burgalossi, Q. Tang, H. Schmidt, and M. Brecht.** 2014. 'Grid-Layout and Theta-Modulation of Layer 2 Pyramidal Neurons in Medial Entorhinal Cortex'. *Science* 343 (6173): 891–96.
- Redish, A. D., A. N. Elga, and D. S. Touretzky.** 1996. 'A Coupled Attractor Model of the Rodent Head Direction System'. *Network: Computation in Neural Systems* 7 (4): 671–85.
- Rich, P. D., H-P. Liaw, and A. K. Lee.** 2014. 'Place Cells. Large Environments Reveal the Statistical Structure Governing Hippocampal Representations'. *Science* 345 (6198): 814–17.
- Rigotti, M., O. Barak, M. R. Warden, X-J. Wang, N. D. Daw, E. K. Miller, and S. Fusi.** 2013. 'The Importance of Mixed Selectivity in Complex Cognitive Tasks'. *Nature* 497 (7451): 585–90.
- Robbins, T. W., J. Semple, R. Kumar, M. I. Truman, J. Shorter, A. Ferraro, B. Fox, G. McKay, and K. Matthews.** 1997. 'Effects of Scopolamine on Delayed-Matching-to-Sample and Paired Associates Tests of Visual Memory and Learning in Human Subjects: Comparison with Diazepam and Implications for Dementia'. *Psychopharmacology* 134 (1): 95–106.
- Robinson, J., F. Manseau, G. Ducharme, B. Amilhon, E. Vigneault, S. El Mestikawy, and S. Williams.** 2016. 'Optogenetic Activation of Septal Glutamatergic Neurons Drive Hippocampal Theta Rhythms'. *The Journal of Neuroscience* 36 (10): 3016–23.
- Rousseau, J-J.** 1762. *Émile, Ou De L'éducation*. Paris: Chez Jean Neaulme.
- Russell, J. C., D. R. Towns, S. H. Anderson, and M. N. Clout.** 2005. 'Intercepting the First Rat Ashore'. *Nature* 437 (7062): 1107–1107.
- Samsonovich, A., and B. L. McNaughton.** 1997. 'Path Integration and Cognitive Mapping in a Continuous Attractor Neural Network Model'. *The Journal of Neuroscience* 17 (15): 5900–5920.
- Sarel, A., A. Finkelstein, L. Las, and N. Ulanovsky.** 2017. 'Vectorial Representation of Spatial Goals in the Hippocampus of Bats'. *Science* 355 (6321): 176–80.
- Sargolini, F., M. Fyhn, T. Hafting, B. L. McNaughton, M. P. Witter, M-B. Moser, and E. I. Moser.** 2006. 'Conjunctive Representation of Position, Direction, and Velocity in Entorhinal Cortex'. *Science* 312 (5774): 758–62.
- Save, E., A. Cressant, C. Thinus-Blanc, and B. Poucet.** 1998. 'Spatial Firing of Hippocampal Place Cells in Blind Rats'. *The Journal of Neuroscience* 18 (5): 1818–26.
- Savelli, F., and J. J. Knierim.** 2010. 'Hebbian Analysis of the Transformation of Medial Entorhinal Grid-Cell Inputs to Hippocampal Place Fields'. *Journal of Neurophysiology* 103 (6): 3167–83.
- Savelli, F., J. D. Luck, and J. J. Knierim.** 2017. 'Framing of Grid Cells within and beyond Navigation Boundaries'. *eLife* 6 (January).

- Schlesiger, M. I., C. C. Cannova, B. L. Boubli, J. B. Hales, E. A. Mankin, M. P. Brandon, J. K. Leutgeb, C. Leibold, and S. Leutgeb.** 2015. 'The Medial Entorhinal Cortex Is Necessary for Temporal Organization of Hippocampal Neuronal Activity'. *Nature Neuroscience* 18 (8): 1123–32.
- Schmidt-Hieber, C., and M. Häusser.** 2013. 'Cellular Mechanisms of Spatial Navigation in the Medial Entorhinal Cortex'. *Nature Neuroscience* 16 (3): 325–31.
- Scoville, W. B., and B. Milner.** 1957. 'Loss of Recent Memory after Bilateral Hippocampal Lesions'. *Journal of Neurology, Neurosurgery, and Psychiatry* 20 (1): 11–21.
- Sharp, P. E.** 1991. 'Computer Simulation of Hippocampal Place Cells'. *Psychobiology* 19 (2): 103–15.
- Simon, A. P., F. Poindessous-Jazat, P. Dutar, J. Epelbaum, and M-H. Bassant.** 2006. 'Firing Properties of Anatomically Identified Neurons in the Medial Septum of Anesthetized and Unanesthetized Restrained Rats'. *The Journal of Neuroscience* 26 (35): 9038–46.
- Skaggs, W. E., J. J. Knierim, H. S. Kudrimoti, and B. L. McNaughton.** 1995. 'A Model of the Neural Basis of the Rat's Sense of Direction'. *Advances in Neural Information Processing Systems* 7: 173–80.
- Skaggs, W. E., and B. L. McNaughton.** 1998. 'Spatial Firing Properties of Hippocampal CA1 Populations in an Environment Containing Two Visually Identical Regions'. *The Journal of Neuroscience* 18 (20): 8455–66.
- Ślawińska, U., and S. Kasicki.** 1998. 'The Frequency of Rat's Hippocampal Theta Rhythm Is Related to the Speed of Locomotion'. *Brain Research* 796 (1–2): 327–31.
- Solanka, L., M. C. W. van Rossum, and M. F. Nolan.** 2015. 'Noise Promotes Independent Control of Gamma Oscillations and Grid Firing within Recurrent Attractor Networks'. *eLife* 4 (July).
- Solstad, T., C. N. Boccara, E. Kropff, M-B. Moser, and E. I. Moser.** 2008. 'Representation of Geometric Borders in the Entorhinal Cortex'. *Science* 322 (5909): 1865–68.
- Solstad, T., E. I. Moser, and G. T. Einevoll.** 2006. 'From Grid Cells to Place Cells: A Mathematical Model'. *Hippocampus* 16 (12): 1026–31.
- Sotty, F., M. Danik, F. Manseau, F. Laplante, R. Quirion, and S. Williams.** 2003. 'Distinct Electrophysiological Properties of Glutamatergic, Cholinergic and GABAergic Rat Septohippocampal Neurons: Novel Implications for Hippocampal Rhythmicity'. *The Journal of Physiology* 551 (Pt 3): 927–43.
- Spiers, H. J., R. M. A. Hayman, A. Jovalekic, E. Marozzi, and K. J. Jeffery.** 2015. 'Place Field Repetition and Purely Local Remapping in a Multicompartment Environment'. *Cerebral Cortex* 25 (1): 10–25.
- Squire, L. R., and S. Zola-Morgan.** 1991. 'The Medial Temporal Lobe Memory System'. *Science* 253 (5026): 1380–86.
- Sreenivasan, S., and I. R. Fiete.** 2011. 'Grid Cells Generate an Analog Error-Correcting Code for Singularly Precise Neural Computation'. *Nature Neuroscience* 14 (10): 1330–37.
- Stemmler, M., A. Mathis, and A. V. M. Herz.** 2015. 'Connecting Multiple Spatial Scales to Decode the Population Activity of Grid Cells'. *Science Advances* 1 (11): e1500816.
- Stensola, H., T. Stensola, T. Solstad, K. Frøland, M-B. Moser, and E. I. Moser.** 2012. 'The Entorhinal Grid Map Is Discretized'. *Nature* 492 (7427): 72–78.
- Stensola, T., H. Stensola, M-B. Moser, and E. I. Moser.** 2015. 'Shearing-Induced Asymmetry in Entorhinal Grid Cells'. *Nature* 518 (7538): 207–12.

- Strange, B. A., M. P. Witter, E. S. Lein, and E. I. Moser.** 2014. 'Functional Organization of the Hippocampal Longitudinal Axis'. *Nature Reviews. Neuroscience* 15 (10): 655–69.
- Strien, N. M. van, N. L. M. Cappaert, and M. P. Witter.** 2009. 'The Anatomy of Memory: An Interactive Overview of the Parahippocampal-Hippocampal Network'. *Nature Reviews. Neuroscience* 10 (4): 272–82.
- Sun, C., T. Kitamura, J. Yamamoto, J. Martin, M. Pignatelli, L. J. Kitch, M. J. Schnitzer, and S. Tonegawa.** 2015. 'Distinct Speed Dependence of Entorhinal Island and Ocean Cells, Including Respective Grid Cells'. *Proceedings of the National Academy of Sciences of the United States of America* 112 (30): 9466–71.
- Sürmeli, G., D. C. Marcu, C. McClure, D. L. F. Garden, H. Pastoll, and M. F. Nolan.** 2015. 'Molecularly Defined Circuitry Reveals Input-Output Segregation in Deep Layers of the Medial Entorhinal Cortex'. *Neuron* 88 (5): 1040–53.
- Swanson, L. W., and W. M. Cowan.** 1979. 'The Connections of the Septal Region in the Rat'. *The Journal of Comparative Neurology* 186 (4): 621–55.
- Swanson, L. W., P. E. Sawchenko, and W. M. Cowan.** 1980. 'Evidence That the Commissural, Associational and Septal Projections of the Regio Inferior of the Hippocampus Arise from the Same Neurons'. *Brain Research* 197 (1): 207–12.
- Swanson, L. W., J. M. Wyss, and W. M. Cowan.** 1978. 'An Autoradiographic Study of the Organization of Intrahippocampal Association Pathways in the Rat'. *The Journal of Comparative Neurology* 181 (4): 681–715.
- Taube, J. S.** 1995. 'Head Direction Cells Recorded in the Anterior Thalamic Nuclei of Freely Moving Rats'. *The Journal of Neuroscience* 15 (1 Pt 1): 70–86.
- . 2007. 'The Head Direction Signal: Origins and Sensory-Motor Integration'. *Annual Review of Neuroscience* 30: 181–207.
- Taube, J. S., R. U. Muller, and J. B. Ranck.** 1990a. 'Head-Direction Cells Recorded from the Postsubiculum in Freely Moving Rats. I. Description and Quantitative Analysis'. *The Journal of Neuroscience* 10 (2): 420–35.
- . 1990b. 'Head-Direction Cells Recorded from the Postsubiculum in Freely Moving Rats. II. Effects of Environmental Manipulations'. *The Journal of Neuroscience* 10 (2): 436–47.
- Tavares, R. M., A. Mendelsohn, Y. Grossman, C. H. Williams, M. Shapiro, Y. Trope, and D. Schiller.** 2015. 'A Map for Social Navigation in the Human Brain'. *Neuron* 87 (1): 231–43.
- Thompson, L. T., and P. J. Best.** 1989. 'Place Cells and Silent Cells in the Hippocampus of Freely-Behaving Rats'. *The Journal of Neuroscience* 9 (7): 2382–90.
- . 1990. 'Long-Term Stability of the Place-Field Activity of Single Units Recorded from the Dorsal Hippocampus of Freely Behaving Rats'. *Brain Research* 509 (2): 299–308.
- Tolman, E. C.** 1948. 'Cognitive Maps in Rats and Men.' *Psychological Review* 55 (4): 189–208.
- Tolman, E. C., B. F. Ritchie, and D. Kalish.** 1946. 'Studies in Spatial Learning I. Orientation and the Short-Cut.' *Journal of Experimental Psychology* 36 (1): 13–24.
- Torrence, C., and G. P. Compo.** 1998. 'A Practical Guide to Wavelet Analysis'. *Bulletin of the American Meteorological Society* 79 (1): 61–78.
- Touretzky, D. S., and A. D. Redish.** 1996. 'Theory of Rodent Navigation Based on Interacting Representations of Space'. *Hippocampus* 6 (3): 247–70.

- Towse, B. W., C. Barry, D. Bush, and N. Burgess.** 2014. 'Optimal Configurations of Spatial Scale for Grid Cell Firing under Noise and Uncertainty'. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* 369 (1635): 20130290.
- Tsuno, Y., N. W. Schultheiss, and M. E. Hasselmo.** 2013. 'In Vivo Cholinergic Modulation of the Cellular Properties of Medial Entorhinal Cortex Neurons'. *The Journal of Physiology* 591 (10): 2611–27.
- Ulanovsky, N., and C. F. Moss.** 2007. 'Hippocampal Cellular and Network Activity in Freely Moving Echolocating Bats'. *Nature Neuroscience* 10 (2): 224–33.
- Vandecasteele, M., V. Varga, A. Berényi, E. Papp, P. Barthó, L. Venance, T. F. Freund, and G. Buzsáki.** 2014. 'Optogenetic Activation of Septal Cholinergic Neurons Suppresses Sharp Wave Ripples and Enhances Theta Oscillations in the Hippocampus'. *Proceedings of the National Academy of Sciences of the United States of America* 111 (37): 13535–40.
- Varga, C., S. Y. Lee, and I. Soltesz.** 2010. 'Target-Selective GABAergic Control of Entorhinal Cortex Output'. *Nature Neuroscience* 13 (7): 822–24.
- Wang, Y., S. Romani, B. Lustig, A. Leonardo, and E. Pastalkova.** 2015. 'Theta Sequences Are Essential for Internally Generated Hippocampal Firing Fields'. *Nature Neuroscience* 18 (2): 282–88.
- Wei, X.-X, J. Prentice, and V. Balasubramanian.** 2015. 'A Principle of Economy Predicts the Functional Architecture of Grid Cells'. *eLife* 4 (September): e08362.
- Welday, A. C., I. G. Shlifer, M. L. Bloom, K. Zhang, and H. T. Blair.** 2011. 'Cosine Directional Tuning of Theta Cell Burst Frequencies: Evidence for Spatial Coding by Oscillatory Interference'. *The Journal of Neuroscience* 31 (45): 16157–76.
- Wells, C. E., D. P. Amos, A. Jeewajee, V. Douchamps, J. Rodgers, J. O'Keefe, N. Burgess, and C. Lever.** 2013. 'Novelty and Anxiolytic Drugs Dissociate Two Components of Hippocampal Theta in Behaving Rats'. *The Journal of Neuroscience* 33 (20): 8650–67.
- Widloski, J., and I. R. Fiete.** 2014. 'A Model of Grid Cell Development through Spatial Exploration and Spike Time-Dependent Plasticity'. *Neuron* 83 (2): 481–95.
- Wills, T. J., F. Cacucci, N. Burgess, and J. O'Keefe.** 2010. 'Development of the Hippocampal Cognitive Map in Preweanling Rats'. *Science* 328 (5985): 1573–76.
- Wilson, M. A., and B. L. McNaughton.** 1993. 'Dynamics of the Hippocampal Ensemble Code for Space'. *Science* 261 (5124): 1055–58.
- Winson, J.** 1978. 'Loss of Hippocampal Theta Rhythm Results in Spatial Memory Deficit in the Rat'. *Science* 201 (4351): 160–63.
- Witter, M. P.** 2007. 'The Perforant Path: Projections from the Entorhinal Cortex to the Dentate Gyrus'. *Progress in Brain Research* 163: 43–61.
- Yartsev, M. M., M. P. Witter, and N. Ulanovsky.** 2011. 'Grid Cells without Theta Oscillations in the Entorhinal Cortex of Bats'. *Nature* 479 (7371): 103–7.
- Yoder, R. M., and K. C. H. Pang.** 2005. 'Involvement of GABAergic and Cholinergic Medial Septal Neurons in Hippocampal Theta Rhythm'. *Hippocampus* 15 (3): 381–92.
- Yoon, K., M. A. Buice, C. Barry, R. Hayman, N. Burgess, and I. R. Fiete.** 2013. 'Specific Evidence of Low-Dimensional Continuous Attractor Dynamics in Grid Cells'. *Nature Neuroscience* 16 (8): 1077–84.
- 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): 2112–26.

- Zipser, D.** 1985. 'A Computational Model of Hippocampal Place Fields'. *Behavioral Neuroscience* 99 (5): 1006–18.
- Ziv, Y., L. D. Burns, E. D. Cocker, E. O. Hamel, K. K. Ghosh, L. J. Kitch, A. El Gamal, and M. J. Schnitzer.** 2013. 'Long-Term Dynamics of CA1 Hippocampal Place Codes'. *Nature Neuroscience* 16 (3): 264–66.
- Zugaro, M. B., A. Arleo, A. Berthoz, and S. I. Wiener.** 2003. 'Rapid Spatial Reorientation and Head Direction Cells'. *The Journal of Neuroscience* 23 (8): 3478–82.