A Functional Role for Hippocampal Non-Local Activity and a Method to Interfere with Functional Circuits

Hauður Freyja Ólafsdóttir

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

I, Hauður Freyja Ólafsdóttir 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

To my supervisor, Hugo Spiers, thank you very much for agreeing to supervise my Ph.D, supporting and inspiring my research and development, and all the interesting conversations over the past few years! I also owe special thanks to my secondary supervisor, Francesca Caccucci, for giving me self-less guidance and advice whenever I needed it and Kate Jeffery for her valuable insights and for allowing me to do research at the IBN. Moreover, I would like to thank Karel Svoboda for welcoming me into his laboratory at Janelia Farm, for giving me the outstanding opportunity to advance my scientific skills, for opening my eyes to different ways of doing science as well as reminding me that rigour is everything. Finally, I am enormously and eternally grateful to my collaborator and mentor Caswell Barry. I'm not sure where to start but I hope you know my Ph.D. would not have been the same without your input, support and consistent and ceaseless guidance. I have hugely enjoyed working with you and am thrilled to know that I will continue working with you in the coming years, yay!

I also owe so many thanks to all my colleagues, past and present, at the IBN, and Janelia – Liz, Martha, Giulio, Yave, Caitlin, Lin, Lorelei, Amir, Martin, Aleks, Bex, Laurenz, Jonas, Josh, Shaz, Robin, Nick, and Simon. Thank you all (and especially Simon) for your inestimable support, you helped made this, sometimes tumultuous, journey so enjoyable!

Lastly, to my friends and family thank you for being there and trying to understand why I want to do what I do……and why I don't have a 'real' job yet!

To my rats, mice and Shog thank you for challenging and uplifting me in proportionate amounts!

Abstract

The hippocampus has been implicated in spatial cognition. Pyramidal cells in the rodent hippocampus show a spatial correlate in their activity, such that they increase their activity when an animal enters their principal firing field ('place field'). However, behavioural state and task demands have been found to influence place cell activity when cells are outside their place field. This thesis will report findings implying a functional role for 'out-of-field' activity.

The first experimental chapter will show ensemble place cell activity is highest at the start of a goal-directed trajectory, and that such enhanced activity is related to navigational performance. These findings accord with hippocampal activity patterns observed in human studies. Moreover, we speculate the enhanced activity may relate to navigational planning. The second experimental chapter shows hippocampal place cell sequences during rest *pre*play a future, motivationally-relevant trajectory through unexplored space. This effect could not be explained by simple preconfiguration in the hippocampal circuitry. Thus, we conclude these findings show that preplay can be under the control of goal-directed influences and may represent the neural signature underlying navigational planning in novel environments.

The aforementioned findings emphasise the importance of studying functional neural circuits. Consequently, the final study describes a method to perform targeted ablations of single-cells in the rodent brain, by means of laser irradiation. We show this lesioning method is spatially precise and has a depth limitation of ~400μm below brain surface. Moreover, we describe one experimental application of the method. We suggest this method is well suited for addressing the causal links between functional neural circuits and behaviour.

The final chapter discusses the data in terms of how out-of-field hippocampal activity may be involved in future planning and the contribution the ablation method can make to the study of hippocampal function.

Table of Content

Index of Figures

Index of Tables

1 Introduction

1.1. Preamble

The hippocampus has long been hypothesised to play a vital role in spatial cognition (O'Keefe and Nadel 1978, Morris, Garrud et al. 1982). More than four decades ago, O'Keefe and Dostrovsky discovered, through electrophysiological recordings in moving rats, the activity of pyramidal cells in areas CA1 and CA3 of the hippocampus could be predicted by the place of an animal in an environment (O'Keefe and Dostrovsky 1971). During locomotion, a *place cell* is silent most of the time, but increases its activity by several orders of magnitude when an animal enters its principal firing field ('place field'). Moreover, spiking activity occurring outside a cell's place field has been generally thought of as noise (Fenton and Muller 1998). However, over the past decade evidence has accumulated implying such out-of-field activity may be functional (Wilson and McNaughton 1994, Wood, Dudchenko et al. 2000, Kentros, Agnihotri et al. 2004, Johnson and Redish 2007). Some of the functions assigned to such activity has been the consolidation of recently formed memory traces (Wilson and McNaughton 1994) and planning of future trajectories (Johnson and Redish 2007).

In the first two data chapters of this thesis (chapters 3 and 4) we will describe findings, obtained from in-vivo electrophysiological recordings of freely moving rats, which add to the growing body of literature implying a functional role of out-of-field activity. Specifically, chapter 3 will describe a study showing non-local hippocampal activity is highest at the start of a goal-directed trajectory, and as an animal makes progress to its destination, population rates drop. Moreover, we found the stronger the described relationship between activity rates and distance to a navigational goal is the better the animal performed on the task. These findings accord with human studies on hippocampal function, and, as such, provide a bridge between human and animal research. Moreover, we speculate these results may be considered consistent with the hypothesised role of place cell sequences observed in theta cycles (i.e. 'theta sequences' (Foster and Wilson 2007)) in planning and guiding navigation (Johnson, Fenton et al. 2009, Foster and Knierim 2012). In the second data chapter, we will describe results showing synchronised population activity observed during rest and sleep (O'Keefe and Nadel 1978, Buzsaki, Horvath et al. 1992) may represent a novel, motivationally-relevant future trajectory. Moreover, we show that

such *pre*play is not a result of pre-existing 'maps' (i.e. 'preconfiguration' (McNaughton, Barnes et al. 1996)) in the hippocampal circuitry. This finding, we believe, shows the population activity during quiescence is not limited to re-capitulating the past (i.e. replay (Wilson and McNaughton 1994)) but may extend to *constructing* future experiences. Moreover, this finding accords with recent studies suggesting replay may support navigational planning (Pfeiffer and Foster 2013, Singer, Carr et al. 2013), yet also extends this hypothesis to planning in novel environments.

The findings described in chapter 3 and 4 highlight the importance of functional micro-circuits within brain regions. Thus, the third and final data chapter of this thesis describes the development of a method to perform targeted laser-induced lesions of single cells in the mouse brain. In this chapter, we will show irradiating single cells with elevated doses of a femtosecond laser reliably leads to ablation of the targeted cell - up to ~400μm below brain surface, without causing damage to surrounding tissue. Moreover, we show preliminary results of one of the many possible experimental applications of this method; namely, the effect on a brain region's sensory tuning after removing a small proportion of its excitatory circuit. Although this method was developed in the Barrel cortex we consider it a *generic* method enabling perturbation of neural circuits underlying behaviour. Consequently, we believe it makes a valuable contribution to the available set of tools designed to probe the relationship between action potentials and behaviour.

In the final chapter of this thesis we will discuss the contribution the foregoing studies have made to understanding the role of hippocampal activity in behaviour, emphasising the functional role out-of-field activity may have for aiding future navigation. Moreover, I will explore some of the possible experimental applications of the laser ablation method to further elucidating the function of this intensively studied brain region. However, before reaching that point I will provide the reader with some background literature of hippocampal anatomy and function relevant to the questions described here.

1.2. Nomenclature

In this thesis the *hippocampal formation* is considered to be composed of six structures: dentate gyrus (DG), the Cornu Ammonis (CA) fields CA1, CA2 and CA3, subiculum, presubiculum, parasubiculum and entorhinal cortex (EC) (Amaral and Witter 1995). The term the *hippocampus* will be used to exclusively refer to CA1, CA3 and CA3 subfields. This nomenclature is based on the predominantly unidirectional pathway that links the 6 structures (Amaral and Witter 1995). The term *place cell* (O'Keefe and Dostrovsky 1971) will be exclusively used to refer to the pyramidal cells of the CA1 and CA3 subfields, that show a distinct spatial modulation in their activity. Moreover, the term *replay* and *preplay* will be used to refer to re-activations of stored memory traces and activations of future memory traces, respectively. Therefore, examples of re-activations that represent future behaviour will still be referred to as replay. Finally, the terms 'non-local' and 'out-of-field' hippocampal activity will refer to any activity of place cells that happens outside a cell's principal firing field (i.e. place field).

1.3. Neuroanatomy of the Hippocampal Formation

The hippocampal formation is one of few brain structures that receives highly processed, multi-modal information from a variety of neocortical as well as subcortical structures. The wealth of intrinsic connections it possesses and its prominent (almost) unidirectional micro-circuitry has made theorists suggest it plays a key role in comparing and integrating information.

The hippocampus draws its name from an analogy the anatomist Arantius made between its shape and that of a seahorse (Arantius 1587). However, its shape is better explained as that of a pair of 'Cs', one in each hemisphere. The open part of the 'C' corresponds to its rostral part. The hippocampal formation is located near the midline of the brain – close to the septal nuclei (rostrodorsally), and stretches over and behind the thalamus into the incipient pole (caudoventrally). It may be noted that the structure is more linear and horizontally oriented in primates. Moreover, there are various specie differences, such as the number of neurons residing in its subfields, the thickness of its cell layers, and extrinsic connection pathways (Amaral and Lavenex 2007). However, for the purpose of this chapter, the anatomy of the rodent hippocampus will be focused on.

One can divide the structure into two interconnected regions – namely, the hippocampus proper composed of the dentate gyrus (DG), and the CA fields (CA1, CA2, CA3). The subiculum (sub), entorhinal cortex (EC), parasubiculum (PaS), and presubiculum (PrS) form a part of the hippocampal formation (Amaral and Lavenex 2007). The hippocampal subfields vary considerably as one travels along its long (septo-temporal) axis. To note, the septal pole is rostral and dorsal to the temporal pole. At the septal pole DG and the CA fields are apparent, with the subiculum appearing about a third way along the axis. Further along, the pre- and parasubiculum appear, and the EC appears at the most caudal position (see Fig 1.1). Moreover, the axis running medial to lateral is often referred to as the transverse axis, with the DG considered to be proximal, and PrS PaS and EC at its distal end.

Figure 1.1 Anatomy of the Hippocampal Formation An image of the rat brain, with regions of the hippocampal formation highlighted in colour. The yellow regions corresponds to the hippocampus proper, dark green to the lateral entorhinal area, light green to the medial entorhinal area, purple to the perirhinal cortex and blue to the postrhinal cortex. **b)** Diagrams of the constituents of the hippocampal formation overlaid on coronal sections taken at different points along the septo-temporal axis (dotted line indicates which portion is being shown). Subfields are colour coded according to the legend at the bottom of the figure. HF = Hippocampal Formation, DG = Dentate gyrus, Sub = Subiculum, PrS = Presubiculum, PaS = Parasubiculum, MEA = medial entorhinal area, LEA = lateral entorhinal area. Diagram adapted from (Sugar, Witter et al. 2011).

1.3.1.   Major Hippocampal Pathways

The fields of the hippocampus proper are linked by a largely unidirectional (excitatory) pathway starting in the entorhinal cortex and ending in CA1 – the so called tri-synaptic circuit (Andersen, Bliss et al. 1971). The first synapse links the entorhinal cortex (from layers II and III, via the performant path) to the dentate gyrus (EC -> DG), the second connects the mossy fibres of the dentate gyrus with CA3 (DG -> CA3) and then finally CA3 projects, via the Schaffer collaterals, to CA1 (CA3 -> CA1). This circuit was thought to largely dictate information flow within the hippocampus and has, as such, influenced many theories of hippocampal function (for example (Marr 1971). However, since then research has accumulated suggesting this circuit represents only one of many functional micro-circuits within the hippocampus. To give an example, it is known a prominent projection from CA1

to subiculum exists which then projects back to the entorhinal cortex (Amaral and Lavenex 2007). Moreover, it is now thought the entorhinal cortex is the main cortical output region of the hippocampal formation (Swanson and Kohler 1986) and the subiculum the main sub-cortical output region (Swanson and Cowan 1975). However, due to its strong theoretical influence the tri-synaptic loop will be given some prominence in this discussion.

Figure 1.2. Major Intrinsic and Extrinsic Pathways of the Hippocampal Formation
a) A diagram of the hippocampal micro-circuit, illustrating the perforant and temporoammonic pathways. **b)** A diagram of projections to and from the entorhinal cortex to the hippocampal subfields. EC = entorhinal cortex. PP = performant path, TA = temporoammonic path. Adapted from (Deng, Aimone et al. 2010).

1.3.2.   *Major Fibre Systems of the Hippocampal Formation*

There are three major fibre bundles by which information travels to and from the hippocampal formation. Each of these will be discussed briefly.

Firstly, the angular bundle is a conduit for information flow between the entorhinal cortex and the hippocampal fields, including the dentate gyrus and the subiculum. It is located between the entorhinal cortex and para- and presubiculum. However, entorhinal projections to CA fields can also travel via the alveus - a thin sheet of myelinated fibres on the ventricular surface of the hippocampus; the so called temporo-ammonic alvear pathway (Amaral and Lavenex 2007).

Second, is the Fimbria-Fornix. This fibre bundle connects the hippocampus with various subcortical regions, such as the basal forebrain, the hypothalamus and the brain stem (Daitz and Powell 1954, Powell, Guillery et al. 1957). At temporal

portions of the hippocampus, fibres to subcortical regions extend in the alveus and collect in a bundle – the fimbria. The fimbria thickens as one progresses septally. The fornix is the continuation of the fimbria. This fibre bundle carries information mostly from the septal third of the hippocampus and subiculum. The fornix targets various subcortical structures such as the septal nuclei, the nucleus accumbens, diencephalon, hypothalamic areas, and the anterior thalamic nuclei (Swanson and Cowan 1975, Canteras and Swanson 1992). This fibre bundle also carries information to the hippocampus.

Finally, the third fibre bundle are the dorsal and ventral commissures that connect the hippocampi of the two hemispheres. The ventral commissure is located just caudal to the septal area and dorso-caudal to the anterior commissure. The dorsal hippocampal commissure crosses the midline just rostral to the splenium of the corpus callosum and carries fibres mainly originating or projecting to the presubiculum and entorhinal cortex.

1.3.3.   Entorhinal Cortex

My discussion of the hippocampal subfields will start with the area proposed as the starting point for information flow in the hippocampal formation; the entorhinal cortex. Its importance to hippocampal processing was noted by Ramon y Cajal (paraphrasing) saying that whatever the rest of the hippocampal formation is doing depends on what the entorhinal cortex has done.

The entorhinal cortex can be split up into two divisions – the medial entorhinal area (MEA) and the lateral entorhinal area (LEA) (largely based on entorhinal-to-dentate projections, (Brodmann 1909). Each region is composed of six layers. The first layer (layer I) is a plexiform layer, relatively cell-free but rich in transversely oriented fibres. Layer II is largely composed of stellate cells and small clusters (islands) of pyramidal cells. The third layer (layer III), is chiefly made up of pyramidal cells. Layer IV is another plexiform layer, while layer V is a cellular layer that is further subdivided into two layers – Va and Vb. A variety of cells reside in layer V, as well as in the deepest layer – layer VI (Amaral and Lavenex 2007).

Perhaps due to the hypothetical role of the entorhinal cortex as the input/output structure in the hippocampal information processing loop, relatively little is known about its intrinsic connections. However, generally speaking, superficial projections tend to terminate superficially (e.g. Layers I-III) (Canto, Wouterlood et

al. 2008). Moreover, numerous studies have shown axons of neurons from the deeper layers project to the superficial layers (Dolorfo and Amaral 1998, Hamam, Kennedy et al. 2000, van Haeften, Baks-te-Bulte et al. 2003). These connections are likely to be excitatory (van Haeften, Baks-te-Bulte et al. 2003). Thus, much research remains to be done to further elucidate the intrinsic circuitry of this brain structure.

Strong projections exists between the entorhinal cortex and the hippocampus. However, due to the emphasis of this thesis on the hippocampus, these projections will be discussed in the sub-section on the hippocampal CA fields below.

In terms of extra-hippocampal connections, one can split input to entorhinal cortex into roughly two categories – those terminating in layers I-III, and those targeting the deeper layers. Superficial input originates predominantly in olfactory structure, such as the olfactory bulb, anterior olfactory nucleus and piriform cortex. Moreover, input from the perirhinal and postrhinal cortices represents another prominent input. Cortical afferents to the deeper layers arise from the agranular insular cortex and various prefrontal regions (Insausti, Herrero et al. 1997). Moreover, the entorhinal cortex sends efferents to olfactory areas as well as the perirhinal, postrhinal, infralimbic and prelimbic cortices, to name a few.

Various subcortical structures target the entorhinal cortex, such as thalamic and hypothalamic nuclei, the amygdala, the ventral tegmental area (VTA) and raphe nuclei of the brain stem (Pikkarainen, Ronkko et al. 1999, Van der Werf, Witter et al. 2002). A prominent cholinergic projection originates in the basal forebrain. The medial septal nucleus and vertical limb of the nucleus of the diagonal band project to the medial entorhinal area, and the horizontal limb of the nucleus of the diagonal band project to the lateral entorhinal area. Finally, the entorhinal cortex returns projections to various subcortical regions – e.g. striatum, nucleus accumbens, and the septum, these projections originate mostly from layer V.

1.3.4. Dentate Gyrus

The first synapse in the hippocampal processing loop is with the dentate gyrus – a hippocampal subfield given much importance in theories of memory formation and recall (Kelsch, Sim et al. 2010, Aimone, Deng et al. 2011), and is strongly implicated in Alzheimer's disease (Gallagher and Koh 2011, Winner, Kohl et al. 2011). The dentate gyrus is a three-layered structure, its layers are the molecular, granular and polymorphic layer. The molecular layer is the most superficial and the polymorphic

the deepest. The molecular and granular layers collectively are referred to as the fascia dentata, which form a U shaped structure enclosing the polymorphic layer.

The granule cells of the granular layer are the most numerous cell type in the hippocampal formation. It is estimated that a total of 1.2×10^6 granule cells reside in the rat dentate gyrus (West, Slomianka et al. 1991, Rapp and Gallagher 1996). The granule cells are the only cells in the dentate gyrus that give rise to axons that leave the dentate gyrus – to innervate CA3 cells. Another cell type resident along the deep surface of the granule layer is the basket cell. Basket cells form pericellular plexuses that surround and form synapses with the cell bodies of the granule cells. These cells are immunoreactive to GABA and are thus classified as an inhibitory interneuron (Ribak, Vaughn et al. 1978, Seress and Ribak 1983).

The superficial molecular layer contains mostly dendrites of granule, basket and polymorphic cells as well as axons and terminal arbors from the entorhinal cortex and other sources. Of the few cells residing in this layer, two types have been identified – molecular layer perforant path–associated (MOPP) and axo-axonic cells (Soriano and Frotscher 1989). These cells innervate the axon initial segment of as many as 1000 granule cells. Both of these cell types are GABAergic.

Finally, little is known about the neurons located in the polymorphic layers, except for the mossy cells that project ipsi- and contralaterally to the granule and molecular layers. These connections are glutamatergic and are thus thought to be the main excitatory source in the dentate gyrus associational projections. Various interneuron types reside in the polymorphic layer as well, such as the fusiform cells, that project to both the granular and molecular layer.

The dentate gyrus receives cortical input from the entorhinal cortex, pre-and parasubiculum (Ramon y Cajal 1893, Kohler 1985). These inputs terminate in the molecular layer. Subcortical input originates from a variety of areas, such as the medial septal nucleus and nucleus of the diagonal band of Broca –projections thought to be largely cholinergic (Mosko, Lynch et al. 1973, Swanson, Wyss et al. 1978). Moreover, the pontine nucleus, the locus coeruleus, the ventral tegmental area and the raphe nuclei of the brain stem all project to the polymorphic layer of the dentate gyrus (Pickel, Segal et al. 1974, Swanson and Hartman 1975).

Until recently, the only extrinsic projection from the dentate gyrus was thought to be to CA3 pyramidal cells. Recently, evidence was obtained indicating that the

dentate gyrus also projects to the CA2 field (Kohara, Pignatelli et al. 2014). However, as the dentate gyrus-CA3 projection has been delineated in considerably more detail, it will be focused on here. Axons projecting to CA3 are called mossy fibres and arise exclusively from the granule cells and terminate on a narrow zone located just above the CA3 pyramidal cell layer (Blackstad, Brink et al. 1970, Gaarskjaer 1978, Gaarskjaer 1978). These fibres give rise to unique complex en passant pre-synaptic terminals called mossy fibre expansions (Amaral and Dent 1981). A single mossy fibre expansion can make up to 37 synaptic contacts with a single CA3 cell. Furthermore, each granule cell contacts about 25 CA3 cells along the full proximal-distal axis. Therefore, each CA3 cell is expected to receive input from around 72 granule cells; potentially exerting powerful control over CA3 activity (Amaral and Lavenex 2007).

1.3.5. The CA Fields

Some early anatomists likened the shape of the hippocampus to a ram's horn. Accordingly, De Garengeout (1742) named the structure 'cornu ammonis', or 'Ammon's horn', after the Egyptian mythological god Amun Kneph. Although this term is rarely used to refer to the hippocampus as a whole it is used to refer to its subfields, namely CA1, CA2 and CA3. Areas CA1 and CA3 are, by far, the most studied areas due to the ease with which one can distinguish them. CA2 is thought to be considerably smaller and its borders hard to precisely localize, although anatomist still contend it is a subfield in its own right. In fact, recently, interest in the CA2 field has gathered momentum (e.g. see (Lein, Callaway et al. 2005, Kohara, Pignatelli et al. 2014). Nonetheless, due to the importance the other fields have been given in theories of hippocampal function these will be focused on here.

The CA fields have a similar laminar organization with the exception of one extra layer in CA3. The principal layer is a tightly packed pyramidal cell layer. Below the pyramidal layer is a layer known as stratum oriens, containing the basal dendrites of the pyramidal cells as well as interneurons. Moreover, CA3-CA3 and CA3-CA1 associations are located in the stratum oriens. Just above the pyramidal cell layer in CA3 is a layer called stratum lucidum, which is occupied by the mossy fibres from the dentate gyrus. Superficial to the stratum lucidum is the stratum radiatum – this layer is immediately above the CA1/CA2 pyramidal cell layer. The most superficial layer of the CA fields is the stratum lacunosum-moleculare – projections from the entorhinal cortex terminate here. The two superficial layers contain interneurons

in addition to apical dendrites of the pyramidal cells (see figure 1.3 for a diagram of the CA fields and their laminar organization).
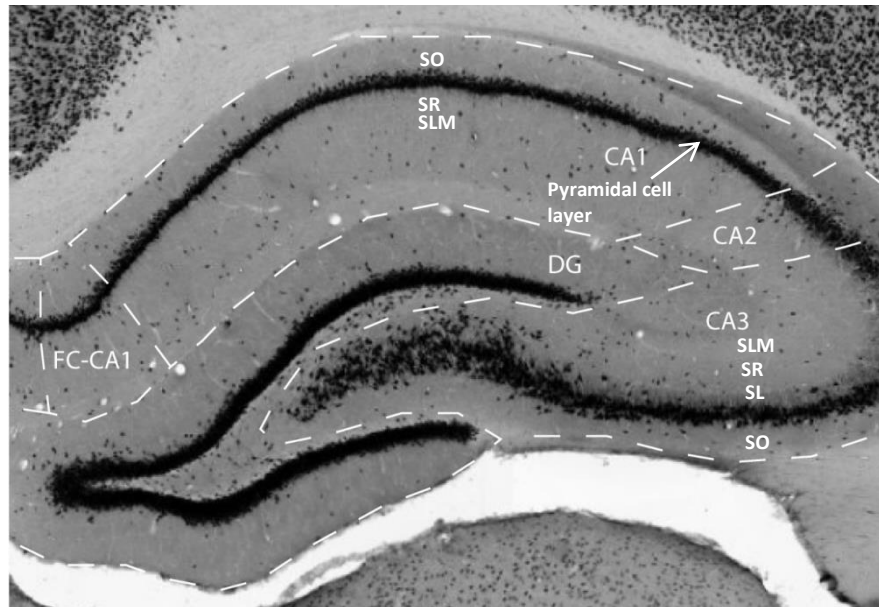
Figure 1.3. **The Hippocampal Fields** A coronal section of the right hippocampus illustrating the four subfields as well as the laminar organization of the CA fields. SLM = stratum locunosum-moleculare, SR = stratum radiatum, SO = stratum oriens, SL = stratum lucidum Adapted from the rat hippocampal atlas.

Various types of interneurons have been identified and studied in the CA fields. The two best known interneurons are the pyramidal basket cells and the axo-axonic cells, although it is thought that many more interneuron types reside in the CA fields (Freund and Buzsaki 1996). Most of these interneurons have been found immunoreactive with GABA (Ribak, Vaughn et al. 1978) and are thus thought to be inhibitory. Freund and Buzsaki (1996) have described the location, dendritic organization and axonal distribution of a number of other interneurons, and I will refer the reader to their excellent review of these.

Most synaptic input in the CA fields is intrinsic - CA3-CA3, or CA3-CA1 associational connections. These connections take place either in the stratum oriens or the stratum radiatum of the CA1 and CA3 fields. It is known that CA3 cells, from any septo-temporal position, distribute their collaterals widely to much of the full septo-temporal axis of CA1. However, what has yet to be estimated is the number of synapses a CA3 cell typical makes with a CA1 cell. The associational network observed in CA3 seems largely missing in CA1.

Another important projection to the hippocampus comes from the entorhinal cortex. Layer II projects to the stratum lacunosum-moleculare of CA3 while layer III project to CA1. Furthermore, fibres from the lateral entorhinal area terminate in more distal portions of CA1, while fibres from the medial entorhinal area terminate in more proximal portions. Projections *to* the entorhinal cortex originate only in CA1 and they target the deeper layers (Amaral and Lavenex 2007). CA1 also projects to the subiculum – the last step in the hippocampal processing loop. The topography of these projections is such that proximal parts of CA1 project to more distal portions of the subiculum while distal parts of CA1 project to proximal portions of the subiculum (Amaral and Lavenex 2007).

The CA fields receive inputs from various subcortical regions. The septum provides major subcortical input to CA3, and to a lesser extent to CA1. It is mostly the medial septum and the nucleus of the diagonal band of Broca that project to the CA fields. The projections terminate mostly in the stratum oriens, but also in the stratum radiatum. Only CA3 projects to the lateral septum (Luo, Tahsili-Fahadan et al. 2011). Moreover, both fields are innervated by amygdala and brain stem projections (Pikkarainen, Ronkko et al. 1999, Pitkanen, Pikkarainen et al. 2000). In terms of cortical connections (beside the one with the entorhinal cortex), CA3 receives projections from the perirhinal cortex. These CA3 cells return input to the perirhinal cortex. Cells from the septal portion of CA1 project to the retrosplenial cortex, and mid-septotemporal portion cells project to the medial frontal lobe.

1.3.6. The Subiculum

The subiculum has long been thought of as the final step in the hippocampal processing loop, and the structure that closes the loop by projecting to the brain region that starts the loop; namely, the entorhinal cortex. However, in recent years findings have gathered suggesting a) it may represent an earlier step in the aforementioned processing loop, and b) it may represent one of the major points at which information flows out of the hippocampus to both subcortical and cortical regions.

The subiculum lies distal and ventral to area CA1, and its border with CA1 is marked by an abrupt widening of the pyramidal cell layer. Moreover, the stratum radiatum of CA1 ends here and is replaced by a subicular molecular layer. The molecular layer can be divided into a deep portion, continuous with the stratum radiatum of CA1 and a superficial portion continuous with the stratum lacunosum- moleculare. The

principal layer of the subiculum, as that of the CA fields, is populated by pyramidal cells. The pyramidal layer starts just underneath the distal end of CA1. Although cells in this layer have received considerably less attention than those of the CA pyramidal layers, one can, broadly speaking, identify two cell types: regular spiking and bursting cells (Greene and Totterdell 1997). Intriguingly, some evidence suggests that only bursting cells contribute to the well-known entorhinal projection (Greene and Totterdell 1997).

Associational connections within the subiculum are abundant, although they appear unidirectional (Harris, Witter et al. 2001). These projections extend from the cells of origin to much of the subiculum lying temporally. In terms of intra-hippocampal projections, a strong projection to all layers of the entorhinal cortex exists, as mentioned above. However, more recent studies suggest that the entorhinal cortex also sends projections to the subiculum, via the perforant path. This projection originates in layer III of the entorhinal cortex – the layer providing input to the hippocampus proper (Baks-Te Bulte, Wouterlood et al. 2005). A meagre projection to the pre- and parasubiculum also exists.

Various neocortical regions are targeted by subicular efferents, such as the medial and ventral orbitofrontal cortex, and prelimbic and infralimbic cortices (Verwer, Meijer et al. 1997) and the retrosplenial cortex (Wyss and Van Groen 1992). The cingulate cortex receives to a lesser extent a subicular projection. Finally, in terms of subcortical connections, the amygdala projects to the subicular/CA1 border (Pitkanen, Pikkarainen et al. 2000), and the subiculum returns fibres to the amygdala. A prominent subicular projection reaches the septal complex, the nucleus accumbens and mammillary nuclei of the basal forebrain. The subiculum also projects heavily to the hypothalamus; namely, the mammillary nuclei, while the supramammillary region returns fibres to the subiculum. Various thalamic nuclei project to the subiculum, such as the nucleus reuniens and the paraventricular nucleus, and the subiculum sends fibres back to the thalamus as well. Finally, the locus coeruleus, the ventral tegmental area and the medial and dorsal raphe nuclei of the brain stem project to the subiculum.

1.3.7. Presubiculum and Parasubiculum

The role of the pre- and parasubiculum in the hippocampal processing loop are not as clear as that of the aforementioned structures. They may provide the hippocampus with sensory input as they project heavily to layers I and III of the entorhinal cortex.   However, their unique connection with the anterior thalamic nuclear complex may imply they play a role in routing thalamic influences over the hippocampus.

The presubiculum can be distinguished from the subiculum by its densely packed external layer made up of pyramidal cells. The dorsal portion of the presubiculum has distinct cell layers, while its ventral components are harder to distinguish from the deeper layers of the entorhinal cortex or the principal layer of the subiculum. The parasubiculum lies adjacent to the presubiculum. Similar to the presubiculum its superficial layer is densely populated by pyramidal cells, while its deeper layers are hard to distinguish from those of the entorhinal cortex.

Both the pre- and parasubiculum are intrinsically connected. The ventral presubiculum projects to more dorsal levels, while projections in the opposite direction arise preferentially from cells in the deep layers. With regards to the parasubiculum, dorsally directed projections are short and weak, whereas those oriented ventrally are denser and stretch over longer distances. Both structures possess numerous intra-hippocampal fibres. The most prominent of these is from the presubiculum to the entorhinal cortex (Shipley 1975, Caballero-Bleda and Witter 1993). The parasubiculum projects to layers I and II of the presubiculum, it has a substantial projection to the molecular layer of the dentate gyrus (Kohler 1985), and also to the stratum lacunosum-moleculare of the hippocampal CA fields as well as the molecular layer of the subiculum.

In terms extra-hippocampal connections, the presubiculum receives and sends back fibres to the retrosplenial cortex (van Groen and Wyss 1992, Wyss and Van Groen 1992). Moreover, visual area 18b also projects to the presubiculum. A minor projection from the prelimbic cortex also exists, and the dorsal portion of the medial prefrontal cortex. The parasubiculum, however, only receives projections from the retrosplenial and visual cortex.

The most prominent subcortical connection of the pre- and parasubiculum is that with the anterior thalamic nuclear complex.   Primarily, the anteroventral and

lateraldorsal nuclei and the anterodorsal nucleus (Kaitz and Robertson 1981, Robertson and Kaitz 1981). Furthermore, both the pre-and parasubiculum receive input from the septal nucleus and the diagonal band of Broca. Reciprocal connections with the medial and lateral mammillary nuclei and the deeper layers of the presubiculum exist (Thompson and Robertson 1987, Allen and Hopkins 1989). Finally, the dorsal and ventral raphe nuclei of the brain stem send fibres to the presubiculum.

1.3.8. Synopsis

The foregoing sections have described the major projection networks within the hippocampal formation. The most prominent, and well-studied, pathway being the one that assumes information arrives at the hippocampus from the entorhinal cortex, where it flows from the dentate gyrus to the CA3 and CA1 fields, finally reaching the subiculum which closes the loop by projecting back to the entorhinal cortex. However, other micro-circuits have been introduced, and the caveats of our current understanding highlighted. For example, the role of the pre – and parasubiculum in hippocampal function, the role of extrinsic connections in the hippocampus proper and finally the role of the CA2 field.

Delineation of the discussed circuits will be useful for discussing possible mechanisms for effects reported in chapters 3 and 4, and the links between the different studies, described in this thesis, as discussed in the final chapter (chapter 6).

1.4. Behavioural and Cognitive Correlates of the Hippocampal Formation

The seminal case of patient H.M. (Scoville and Milner 1957) was the first to indicate the role of the human hippocampus in memory formation and consolidation. In animals, the role of the hippocampus in another (perhaps related) mnemonic function has been extensively studied; namely, that of encoding a spatial representation of the environment (O'Keefe and Dostrovsky 1971, O'Keefe and Nadel 1978). O'Keefe and colleagues termed the principal cells in the pyramidal layer of CA1 and CA3 subfields 'place' cells as their firing was best predicted by the location of an animal in the environment (see fig. 1.4a) over other factors such as on-going behaviour, motivational states, and so forth (O'Keefe 1976, O'Keefe and Nadel 1978). Since their discovery a wealth of research effort has been devoted to understanding their mechanism, and an accumulation of evidence implies they indeed provide the animal with a spatial representation of their environment (i.e. a so called 'cognitive map' see (Tolman 1948)).

The cognitive map hypothesis of hippocampal function has been reinforced by a series of findings showing other 'spatial' cells exist in and around the hippocampal formation that form a part of the greater hippocampal circuitry. The second spatial cell to be discovered was the so called head direction (HD) cell (Taube, Muller et al. 1990, Taube, Muller et al. 1990, Taube 1995). The response of a HD cell is tuned to a particular allocentric (i.e. world-centred) head direction, in such a manner that a single cell will fire maximally whenever the animal's head faces its preferred direction (fig 1.4d). HD cells are found in a handful of regions that either provide direct or indirect input to the hippocampus, such as the post-subiculum, medial entorhinal cortex, lateral mammillary nuclei and the thalamus (Taube 2007). The third cell discovered was the grid cell, found in the medial entorhinal cortex mainly (Hafting, Fyhn et al. 2005, Sargolini, Fyhn et al. 2006). The firing pattern of grid cells is best described as a hexagonal grid spanning the extent of an environment (Fig 1.4b). Finally, the most recent ingredient to the brain's spatial cognition system is the so called border cell (Barry, Lever et al. 2006, Solstad, Boccara et al. 2008, Lever, Burton et al. 2009) found in the subiculum and medial entorhinal cortex that seems tuned to a boundary in an environment, such as a wall, of a particular allocentric direction (see fig.1.4c).
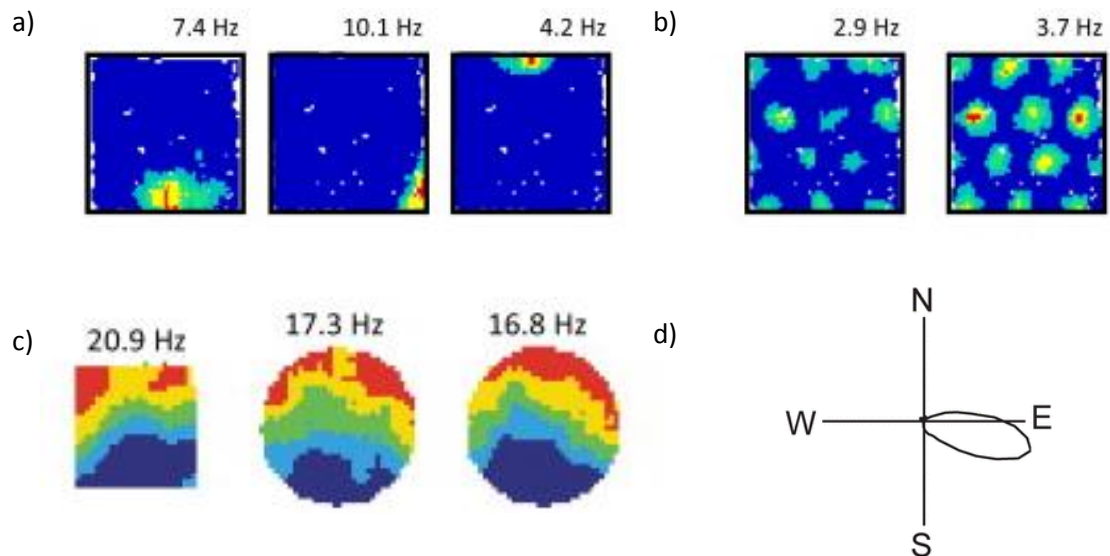
Figure 1.4. Spatial Cells in the Hippocampal Formation Four types of spatial cells recorded in the hippocampal formation, described in the foregoing text. **a)** Heatmaps of three different place cells, hotter colouring indicating higher firing rate. For all three cases, the cells have a clear location where they prefer to fire. **b)** Example of heatmaps for two grid cells recorded in a square environment, both cells demonstrate a clear hexagonal structure in their firing which tessellates the environment. **c)** Examples of ratemaps for a subicular boundary cell, whose firing is clearly tuned to the presence a north boundary. **d)** An example poplar plot of a head direction cell, whose firing is maximal when animal faces east. **a-c** adapted from Bush et al., (Bush, Barry et al. 2014), **d** adapted from Hartley et al (Hartley, Lever et al. 2014).

Furthermore, numerous recent studies imply the *human* hippocampus plays an equal role in navigation and memory (Spiers, Burgess et al. 2001, Maguire, Nannery et al. 2006, Spiers and Maguire 2007), although the spatial ability may be lateralized (Spiers, Maguire et al. 2001). Moreover, Spiers and colleagues in a thorough analysis of the brain network underlying complex navigation found that the hippocampus was preferentially involved at the start of navigation (Spiers and Maguire 2006).

Studies carried out over the past decade support the hypothesised role of the hippocampus in consolidation. As mentioned earlier, the case of patient H.M. suggested that the hippocampus plays a time-limited role in memory storage (Scoville and Milner 1957). This is because the retrograde memory impairment patient H.M. demonstrated seemed to have a temporal slope such that remote memories were retained while more recent ones were lost. Moreover, findings showing hippocampal lesions are ineffective at impairing learning if sufficient time

lapses between learning and surgery (Kim, Fanselow et al. 1992) support this hypothesis. Furthermore, a number of studies have found recent place cell activity is re-activated (*replayed)* during sharp-wave-ripple (SWR) -associated population bursts (Wilson and McNaughton 1994); providing scientists with a possible hippocampal mechanism for consolidation.

It is worth noting, the functional implications of the replay phenomenon has recently been extended to representing, and perhaps *planning,* future trajectories (Dragoi and Tonegawa 2011, Pfeiffer and Foster 2013). In agreement with such a 'constructive' role in cognition, Hassabis and colleagues have found that hippocampal patients are impaired at imaging novel scenarios (Hassabis, Kumaran et al. 2007, Hassabis, Kumaran et al. 2007, Maguire and Hassabis 2011). This hypothetical extension of the role of the hippocampus to planning/simulations of future experiences will be discussed thoroughly at later stages in this document.

To summarise, an amalgamation of studies suggest the hippocampus plays a special role in spatial cognition, the consolidation of memories and perhaps even construction of future scenarios. In this chapter, I will describe some of the electrophysiological characteristics of the hippocampus, mainly the dominant oscillatory states and the behavioural correlates of its pyramidal cell.

1.4.1.   Oscillatory Dynamics of the Hippocampus

The local field potential of the hippocampus is dominated by two distinct oscillations, associated with different behavioural states. Namely, the theta (6-12Hz) oscillation associated with any behaviour where an animal moves (Vanderwolf 1969, Buzsaki 2002), and the so called sharp-wave ripple (SWR) complexes associated with slow-wave sleep as well as periods of behavioural immobility (e.g. eating, grooming) (Buzsaki 1986, Buzsaki, Horvath et al. 1992). I will give a brief overview of the physiological characteristics of these two oscillations as well their hypothetical contribution to behaviour.

Theta

The physiology of hippocampal theta has been extensively studied by Buzsaki and colleagues (Buzsaki, Leung et al. 1983, Buzsaki, Rappelsberger et al. 1985, Buzsaki 2002). Buzsaki et al. have found theta to be most regular in frequency and largest in amplitude in the stratum lacunosum-moleculare of CA1. Both amplitude and

phase change as a function of depth whereas in the same layer amplitude and phase are robustly similar along the septo-temporal axis of the hippocampus (Bullock, Buzsaki et al. 1990) (Fig. 1.5a). Moreover, a strong relationship between theta frequency and speed has been observed (Whishaw and Vanderwolf 1973). Although predominantly associated with the hippocampus, the theta oscillation and phase-locking has been observed in many regions such as the subiculum (Adey 1967), entorhinal cortex (Mitchell and Ranck 1980), cingulate cortex (Borst, Leung et al. 1987, Leung and Borst 1987) and amygdala (Pare and Collins 2000). Finally, the medial septum-diagonal band of Broca (MS-DBB) is assumed to be the source (pacemaker) of theta, as lesioning/inactivating MS-DBB abolishes theta in all cortical targets (Petsche and Stumpf 1962). For a detailed description of prominent models of theta generation see an excellent review by Buzsaki (2002).

A plethora of hypotheses regarding the behavioural function of theta have been put forward over the past decades. Some have suggested theta is involved in comparing sensory activity with previously stored information (Pickenhain and Klingberg 1967), positive emotions and motivation (Grastyan, Karmos et al. 1966), short term memory (Jensen, Idiart et al. 1996) and response inhibition (Sainsbury 1998). In fact, to date a universal consensus regarding the behavioural correlate of theta has yet to be reached, although scientists agree that during rest or periods of immobility theta is absent.

With respect to neuronal activity of the hippocampus, O'Keefe and Recce (1993) hypothesised that theta might endow place cells with a phase code. As an animal passes through the place field of a given place cell, spikes are fired at progressively earlier phases of theta, perhaps providing the animal with a distance-through-field measure (O'Keefe and Recce 1993). Phase precession is thought to derive from the interference of two oscillators of different frequencies. One oscillator being the LFP theta frequency, and the second a theta-band membrane potential oscillation (MPO) whose frequency increases above baseline when the animal enters its place field, resulting in phase precession (Fig. 1.5b). Some experimental evidence supports this model (Harvey, Collman et al. 2009, Epsztein, Lee et al. 2010), however other theories have also been posited (Tsodyks, Skaggs et al. 1996, Losonczy, Zemelman et al. 2010). Others have interpreted the phenomenon of phase precession as a temporal code, lending support to ideas that hippocampal pyramidal cells underlie an episodic memory function (Lisman and Idiart 1995,

Eichenbaum, Dudchenko et al. 1999, Eichenbaum 2000, Gupta, van der Meer et al. 2012).
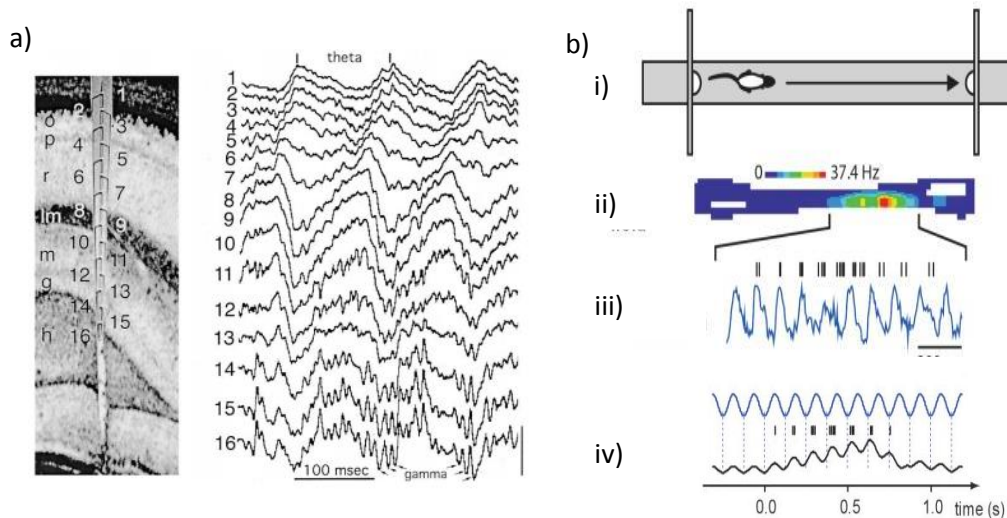
Figure 1.5. Theta and Phase Precession in the Hippocampus a) Examples of LFP recordings of theta at 16 different depths (right), starting in the stratum oriens of CA1 and ending in the hilus. From these traces it is clear that with increasing depth the amplitude of theta increases and phase shifts. Left - the 16 recording-site silicon shank used. **b)** An example of phase precession and oscillatory interference. i) A linear track an animal runs on, parts of the maze in between vertical lines are used for analysis. ii) An example of a place field from a given cell, iii) spikes (black vertical lines) plotted as a function of location in the on-going theta oscillation. From this trace it is clear that spikes fire at progressively earlier phases in the theta cycle as the animal passes through the place field. iv) An example theta trace (blue) and a MPO trace (black) increasing in frequency as the animal enters a place field. The spikes (black vertical lines) occur at the peak of the MPO, which precesses in relation to the theta oscillation. Panel **a)** adapted from Buzsaki (2002), panel **b)** adapted from Burgess and O'Keefe (2011).

Irrespective of the behavioural function of theta, it is one of the most robust electrophysiological markers of the hippocampus during active movement. When the animal is immobile, on the other hand, a distinct, more irregular, oscillation prevails in the hippocampus; namely, sharp wave ripples (SWRs) (Buzsaki 1986). Below, I will give an overview of the physiological properties of this oscillation and its proposed role in hippocampal function.

Sharp-Wave Ripples (SWRs)

For decades, it has been known that during slow wave sleep and period of behavioural immobility (such as eating, drinking, grooming) the electrophysiological state of the hippocampus changes drastically. One of the most obvious changes is the strong theta oscillation ceases and it is replaced by an irregular, high in amplitude but low in frequency, oscillation known as 'sharp waves (Buzsaki 1986). Moreover, during sharp-waves a high frequency oscillation ('ripples') co-occurs, although its occurrence is almost limited to the CA1 field (O'Keefe and Nadel 1979, Buzsaki, Horvath et al. 1992). Together, the two oscillations are referred to as sharp-wave ripple (SWR) complexes.

As noted above, sharp waves most frequently occur during slow wave sleep (SWS) and behavioural immobility (Jouvet, Michel et al. 1959, Buzsaki, Leung et al. 1983, Suzuki and Smith 1985, Buzsaki 1986). In fact, Jouvet (1959) regarded sharp waves (SPWs) as the electrophysiological marker of SWS in the cat. SPWs are seen in different species, such as in cats, rodents, chimpanzees and even humans (Freemon, McNew et al. 1969, Freemon and Walter 1970). Buzsaki (1986) carried out a thorough analysis of their properties and behavioural correlations. A single sharp wave varies in duration from 30-120ms, they have a low incidence rate (0.01 – 3/sec), and are suppressed by movement and sensory stimulation during immobility. Moreover, SPWs are invariably associated with a population burst in all hippocampal regions. Although a single pyramidal cell may not fire in every single sharp wave, its discharge likelihood is substantially higher during SPWs than during comparable periods outside SPWs. Interneurons as well as granule cells increase their discharge rate also during SPWs. Cellular activity is generally suppressed following a population burst for 100-500ms. SPWs are synchronous between hemispheres and different ipsilateral hippocampal fields. Administering atropine (a drug that interferes with theta) or a GABA antagonist increases the amplitude of SPWs, while GABA agonists significantly reduce the incidence of SPWs. Interestingly, neo-decortication or lesions to the fimbria/fornix fail to abolish SPWs, in fact SPWs are larger in amplitude in the neo-decorticated/lesioned brain (Buzsaki, Leung et al. 1983, Suzuki and Smith 1985). This finding has led scientists to hypothesise that SPWs are an intrinsic characteristic of the hippocampus and are generated when inputs to the hippocampus are weakened or blocked. The CA3

field has been hypothesised as the triggering zone for SPWs, due to its recurrent circuitry.

Ripples, a high frequency oscillation (100-250Hz), were first document by O'Keefe and Nadel in 1978 (O'Keefe and Nadel 1978), but were then comprehensively studied by Buzsaki and colleagues (1992). Buzsaki et al. (1992) noted ripples occurred in conjunction with SPWs and were most prominent in the CA1 pyramidal layer. Ripples are a spindle-shaped oscillation consisting of 5-15 sinusoid waves with ~200Hz intraburst frequency. The probability of a spike during ripple epochs has been estimated to be 3-8 times higher than during other comparable periods, and overall about 15% of recorded neurons are thought to be active during any given oscillatory epoch. As the amplitude of ripples is highest in the CA1 pyramidal layer and their phase reverses about 100μm below the layer it has been suggested that the cell bodies of pyramidal cells are the source of the oscillation. Ripples are synchronous across hemispheres, although their phase differs. Ripples have also been found in CA3, but the dominant frequency is lower - ~100Hz. Finally, interneurons also fire during ripples, in fact interneurons fire rhythmically at the fast population frequency. The phase locking of interneurons to ripples and their co-occurrence with sharp waves has led scientists to believe that ripples are caused by strong depolarisation in the CA1 layer caused by a dramatic increase in the CA3 recurrent circuitry during SPWs. The powerful excitatory input causes all cell types to fire, but particularly the interneurons, which fire at an exceptionally high frequency. The interneurons then also control the timing of pyramidal cell spiking during a ripple and finally determine the cessation of the ripple (Buzsaki, Horvath et al. 1992). Csicsvari and colleagues (2000) indeed found that 10% of CA3 cells firing within 100ms exerted a detectable effect on the CA1 network (Csicsvari, Hirase et al. 2000).
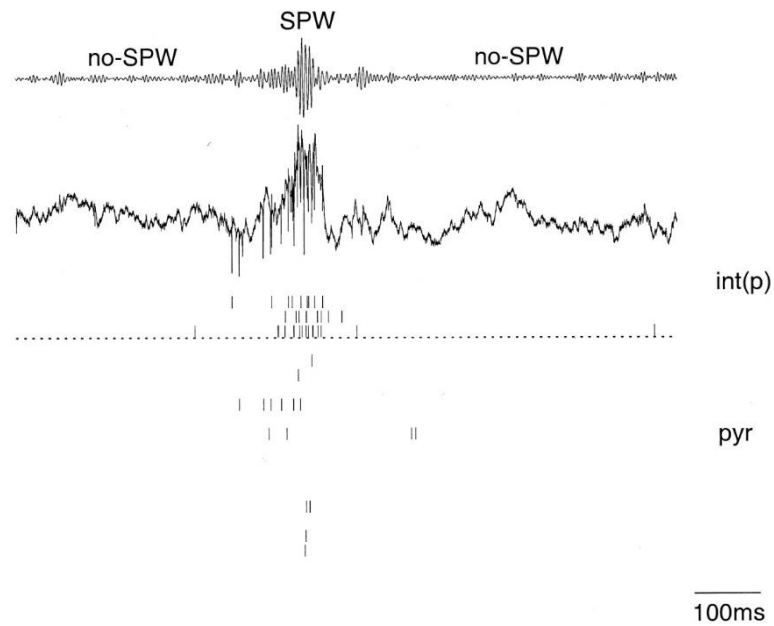
Figure 1.6. Sharp-Wave Ripple Population Bursts in area CA1 Example of a SWR (upper trace: filtered, lower trace: unfiltered), and concomitant population burst (below). Vertical black lines represent spikes. Int (p) = interneurons, pyr = pyramidal cells. Adapted from Csicsvari et al. (Csicsvari, Hirase et al. 1999).

Over the past 20 years, interest in the contribution of SWRs to spatial cognition, and memory generally, has grown considerably. In fact, now the remarkably synchronous population bursts observed during these oscillatory epochs are thought to represent the mechanism for consolidating new learning (Wilson and McNaughton 1994). Moreover, research over the past few years implies the function of SWRs may extent to navigational planning and representations of future trajectories (Dragoi and Tonegawa 2011, Dragoi and Tonegawa 2013, Pfeiffer and Foster 2013). The behavioural correlates of the SWR-associated population bursts will be discussed in more detail later in this chapter.

The foregoing discussion has given the reader a brief overview of the dominating oscillations prevalent in the local field potential of the hippocampus. For more detailed descriptions of both theta and SWRs I refer the reader to some excellent reviews by Buzsaki and colleagues (Buzsaki 1986, Buzsaki, Horvath et al. 1992, Buzsaki 2002). Now, we turn to the principal cells of the CA fields of the hippocampus.

1.4.2. Place cells: A Very Brief Introduction

As mentioned above, place cells were discovered by John O'Keefe and colleagues (O'Keefe and Dostrovsky 1971, O'Keefe and Nadel 1978), and are found in the pyramidal layers of the CA1 and CA3 subfields. Since their discovery, place cells have been studied in a range of species such as mice (Dragoi and Tonegawa 2011) , rats (O'Keefe 1976), bats (Yartsev and Ulanovsky 2013) and even humans (Ekstrom, Kahana et al. 2003) . From rodent studies, it is known that place cells fire in complex bursts (1.5-6ms spike intervals) (Ranck 1973), their firing rate is modulated by speed (McNaughton, Barnes et al. 1983), and when a place cell enters its firing field its discharge rate increases several-fold (Hill 1978, Olton, Branch et al. 1978) – although their firing is also predicted by the onset of SWRs (O'Keefe and Nadel 1978, Buzsaki, Horvath et al. 1992). Most place cells have one firing field (place field) per environment, although with increasing environmental size fields are generally bigger and a single cell is more likely to have multiple fields (Fenton, Kao et al. 2008). Place cells are often active in multiple environments, although the location and/or firing rate will differ between environments (O'Keefe and Conway 1978, O'Keefe and Speakman 1987). In other words, a place cell *remaps* when an animal enters a different environment or when the current environment is changed substantially (O'Keefe 1979, Bostock, Muller et al. 1991). Place cells develop quickly (i.e. within the first few minutes) upon entry into a new environment, and in any given environment it is estimated that 15-50% of all place cells are active (Wilson and McNaughton 1993). Finally, once a place field is established it can remain stable for months (Lever, Wills et al. 2002).

The cues influencing a place cell's activity in an environment are manifold. However, the role of sensory cues – such as visual landmarks and boundaries of an environment, are thought to play a particularly powerful role. Consequently, I will now describe in detail studies implicating the role of these cues in place cell firing.

1.4.3. The Role of Sensory Cues in Place Cell Activity

Early studies on place cell activity all converged on the role of sensory cues in determining place cell firing (O'Keefe and Conway 1978, Muller and Kubie 1987, O'Keefe and Speakman 1987). A range of sensory cues have been found to exert control over place cells, such as visual (Muller and Kubie 1987), geometric (O'Keefe and Burgess 1996), colour (Bostock, Muller et al. 1991) and olfactory (Anderson and Jeffery 2003) cues, and it appears place cells can make flexible use of available cues.

Place cells maintain their firing fields with partial removal of cues, yet with substantial changes place cells will either re-map or completely cease activity (O'Keefe 1976, O'Keefe and Conway 1978).

Although place cells have been found to use a diverse set of cues to set their place field, many studies have found prominent visual cues, such as distal landmarks, exert a particularly strong influence over place cell discharge (O'Keefe and Conway 1978, Muller and Kubie 1987). In an important study carried out but Muller and Kubie (1987) it was found that rotating a large distal cue card in a cylindrical environment resulted in concomitant (and equal in magnitude) rotation of place fields (Muller and Kubie 1987)(fig.1.7a). However, Jeffery & O'Keefe (1999) noted that such field rotation was dependent on the animal 'trusting' the visual landmark. That is, if the animal had learned that the cue card was not reliable (by observing the experimenter moving it on a previous occasion) the place cells failed to rotate (Jeffery and O'Keefe 1999).

Another prominent influence over place cell firing seems to be the presence of boundaries/barriers in an environment. For example, O'Keefe and Burgess (1996) found that if in a rectangular enclosure the animal foraged in was stretched along one axis, place cells having place fields along that axis stretched (or were split in two) by the same amount (O'Keefe and Burgess 1996). Moreover, Muller and Kubie (1987) inserted a wall into a familiar environment, cells whose place field was bisected by the wall ceased to fire after the insertion, while cells with place fields far from the novel wall seemed unaffected (Muller and Kubie 1987) (fig. 1.7b). Moreover, in a systematic study by Barry and colleagues (2006) it was observed that reducing the number of boundaries in an environment caused a gradual degradation in place cell firing (Barry, Lever et al. 2006). However, isolated objects that don't impose an impediment to movement do not seem to exert the same influence (Cressant, Muller et al. 1997).
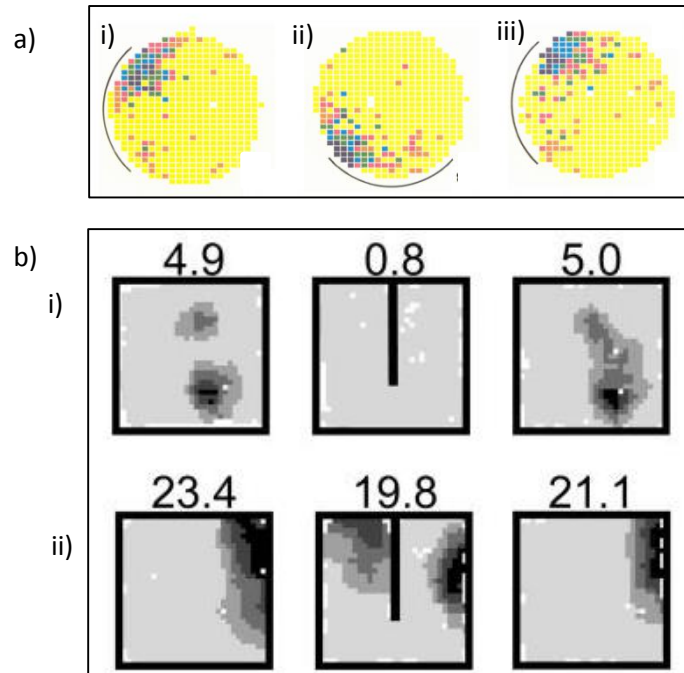
Figure 1.7. Role of Landmarks and Boundaries in Place cell Firing a) An example of a place cell whose firing field is controlled by the presence of a distal cue (curved line). In i) the cue is in its original location (west) with the place field located in the north west of the platform, then the cues moves 90 degrees south the place field rotates in the same direction by the same amount ii). Finally in iii), the cue is back in its original location as well as the place field. **b)** Example of place cells modulated by the insertion of a north-south barrier in the middle of the square recording box. For cell i) whose place field is bisected by the barrier, it ceases to fire with the presence of the barrier, while for cell ii) whose place field seems related to the west wall duplicates its field along the novel barrier. Panel **a)** adapted from Muller and Kubie (Muller and Kubie 1987), panel **b)** adapted from Barry et al. (Barry, Lever et al. 2006).

To sum, many studies imply a strong involvement of sensory cues in place cell firing, with a particular strength given to landmarks and boundaries. However, early studies also noted that many place cells remained remarkably stable when lights were turned off while an animal foraged in an environment (O'Keefe 1976). Such apparent independence of visual input for place cell firing led researcher to study another possible source of place cell firing – self-motion cues.

1.4.4.   The Role of Path Integration in Place Cell Activity

One way to locate oneself is to integrate information regarding how far one has moved and in what direction from a previous location; in other words, to path

integrate. Path integration was first described by Darwin (1873) as the integration of inertial signals (Darwin 1873). These inertial (internal) signals can derive from different sources, however much focus has been given to vestibular signals. Mittelstaedt and Mittelstaedt (1980), in the first study conclusively showing path integration as a navigation strategy, had gerbils retrieve pups from the centre of a circular arena and return back to their home cage located on the boundary of the arena. When the gerbils were rotated at a speed below their vestibular threshold, their home-bound trajectory deviated by the amount that they had been rotated by; as if vestibular information had been the basis for navigation (Mittelstaedt and Mittelstaedt 1980). The role of vestibular signals has since then been shown to play an important role in human and animal navigation (Cohen 2000, Stackman, Clark et al. 2002, Wallace, Hines et al. 2002)

O'Keefe and Nadel (1978) commented on role of path integration in place cell firing, however the idea only gained prominence after an accumulation of evidence implied place cell firing was seemingly impervious to removal of cues, especially visual cues (Hill 1978, O'Keefe 1979). Moreover, findings showing that place cell firing correlates with running speed (McNaughton, Barnes et al. 1983) and that place cells quickly stabilised in a new environment (Wilson and McNaughton 1993) afforded credence to the idea.

Some direct evidence for the role of path integration in place cell activity also exists. An important study by Gothard and colleagues (1996) found that place fields located close to the start of linear track fired in relation to distance from a moving start box (fig. 1.8), as if coding distance travelled on track. On the other hand, fields located near the end seemed unrelated to distance (Gothard, Skaggs et al. 1996). Another study by Jeffery (1997) followed the logic of the Mittelstaedt and Mittelstaedt study. Rats familiar with a given environment were slowly rotated so their vestibular system did not detect the rotation. When the animals were placed back in the environment the place fields seemed to rotate by the same amount the animals had been rotated by (Jeffery, Donnett et al. 1997). Furthermore, rats with hippocampal or fimbria/fornix lesions have been found to be impaired at navigational tasks requiring path integration (Whishaw and Maaswinkel 1998, Maaswinkel and Whishaw 1999). Finally, the recent discovery of grid cells (Hafting, Fyhn et al. 2005) – a cell type assumed to perform path integration and providing the major projection to the hippocampal system, as well as the delineation of the pathways from the vestibular system to the hippocampus (Bassett and Taube 2001)

only serve to reinforce the hypothesised role of path integration in place cell activity.
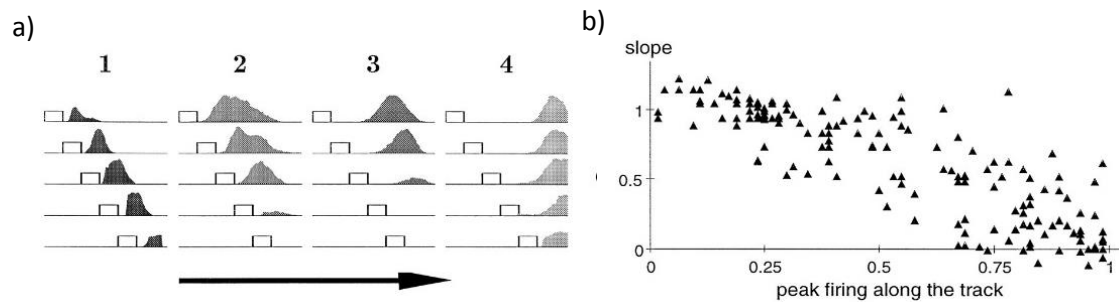
Figure 1.8. Example of Path Integration Control over Place Cell Firing a) Examples of place fields of four cells (1-4) on a linear track (horizontal line) for four different locations of the start box (rectangle on track). Cells 1-3 notably shift their place field in response to shifts in the location of the start box. Whereas cell 4 maintains the location of its place field. **b)** Relationship between location of peak firing rate along track (0 = start box, 1 = food cup, end of track), and displacement slope indicating to what extent a cell was controlled by the location of the start box (1 = fully controlled by start box, 0 = fully controlled by food cup). It is clear that cells closer to the start box seem to be influenced by the location of the start box, whereas cells further down the track are less affected by the location of the start box. Figure adapted from Gothard et al (Gothard, Skaggs et al. 1996)

However, a known problem with path integration is that over distance errors accumulate (Etienne, Maurer et al. 1996). Thus, it is unlikely that place cells solely rely on path integration mechanisms to dictate their firing. Rather, it is more likely place cells use a combination of external and internal cues. Indeed, the wealth of evidence suggesting a role for both cue types has influenced models of place cell activity. I will briefly discuss two of these models, one emphasising the role of sensory cues and the other path integration.

1.4.5. Models of Place Cell Firing

As discussed above, prominent visual landmarks seem to play an influential role in determining where a place cell fires in a particular environment (O'Keefe and Conway 1978, Muller and Kubie 1987). As such, early models of place cell firing derived the activity of place cells from the location and orientation of environmental visual landmarks (Zipser 1986). Moreover, O'Keefe and Burgess (1996) commented on the importance boundaries, such as walls and barriers, seemed to play in place field shape and location (Burgess and O'Keefe 1996).

Together, these findings lead Barry and colleagues (2006) to put forward the so called boundary vector cell model (BVC model) of place cell firing.

The boundary vector cell model of place cell activity posits that place cell firing is a continuous function of the relative location of a barrier in and around the animal's environment. Thus, place cell firing is derived from feed-forward connections from putative BVC inputs, whose firing is determined by the presence of extended barriers. A BVC cell would behave in the following way: it would fire optimally when a barrier is encountered at a defined distance and allocentric direction from the rat. To give an example, a place cell receiving input from a BVC tuned to respond to nearby barriers to the north and east would fire whenever the rat is in the north-east corner of its environment. Simulations of data based on this model are able to replicate various experimental findings (Fig. 1.9). Such as place fields having a crescent shape in a circular environment but an elongated shape along a particular axis in a rectangular environment (Muller, Kubie et al. 1987), that insertion of a novel barrier leads to either cessation of firing (if a place field is bisected by the novel barrier) or the addition of new fields (Muller and Kubie 1987, Barry, Lever et al. 2006), and that with decreasingly fewer boundaries place fields deteriorate (Barry, Lever et al. 2006). Moreover, the recent discovery of boundary cells (Barry, Lever et al. 2006, Solstad, Boccara et al. 2008, Lever, Burton et al. 2009) whose firing characteristic strongly resembles the description by Barry and colleagues of BVC cells has only reinforced this model. However, a caveat with this model is that it fails to explain why place cells are often unaffected when visual cues are removed, for example when the lights are turned off (e.g. O'Keefe, 1976). In order to explain such results, theorists have modelled place cell activity based on a so called continuous attractor.
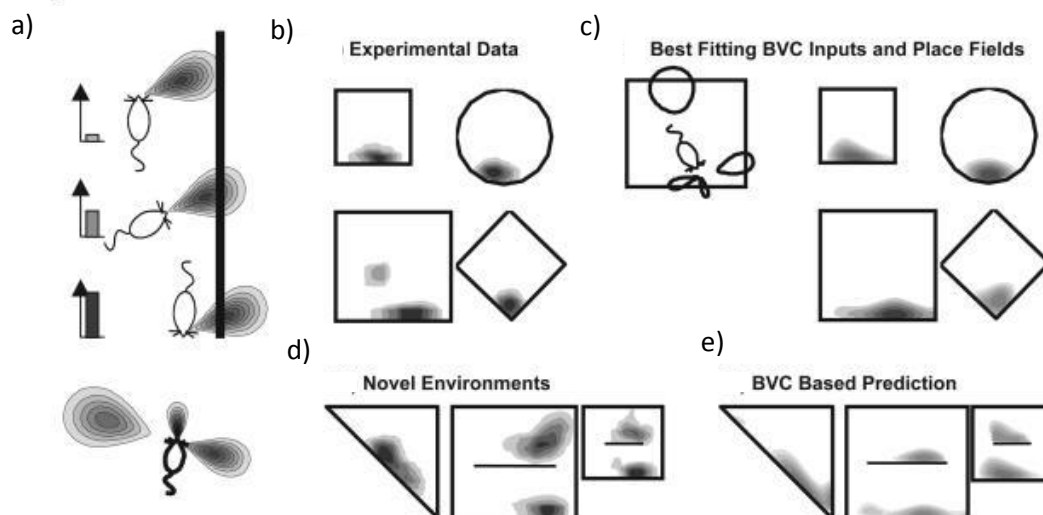
Figure 1.9. The Boundary Vector Cell Model of Place Cell Firing (previous page) **a)** A boundary vector cell that responds to a boundary at a preferred direction and distance. The firing rate decreases (shown in bar charts) as the animal's direction and distance deviate from its preference. **b)** Example experimental data, one place cell was recorded in four environments, differing in geometry. **c)** Modelled firing of place cell based on BVC inputs (inputs shown on left). **d)** The response of the cell from panel b) in three novel environments. **e)** The predicted firing of the modelled place cell in the three novel environments. Figure adapted from Barry et al. (2006).

Various experimental findings have pointed to the mechanism of a continuous attractor in determining place cell firing (McNaughton, Battaglia et al. 2006). For example, Wills et al. (2005) found the place cell ensemble code abruptly changed between two environments that only differed in geometry (Wills, Lever et al. 2005). The mechanism of a continuous attractor is best illustrated by the example of the one-dimensional head-direction (HD) system. If one imagines HD cells for different allocentric directions are symbolically arrayed in a circle, ordered by their HD preference. Each cell connects to a nearby cell with a synaptic strength that declines as a function of distance. That is, cells with a similar HD preference will be strongly connected, while cells with distinct HD preferences will be weakly connected. To note, the network is subject to global feedback inhibition to limit total activity. Activity is focused at one point and declines with distance from that point, resulting in an activity 'bump'. Such a network would keep track of head direction if the bump could be made to rotate around the ring corresponding with changes in head direction. To achieve rotation, a second layer would be required which, let's say, integrates information regarding vestibular input (Fig. 1.10a). Such a continuous attractor can easily be extended to account for 2-dimensional place cell activity. Instead of neurons being arranged in a circle they could be arranged in a plane having interconnections that decline in strength with distance. The bump can then be made to move in correspondence with a rat's motion using an intermediate layer of cells that conjunctively code for position and head direction (Fig. 1.10b). The continuous attractor model is not only supported by findings showing place cells can maintain their firing in the dark but also by the recent discovery of grid cells – cells that can integrate both head direction and position (as well as speed) (Sargolini, Fyhn et al. 2006). In fact, recent theoretical models of place cell formation assume that place cells are formed from the summation of grid cells of different scales (Burgess and O'Keefe 2005, Fuhs and Touretzky 2006, Solstad, Moser et al. 2006).
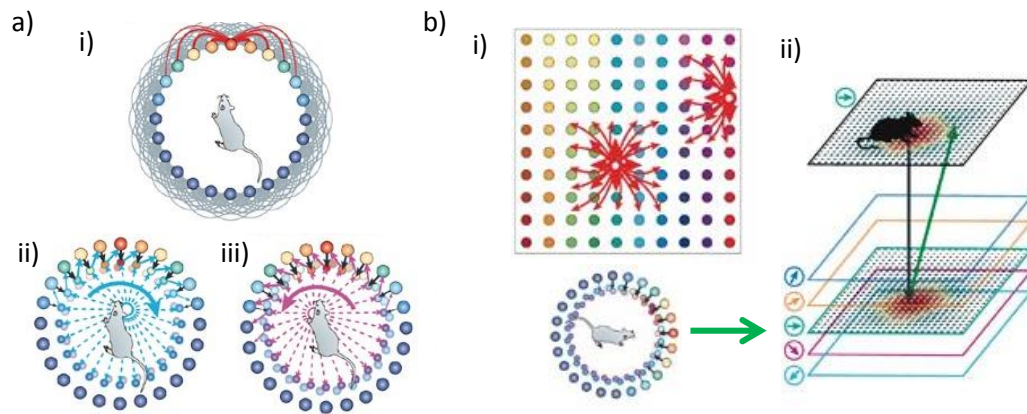
Figure 1.10. Illustration of a 1-Dimensional and 2-Dimensional Continuous Attractor a) An example of a continuous attractor for a 1-dimensional head direction system. i) Cells with similar preferred head directions (closer together on ring) are strongly connected, while cells with dissimilar preferred head directions (far apart) are weakly connected. ii) A hidden layer receiving vestibular input indicating head movements to the east project to head direction cells to the right so to move the bump of activity clockwise. Alternatively, iii) if the hidden layer cells receive information that the animal has moved its head to the west, cells will project to head direction cells to the left, moving the activity bump anti-clockwise. **b)** An extension of the continuous attractor to 2-dimensional coding of place cells. Place cells are organised in a plane depending on the location of their place field (i), as the animal moves in the plane, a hidden layer of cells coding for direction and position will project to the place cell plane, moving the activity bump in accordance with the animal's movement (ii). Figure adapted from McNaughton et al. (2006).

However, a problem with assuming place cells are underpinned by path integration mechanisms is that a path integrator will inevitable accumulate error with distance, thus other cues, such as landmarks and boundaries need to be used to correct the integrator. Therefore an interaction between path integration and sensory cues, and consequently a continuous attractor and BVC inputs, is perhaps a more comprehensive way of accounting for place cell firing. In fact, mechanisms involved in a continuous attractor and a BVC model may not be mutually exclusive. A place cell may flexibly switch between the two mechanisms depending on available cues, or, simply, cells will be arranged in the continuous attractor plane based on their BVC inputs.

Yet, it is also clear that the variability of a place cell's activity cannot be entirely explained by path integrative or sensory cues. Muller and colleagues (1987) found

the activity of a place cell for two very similar runs through its place field varied dramatically. In fact, on one run a cell fired 18 spikes whereas on another it did not fire at all (Muller, Kubie et al. 1987). Moreover, Fenton and Muller (1998) comprehensively analysed the variability of place cell activity, and found it exceeded the variance assumed by an inhomogeneous Poisson process (IPP) (Fenton and Muller 1998). In recent years, the role of cues that do not fall in either the sensory or path integrative cue category have been given increment attention.

1.4.6. *A Role for Non-Spatial Cues in Determining Place Cell Activity*

The foregoing discussions have discussed the role of 'spatial' cues in dictating where a place cell fires in an environment. That is, both path integrative and sensory cues relate to cues that refer to some kind of spatial information, such as distance or direction. However, recently findings have accumulated indicating that non-spatial cues can influence place cell activity. The location of goals (Hok, Lenck-Santini et al. 2007, Dupret, O'Neill et al. 2010), future trajectories (Frank, Brown et al. 2000, Wood, Dudchenko et al. 2000) and time (MacDonald, Lepage et al. 2011), are examples of some of these non-spatial cues. Based on the studies described later in this thesis, I will focus this discussion on the literature implicating an influence of goal locations and future trajectories over place cell activity.

Ranck (1973) was the first to note a correlation between hippocampal complex spikes and approach and consummation of reward (Ranck 1973). Moreover, place cell experiments carried out on linear tracks generally reveal a U-shaped distribution of place field locations; that is, an over-representation of the ends of the track, where reward is generally located (for example see fig 4.6 in Chapter 4). Until recently, scientists often assumed such an over-representation was due to a behavioural confound, as sharp-wave ripple complexes (Buzsaki, Horvath et al. 1992) are frequently observed during food consumption which often happens at the end of tracks. However, some recent studies have called this assumption into question (Hollup, Molden et al. 2001, Fyhn, Molden et al. 2002, Hok, Lenck-Santini et al. 2007, Dupret, O'Neill et al. 2010).

Hollup and colleagues (2001) trained rats to navigate towards a hidden platform in an annular water maze. Throughout trials, CA1 place cells were recorded and the location of the place field in the annular maze analysed. Hollup et al observed a striking clustering of place fields around the (goal) platform location. About 30% of place cells had their place field at the platform location. Importantly, such

clustering was maintained when animals were trained on a version of the task where the platform changed location on different trials, and when behavioural factors such as swimming speed were controlled for (Hollup, Molden et al. 2001). Moreover, Fyhn and colleagues (2002) observed such goal clustering was accentuated when a learned platform location was changed to a novel one. However, a possible caveat with these findings is that these water maze tasks are stressful for the animals, thus perhaps the clustering at the platform location may be a result of the animal feeling relieved to have found the platform. In fact, some evidence suggests emotions such as fear influence place cell activity (Moita, Rosis et al. 2004, Wang, Wann et al. 2012).

a)
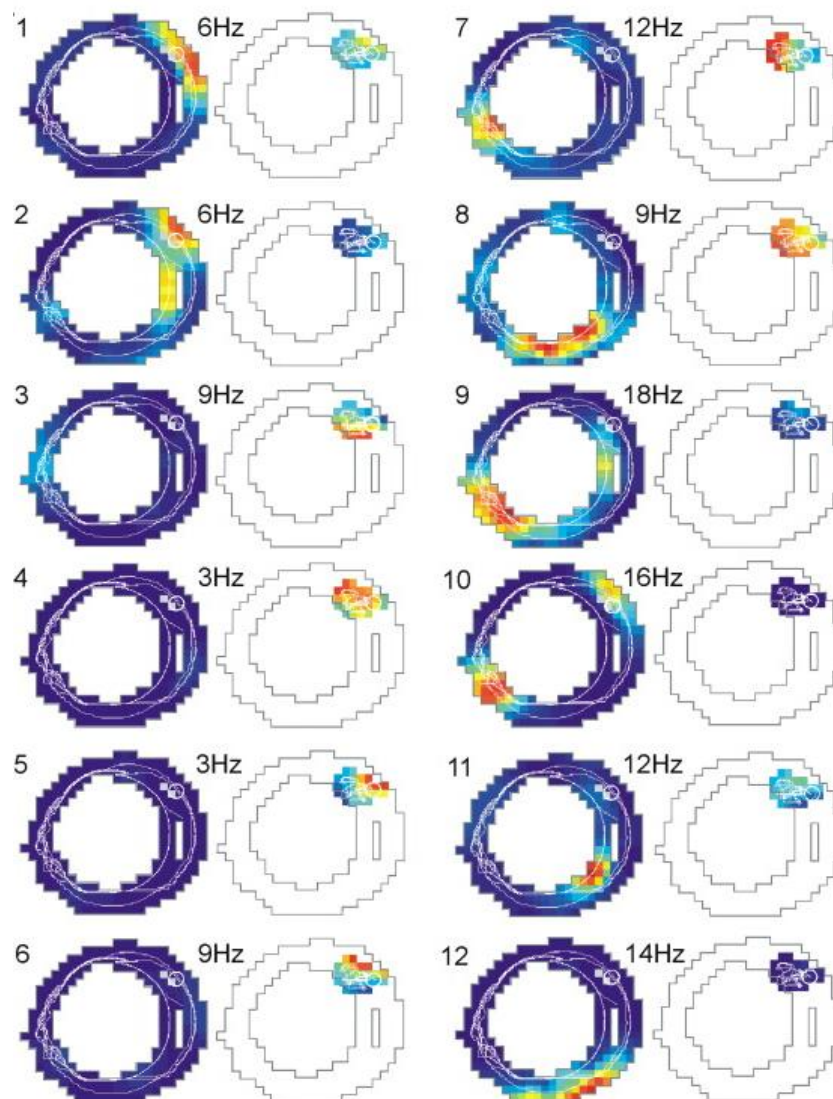
b)
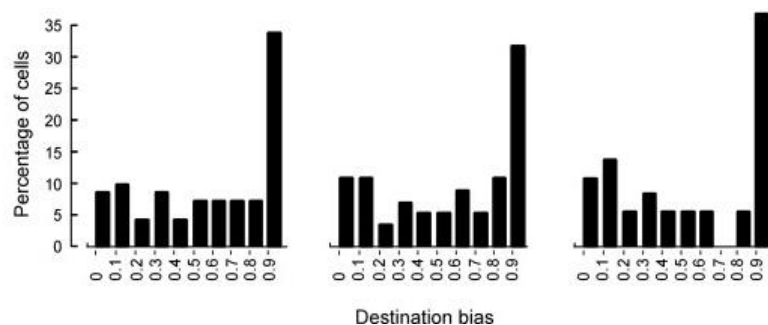
Figure 1.11. Clustering of Place Cells around Goal Locations (previous page) **a)** Ratemaps for 12 CA1 place cells recorded during an annular water maze task. Rats learned to search for a hidden platform in order to escape the maze. On a probe trial the location of the platform changed. The ratemaps are from the probe trial. Hotter colours represent higher firing rates. For each cell the activity during the swim (left) and goal (right) phase is plotted separately. Goal phase refers to the period after the animal has reached the platform. Peak firing rate is given above each ratemap. These examples show many cells fired during the goal phase of the trial. **b)** Distribution of 'destination biases' of place cells recorded. Destination bias was estimated by dividing the firing rate during the goal phase by the sum of the firing rates during the swim and goal phase. On this measure, 0 indicates a cell is just active during swim phase and 1 that is just active during the goal phase. The data was split by the firing rate at the platform (left: < 0.5Hz, middle: >1Hz, right: >3Hz). For all three data subsets it is clear that a disproportionate amount of cells are active at the goal location. Adapted from Fyhn et al. (2002)

However, two studies carried out by Hok and colleagues (2007) and Dupret and colleagues (2010) rule out an effect of emotion. Hok and colleagues (2007) recorded place cells while animals ran in a circular arena. The animal learned an area of the arena was associated with reward. In order to receive rewards, animals had to enter the reward area and wait two seconds. Hok et al. observed increased place cell activity in the goal area, prior to the animal receiving the reward (Hok, Lenck-Santini et al. 2007). Yet, a possible caveat with this finding is that the animals might have been relatively immobile while in the goal area, waiting for food, thus the hippocampus might have been in a SWR state, where it is known that place cells increase their activity (Buzsaki, Horvath et al. 1992). However, such a behavioural confound is unlikely to explain a similar finding by another group. Dupret and colleagues (2010), similarly, trained rats to run to a reward location on a circular platform ('the cheeseboard maze' (Gilbert, Kesner et al. 1998)). Place cell activity was recorded over several days, and on each day the location of the reward was changed. Dupret et al. observed a 'reorganisation' of the hippocampal CA1 pyramidal cell ensemble between days, such that place fields seemed to cluster around the location of the goal on a particular day (Dupret, O'Neill et al. 2010). Moreover, Dupret and colleagues also recorded from place cells in the CA3 field, and found no evidence for such re-organisation; thus, importantly ruling out possible behavioural confounds.

A few studies have found evidence for prospective and retrospective coding in the activity of hippocampal place cells. That is, the location of the firing field of a place

cell is modulated by the future or past trajectory of the animal (Frank, Brown et al. 2000, Wood, Dudchenko et al. 2000, Ferbinteanu and Shapiro 2003). Wood and colleagues (2000) had rats perform a continuous alternation task on a T-maze, where animals ran up a central stem and then had to turn either left or right in order to receive reward on their return to the central stem. Wood et al., found place cell activity on the central stem varied as a function of the future path taken; that is, some place cells would be active on the central stem only if the animal would then turn left (Wood, Dudchenko et al. 2000). Moreover, Frank and colleagues (2000) observed both prospective and retrospective coding in the activity of place cells on a 'W' maze – a maze where an animal either needs to make 'inbound' trajectories from the right or left arm of the maze to the middle arm, or 'outbound' trajectories from the middle arm to the left or right arm. Some place cells active on the central arm modulated their activity based on where the animal was coming from (i.e. on inbound trials) or where it was going (on outbound trials). The modulation was manifested either as a cell turning 'on' or 'off' or changing the firing rate of its place field. Frank et al. found about 30% of place cells showed either type of trajectory coding – although this estimate was lower when speed, path and head direction variations between trajectories were controlled for (Frank, Brown et al. 2000). Finally, Ferbinteanu and Shapiro (2003) found trajectory coding to be diminished on trials where the animal made an error, especially for cells showing prospective coding; consistent with such coding having a task-related function.

a)

i)

b)

ii)
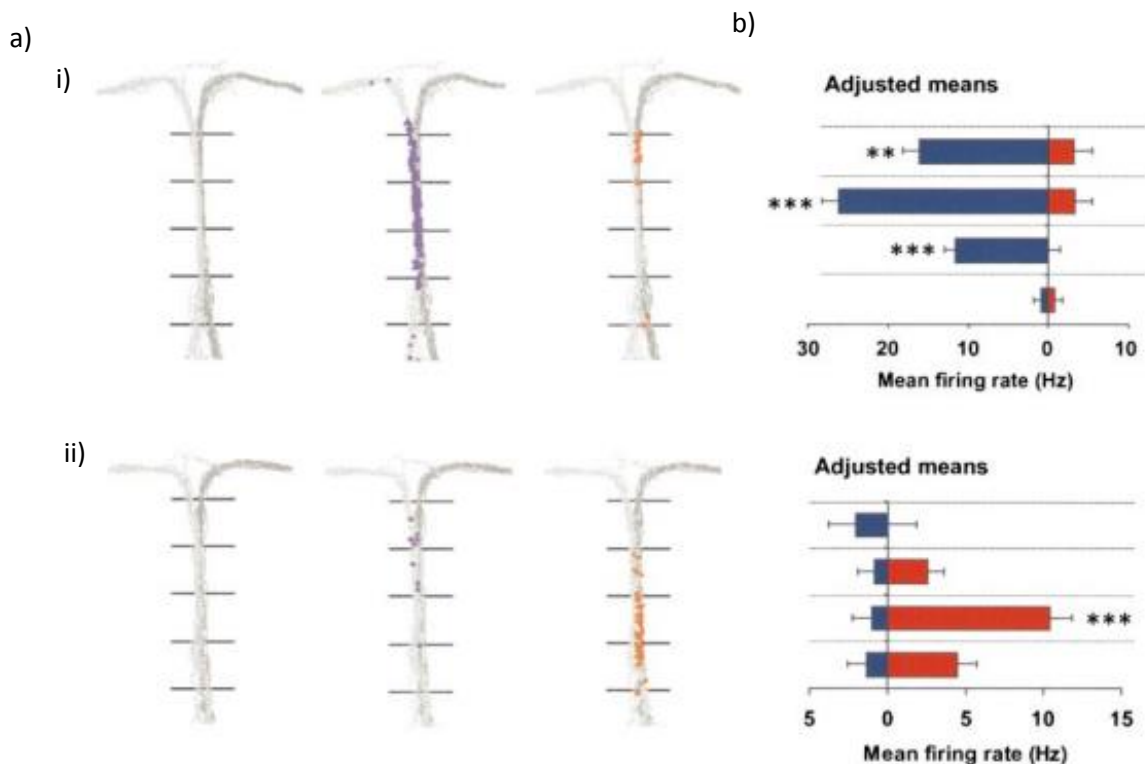
Figure 1.12. Prospective Coding in CA1 Place Cells (previous page) Example of two CA1 place cells (I and ii) whose activity is modulated by the future turn the animal will take. **a)** Left most panel shows trajectory of animal (light grey = left trials, dark grey = right trials), middle panel shows location of spikes (in blue) for all left trials, and right panel shows location of spikes (in red) on right trials. The central stem is divided up into four sectors as indicated by horizontal lines. **b)** Mean firing rates for each sector (red = right trials firing rate, blue = left trials firing rate). It is clear that cell (i) fires only the central stem when the animal will turn left on the arms, and that cell (ii) only fires on the arm when the animal will turn right. Adapted from Wood et al. (Wood, Dudchenko et al. 2000).

To summarise, evidence is gathering implying hippocampal place cells are influenced by a set of non-spatial cues, such as future/past trajectories and goal locations. Moreover, later in this thesis I will describe data that may be interpreted as supporting the findings described in this section; namely that the hippocampus represents future motivationally-relevant trajectories during rest. However, before describing these (and other) findings I will give an overview of the fast-growing literature dedicated to understanding the role and determinants of place cell firing occurring outside a place cell's main firing field – another firing characteristic not explained by sensory or path-integrative cues.

1.4.7.   The Function of Hippocampal Out-of-Field Firing

O'Keefe and Nadel (1978) noted one of the strongest determinant of a place cell's activity, beside the spatial location of an animal, is an animal's behavioural state. Namely, during slow-wave-sleep (SWS) SWR complexes dominate the hippocampal local field potential, and the likelihood of a hippocampal pyramidal cell firing is increased manifold (also see (Buzsaki, Horvath et al. 1992)). The activity during SWRs seems independent of the location of the animal; thus, this activity represents so called 'out-of-field' firing (also known as 'non-local' firing). The activity pattern of place cells observed during SWR events (both during sleep and behavioural immobility) has received much attention in recent years, and data has accumulated implying such activity may play a role in consolidation of recently formed memories (Wilson and McNaughton 1994) and perhaps planning of future trajectories (Pfeiffer and Foster 2013). A similar function has been assigned to the out-of-field firing sometimes observed during wakefulness when an animal is at choice points in a navigational task (Johnson and Redish 2007). The literature covering both types of out-of-field firing will be discussed in detail below.

Wilson and McNaughton (1994) recorded from large (100+) ensembles of CA1 pyramidal cells while an animal foraged on a novel circular platform, and during a preceding (PRE) and subsequent (POST) sleep session. Remarkably, Wilson and McNaughton found that place cells that tended to fire together during wakefulness (i.e. had place fields near to each other) also tended to firing during the POST sleep session. Importantly, the correlated firing was weaker during the PRE sleep. In other words, the CA1 network seemed to re-activate (or *replay)* recently formed memory traces during sleep (Wilson and McNaughton 1994). Wilson and McNaughton hypothesised such replay represented the mechanism underlying strengthening (or consolidation) of memory traces through a Hebbian learning process (Hebb 1949). Since Wilson and McNaughton's seminal study, replay has been studied by many different research groups (Diba and Buzsaki 2007, Karlsson and Frank 2009, Gupta, van der Meer et al. 2010), and has been found to occur both during sleep and periods of behavioural immobility, such as when an animal is eating or grooming (Diba and Buzsaki 2007, Gupta, van der Meer et al. 2010). Replay is best described with the example of a linear track. In the hippocampus, a path down a linear track is represented as a sequence of place cells becoming active as an animal progresses down a track – e.g. cells A-B-C-D-E. During sleep or wakeful quiescence the place cell sequence underlying the path is then re-activated multiple times. Moreover, such replay can either recapitulate the path in the original running direction (i.e. A-B-C-D-E) or in the reverse direction (E-D-C-B-A).

As mentioned above, Wilson and McNaughton (1994) hypothesised replay might represent the mechanism underlying consolidation in the hippocampus – a function assigned to the hippocampus decades earlier (Scoville and Milner 1957). In fact, various studies on replay have found evidence consistent with such a function. For example, Cheng and Frank (2008) found activity of place cells representing a novel arm to be more coordinated during ripples than those for a familiar arm. Specifically, Cheng and Frank recorded large ensembles of place cells during an alternation task where one arm was familiar and one was novel. Cheng and Frank analysed cross-correlations of place cells with overlapping fields on either am. Cell-pairs on the novel arm had more 0-lag correlations than cell pairs on the familiar arm. Moreover, the enhanced synchrony observed for the novel arm seemed to diminish with experience (Cheng and Frank 2008). If replay underlies consolidation, one would indeed expect it to be stronger for novel experiences than familiar ones. In a similar study, Singer and Frank (2009) found SWRs following a rewarded

trajectory had significantly more CA3 place cells active than SWRs following an unrewarded trajectory (Singer and Frank 2009). Furthermore, one would expect replay to correlate with learning if involved in memory consolidation. Dupret and colleagues (2010) found replay of goal-locations in the 'cheeseboard' maze during learning and rest to predict memory performance on the task. Finally, the causal links between task-related ripples and learning was assessed by Girardeau and colleagues (2009). Girardeau et al. selectively interfered with SWRs during rest following learning and found rats were significantly impaired on a hippocampal-dependent spatial memory task compared to controls (Girardeau, Benchenane et al. 2009).

a)

b)

c)
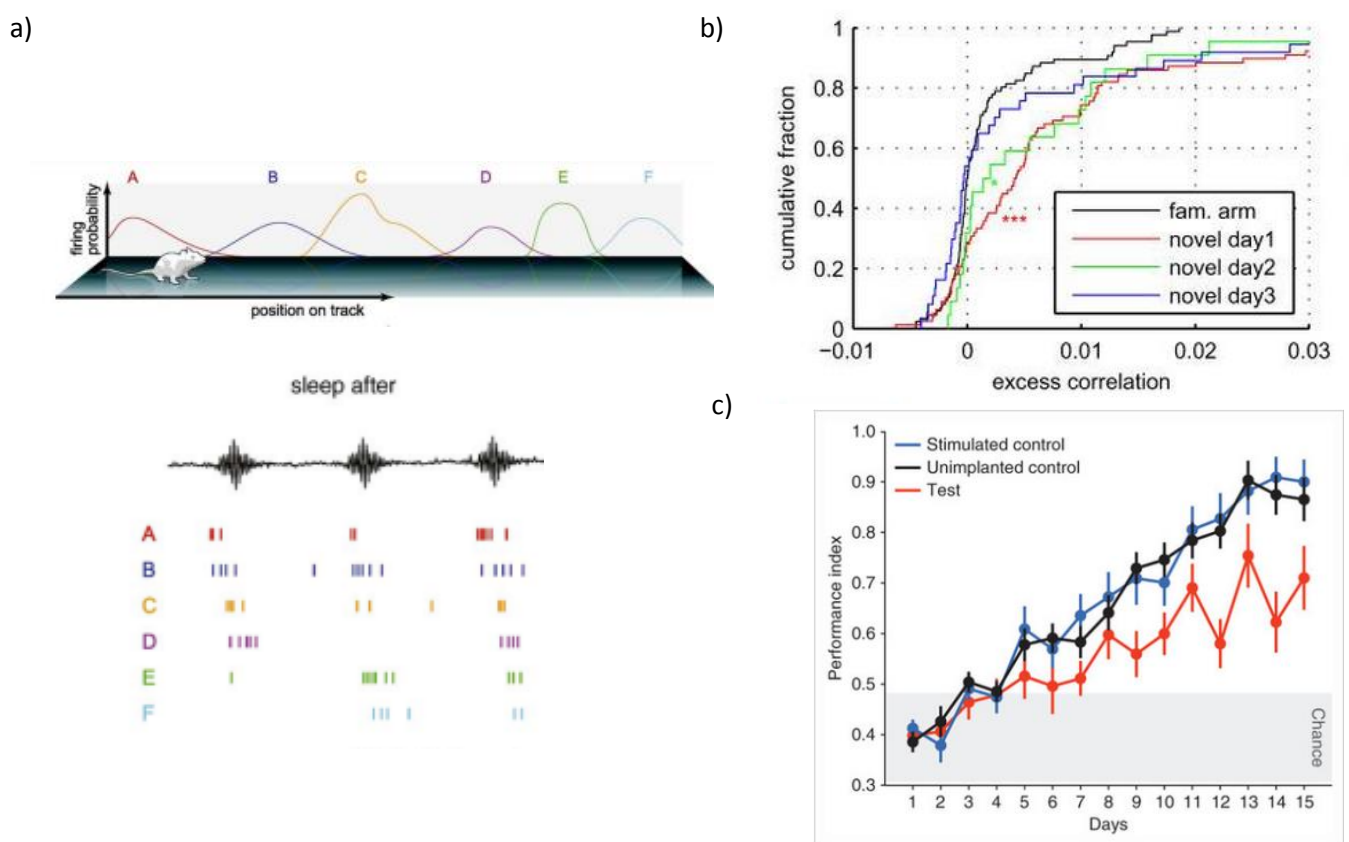
Figure 1.13. Reactivations of Place Cell Sequences during Rest a) An example of a place cell sequence (top) for a run down a linear track and forward replay of the place sequence (bottom) during SWRs. **b)** Enhanced synchrony of place cells with place fields on a novel arm in a maze. The graph shows place cell synchrony reduced across days, and approaches that observed for a familiar arm (black) line. X-axis shows the excess correlation and Y-axis the cumulative frequency distribution of excess correlations. **c)** Interfering with sharp-wave ripples interferes with learning on a spatial task. Orange line shows performance for animals where SWRs were selectively interfered with during sleep following training. The blue and black lines show performance for control groups. Panel a)

is adapted from O'Neill et al (2010), panel b) from Cheng and Frank (2008) and c) from Girardeau et al. (2009).

To summarise, evidence has gathered supporting the proposition that replay may represent the mechanism underlying the consolidation of recently formed memory traces. Yet, recently scientists have speculated that replay may also serve a prospective function. Namely, it may represent the neural signature underlying future planning.

As noted earlier, replay may either recapitulate place cell sequences in the order they were activated during experience or in the reverse order (Foster and Wilson 2006). Theorists have speculated reverse replay may serve an important role for consolidating goal directed, future-relevant, trajectories (Foster and Knierim 2012). Moreover, Diba and Buzsaki (2006) found the order in which place cells were active during replay depended on whether an animal was concluding or commencing a run. Specifically, Diba and Buzsaki found that as an animal reached the end of a track, place cell sequences were replayed in the reverse order while replay at the start of a run seemed to activate place cells in the original ('forward') order. That is, at the start of the run, the hippocampus seemed to anticipate the trajectory an animal was about to take. Furthermore, Singer and colleagues (2013) found SWR activity predicted whether an animal would make the correct turn on an alternation task. That is, on the 'W' maze described above (Frank, Brown et al. 2000), the number of cell pairs active during SWRs on the central arm during outbound trials predicted whether the turn at the end of the arm would be correct or incorrect. Perhaps the synchronised activity on the central arm represented route planning? Yet, to convincingly show replay serves a planning function a relationship between a replayed trajectory and the future path of the animal, on a task where an animal can take multiple trajectories, needs to be demonstrated. In an ambitious study by Pfeiffer and Foster (2013) rats learned to flexibly navigate to a goal-location in a square environment. During periods of immobility, such as when an animal consumed food at the end of a trajectory or just prior to the animal initiating a goal-directed trajectory, Pfeiffer and Foster (2013) analysed trajectories represented during SWR events. They found, first of all, replay events over-represented goal-directed trajectories and that replayed trajectories seemed to represent the *future* goal-directed path of the animal; consistent with the suggestion that replay represents a navigational planning mechanism. Furthermore, some recent studies have shown future place cell sequences of a *novel* environment are *pre*-played in

sleep and/or rest prior to an animal entering the novel environment (Dragoi and Tonegawa 2011, Dragoi and Tonegawa 2013). Although, such preplay might be interpreted as preparation for future navigation, the studies by Dragoi and Tonegawa do not differentiate between goal-directed influences and influences derived from supposed pre-existing maps in the hippocampus (i.e. pre-configuration (McNaughton, Barnes et al. 1996). In this thesis, I will describe a study which attempts to tease apart these two influences.
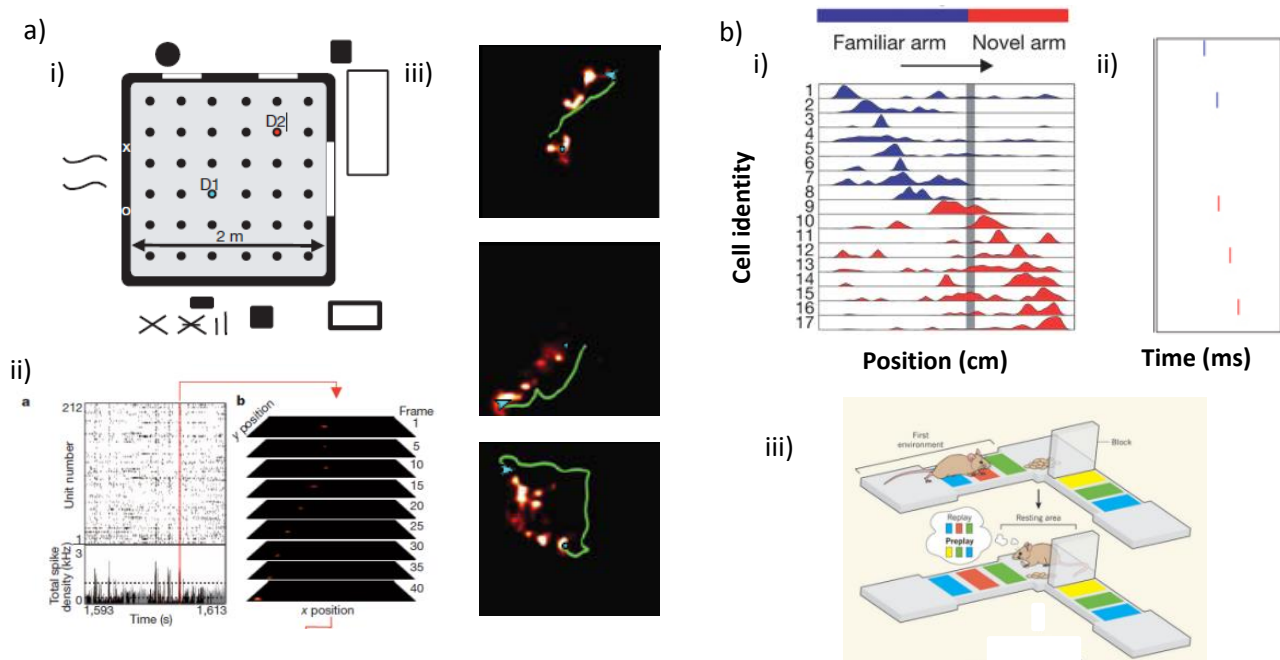
Figure 1.14. Sharp-Wave Ripple Activity Represents Future Trajectories a) During a goal-directed navigation task hippocampal cell ensembles replay the future path of the animal. i) Schematic of apparatus used, D1 and D2 mark locations of goals on two different days. ii) Left: raster plot and spike density plot of neurons recorded from CA1 during running and behavioural immobility. The activity during behavioural immobility (during SWRs) was used to estimate the trajectory represented (right). iii) Decoded trajectories from SWR events, lighter colours represent more probable locations, the green line the actual future path of the animal and the blue triangle the location of the animal at the start of the event. **b)** Future hippocampal sequences for a novel environment are *preplayed* during rest before an animal enters the environment. i) Tuning curves for place cells active on familiar (blue) and novel (red) arm of the environment, x-axis shows position on arms (cm) and y-axis cell ID, ii) an example of a preplay event of the novel environment observed while the animal is immobile in the familiar environment, not having entered the novel environment. Cells in spiking event are ordered according to the location of their peak on both arms, x-axis shows time (ms), which is in this case 300ms iii) A schematic of the preplay phenomenon, during rest in a familiar environment both place cell sequences for the familiar and novel

environment are activated. Panel a) is adapted from Pfeiffer and Foster (2013) and b) from Dragoi and Tonegawa (2011).

In sum, some preliminary evidence suggest hippocampal replay, observed during sleep and rest, may represent a navigational planning mechanism. Yet, some evidence also exists showing such planning may not be limited to quiescent periods.

In an important study carried out by Johnson and Redish (2007) it was observed that ensembles of CA3 place cells 'swept' ahead of the animal, while an animal was located at a choice point in a navigational task. Specifically, animals made several 'T turns' on the central stem of a multiple T-maze. When animals reached the end of the stem they had to make an important decision as to whether to turn left or right on the return arm, in order to receive a food reward. As the animal approached the choice point Johnson and Redish observed the CA3 ensemble seemed to represent trajectories down the left and right return arms, separately. Such 'sweeps' were not reliably seen at any other point in the task, and they preferably represented future paths rather than past paths. Moreover, Johnson and Redish found sweeps to be associated with *theta* not SWRs (Johnson and Redish 2007). The authors speculated such sweeps represented decision-making regarding which path to take in the task. However, an alternative interpretation is that such sweeps are a mere consequence of phase precession, described earlier. That is, place cells fire at a progressively earlier phase of a theta cycle as an animal passes through their place field. This precession then results in 'micro sequences' of future (and past) paths within a theta cycle. However, Foster and Wilson (2007) analysed such theta sequences during spatial experiences and found they could not be merely predicted by phase precession (Foster and Wilson 2007). Moreover, in a follow-up study by Gupta and colleague (2012) it was shown that sweeps seemed to *segment* an environment animals performed a spatial decision-making task in. The segmentation seemed to be based on landmarks in the environment, for example turns and choice points, such that sweeps represented paths between landmarks rather than paths across them (Gupta, van der Meer et al. 2012). Segmentation would not be predicted by phase precession. Moreover, the segmentation seemed akin to mnemonic 'chunking'; consistent with the suggestion that sweeps during theta cycles are *functional*, and relate to performance on a navigational task. However, to show sweeps have a functional role to play in navigation a relationship between sweeps and future choices must be shown, similar to that shown by Pfeiffer and Foster (2013) for replay.
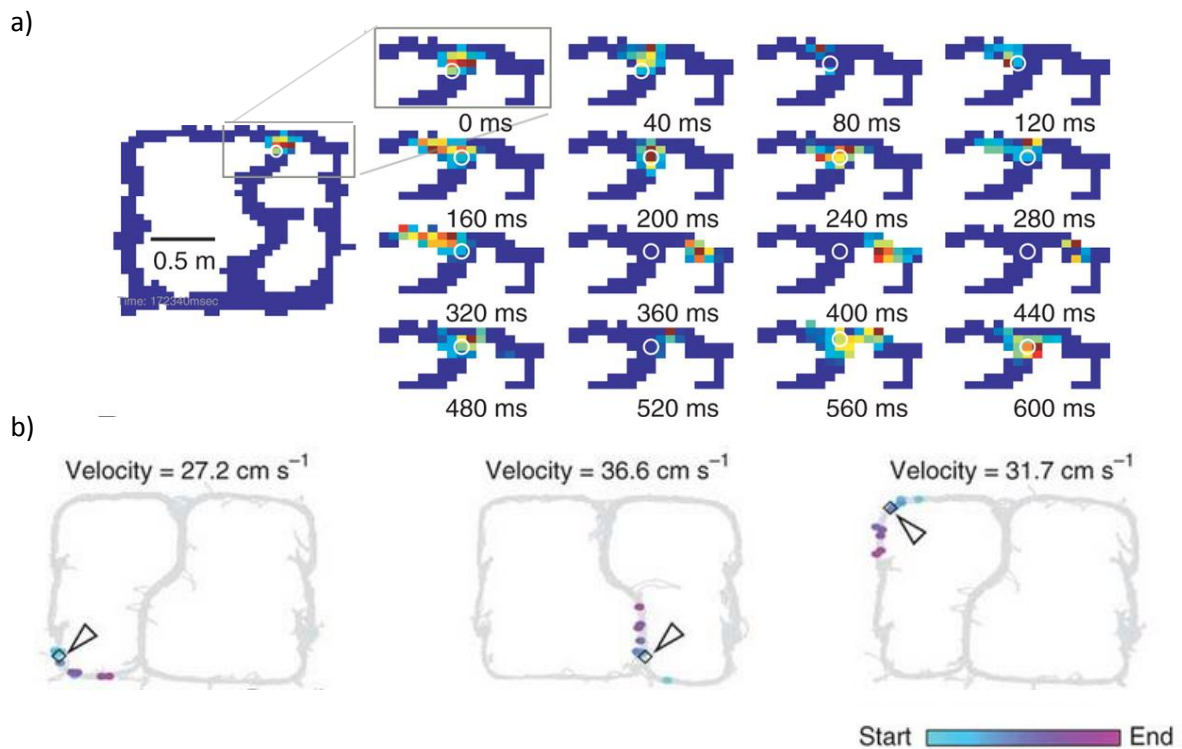
Figure 1.15. Hippocampal Ensembles 'Sweep' ahead of the Animal during Theta States a) On a spatial decision-making task, CA3 place cell ensembles represent future trajectories of the animal when the animal is approaching a choice point (white circle). Heatmaps represent the probability of position being represented, based on decoding place cell ensembles. Hotter colours indicate higher probabilities. Number below each heatmap shows duration of 'sweeps'. **b)** During behaviour on a double T-maze hippocampal ensembles represent trajectories between prominent physical landmarks. For each plot, coloured dots represent positions represented, cyan colours represent start of sweep and magenta colours the end. Triangle represents location of the animal at the time of the sweep. Panel a) is adapted from Johnson & Redish (2007) and panel b) from Gupta et al. (2012)

To conclude, a wealth of research implies out-of-field hippocampal activity, such as replay observed during SWRs and sweeps during theta states, serve an important function in behaviour and cognition. Replay seems like a plausible candidate mechanism for memory consolidation and, perhaps, also navigational planning. Moreover, sweeps of future place cell sequences during theta states may serve an important decision-making function during navigation. In this thesis I will describe two studies that provide further support to the suggestion that out-of-field hippocampal activity is functional; namely, out-of-field activity may underlie navigational planning in novel environments and that it may be important for successful goal-directed navigation.

2 General Methods

During my Ph.D., I used two different neuroscience techniques, namely, in-vivo chronic extracellular electrophysiological recording in rats and two-photon calcium imaging in head-fixed mice. Due to the distinct methods used for these techniques, this chapter will be divided into two subsections; the first covering methods used for electrophysiological recording and the second methods for calcium imaging.

2.1. In-Vivo Chronic Electrophysiology in Freely Moving Rats

2.1.1. Animals

Male lister-hooded rats (weighing 330–450g and 3-6 months of age at implantation) were used in all studies. Prior to surgery, animals were housed communally, 2-4 to a cage, with free access to water and food. After surgery the rats were transferred to individual Perspex cages (70cm x 45cm x 40cm), with free access to water and food sufficient to keep them at ~90% of their free-feeding weight. Animals were maintained on a 12:12 light-dark cycle, with two, one hour half-light periods simulating dawn and dusk. All procedures were approved by the UK Home Office, subject to the restrictions and provisions contained in the Animals (Scientific Procedures) Act of 1986.

2.1.2. Surgery and Electrodes

A total of 7[1] animals with either a single or dual, bilateral, hippocampal implants participated in the two electrophysiological studies. For the data presented in chapter 4, one surgery was carried out by Caswell Barry (CB). This animal had a dual hippocampal and medial entorhinal cortex (MEC) implant. The procedure for this implant was almost identical to that for the hippocampal implants, thus will not be described specifically. For the study described in chapter 4, all hippocampal implants were carried out by collaborators of the author (HFO). Namely, Hugo Spiers (HS), Colin Lever (CS) and Aleksander Jovalekic (AJ) (for details of implants see table 2.1 below).

---

[1] Four animals were used for the study described in chapter 3, and three for the study described in chapter 4. A considerably higher number of animals were implanted for the study in chapter 4 (n = 12). However for a rat to be able to participate in the study we required a minimum of 35 place cells to be held in parallel.

Table 2.1. Microdrive Implants and Operating Surgeons for all Implants used in Electrophysiology Studies.

Rat ID	Implant location	Hemisphere	Surgeon
R296	Hippocampus	Right and left	CL/HS
R346	Hippocampus	Right and left	CL/HS
R351	Hippocampus	Right	AJ
R1838	Hippocampus and MEC	Right and left	CB
R504	Hippocampus	Right and left	HFO
R505	Hippocampus	Right and left	HFO
R584	Hippocampus	Right	HFO

Each animal received one or two 'poor-lady' microdrives (Axona Ltd., St. Albans, UK) loaded with four-eight tetrodes (Recce & O'Keefe, 1989) of twisted 12-25μm HM-L coated platinum-iridium wire (90%-10%) (California Fine Wire, USA). The 'poor-lady' design allows tetrodes to be advanced through the brain in steps as small as 25μm, though tetrodes cannot be moved independently of one another. Prior to surgery tetrodes were cut level with one another.

Anaesthesia was induced and maintained with an isoflurane-oxygen mix (1.5-4l/min). The dose of isoflurane delivered varied depending on the animal and time point of surgery. Induction occurred in a chamber for around 3-5 minutes. After a surgical level of anaesthesia was achieved, animals were positioned via ear bars in a Kopf stereotaxic frame with Lambda and Bregma in the same horizontal plane. Anaesthesia was delivered during surgery via a mouthpiece in the stereotaxic frame. Prior to surgery animals received Baytril (0.1mL/100g, enrofloxacin) as a microbial prophylaxis and continued to receive Baytril in their water for 5-7 days following surgery (4.0mL enrofloxacin/100mL water). Moreover, all animals received a pre-operative analgesic, either Vetergesic (0.2mL, Buprenorphine 0.3mg/mL) or Rimadyl/Caprieve (0.1mL/100g, Carpofen 5.0%). Animals continued to receive the analgesic, delivered orally in the form of jelly, post-operatively for 2-3 days. During surgery, eyes were covered with Viscotear to prevent corneal damage, and the animal's body temperature was controlled via a heating pad the animals lay on throughout surgery.

Prior to incision, Betadine was used to disinfect the skin on and around the incision area. An anterior-posterior incision was made with a scalpel along the midline of

the animal's head. Skin and muscles were retracted with sterile surgical haemostats. The skull surface was then cleaned with sterile cotton buds and saline. Six to seven holes were drilled into the skull using a burr drill and a jeweller's screw was screwed into each of them. One of the screws, soldered to an electrical wire, was later connected to the microdrive to provide electrical grounding. One or two larger (1mm) holes were drilled using a trephine drill. Into these larger holes, the microdrives were implanted. For right hemisphere implants, the drive was inserted at 3.2mm posterior and 2mm lateral to Bregma. For left hemisphere implants (i.e. dual implant surgeries), drives were implanted 4.8mm posterior and 3mm lateral to Bregma (for the dual hippocampal-MEC implant surgery, the MEC drive was implanted at 1.2mm posterior and 4.5mm lateral to lambda over the right hemisphere, and the hippocampal drive was implanted at 3.5mm posterior and 2.2 mm lateral to Bregma over the left hemisphere, the hippocampal coordinates used for these surgeries were optimised so to fit two implants snuggly). The general procedure for implanting microdrives was similar for all surgeries. Briefly, the dura and pia were removed around the area where tetrodes were to be inserted. The tetrodes were then inserted into the brain with care being taken to ensure they did not bend. Once the tetrodes were in place in the brain (1.3-1.5mm below brain surface), a metallic guide cannula was lowered over the implant in order to protect the tetrodes, the brain surface was then covered with sterile Vaseline. The guide cannula and the feet of the Microdrive were then secured permanently to the skull using dental acrylic. The edges of the incision site were also sealed with dental acrylic to prevent infection.

All surgeries were conducted under sterile conditions – sterility was maintained by autoclaving surgical instruments prior to surgery, keeping tools in an ethanol solution during the procedure and placing them in a hot-bead steriliser for 15 seconds before use. The animals' body temperature and breathing was checked regularly and adjustments made as necessary.

2.1.3. Single-Unit Recording

Screening was performed post-surgically after one week recovery period post-surgery. An Axona recording system (Axona Ltd., St. Albans, UK) was used to acquire the single unit and positional data. The local field potential recorded from each of the 32 channels were passed through a RC-coupled, unity-gain operational amplifier mounted on the animal's head and fed to the recording system using lightweight wires. Each channel was amplified 20,0000 – 40,000 times, band pass filtered (500Hz-7kHz) and recorded differentially against a channel on a separate tetrode. Spikes exceeding a user-defined threshold were sampled at 48 kHz (50 samples from each of four channels) and time stamped with a 96 kHz clock signal. The position of a rat in the apparatus was captured using an overhead video camera to record the position of one LED on the animal's head-stage. The image was digitized and sampled at a rate of 50Hz to identify the rat's position. Head direction was inferred from the animal's trajectory.

Prior to recording, tetrodes were advanced in 50-200μm steps. Electrodes were advanced until ripples, high frequency (~200Hz) oscillations, were seen in the EEG (O'Keefe and Nadel 1978). At this point tetrodes were moved more cautiously, in smaller steps (25-50μm) until complex spikes were detected on the oscilloscope. To confirm complex spikes were emitted by place cells, single-units were recorded while the animal foraged in a 60cm x 60cm square box. Screening generally took place in the experimental room, on an elevated holding area. In between screening sessions, animals were returned to their home cage for at least 4 hours.

2.1.4. Spike sorting and Binning

Spike-sorting took place offline. Units were isolated manually using a custom-designed spike-sorting software (Tint, Axona Ltd, St. Albans, UK). Action potentials were assigned to putative cells based on plotting peak-to-trough amplitude of different tetrode channels against each other as well as temporal auto-correlation characteristics (see figure 2.1). Moreover, pyramidal cells were distinguished from interneurons based on waveform properties. Interneurons, identified by narrow waveforms (peak-to-trough duration < 0.4ms) and high firing rates (mean firing rate > 5Hz), were excluded from all analyses. Binning of spikes and position for ratemap generation differed for the two electrophysiology studies described in this thesis. Thus, I will refer the reader to the methods section of chapter 3 and 4 for details of this.
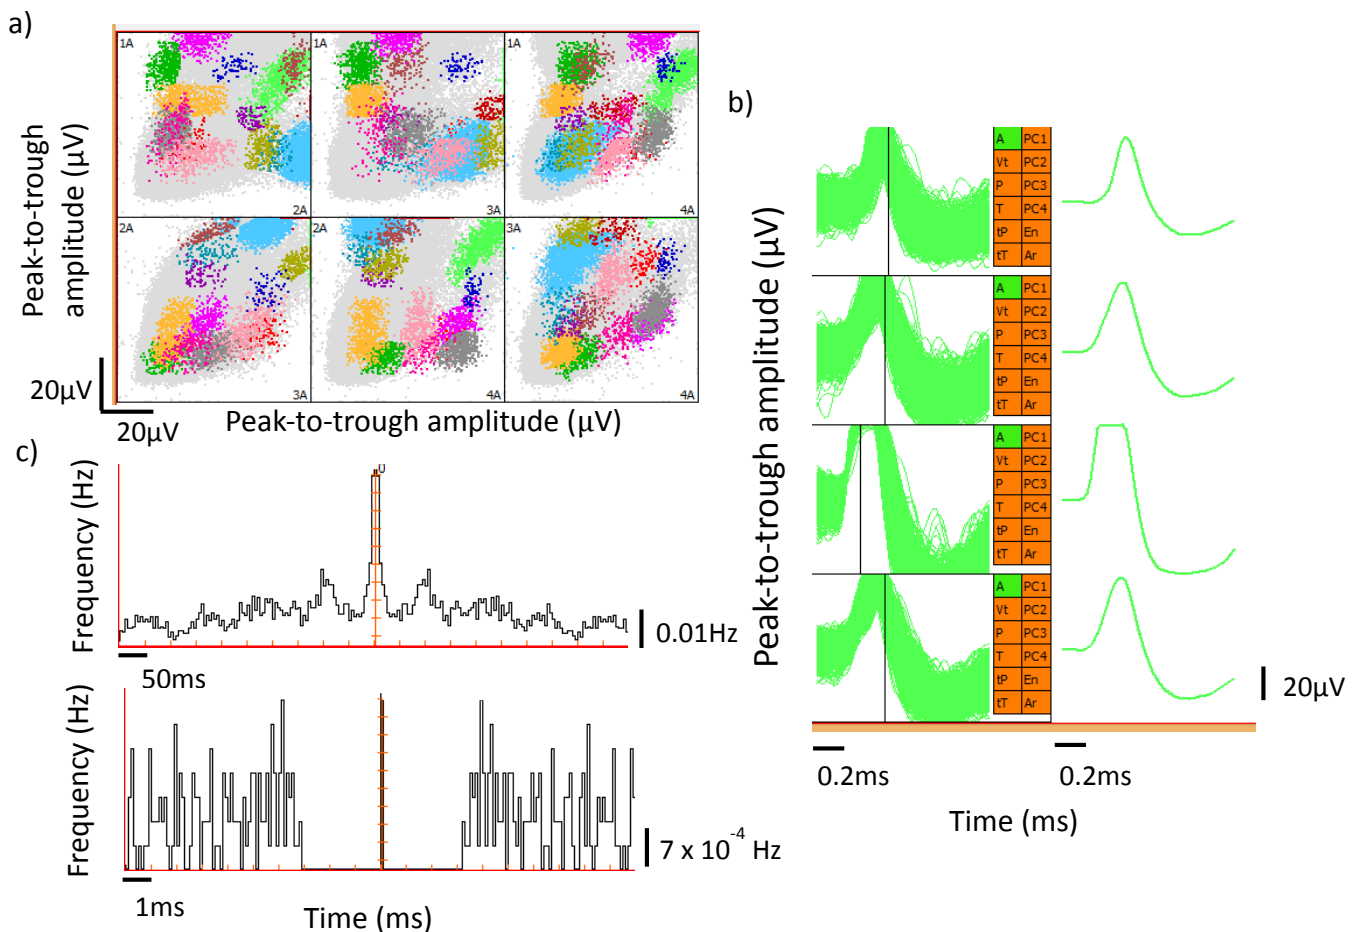
Figure 2.1. Examples of Spike Sorting from Tint Cluster Cutting Software. **a)** Peak to trough amplitude(µV) of each spike on each channel is plotted against its amplitude on each of the other channels. **b)** Waveforms (left) corresponding to multiple action potentials captured on four spatially separated electrodes (rows). Mean waveform for each channel is shown on the right. The waveform shown corresponds to the green cluster in panel a). **c)** Temporal autocorrelation corresponding to the green cluster in panel a). The autocorrelation is calculated for each pair of spikes and indicates the frequency with which a spike pair at a given lag (ms) is encountered. Upper panel: time range of -500:500ms, lower panel: time range of -10:10ms, showing the cell's refractory period. Theta modulation of approximately 8Hz is apparent in the upper panel.

2.1.5. Histology

After data collection, rats were anaesthetised (4% isoflurane and 4L/min O_2), injected intra-peritoneal with an overdose of Euthatal (sodium pentobarbital) after which they were transcardially perfused with saline followed by a 4% paraformaldehyde solution (PFA). Brains were carefully removed and stored in PFA which was exchanged for a 4% PFA solution in PBS (phosphate buffered saline) with 20% sucrose 2-3 days prior to sectioning. Subsequently, 40-50µm frozen coronal sections were cut using a cryostat, mounted on gelatine-coated glass slides and stained with cresyl violet. Images of the sections were acquired using an Olympus

microscope, Xli digital camera (XL Imaging Ltd.). Sections in which clear tracks from tetrode bundles could be seen were used to determine the location of cells recorded.

2.2. Two-Photon Calcium Imaging in Head-Fixed Mice

2.2.1. Animals and Transgenic Lines

Prior to surgery, mice were housed communally with free access to food and water. Mice were maintained on a 12:12 light-dark cycle. After surgery, animals were housed individually with free access to food. For experiments, animals were water restricted, generally receiving 1mL of water per day. All procedures were approved by the Janelia Farm Research Campus Institutional Animal Care and Use Committee.

Three different transgenic lines were used – namely, LSL-H2B-mCherry x emx1-Cre, gad2-NLS-mCherry x six3cre, PV-ires-Cre X ai9. The first, expresses mCherry in the nucleus of pyramidal cells, the second expresses mCherry in inhibitory neurons and the third tdTomato in parvalbumin expressing interneurons (tdTomato was not imaged). mCherrry was trafficked to the nucleus by engineering a mouse where mCherry was fused to the HSB histone tag. The expression of mCherry was under the control of Cre. Animals were typically crossed with emx1-Cre - resulting in expression in pyramidal cells, while for some they were crossed with six2cre which resulted in expression in inhibitory neurons. Moreover, mCherry was chosen due to its low toxicity and far red excitation wavelength that does not interfere with GCaMP imaging.

2.2.2. Surgery and Virus Selection

At 6-8 weeks of age mice underwent a chronic imaging window surgery[2]. Mice were anaesthetized with 1-2% isoflurane (by volume). A circular piece of scalp was removed and the underlying bone was cleaned and dried. A titanium head post was attached to the skull using cyano-acrylic primer (crazy glue). A craniotomy was made over the barrel cortex and 9-12 viral injections (20nL each) of AAV2/1 syn-GCaMP6s were delivered 450μm below the dura. Injections were made using a custom, piston-based, volumetric injection system. Glass pipettes were pulled and beveled to a sharp tip (outer diameter = 30μm). Pipettes were back-filled with mineral oil and front-loaded with viral suspension immediately before injection. Injections were laid out in a 300μm spaced grid centered 3.7mm lateral and 1.7 mm posterior to Bregma. Viral suspension was delivered at 15nL/min, followed by a 1-minute pause before withdrawal. After virus injection, the glass window was lowered into the craniotomy. The imaging window was constructed from two layers

---

[2] All surgeries were carried out by Simon Peron

(4.5mm and 3.5mm diameter) of standard microscope cover glass, joined together with an ultraviolet optical glue. The larger piece was made flush with the bone (by thinning the bone) and the smaller fit fitted tightly into the craniotomy. The window and head post were affixed with dental acrylic.

Prior to surgery animals received subcutaneous injection of 6mg kg^{-1} ketaprofen and a local injection of 0.5% Marcaine solution (injected under the scalp). Finally, for two days post-surgery animals received a general analgesia – a subcutaneous injection of 0.1mg kg^{-1} buprenorphine.

The reason for using the GCaMP6 virus is that the virus reports cytoplasmic changes in calcium ions by increasing its fluorescence in the presence of calcium (Chen, Wardill et al. 2013) – rise time of GCaMP6s ~0.2s, decay time ~0.6s. Neural activity causes rapid changes in intracellular free calcium (Baker, Hodgkin et al. 1971, Tank, Sugimori et al. 1988). Thus, GCaMP can track neural activity of all infected cells by reporting changes in intracellular calcium (i.e. often thousands of cells, depending on the size of the infection site). Moreover, as GCaMP6 is a genetically encoded calcium indicator (GECI) it allows for cell-specific targeting (e.g. only pyramidal cells) (Tian, Hires et al. 2009, Zariwala, Borghuis et al. 2012) and reliable imaging of cell populations over weeks (Huber, Gutnisky et al. 2012, Chen, Wardill et al. 2013); advantages that synthetic calcium indicators do not have. Moreover, the recent GCaMP variant – GCaMP6, has kinetics and sensitivity that rivals synthetic indicators while still maintaining a similar baseline fluorescence to previous variants (e.g. GCaMP5 and GCaMP3), allowing single-spike resolution and stable imaging over weeks (Chen, Wardill et al. 2013).

2.2.3. Imaging

Imaging carried out during laser-induced ablations was controlled by an older version of ScanImage (SI 3, scanimage.org), while imaging for all other experimental purposes were controlled by ScanImage 4 (SI4, scanimage.org). The system parameters differ slightly between the two versions, thus each will be described separately.

Imaging using ScanImage 4

GCaMP6s was excited using a Ti-Sapphire laser (Chameleon, Coherent) tuned to λ = 1000nm. Images were acquired using GaAsP photomultiplier tubes (PMTs, Hamamatsu) and a 16X 0.8 NA objective (Nikon). Green (GCaMP) and red (mCherry)

fluorescence channels were collected simultaneously. The voltage gain of the PMTs was set to 0.6V throughout all imaging sessions, to enable legitimate comparison of fluorescent traces from different imaging sessions. The gain setting was chosen to maximize signal-to-noise ratio (i.e. minimizing the influence of possible sources of noise, such as shot noise and dark current) while not causing the PMTs to trip. Horizontal scanning was accomplished using a resonant galvanometer (Thorlabs; 16 kHz line rate, bidirectional). Either four 600 x 600μm (512 x 512 pixels, 0.5μs pixel dwell time) imaging planes were acquired at 7Hz each (with axial motion controlled by a piezo collar (Physic Instruments)) or one 600 x 600μm plane imaged at 28Hz (512 x 512 pixels, 0.14 pixel dwell time). The system was controlled using ScanImage (scanimage.org). For volumetric imaging, three imaging planes 20-40μm apart were imaged sequentially, and the fourth fly-back frame was discarded. At the start of an imaging session, reference images was generated (by imaging 100 frames) and these were displayed during the imaging session, and positional adjustments were made to ensure that the same neurons were imaged. The scanning parameters were chosen to optimize speed and spatial extent of imaging while also diminishing the effect of possible noise artifacts, such as shot noise and dark current from the PMTs.

Imaging using ScanImage 3 System

GCaMP6s was excited at λ = 880/900nm using the same Ti-Sapphire laser as well as PMTs and objective. Moreover, green and red channels were collected simultaneously. Scanning was accomplished by the use of scanning mirrors. 350 x 350μm (256 x 512 pixels, 1.9μs pixel dwell time) planes were acquired at 4Hz. Images were acquired before and after ablation, and in some cases a day after ablation. As cell ablations resulted in dramatic increases in fluorescence from the stimulated region the gain of the PMTs was kept lower during ablations (0.3-0.4V) so to prevent tripping of the PMTs.

2.2.4. Image registration

To align frames from an imaging session, each experimental session underwent a motion correction pipeline. Firstly, images were registered on individual trials – rigid registration was performed using a down sampled fast Fourier transform (dFFFT; reference image, five consecutive frames of trial with minimal luminance change), followed by registration using a custom line-by-line algorithm. Second, registration was carried out across trials within a session. The mean of a trial with

stable luminance toward the middle of the session was used as the reference image for the imaging session. All other trials were aligned to this reference image using a rigid transform (dFFt), followed by a non-rigid interpolated warp field transform.

2.2.5. Selection of Regions of Interest (ROIs)

ROIs were drawn using a semi-automated algorithm based on a master reference image (the mean of all trial images). ROIs were matched across days using the warp-field transform. Drawing employed a custom user interface (MATLAB). First, we selected a point near the cell center. Then the algorithm constructed a matrix of intensities from the mCherry image (red channel); columns represented a range of distances from the selected point and rows spanned all angles around the point. The border of the nucleus was defined as the sharpest intensity contrast along varying distances that varied minimally. This process was repeated for the GCaMP image, but starting at the nuclear border in order to identify the border of the cytoplasm. Pixels between the two identified borders were assigned to the ROI. Once ROIs were drawn, corresponding peri-somatic neuropil ROIs were generated consisting of an annulus 3-13μm away from the outer edge of the ROI. Pixels with a correlation above 0.2 to an adjacent pixel as well as pixels belonging to non-neuropil ROIs were excluded from the annulus.

2.2.6. Fluorescence Extraction

In the presence of calcium (i.e. when a neuron fires an action potential) GCaMP6 increases its fluorescence. Moreover, the fluorescence increase is proportional to the number of action potentials emitted by the cell (Chen, Wardill et al. 2013). Thus, by measuring fluorescence from cells one can indirectly infer the activity of cells. Raw fluorescence was extracted for each ROI, and ROI fluorescence transients were neuropil-corrected ($F_{ROI} = F_{ROI} - F_{neuropil}$). If the corrected trace fell below a baseline fluorescence calculated for that ROI, the value was set to the baseline value, so as to prevent over-correction. Baseline fluorescence (F_0), was calculated using a 3-minute sliding window. For highly active cells, the 5^{th} percentile of raw fluorescence within the window was used as F_0. For inactive cells, the median was used. For cells displaying intermediate activity, an intermediate percentile was used. The reason for estimating baseline fluorescence of a cell was to account for differences in baseline fluorescence between cells and changes in baseline fluorescence/resting calcium levels for given cells between days. $\Delta F/F = (F-F_0)/F_0$

was then calculated for each ROI and used as a measure of neural activity in all activity-based analyses.

3. Ensemble Hippocampal Activity Varies as a Function of Distance to Goals and Predicts Navigational Performance

3.1. Introduction

Early human neuropsychological evidence implicated the hippocampus in the formation and consolidation of episodic memories (Scoville and Milner 1957). Patient H.M. had his hippocampi bilaterally removed as an attempt to relieve his epilepsy. Although his epilepsy symptoms improved as a result of the surgery, he was left with a profound memory impairment; he could neither form new memories nor retrieve the recent past. Studies using rodent electrophysiology have focused on the role of the hippocampus in spatial memory rather than episodic memory (O'Keefe and Nadel 1978). Spatial memory shares some core features with episodic memory, for example both consist of sequence encoding (Lisman 1999). Moreover, in some respects, spatial memory could be conceived of as a rudimentary form of episodic memory (O'Keefe and Nadel 1978); where places are tied together in a meaningful way to encode paths while episodic memories involve binding *both* space and time together in a unified construct (Buzsaki and Moser 2013). Following such a proposition some theoreticians have speculated that an episodic memory system evolved from a spatial memory system (Buzsaki and Moser 2013). Consistent with these ideas, various recent human neuroimaging and neuropsychology studies have shown that the human hippocampus is crucial for successful completion of spatial memory/navigational tasks (e.g. Spiers, Burgess et al.2001; Spiers, Burgess et al. 2001; Maguire, Burgess et al. 1998; Hartley, Maguire et al. 2003).

Early evidence for the role of the human hippocampus in spatial memory comes from the symptoms often seen in neuro-degenerative disorders. Namely, one of the important markers of Alzheimer's Disease (a neuro-degenerative disease involving the hippocampus) is impairments to both episodic and spatial memory (Kolb and Whishaw 1996). Moreover, neuropsychological studies of hippocampal patients have generally observed a deficit in topographical memory (Abrahams, Morris et al. 1999, Spiers, Burgess et al. 2001, Spiers, Burgess et al. 2001). Spiers and colleagues (2001) presented hippocampal patients with scenes from a fictitious town, using virtual reality (VR). The patients were then tested on their ability to recognize scenes from the town and their ability to construct accurate maps of routes between town destinations. Hippocampal patients performed worse than

controls on measures of route length, quality of drawn maps and recognition accuracy. However, the topographical impairment was accentuated in patients with damage to their right hemisphere (RTL). Moreover, on episodic memory tests, although the performance of all patients was deficient compared to controls, the deficit was greater in patients with damage to the *left* hippocampus (LTL). However, it is not clear whether the deficit in topographical memory performance was due to an impairment in semantic or episodic memory for spatial information. The initial presentation of the town might have enabled the formation of a semantic memory trace (due to formation of multiple overlapping memory traces), which could have supported the later navigation and map construction of the town. If semantic memory for the town was deficient in RTL patients this might have resulted in the poor performance on the topographical memory tests observed. However, such an explanation does not conceivably account for the deficit the RTL patients showed on some episodic memory tests, such as remembering the location of an object only experienced once. Such a one-shot memory task requires the retrieval of an episodic memory trace. Thus, although it is conceivable that topographical tasks can be solved by using semantic memory at least some tasks require the retrieval of an episodic memory trace; highlighting, again the overlap between spatial and episodic memory.

More recently, human studies have assessed the role of the hippocampus in navigation in healthy human subjects. Most of these studies have used functional-MRI (fMRI) as an imaging technique while subjects carry out a navigational task in virtual reality (VR). Although some early studies did not find the hippocampus to be significantly more engaged in spatial tasks requiring way-finding compared to route-following (Hartley, Maguire et al. 2003), more recent studies have shown hippocampal activity during navigation is significantly modulated by navigational performance and phase (Spiers and Maguire 2006, Brown, Ross et al. 2010, Howard, Yu et al. 2011, Sherrill, Erdem et al. 2013).

Hartley and colleagues compared hippocampal activity between a way-finding condition requiring flexible navigation and a condition where they had to simply follow a prescribed route. The hippocampus was not found to be significantly more active in the way-finding condition compared to the one involving route-following. However, the activity in the right posterior hippocampus significantly predicted trial performance as well as discriminated between good and poor navigators. Other studies have also found hippocampal activity to predict navigational performance

(Maguire, Burgess et al. 1998, Cornwell, Johnson et al. 2008, Cornwell, Arkin et al. 2012, Sherrill, Erdem et al. 2013). Cornwell and colleagues (2008, 2012), in a MEG study, found hippocampal theta during a virtual water maze task to predict navigational accuracy, indexed by path length and heading accuracy. Moreover, Cornwell (2012) found theta at the *start* of navigation to significantly predict heading accuracy. Recent fMRI studies have also found pronounced hippocampal activity at the start of navigation, for example (Spiers and Maguire 2006, Brown, Ross et al. 2010, Howard, Yu et al. 2011). Spiers and Maguire (2006) had participants perform a taxi-driver game while undergoing fMRI imaging. The only time point when the hippocampus seemed to be involved was at the start - after participants received a navigational request from a fictitious passenger (Spiers and Maguire 2006). A similar findings was reported by Brown and colleagues (Brown, Ross et al. 2010). Furthermore, Howard and colleagues (2011) found right posterior hippocampal activity to be linearly correlated with path distance to a navigational goal (more activity further away from goal). Moreover, when participants were forced to take a detour, the hippocampus was found to be engaged if the de-tour resulted in a large change in path distance (Howard, Yu et al. 2011).

How does one reconcile these performance and start-related effects? Perhaps, the reason the hippocampus is active at the start of navigation is because that is when subjects are engaged in route-planning. Spiers and Maguire (2006), in their taxi-driver study, indeed found hippocampal activity to correlate with customer-driven route-planning. Moreover, quality of route-planning presumably relates to navigational performance.

On the face of it, the pattern of hippocampal activity in humans appears to be quite different to that observed in rodents. Place cells are active irrespective of the navigational contexts, even when the animal simply needs to run up and down a linear track (Wilson and McNaughton 1993), or forage for food in an open arena (O'Keefe and Dostrovsky 1971). However, as mentioned in chapter 1, some recent theories have proposed that synchronized population firing observed during periods of awake immobility play a particular role in navigational planning (Johnson and Redish 2007, Foster and Knierim 2012, Pfeiffer and Foster 2013). At choice-points in a double T-maze task Redish and colleagues have found hippocampal ensembles to transiently represent routes ahead of the animal (Johnson and Redish 2007, Gupta, van der Meer et al. 2012). Gupta and colleagues found such 'sweeps' represented trajectories extending up to proximal salient navigational landmarks in

an environment. Moreover, Pfeiffer and Foster (2013) found replay during awake immobility to anticipate the future goal-directed path of an animal.

Thus, two hypotheses can be drawn from the discussed findings: 1) at navigational time points requiring route planning the average population rates of CA1 place cells should be higher than at other points during active navigation, 2) such population activity might be predictive of navigational performance. In this chapter, I will describe results from an analysis carried out on data collected from a study where CA1 place cells were recorded from rodents carrying out a goal-directed navigational task requiring the hippocampus (i.e. 'event arena' (Day, Langston et al. 2003))[3]. During this task, population single-unit activity was found to vary positively and linearly with distance to goal. Moreover, the degree to which population activity was modulated by goal distance was found to significantly predict navigational performance. These findings bridge rodent and human hippocampal studies. Moreover, they may lend support to the idea the hippocampus may be involved in route planning, but highlight the need to study the content of population activity at navigational points where route-planning is likely to occur.

---

[3] This data was collected by Dr. Hugo Spiers, but spike-sorting and all analyses were carried out by the author

3.2. Method

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

3.2.1. Animals

Three male Lister Hooded rats (300-450g in weight) at three to six months of age at implantation were used in this study. Two animals had dual hippocampal implants, each hemisphere containing a 16-channel microdrive. One animal had a single 16-channel implant over the right hemisphere. The number of place cells recorded in a session varied greatly (range: 7 – 45). However, on average 26 cells were active in a session (see table 3.1 for details of maximum number of cells recorded per session)[4].

3.2.2. Apparatus

The experiment took place in the 'event arena' apparatus (Tse, Langston et al. 2007). The event arena was made of wood painted with matte black paint and was surrounded by a set of fixed distal cues (see Fig 3.1). The 80cm x 80cm platform in the event arena was at a height of 100cm in the room and covered with sand approximately 1cm deep. The edges of the platform had 14 cm high walls such that the distal cues were clearly visible to the rats. The arena floor had 4 locations in which sand-filled tubs ('goal-wells', diameter 11cm, 6cm deep) were placed into holes cut in the arena floor such that the top of the open tub was flush with the floor of the arena. Two steps were taken to reduce the possibility that rats were choosing the well by using olfactory cues. Each had a wire mesh grid half way down, below the grid the wells were filled with sand and a ground up mix of the two food types used in the experiment. In addition, the sand in the wells was mixed with an equal amount of the two foods. Above the grill the tub was filled with sand to the top. The sand in the wells was at a level with the sand in the arena floor, such that they were not visually identifiable (see Fig 3.1). Four 'start boxes' were located at each corner of the arena (see Fig 3.1). These were 30cm high, 10cm wide and 30cm in length. Each one contained a distinct visual cue (white card, white card with a

---

[4] Other place cell data collected from these animals has been published in Spiers, H. J., R. M. Hayman, A. Jovalekic, E. Marozzi and K. J. Jeffery (2013). "Place Field Repetition and Purely Local Remapping in a Multicompartment Environment." Cerebral cortex.

diagonal, black card, or white card with a large back circle, see Fig 3.1). Each had a door with a grill cut with 1cm wide bars, so that the rat could view the arena when in the start box. Between every session sand was moved thoroughly around the apparatus and mixed with new sand and the wooden painted parts of the apparatus wiped down with water.   A holding platform (10cm x 10cm) was located 1m east of the arena, at the same height as the arena.

Figure 3.1 The Event Arena Apparatus A picture of the experimental apparatus. The arena itself is covered with sand, obscuring the goal wells. Each corner of the arena has a start-box with a unique cue clipped inside of it.

3.2.3 Training

Habituation

The experiment was initially designed to adapt the protocol of Tse and colleagues (2007) for place cell recording. Specifically, Tse at al. trained animals to associate six different flavours with different behavioural goals in an open arena. However, initial pilot work showed the animals took long (> 3 weeks) to perform reliably above chance thus the task was simplified to a delayed-match-to-sample task where animals had to learn to associate two different flavours with two different navigational goals.

For one rat habituation started prior to surgery (R296), for the other two it occurred after surgery (R346, R351). Each rat was tested from start of the protocol to the end before the next rat began the experiment (order R296, R346, R351). On the first day, rats were allowed to explore the arena for 15 minutes freely. At the end of the session they were placed in each of the four start boxes for 30 seconds each. The next session involved placing a sand-filled tub (11cm diameter, 6cm high) on

top of the middle of the arena floor. A food not used in the recording session ('Curiously Cinnamon' squares, Nestle ©) was placed in the tub. The rats were allowed to learn to dig from the tub for the food for 15 minutes. After this the tub was removed and the food was placed protruding from the sand in the four goal-wells. The rat was allowed to travel between the four locations to dig for the food, with the food being submerged in subsequent visits, such that the rat learned to visit each of the goal-wells. Rats were habituated to the food used for training and the other two types of food (chocolate pellets and banana flakes), used in the experiment, by placing a small quantity of each (< 2g) in their home cage. Chocolate pellets and banana flakes were chosen because pilot work showed they were equally desired.

Acquisition

Prior to performing the task rats were initially trained to a high level of performance (>80% correct) in sessions with a single goal, before introducing sessions combining two different goal locations.

3.2.4 Experimental Procedure

During a single session (Figure 3.2), there were two goal locations, only one of which was correct in a given series of trials. A 'goal trial' began with the rat being placed in the start box and a particular flavoured food being delivered into the start box. After the rat consumed the food a 30-second delay began, after which the experimenter opened the door of the start box (see figure 3.3). In order to obtain the reward, the rat had to proceed to the correct goal-well and dig there to uncover more of the same food that was delivered into the start box (as in Tse et al., 2007). The other three goal-wells were not baited in a given goal trial. Navigation to two different goal locations was tested during a session (e.g. South, West, South, and West). A trial was scored 'Correct' when the rat's first dig occurred at the goal-well. Start box location was randomized with each set of four trials (e.g. 1,2,3,4; 4,3,2,1). Each session began randomly with the south goal or west goal being baited. The south goal was associated with chocolate and the west goal with banana, for all rats. After rats had successfully navigated from each of the start boxes to the baited goal-well, that goal-well was no longer baited and the other goal-well was baited. In between trials the rat rested on the holding platform for approximately 15-30 seconds.

Two 15-minute 'baseline' foraging trials were given, one BEFORE the navigation task and one AFTER (see figure 3.2), during which rats searched for rice scattered throughout the arena. CA1 place cells were recorded during goal trials and during these baseline foraging trials.
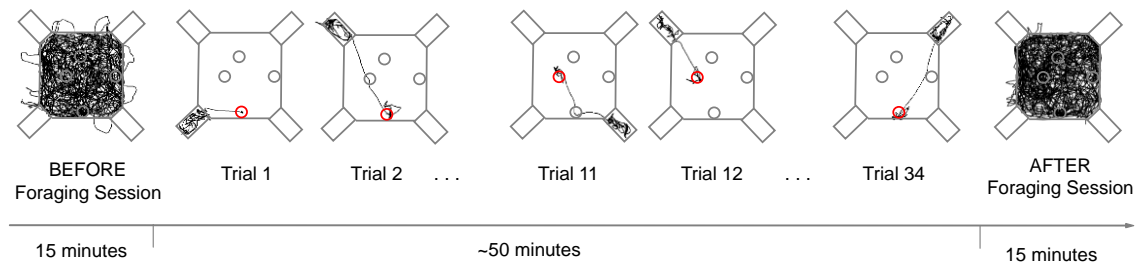
Figure 3.2 Session Timeline Each session began with a 15-minute foraging session, followed by 14-40 goal-trials. During the goal trials animals ran from one of four start boxes to either the south or west goal wells (marked in red), depending on the food they ate while in the start box. Following goal-trials, the animals completed another 15-minute foraging session. Event arena walls are marked in grey. The rat's trajectory is marked in black. Circles indicate the position of goal-wells. The red circle indicates the rewarded goal in each trial.
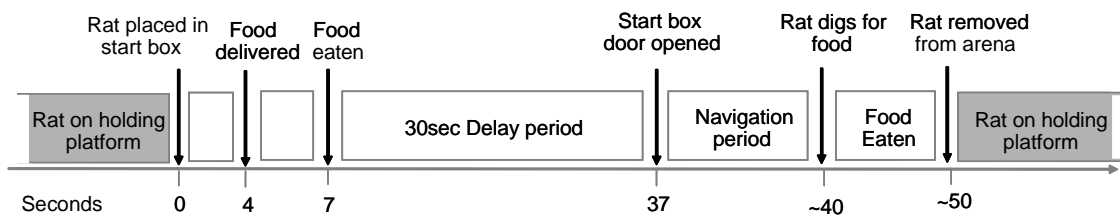
Figure 3.3 Trial Timeline A schematic of a trial time line. Before a trial began, the rats sat on a holding platform. From there they were moved to one of four start boxes, in which they received food (chocolate or banana). It generally took the animals 3-4 seconds to realise the food had been delivered into their start box and thus, to start eating it. Upon food consumption, a 30second delay period began, after which the start box was opened and the animal had to navigate to the goal-well associated with the food it received in the start box ('Navigation period'). Once the animal had dug for food, the rat was removed from the arena and put back on the holding platform. Each trial lasted about one minute.
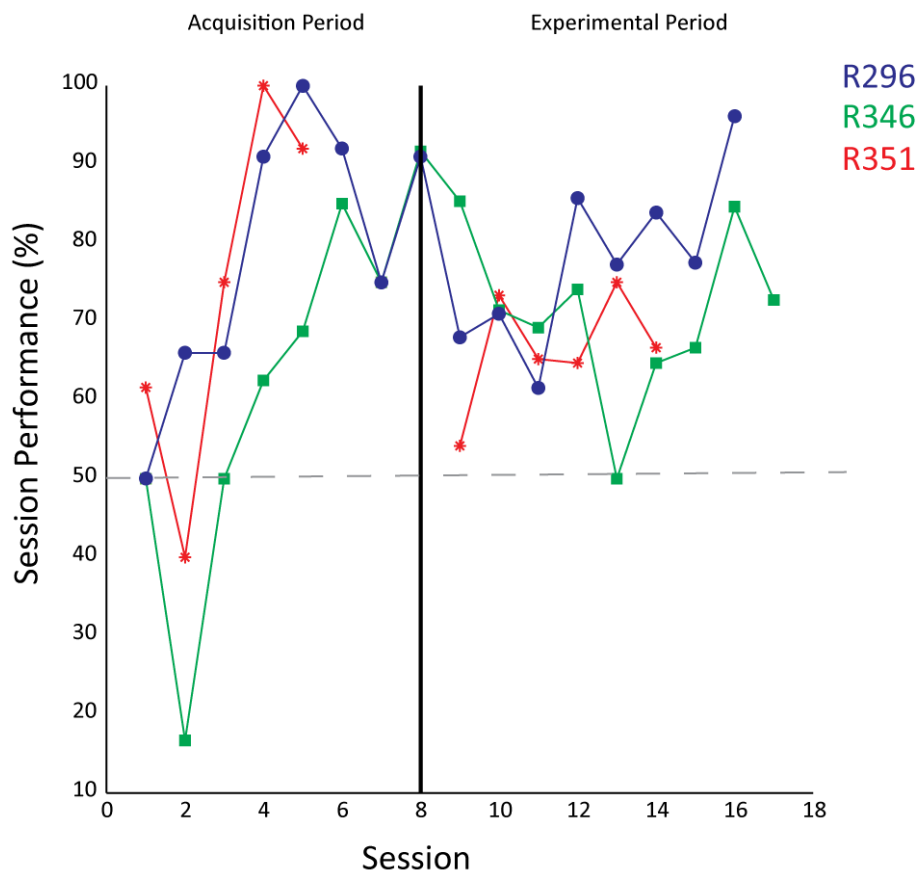
Figure 3.4 Learning Curves Learning curves for the three rats used for this study. The session performance (number of trials where animal ran to correct goal well divided by total number of trials in a session) is plotted against session number. The first eight sessions represented acquisition sessions where animals only had to run to one goal well (either South or West). For R351, the acquisition period was shorter (five session). Following acquisition (>80% performance), the animals carried out the experimental protocol where they either ran to the South or West goal well, depending on the food they received in the start box. The grey dashed line represents chance performance, the black vertical line the separation between the training and experimental period. Y-axis shows session performance (%) and x-axis the session ID.

3.2.5 Session and Trial Selection

We only considered sessions where the animal performed above chance levels (50%, assessed using a binomial test). Here, session performance was defined as the number of times the animal ran to the correct goal well (either west or south) divided by the total number of trials in a session. Animal R351 only had five sessions where it performed above chance. To ensure each animal contributed equally to the analysis, we chose five sessions for the other two animals. The criteria we used to select these sessions were a) the animal completed at least 20 trials, b) the sessions chosen needed to represent a range of performance scores.

Only trials where the animal's trajectory concluded at one of the two goal wells (south and west) and where it reached its destination within 15 seconds were included in the analysis. Following this criterion, on average, 85% of all trials in a session were included in the final analysis. A total of 364 trials were included in the analysis. See table 3.1 below for information regarding number of cells, performance and trials for each session.

Table 3.1 Session Statistics

Rat	Session	Max cells	Performance	# Trials	# Trials included	South trials included	West trials included
296	9	23	68	30	23	9	14
296	10	21	70.97	31	27	17	10
296	12	18	85.71	21	19	11	8
296	15	10	77.5	40	38	16	22
296	16	7	96.16	26	24	8	16
346	9	24	85.29	34	32	9	23
346	12	44	74.07	27	25	10	15
346	15	43	66.67	39	33	10	23
346	16	45	84.62	27	27	10	17
346	17	42	72.73	19	19	5	14
351	10	22	73.33	30	16	7	9
351	11	25	65.22	23	16	7	9
351	12	23	64.71	34	25	13	12
351	13	19	75	16	13	2	11
351	14	19	66.67	36	27	14	13
mean		26	75.11	29	24	10	14

3.2.6 Analysis

The principal analysis is based on the goal trial data – when the animal is actively navigating to its goal ('Navigation period'). We did not analyse activity during the delay period in the start box. Although one might hypothesise that place cell activity during this period might relate to the animal's future trajectory (i.e. sharp-wave ripple activity), we were unable to analyse position representation during this period due to low cell yield. Namely, to decode position from SWR activity for an open environment one needs 150-200 simultaneously recorded neurons (Pfeiffer and Foster, 2013). Thus, we focused the analysis on the 'Navigation period'.

Trial Definition

Trials started 250ms before the animal crossed the boundary of the arena from one of the four start boxes. We chose this definition for the start of a trial as the data showed the animals started accelerating while they were still in the start box, immediately preceding their entry into the arena. Moreover, we found 250ms captured the initial acceleration while not including many time points when the animal was immobile. We tried using a longer time window (i.e. 500ms), however this resulted in the inclusion of slow movement samples. A trial ended 500ms after the animal entered the goal zone of its destination goal well. The goal zone was a 20cm x 20cm square centred on the goal well (diameter of goal wells = 11cm). 500ms was chosen to limit the inclusion of samples where the animal was moving slowly, but also long enough to include samples when the animal had reached its destination. Position samples when the animal's speed was less than 3cm/sec were excluded. The average trial duration was 2.44sec (SD = 1.58sec). See figure 3.5 below for two examples of representative trial trajectories. To note, position estimates were smoothed with a boxcar kernel (400ms long).
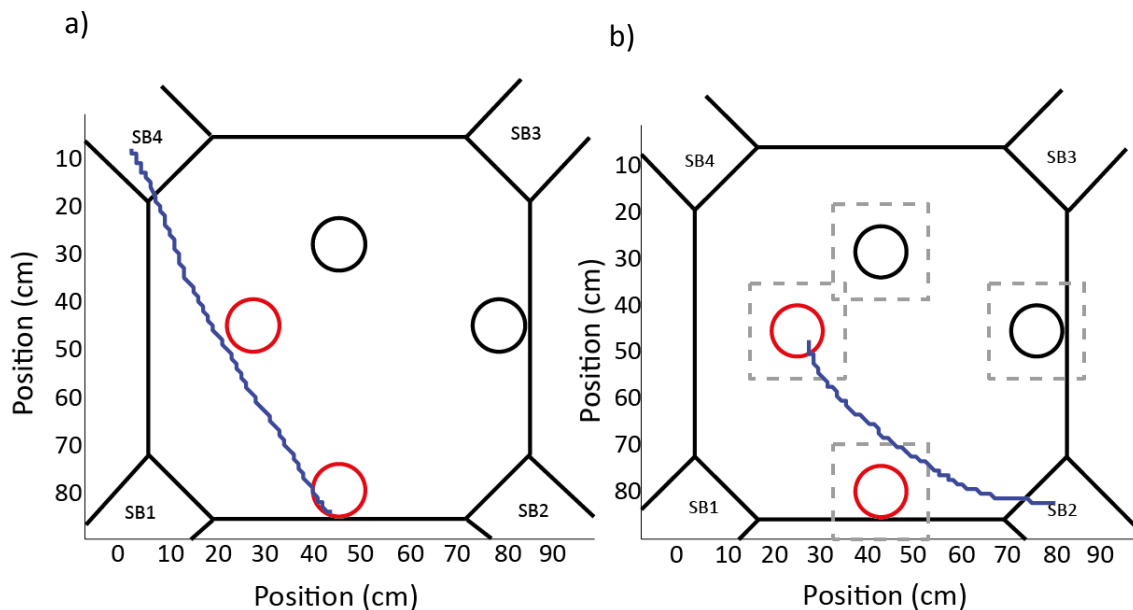
Figure 3.5. Example Trial trajectories a) Example of a trial trajectory (blue) that ends at south goal well. Red circles represent goal wells, black circles unbaited wells. Dotted lines around goal wells represent goal zones. Trial start was defined as 250ms prior to the animal crossing any of the entries into the arena (diagonal corners). A trial end was defined as 500ms after entry into goal zone. **b)** Same as a) except for a different goal trial ending at the west goal well. X and Y axis represent dimensions of arena in centimetres. SB = start box

Spatial Binning

To analyse the relationship between firing rates and distance to goal, the position vector of a trial trajectory was first normalised so that 0 represented positions at goal and 1 the start of the trajectory. The normalised position vector was then divided up into 20 equally sized distance bins (bin size = 0.05 normalised distance units). See figure 3.6 for a demonstration of binning. The indices of the position samples in each distance bin were used to bin the rates, speed and acceleration samples. As bins were defined in terms of distance travelled to goal, the time the animal spent in each bin varied. The average time spent per bin was 116.32msec (SD = 101.12msec).
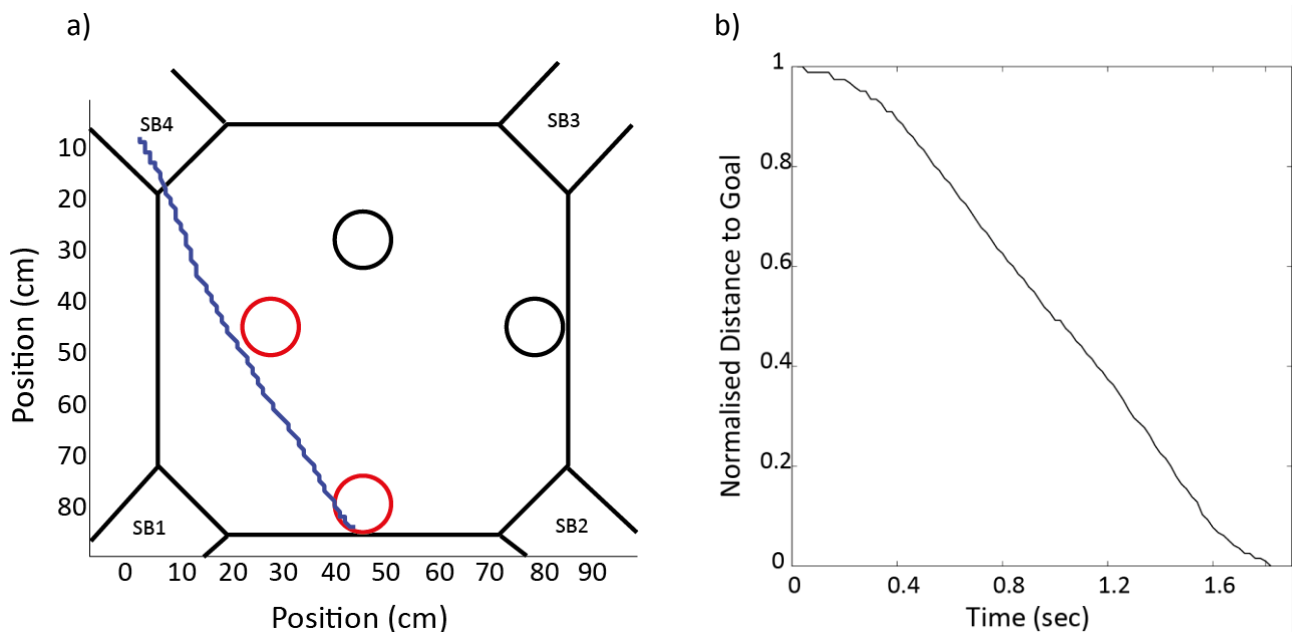
Figure 3.6 Spatial Binning of Position a) A representative trial trajectory (blue) showing in figure 3.5a overlaid on event-arena. **b)** Normalised trajectory shown in a), now position farthest from goal has a value of 1, and position closest a value of 0. X-axis shows time in seconds over a trial, y-axis normalised position.

Goal Trial Analysis

In the goal trial analysis we wanted to estimate population activity while an animal navigated to a goal well. Consequently, estimates of population activity in each goal trial included both cells active in that trial as well as cells that may have been silent in that trial but active in other trials. However, to ensure silent cells did not exert too much weight we excluded cells that fired less than 10 spikes in a given goal trial session (on average 29% of all cells). To note, the results of the principal analysis were the same with or without cells firing less than 10 spikes (population

rates x distance to goal correlation: r = 0.74, including inactive cells: r = 0.75).

Moreover, any samples where the rat's velocity was less than 3cm/sec were not included in the analysis. We estimated firing rate (Hz) in each 20ms position bin from trial start to trial end. Rates were re-binned into 20 spatial bins, according to the procedure described above. Namely, the indices for the positions in each spatial bin were used to assign rate samples to different spatial bins. We then computed a population activity vector by taking the mean rate at each spatial bin. Finally, to ensure trials contributed equally, we divided the rate in each distance bin by the maximum firing rate in a goal trial. To assess the relationship between population rates and distance, population rate vectors were correlated with distance to goal using the Pearson correlation coefficient (r). A correlation was deemed significant if it had an associated probability below 0.05. For the principal analysis, population vectors for all 364 trials were averaged to get one activity vector which was correlated with distance to goal. This analysis was then broken down to look at animal/session differences, and goal-well differences (i.e. trials ending at west vs south goal). See figure 3.7 for an example of population activity for a single trial, before and after normalisation of distance to goal.
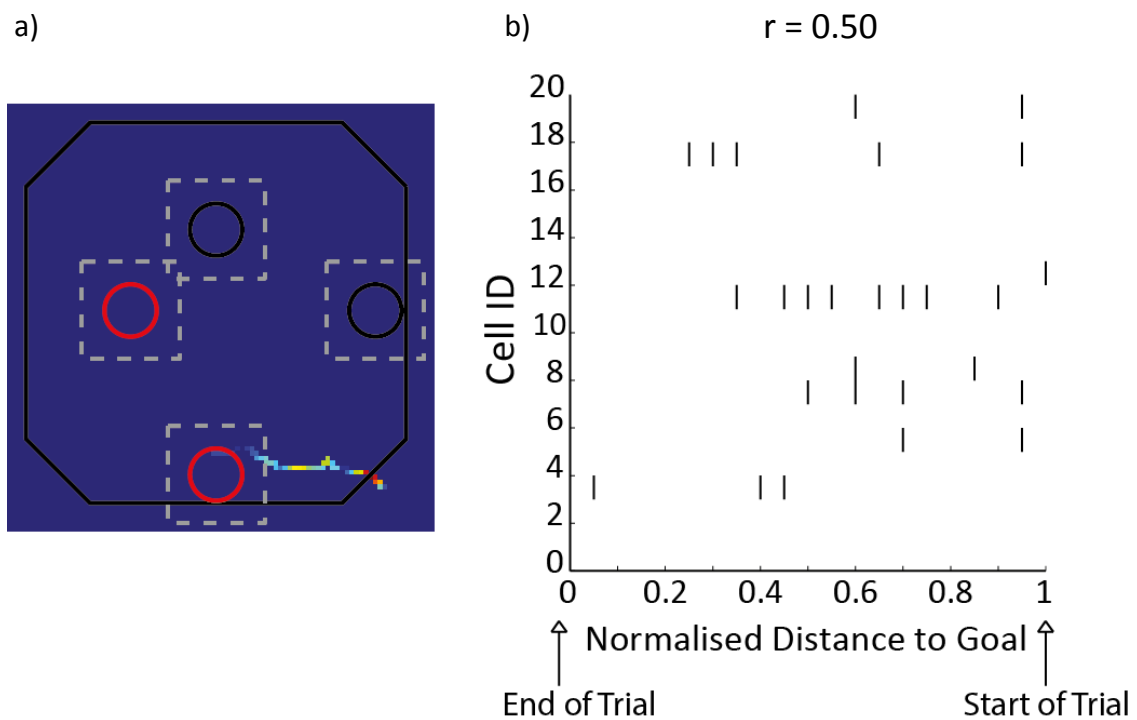
Figure 3.7 Example of Population Rate Binning a) Population activity heatmap for a goal trial. Hotter colours represent higher activity rates. In this example the rates are clearly highest at the start of a trajectory. **b)** A raster plot of cells recorded from the trial shown in panel a), note spikes are plotted against their normalised distance to goal not time. The raster shows that more cells tend to be active at longer distances from the goal. The

correlation between population rates and distance to the goal is stated above the plot. Data from rat R346, session 9, trial 2.

Speed and Acceleration

For each trial instantaneous speed and acceleration were estimated at every 20ms time bin (based on smoothed position estimates). Acceleration was defined as the change in speed between consecutive speed samples. Similar to rate, speed and acceleration were re-binned into 20 spatial bins, and correlated with distance to goals using the Pearson correlation coefficient.

Foraging Trial Analysis

To assess whether population rates during goal trials reflected the distribution of place fields in the recording environment or 'out-of-field' place cell activity we analysed activity rates during the foraging sessions. The data from the foraging sessions that preceded and followed the goal trials were combined and used to control for position in the analysis looking at the relationship between population rates and distance to goal. That is, if population rates during trial trajectories vary as a function of distance to goals, one might suggest that this could result from place cells clustering around the start boxes. This would then create an artificial positive correlation with goal distance. Ratemaps for all recorded cells were generated by binning positions and spikes into 2cm spatial bins, smoothing spikes and position (i.e. dwell time) separately with a Gaussian kernel (σ= 2cm). Although low firing cells were excluded from the goal trial analysis we included them in the foraging analysis since the majority of cells met the activity threshold criteria for the goal trial analysis and the principal results of that analysis were the same with them included (r = 0.75 vs r = 0.74). Firing rate for each ratemap bin was obtained by dividing the smoothed spikes by the smoothed dwell time. Once ratemaps for every cell had been generated, they were averaged (for both before and after foraging sessions) to create one population ratemap per session (see Figure 3.8). For each goal trial, the positions sampled were extracted from the population ratemap, for the session the goal trial belonged to, maintaining the order in which positions were sampled in the goal trial (i.e. 'foraging trials'). This gave the *expected* population vector (see figure 3.9a-b). The expected population rate vector from these foraging trials was re-binned into 20 distance bins, similar to the trial rates. We then estimated the relationship between distance to goal and the expected population rates, obtained from the foraging trials, using a Pearson correlation. If

the correlation between distance and rates in the foraging trials is similar to the one we obtain for that of the goal trials this would indicate the relationship between the population rates and distance in the goal trials is confounded by place field distribution.
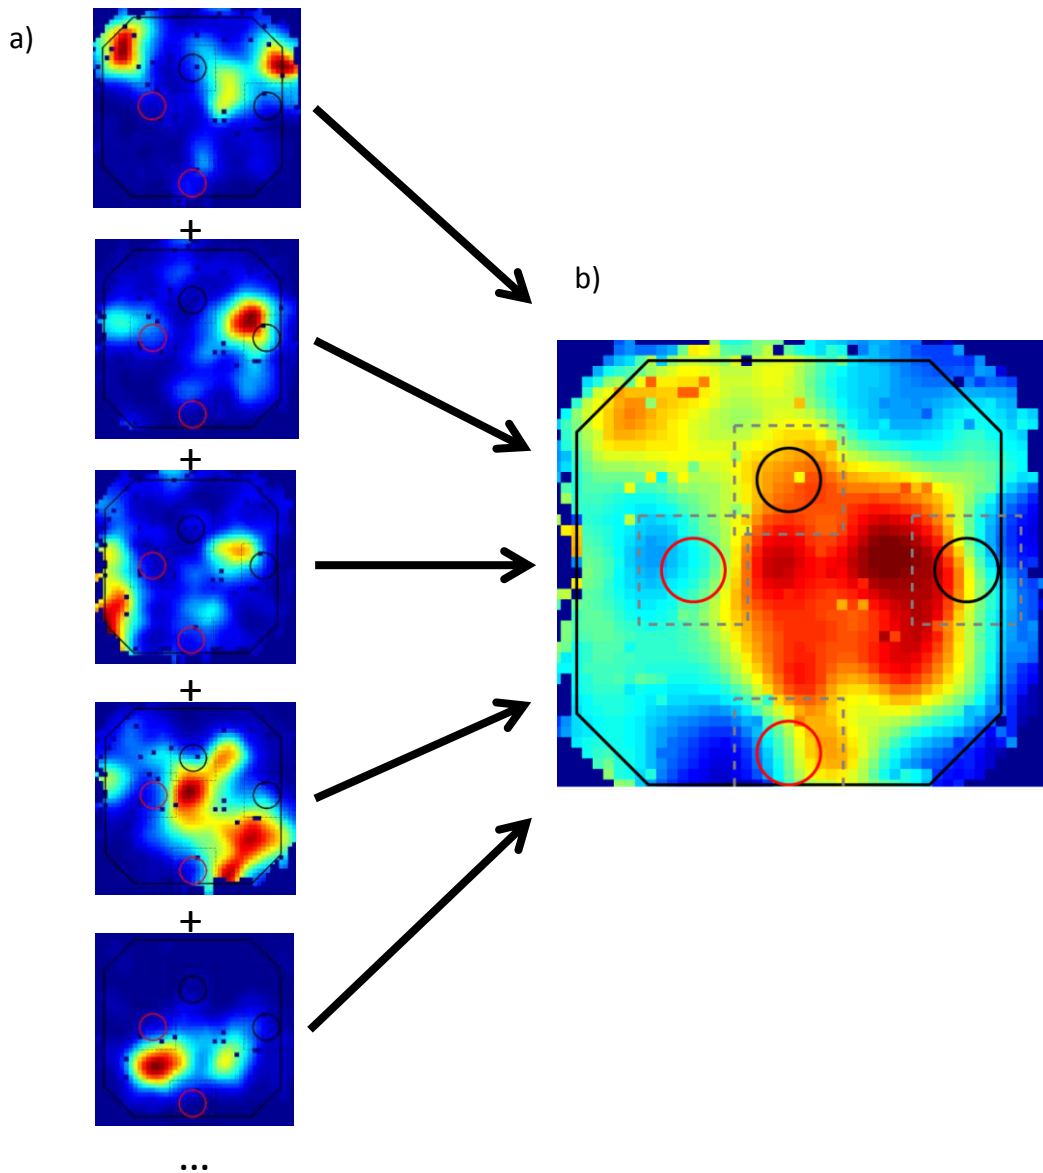
Figure 3.8 Population Ratemap a) Ratemaps for five place cells recorded in one foraging session (session 9, rat R346). A total of 24 cells were recorded in this session. (**b**) A population ratemap for session 9, from rat R346. The population ratemap is obtained by averaging the ratemaps of all cells recorded in the environment in that session. The five ratemaps shown in panel a) contribute to the ratemap in panel b.
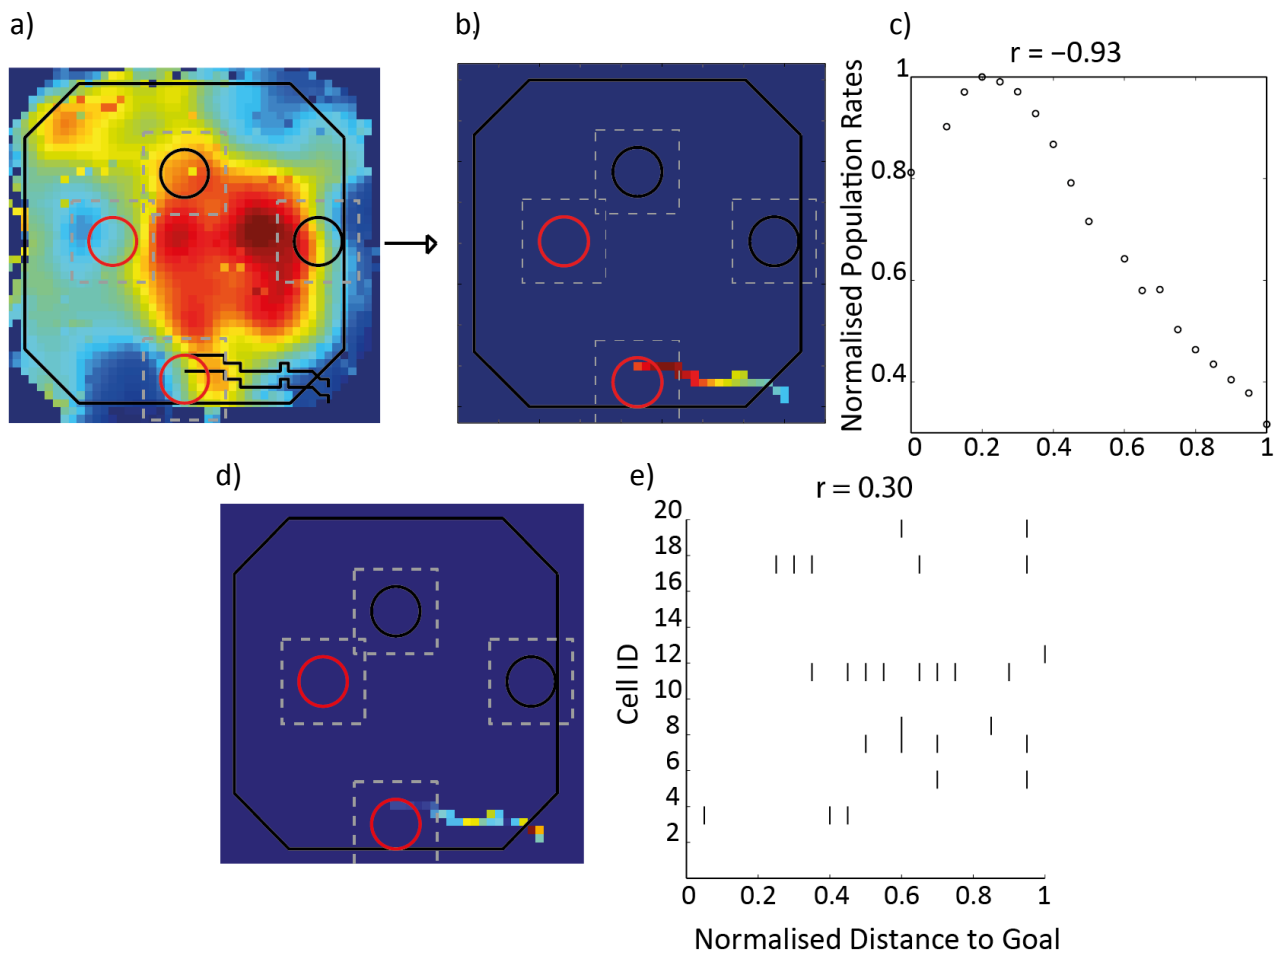
Figure 3.9 Predicted Modulation of Population Rates by Distance to Goal a) An example of a population ratemap from rat R346, session 9. Positions sampled for a given goal trial in a session were extracted from the population ratemap (black lines drawn in the bottom right corner). **b)** A predicted population rate vector was extracted, and rebinned into 20 distance bins. **c)** Expected population rates for a trial plotted against normalised position to goal. The title of the scatterplot shows the correlation between distance to goal and expected rates. The relationship is strongly negative. X-axis shows normalised distance to goal, y-axis average population rates. **d)** Heatmap for population rate obtained for the goal trial in a-c. **e)** A raster plot of cells recorded from the goal trial shown in panel d), note spikes are plotted against their normalised distance to a goal not time. The raster shows that more cells tend to be active at longer distances from the goal. The correlation between population rates and distance to the goal is stated above the plot. In the goal trial**,** the relationship is clearly positive, and quite different from the predicted relationship based on place field distribution (show in panel c). A raster plot of foraging trial rates could not be generated as rates in each spatial bin are based on *average* activity in a spatial bin, over multiple runs through a bin. Data from rat R346, session 9, trial 2.

Performance Analysis

To examine whether the relationship between population rates and distance to goal is modulated by navigational performance, each session was divided in two – based on which goal was baited in each trial (i.e. 'session blocks'). That is, all trials when the south goal was baited belonged to one session block and all trials when the west goal was baited belonged to another session block. Since a total of 15 sessions were included in this analysis, dividing sessions into session blocks resulted in 30 blocks. We then estimated performance in each session block by dividing the number of correct trials in a session block by the total number of trials in a block. Moreover, we estimated the correlation between population rates and distance to goal for each session block. To do this, we averaged the re-binned population rates for all trials in a session block to obtain a session block population vector and then correlated this with distance to goal, using a Pearson correlation, as before. We then correlated the obtained Pearson correlation coefficients for each session block against performance. We repeated this procedure to estimate the relationship between performance and speed-by-distance correlations and acceleration-by-distance correlations.

3.3. Results

3.3.1 Behavioural Observations

Once animals had demonstrated learning during the acquisition period (when only one goal well was baited in a session) they consistently performed above chance when they carried out the alternation version of the task where either of the two goal wells could be baited in a trial ('delayed-matching-to-sample'), during the experimental period. Data from trials completed during the experimental period were used in all analyses.

A total of 364 trials were included in this analysis. As described in section 3.2, only sessions where the animals performed above chance level were included (50% correct). When performance for the two goal wells – south and west, are analysed separately (i.e. data divided into 'session blocks') we observed a slightly higher mean performance for the south goal well - 85.35% (SD = 12.18), than the west goal well - 73.26% (SD = 18.35) - $t_{(28)}$ = 1.95, p = 0.06. Although the difference in performance for the two goal wells was not *statistically* significant, the performance is numerically higher for the south goal. Perhaps the performance for the south goal was a bit higher as it was located close to the exit of two of the start boxes. To ensure this difference did not confound our results, we carried out our main analyses for south and west goal trials separately.

3.3.2. Population Rates Vary Positively with Distance to Goal

To address the experimental hypothesis of whether population activity in the hippocampus is higher at navigational points where route planning is likely to occur, we assessed the relationship between population rates in a goal trial and distance to the destination goal well. Namely, distance to a goal was normalised and normalised population rate at each 0.05 distance unit estimated (see section 3.2.6. for description of methods). Figure 3.10 below shows examples of population activity for several trials recorded, revealing a positive relationship between population rates and distance to goal. Moreover, we observed a significant, strongly positive, correlation between population rates and distance to goal when data from all trials was combined - r = 0.74, p = 1.40 x 10^{-4}. Indicating rates are higher early in goal-directed navigation, and that they drop progressively as the animal moves closer to its goal (see figure 3.11a). Importantly, a positive linear relationship between population rates and distance to goal was observed when

trials going to south and west goals were analysed separately (see Figure 3.11b-c) – south: r = 0.65, p = 0.0016, west: r = 0.68, p = 7.09 x 10^{-4}. Furthermore, when data for each animal was analysed separately we obtained a positive correlation for all three animals, although only the correlation for animal R346 reached statistical significance (R296: r = 0.42, p = 0.057, R346: r = 0.76, p = 6.32 x 10^{-5}, R351 = 0.21, p = 0.37, see figure 4.7d-f)

a)

b)

R296,S12,T17
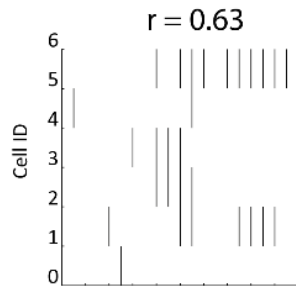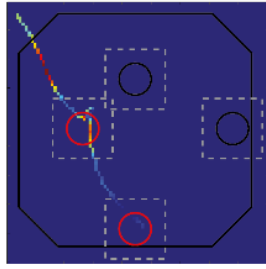
R346,S12,T18
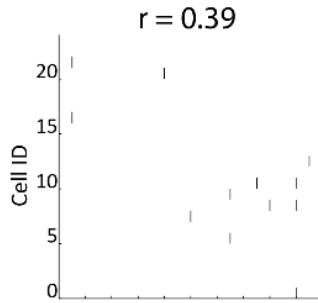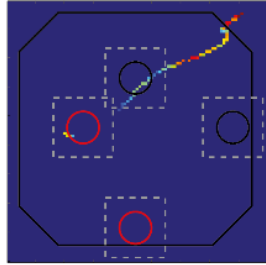
R346,S16,T
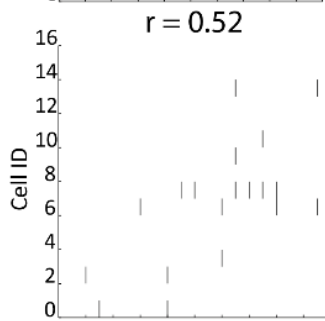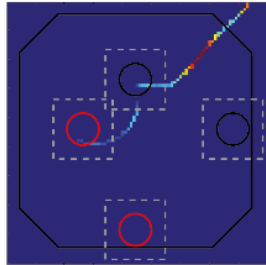
R351,S10,T12

R351,S11,T27

R346,S9,T2

r = 0.63

r = 0.39

r = 0.52

r = 0.53

r = 0.32

r = 0.50

Cell ID

Normalised Distance to Goal
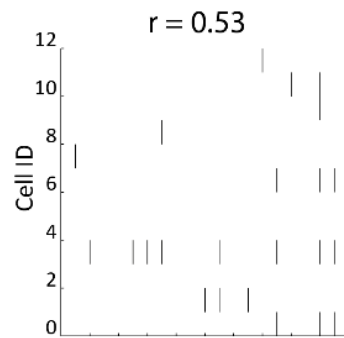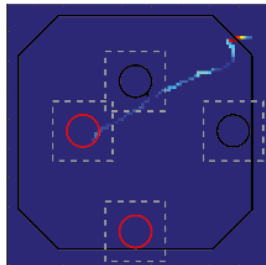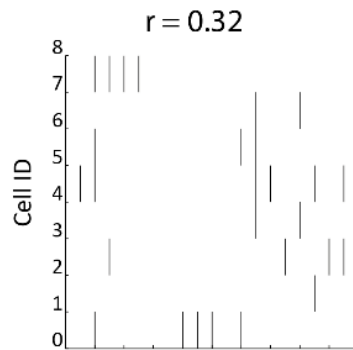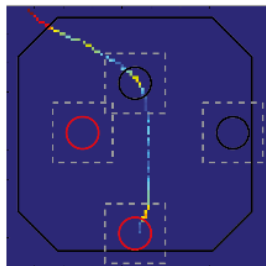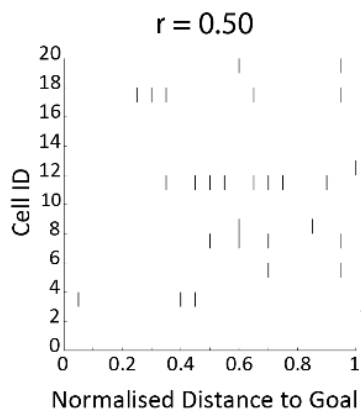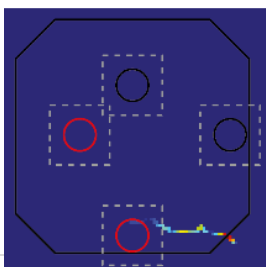
Figure 3.10 Trial Examples of Modulation of Population Firing Rate by Distance to Goal (previous page). Examples from all three animals where population activity during a trial correlated positively with distance to goal. **a)** Heatmaps of population activity in a trial. Hotter colours indicate higher population activity. Arena where animal did not traverse during a trial are coloured in blue. All examples show that population activity is higher early in a trial **b)** Raster plots of all cells recorded in a trial. Note spikes are plotted against normalized distance to goal (x-axis). Plots clearly show more cells are active early in a trial then later. Title shows correlation between population rates and normalised distance to goal.
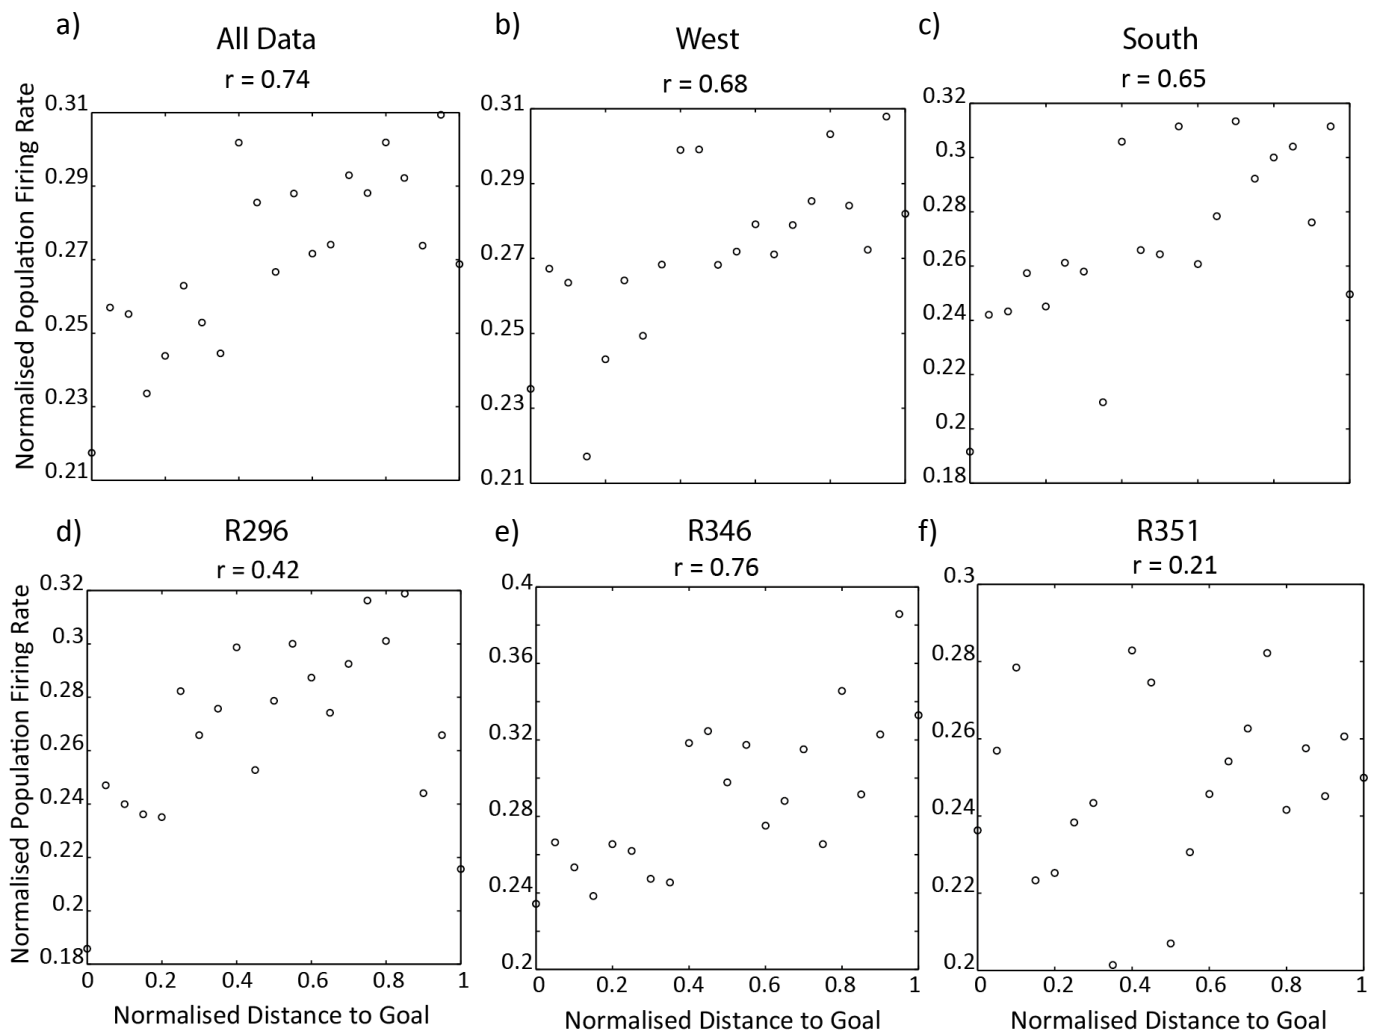
Figure 3.11 Distance-to-Goal Modulation on Population Rates Scatterplot of population rates against normalised distance to goal. X-axis shows distance to goal, and y-axis the normalised population rates. The title of each plot shows the correlation between rates and distance. **a)** Average population rates across goal distance for all data. **b)** Average population rates across goal distance for all data where animal goals to west goal well, **c)** same as b) but for south goal well. **d-f)** Population rates across distance to goal for R296, R346 and R351, respectively.

In sum, population rates vary as a positive linear function of distance to goal. Such that rates are higher at the start of a navigational trajectory than at the end. This is in agreement with our experimental hypothesis. However, one might argue that the reported relationship between population rates and distance to goal may be a simple result of place field distribution; perhaps, place fields are clustered around the start boxes resulting in an overall decreasing firing rate as a rat moves away from them and approaches its destination goal well. To address this possible confound we estimated what the predicted relationship between population rates and goal distance should be using ratemaps obtained from the foraging sessions.

3.3.3 Modulation of Goal Distance over Population Activity cannot be explained by Place Field Distribution

Before and after each goal trial session, animals completed a 15-minute foraging session. We used the ratemaps from these foraging sessions to estimate the expected relationship between population rates and distance to goal. Specifically, for every goal trial we extracted from the population ratemap, for the session the goal trial belonged to, the spatial bins sampled during the goal trial, maintaining the order in which the bins were visited and the dwell time in each bin (i.e. 'foraging trials'). Similar to before, we re-binned the expected population vector into 20 distance bins and computed correlations between the rate vector for the foraging trials and distance to goal. We compared the correlation obtained from the foraging trials to those obtained from the goal trials. If the foraging trials correlations were positive like the correlations from the goal trials this would indicate the relationship between goal distance and place cell activity in the goal trials reflects the distribution of place fields in the environment. On the other hand, if the foraging trial correlations were different from those obtained in the goal trials this would imply the goal trial correlations reflect activity of place cells outside their main firing field. Figure 3.12 shows examples of these 'foraging trials', and the relationship between population rates and distance to a goal for each trial. These examples clearly show that the observed relationship cannot be explained by place field distribution. The average correlation between distance to goal and predicted population rates was found to be strongly negative (figure 3.13a) - r = -0.83, p = 3.19 x 10^{-6}, different to the robust positive relationship observed for the goal trials (r = 0.76). Moreover, when foraging trials concluding at the different goals were analysed separately, we obtained a strongly negative correlation for trials

concluding at the West goal (r = -0.90, p = 1.80 x 10^{-8}, figure 3.13b), and a non-significant positive correlation for those ending at the South goal (r = 0.15, p = 0.51, figure 3.13c). Although, the correlation between distance to goal and population rates for foraging trials ending at the South goal was positive, this correlation is considerably weaker, *and* non-significant, than the one we obtained for the goal trials (i.e. r = 0.66). Moreover, the negative foraging trials correlation between population rates and distance to goal might imply place fields were clustered around goals, as found by previous work (Dupret, O'Neill et al. 2010). However, the weak correlation between population rates and distance to the south goal well might suggest that rather place fields clustered around the centre of the arena.
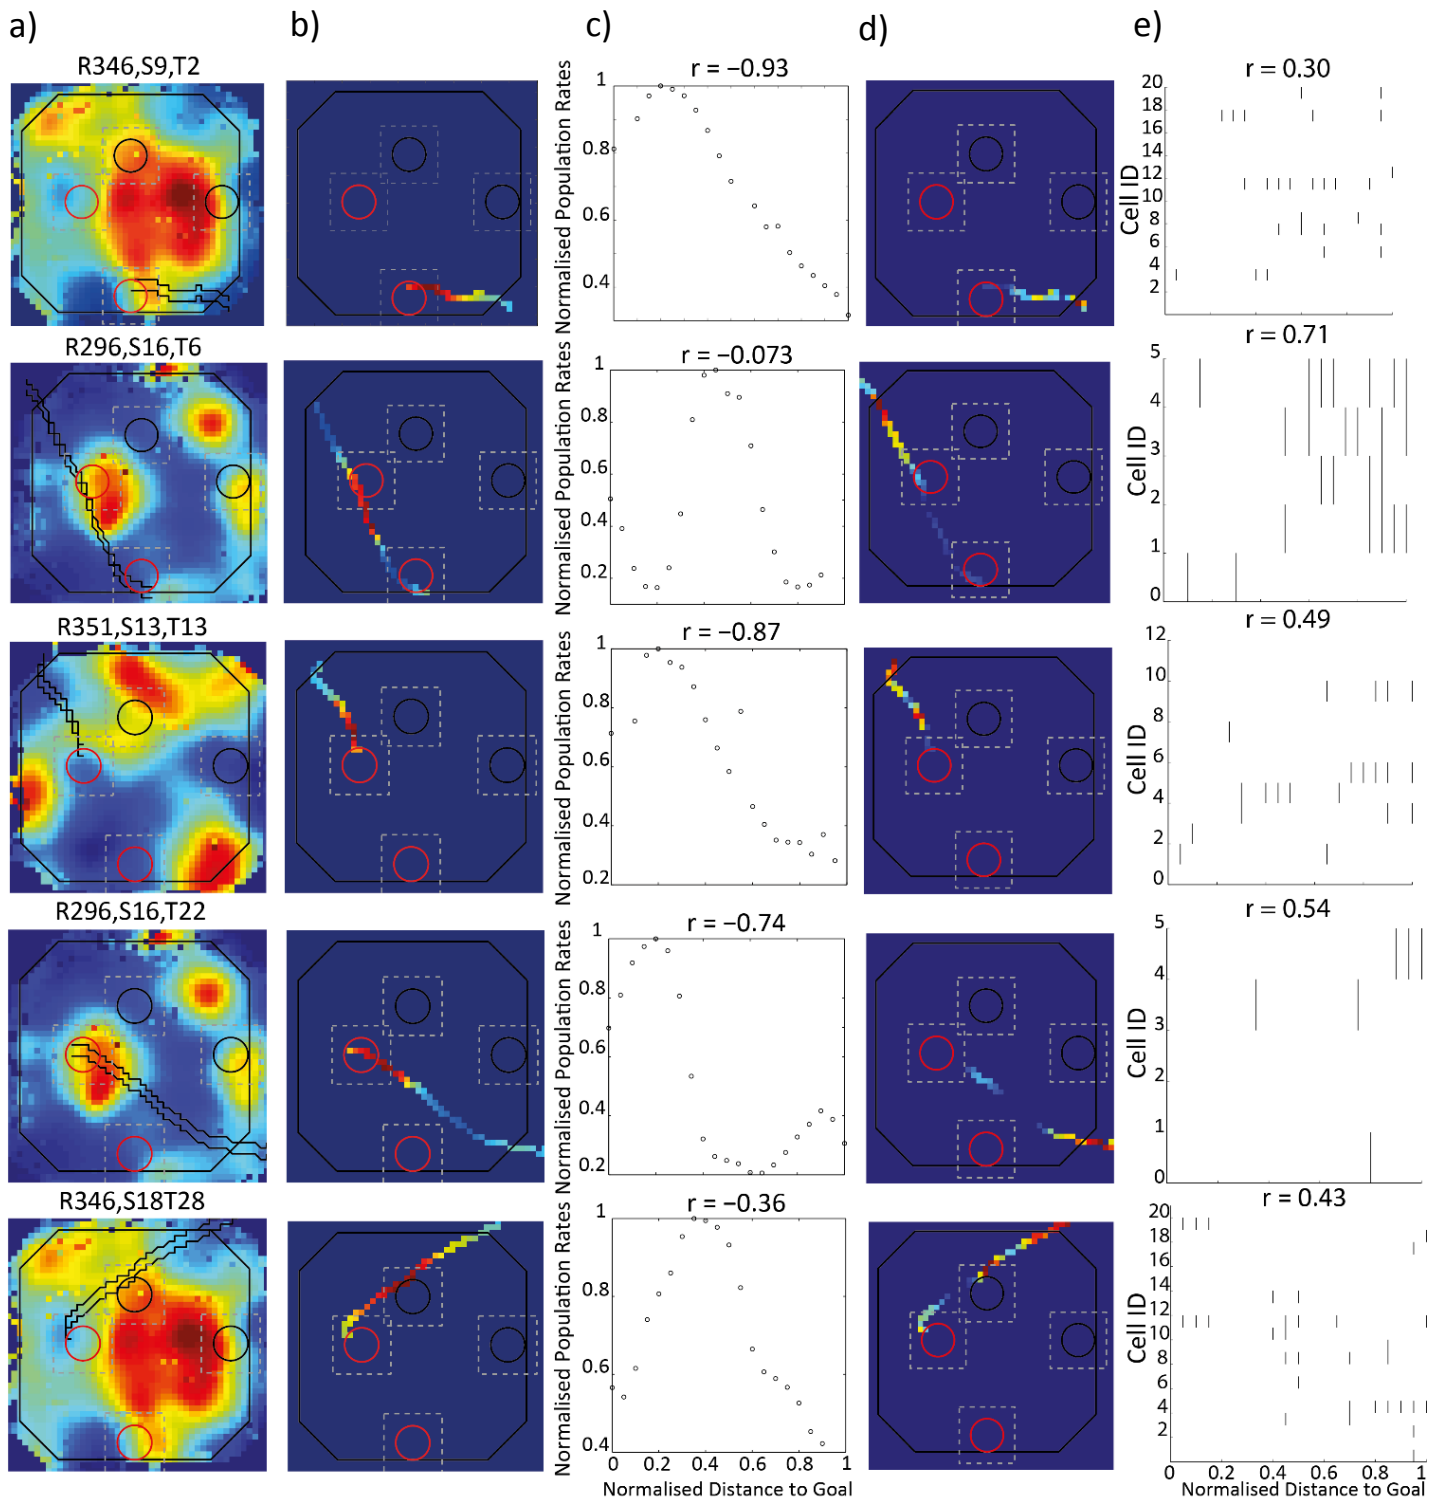
Figure 3.12 Distance Modulation of Population Activity is not a Result of Place Field Distribution Representative examples of population ratemaps (a), foraging trials (b), expected relationship between rates and goal distance (c), goal trial heatmaps (d) and observed relationship between rates and distance to goal (e). **a)** Population ratemaps derived from foraging sessions, with an example goal trial trajectory shown in black. **b)** Population ratemaps of a goal trial trajectory (shown in black in a) based on rates in the

foraging ratemap ('foraging trial'). **c)** Expected population rates (derived from panel b) plotted against distance to goal. Title shows the correlation between rates and distance. Y-axis shows distance to goal and y-axis normalised, expected population rates. **d)** Heatmaps for goal trials. **e)** Raster plots of goal trials shown in panel d), title shows correlation between population rates and distance to goal (x-axis). Labels on top of ratemaps in panel a show rat, session and trial ID. A raster plot for foraging trials could not be constructed as each foraging trial is based on average activity over spatial bins, rather than instantaneous rate over a continuous run to a goal, as is the case for the goal trials.
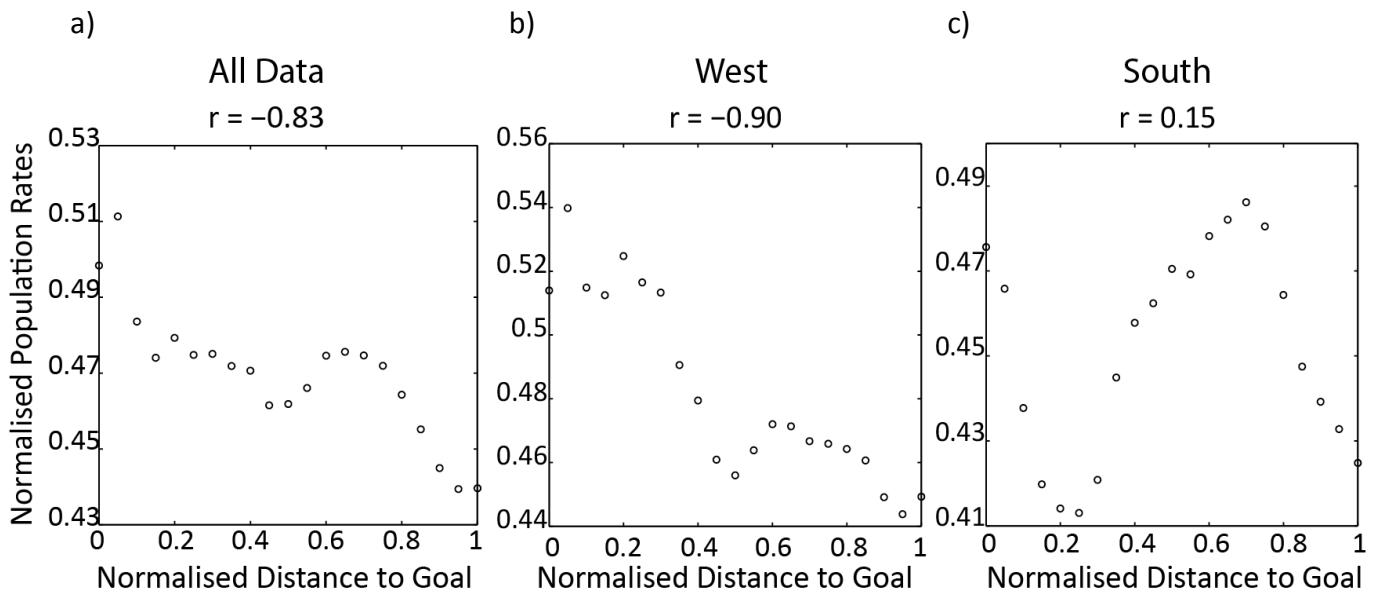
Figure 3.13 Relationship between Distance to Goal and Population Rates in Foraging Trials Normalised population rates obtained from foraging trials plotted against distance to goal. **a)** Rates for all foraging trials, **b)** rates for all foraging trials concluding at the West Goal, **c)** rates for all trials concluding at the South goal. X-axis shows distance to goal, Y-axis normalised population rates, and title the correlation between the two axes.

To conclude, the foregoing analysis has shown that the positive linear relationship between goal trial population rates and distance to goal is unlikely the result of place field distribution in the environment. Rather, the effect may reflect the activity of place cells when the animal is located outside their place field. Based on these results, one might speculate the observed distance modulation is a result of task demands and thus when the animal performs the task better the degree to which distance to goal modulates population rates is stronger. To address this question, we analysed the influence of performance over the distance to goal modulation on rates.

3.3.4. Navigational Performance Modulates the Relationship between Goal Distance and Population Activity

In order to assess whether the performance on the navigational task influences the reported positive relationship between population activity and distance to goal we computed performance for each session block; each session was divided in two blocks based on which goal the animal ran to in a trial. Then the relationship between population rates and distance to goal was estimated, the obtained correlation coefficients were then correlated with the block performance. Overall, we observed a modest positive relationship between performance and distance modulation over population rates in session blocks ($r = 0.31$, $p = 0.046$); implying the better the animal's performed the more positive the relationship between goal distance and population rates in a session block. Thus, when an animal performed well, rates were highest early on in a trial (see figure 3.14a). Moreover, when we analysed the south and west session blocks separately we observed a similar trend (South: $r = 0.51$, $p = 0.026$; West: $r = 0.22$, $p = 0.22$, see figure 3.14c-d). Although the performance modulation appears quite weak for the West session blocks this may be influenced by one or two outliers (see figure 3.14a and 3.14d) where the animal performs at ceiling (i.e. 100%). The inconsistent correlations observed for the high performing blocks alerted us to consider the influence the number of trials in a block can have on performance measures. Namely, high performance scores may be more likely for blocks consisting of only a few trials. Consequently, we imposed a filter on the session blocks, requiring each block to be made up of at least 2 trials from each start box (i.e. 8 trials in total). We then re-estimated the correlation between block performance and the distance modulation over population activity. This increased the correlation considerably – $r = 0.37$, $p = 0.036$, see figure 3.14b.
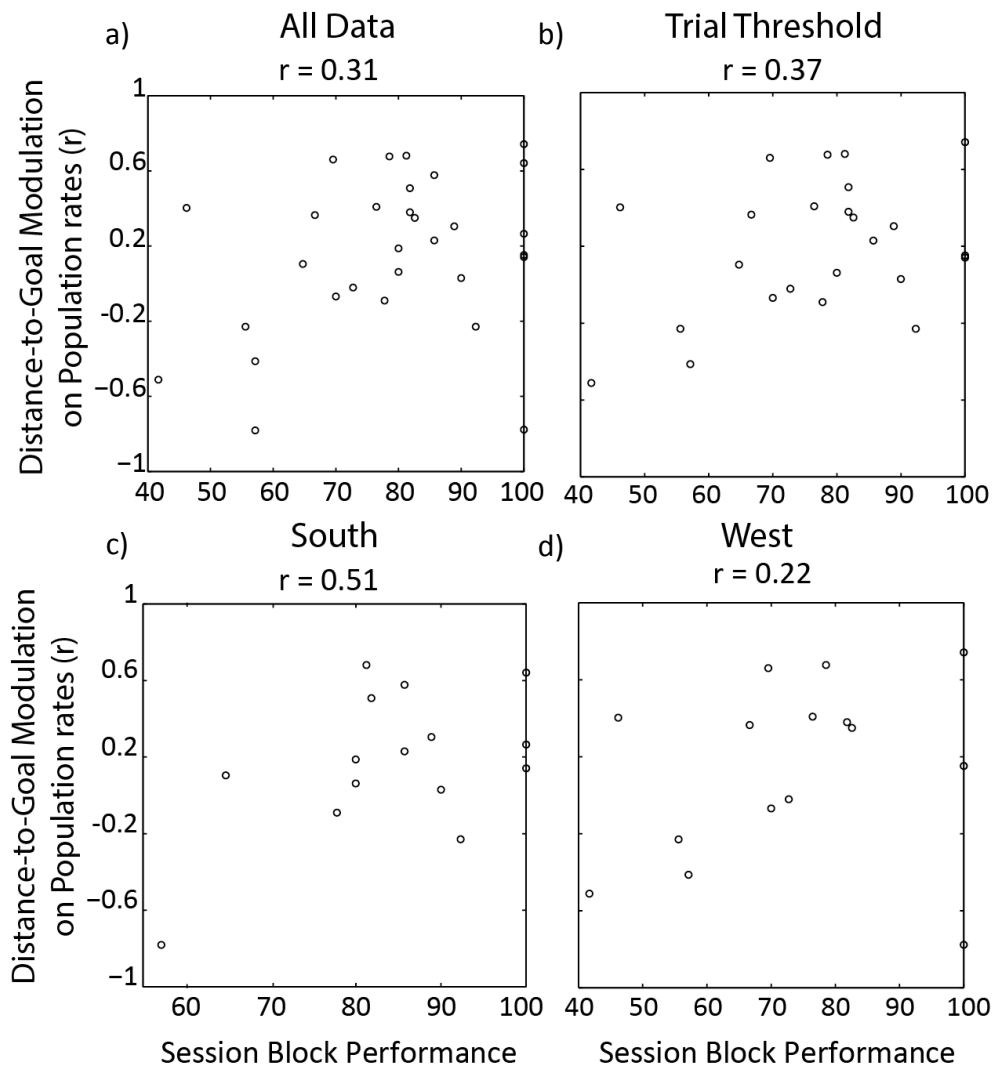
Figure 3.14. The Influence of Navigational Performance over the Relationship between Population Rates and Distance to Goal a) Correlation coefficients, derived from correlating distance to goal and population rates in each session block, plotted against session block performance. **b)** Same as a) but only displaying data point from session blocks consisting of more than eight trials. **c-d)** Same as a) but where South and West session blocks are analysed separately. The X-axis shows session block performance, and y-axis correlation coefficients. The title displays the observed correlation between the two axes.

In summary, not only do population rates in CA1 vary as a function of distance to goal, this relation seems to also be modulated by how well the animal performs. Although this correlation is modest, it is consistent across the West and South session blocks. These results serve to reinforce the suggestion that the heightened activity observed at the start of a trial play a functional role in the navigational task; similar to the increased hippocampal activity observed at the start of navigation in humans (Spiers and Maguire 2006). Perhaps the heightened activity serves a prospective function, such as planning of a future trajectory.

Before concluding the analysis, it is important to assess how running speed and acceleration vary as a function of distance to goal. It is known that place cells are positively correlated with running speed (McNaughton, Barnes et al. 1983). With respect to acceleration, direct evidence for a relationship with place cell activity has not been reported. Yet as speed and acceleration are intrinsically related, and the power of the theta rhythm has been found to correlate with acceleration (Long, Hinman et al. 2014), one might expect a correlation between place cell activity and acceleration. Perhaps at the start of a trial the animal's velocity is at its highest and then progressively drops as it approaches its destination goal. This would consequently lead to a positive relationship between cell activity rates and distance to goal, yet confounded by running speed. Moreover, one might even speculate that perhaps when the animal is performing well the positive relationship between running speed and distance to goal is even stronger; in other words, the correlation between velocity/acceleration and distance to goal may be modulated by performance. This would consequently also confound the reported performance correlations with the distance-to-goal modulated population rates. To address these possible confounds we assessed the relationship between running speed and acceleration and distance to goal as well as the correlation between performance and the distance-to-goal modulation over speed/acceleration.

3.3.5 The Behavioural Correlates of Population Rates, Distance to Goal and Navigation Performance

To assess whether the reported effects are a simple result of systematic variations in running speed/acceleration as an animal progresses through a navigational trial we performed a series of correlations. Firstly, we estimated the relationship between population rates and running speed and acceleration. We observed a positive correlation for both running speed and acceleration (speed vs rates: $r = 0.51$, $p = 0.018$, acceleration vs rates: $r = 0.67$, $p = 4.24 \times 10^{-4}$, figure 3.15a-b). However, running speed did not predict distance to goal significantly – $r = 0.40$, $p = 0.07$, see figure 3.15c). Moreover, when we performed a partial correlation analysis between population rates and distance to goal while controlling for running speed, population rates still had a strongly positive and significant relationship with goal distance – $r = 0.67$, $p = 0.0011$. However, the correlation between acceleration and goal distance was found to be strongly linear and positive – $r = 0.94$, $p = 1.48 \times 10^{-10}$, see figure 3.15d. Moreover, performing a partial correlation between activity rates and goal distance while partialing out the effect of acceleration did reduce

the correlation between rates and goal distance such that it no longer reached statistical significance – r = 0.33, p = 0.16.

This result could imply the relationship between population rates and goal distance is confounded by changes in acceleration during navigation. However, before coming to such a conclusion the relationship between acceleration and place cell activity should be established. Preliminary attempts were made to separate the influence of acceleration and goal distance on population activity in the current study. However, we found the two variables to be too inter-linked in our data. Thus, at this point, one cannot say for certain that the reported relationship between population place cell activity and goal distance during navigation is independent of acceleration; hopefully, future research will be able to tease apart the influence of different behavioural variables over place cell activity. However, what one can conclude is that the modulation of goal distance over population rates is unlikely to be the product of systematic changes in velocity during navigation or place field distribution in the environment.
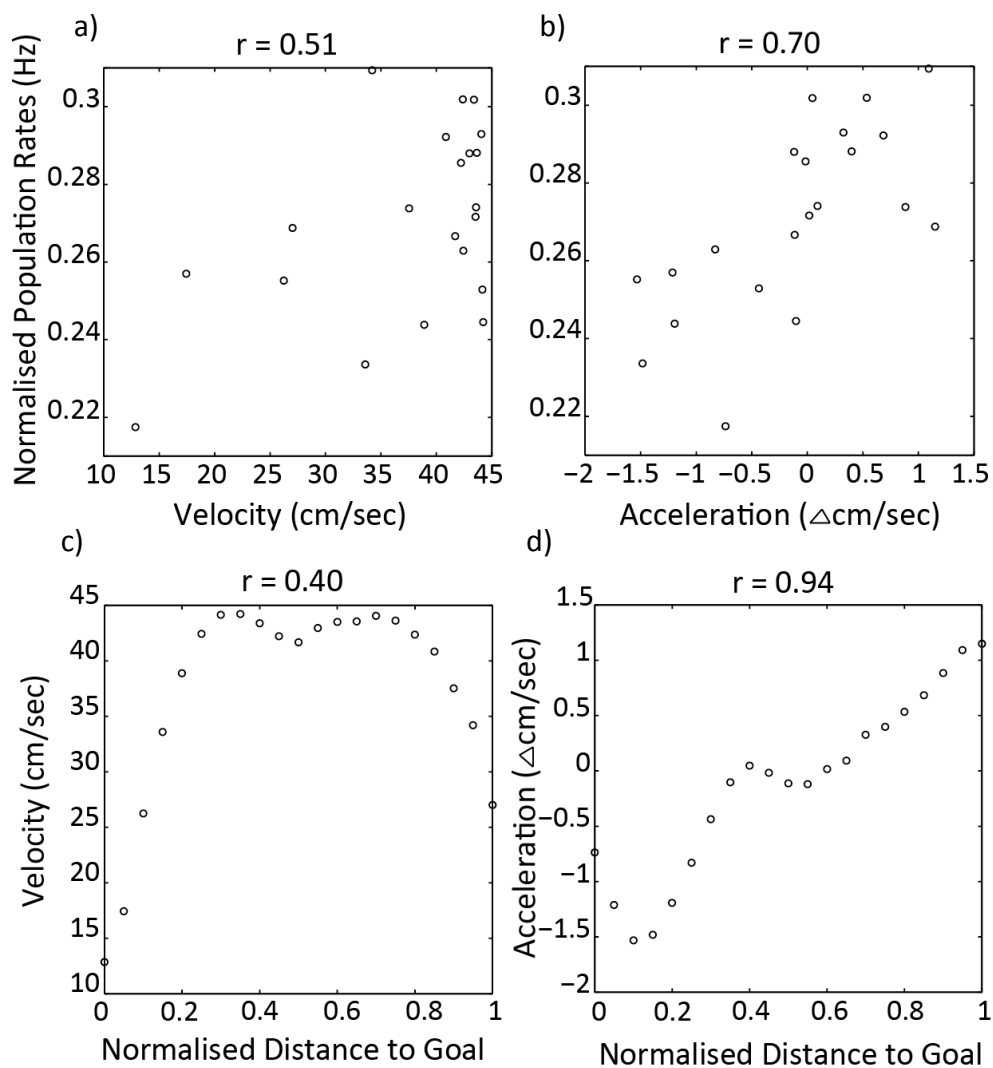
Figure 3.15. The Relationship between Velocity/Acceleration and Population Activity and Distance to Goal (previous page) **a)** population rates plotted against velocity, **b)** same as a) but where rates are plotted against acceleration. **c)** Running speed, across all trials, plotted against distance to goal, revealing a non-linear relationship. **d)** Same as a) but for acceleration. X-axis shows velocity and acceleration for panels a and b, respectively and normalised distance to goal for panels c and d. Y-axis shows normalised population rates for panels a and b and the velocity/acceleration for panels c and d, respectively. The title demonstrates the estimated correlation between the two axes.

To address whether the performance-based effects are confounded by speed and acceleration we correlated session block performance with the correlation coefficients between speed/acceleration and distance to goal for each session block. For velocity we obtained a weak, positive relationship between performance and the extent to which goal distance modulated running speed in a session block, which did not reach statistical significance – $r = 0.26$, $p = 0.083$, see figure 3.16a. The relationship between acceleration and goal distance did not appear to be influenced by session performance – $r = 0.095$, $p = 0.31$, figure 3.16c. Therefore, it is unlikely that the reported relationship between session block performance and the modulatory effect of distance to goal on population rates is confounded speed or acceleration. Moreover, in our performance analysis described above we reported stronger effects after removing blocks containing eight or less trials. Consequently, we applied the same criterion here and re-estimated the correlation between session block performance and the relation between speed/acceleration and goal distance. Although this had the effect of increasing slightly the correlations, the increase was modest and not statistically significant (speed: $r = 0.31$, $p = 0.073$, West: $r = 0.19$, $p = 0.19$, see figure 3.16b and d).
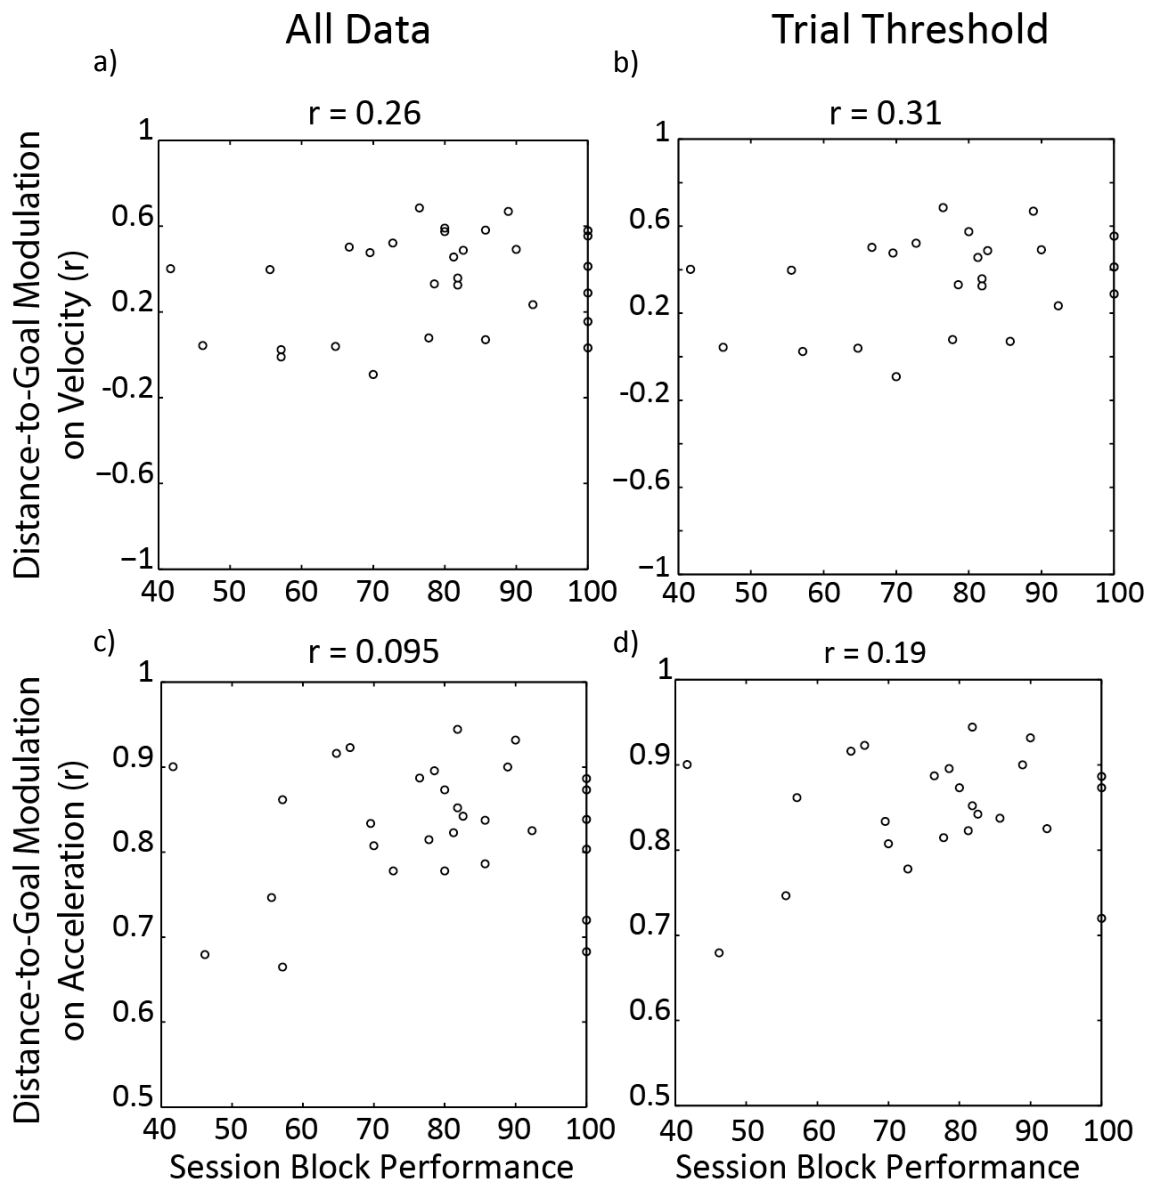
Figure 3.16. The Influence of Navigational Performance over the Relationship between Running Speed/Acceleration and Distance to Goal a) Session block correlation between velocity and distance to goal plotted against session block performance. **b)** Same as a) but excluding session blocks with eight or less trials. **c)** Same as a) but for the session block correlations between acceleration and distance to goal. **d)** Same as c) but excluding session blocks with eight or less trials. X-axis shows session block performance, y-axis correlation between running speed/acceleration and distance to goal. The title shows the correlation between the two axes.

To conclude, the reported correlation between population rates and distance to goal does not seem to be a mere function of speed. Running speed did not have a significant correlation with distance to goal, and only seemed to account for a small portion of the variance between activity rates and goal distance. Acceleration, on the other hand, did reveal a strongly positive relationship with goal distance, which did account for considerable variance between the place cell activity and distance to goal relation. Future studies hopefully will be able to tease apart the correlation

between acceleration and population rates. With regards to navigational performance, the session block correlation between goal distance and running speed/acceleration did not vary significantly with session block performance. Thus, the reported correlation for that of population rates is unlikely confounded by systematic variation in speed/acceleration for session blocks of different performance levels.

3.4. Discussion

This chapter has described the results from a goal-directed navigational study where large ensembles of hippocampal place cells were recorded simultaneously. We found that population activity rates varied as a positive function of distance to a navigational goal. That is, on the whole, CA1 place cells displayed the highest activity rates at the start of navigation and as an animal progressively approached its goal, rates dropped linearly. We have ruled out the possibility that the results are a mere product of place field distribution or the velocity profile of the animal during the task. Interestingly, the degree to which goal proximity modulated population rates seemed to be influenced by how well the animal performed in a session. Together, we think these findings suggest the increase in place cell activity at the start of navigation has a functional role to play in navigation.

A possible caveat with the reported findings deserves some discussion. We observed acceleration during a navigational trial to both correlate significantly with population rates as well as goal distance. Moreover, when the relationship between population rates and distance to goal was estimated while controlling for the effect of acceleration we observed a considerable amelioration of the goal proximity-population rate correlation. Thus, it is unclear whether, in our data, goal distance can influence place cell activity independent of changes in acceleration during navigation. This caveat highlights the need to study *directly* the relationship between place cell activity and acceleration. Studies have shown place cell activity correlates positively with running speed (McNaughton, Barnes et al. 1983), and that the power of the theta oscillation is modulated by acceleration (Long, Hinman et al. 2014). Yet, a direct relationship between place cell activity and acceleration has yet to be established. Considering findings from the human neuroimaging literature, it seems conceivable that hippocampal activity could predict distance to a goal independent of changes in behavioural variables. Howard and colleagues (2011) found the hippocampus was preferentially involved at the start of a goal-directed trajectory in a virtual reality study where speed and acceleration are held constant. Perhaps future studies will be able to separate the influence of different behavioural variables over place cell activity. For example, if attempts are made to control the velocity and acceleration of an animal at different navigational phases. Such controls would allow one to establish whether place cell activity can be independently modulated by distance to a navigational goal.

How do these findings fit with what we know about the role of the hippocampus in navigation? In relation to human neuroimaging findings these findings appear to accord with the few reported findings of hippocampal activity being higher at the start of navigation compared to elsewhere in navigation (Spiers and Maguire 2006, Brown, Ross et al. 2010, Howard, Yu et al. 2011). Howard and colleagues (2011) found a linear relationship between hippocampal activity and the euclidean distance to a navigational goal, similar to the current study. Moreover, Howard et al. also found the hippocampus to be significantly involved at important navigational decision points, such as when participants were forced to take a detour. Furthermore, Brown and colleagues (2010) found the hippocampus was selectively engaged at the start of a goal-directed trajectory that partially overlapped with another trajectory. In the current study, trajectories to the two goal wells often overlapped. This overlap between trajectories concluding at different locations might play an important role in regulating the high place cell activity at the start of navigation; future studies could investigate this question further.

However, the findings described here do not accord with recent findings reported by Sherrill and colleagues (2013). Sherill et al. had participants perform a goal-directed navigation task in an environment deprived of visual cues and prominent landmarks; thus, encouraging participants to path integrate. They found hippocampal activity was higher at the *end* of navigation, rather than at the start. The dearth of research dedicated to understanding distance coding in the hippocampus during navigation makes it hard to conclusively reconcile these discrepant findings. However, one might speculate that, perhaps, the findings described here and those of Howard and colleagues reflect decision making/planning during navigation, whereas those reported by Sherrill and colleagues might rather reflect coding of navigational goals. Although we did not observe an increase in place cell activity when the animal reached its goal, this could, partly, be due to low sampling at the goal location (we only included 500ms of time in the goal zone in the analyses). Moreover, perhaps in Sherrill et al.'s study hippocampal activity related to navigational planning occurred prior to the onset of navigation; participants were presented with a 'survey' like presentation of their location and the goal before the navigation phase began, this initial presentation might represent the time point where participants planned their path. As mentioned earlier, one can merely speculate at this point the reasons for the

diverging findings for distance coding in the hippocampus. Hopefully, future studies, both human as well as animal, will be able to clarify the role of the hippocampus in this particular function.

Human studies have, moreover, reported that the hippocampus significantly predicts navigational accuracy (Hartley, Maguire et al. 2003, Cornwell, Johnson et al. 2008, Sherrill, Erdem et al. 2013) – such findings fit with the findings reported here. Findings from some fMRI studies also imply the hippocampus is particularly involved in route planning (Spiers and Maguire 2006). Although the current results cannot speak directly to these findings, intuitively one might expect route planning to occur at the start of navigation; thus, our findings might be considered consistent with these results. However, future studies need to study the 'content' of ensemble hippocampal activity at the start of navigation and how it may relate to future paths taken in order to understand whether this increase in population activity could reflect route planning. Specifically, the place-selective nature of hippocampal place cells allows one to assess whether at the start of navigation place cell activity depicts a route to the animal's destination. Such a 'decoding' analysis was not possible with the current data, due to low cell yield. Hopefully, future studies will be able to address this question.

As discussed in chapter 1, some theorists have proposed coordinated ensemble activity observed at choice points in a navigational tasks, as well as prior to an animal initiating a goal-direct trajectory, as a neurological marker for navigational planning (Johnson and Redish 2007, Foster and Knierim 2012, Pfeiffer and Foster 2013). Pfeiffer and Foster (2013) found replay in a goal-directed navigational task anticipated the future path of the animal. Yet, replay was associated with behavioural immobility and sharp-wave ripples (SWRs) – a behavioural and oscillatory state quite different from the current study. On the other hand, Johnson and Redish (2007) observed CA3 place cell activity at choice points, when the animal was moving and theta dominated the local field potential (LFP) to represent possible future trajectories (i.e. sweeps of future place cell sequences). Moreover, such 'sweeps' have been observed in other studies (Gupta, van der Meer et al. 2012). Within a theta cycle, upcoming trajectories are bound (O'Keefe and Recce 1993) and some claim such 'theta sequences' have a navigational function, helping the animal predict what's ahead (Foster and Wilson 2007). Perhaps, at the start of navigation, in the current study, the future trajectory of the animal is activated in theta cycles and as the animal approaches its goal progressively fewer cells are

active, resulting in a positive relationship between neural activity and distance to goal. Hopefully future studies will be able to address this possibility.

The current findings might be seen to contradict reported goal-related activity of hippocampal neurons. Dupret and colleagues (2010) recorded from large ensembles of place cells while an animal searched for food on a circular platform. Between days, the location of the food reward changed. Dupret et al. found CA1 cell ensembles to re-organise in response to the change in reward location. Namely, place fields seemed to cluster around the rewarded location, and as the reward location changed so did the place fields (Dupret, O'Neill et al. 2010). Similarly, Hok and colleagues found ensemble CA1 activity increased while an animal was positioned in an area of a platform where it would receive reward (Hok, Lenck-Santini et al. 2007). The current study did not find evidence for an increase in place cell firing at reward locations in the goal trials- in fact, population activity seemed lowest at goal wells. How does one reconcile these findings?

In the current study it was shown the increased place cell activity at the start of a goal directed trajectory was unlikely to be a result of place field distribution in the navigational environment. Thus, this effect might reflect functional out-of-field activity rather than the distribution of place cell firing in the environment. Consequently, this effect may not be interpreted as being at odds with the findings reported by Dupret and colleagues. Moreover, we did not directly assess whether place fields in the arena clustered around goal locations. However, the foraging trial analysis showed the expected variation in population rates with proximity to goal location was the reverse to what we saw in the goal trials; rates were higher when the animal was further away from the start. Thus, one might, infer the hippocampal spatial representation was stronger near the goals, similar to what Dupret et al. and Hok et al. reported. Yet, this effect needs to be investigated in more detail as further analyses of the foraging trials showed the relationship between place cell activity and goal distance was different for the two goal wells. Moreover, it is unclear whether any such clustering of place cells around goal locations in our study is due to the motivational relevance of these locations or difference in behaviour when the animal was located at the goal wells. In general, the role of navigational goals in determining place cell firing has been unclear, as rewarded locations are often associated with the occurrence of ripples (Buzsaki, Horvath et al. 1992). In fact, it is unclear whether the increased place cell activity observed at reward locations in Hok et al.'s study is simply due to quiescent

behaviour observed at the rewarded site. Hopefully, future research, where behaviour at goal locations is better controlled for, will establish the influence goal locations can exert over place cell activity.

To summarise, the current study reports findings consistent with those observed in numerous human neuroimaging studies – providing an important bridge between human and animal studies. Moreover, these findings are consistent with ideas suggesting a role for the hippocampus in route planning during goal-directed navigation yet also highlight the need to study the possible modulatory effect of acceleration over place cell firing. Finally, the study calls attention to theta sequences observed during active navigation, and the possibility that these are functional and may help the animal prepare upcoming routes. Future studies should further investigate cells active within theta cycles while an animal carries out a goal-directed navigational task. Preferably, such a study would be designed so that the relationship between theta sequences and the navigational choice of the animal can be confirmed – for example, in a study where an animal can take multiple, prescribed, routes to a goal. Finally, the use of virtual reality would help with this quest as the speed and acceleration profile of the animal can be better controlled than in studies where the animal is freely moving. This would make a significant contribution in our quest to understand the neural mechanisms underlying successful navigation and clarify the role of the hippocampus in such tasks.

4. The Hippocampus Simulates Desired Paths through Unexplored Space during Rest

4.1. Introduction

As discussed in chapter 3, recently theorists have suggested the contribution of the hippocampus to cognition may not be limited to recollecting the past (Scoville and Milner 1957) but may extend to preparing and perhaps planning future behavior (Johnson and Redish 2007, Foster and Knierim 2012, Pfeiffer and Foster 2013). Theta sequences (Foster and Wilson 2007) discussed in the foregoing chapter, could represent the mechanism for such a planning function (Johnson and Redish 2007, Erdem and Hasselmo 2012, Gupta, van der Meer et al. 2012). However, as mentioned above, this is a nascent area of research. Another candidate mechanism is the synchronized population bursts associated with sharp-wave ripple (SWR) oscillatory states, observed during rest and sleep (O'Keefe and Nadel 1978, Buzsaki, Horvath et al. 1992); a phenomenon hypothesized to serve the long hypothesized consolidative function of the hippocampus (Wilson and McNaughton 1994).

In the seminal study by Wilson and McNaughton (1994) place cell activity patterns underlying wakeful activity were re-activated during SWRs in subsequent sleep. That is, place cells that tended to fire together while an animal foraged on a novel circular platform (i.e. had place fields close to each other) also tended to fire together in the sleep session following the wakeful activity. Importantly, the coordinated activity was not as strong in sleep preceding the run on the novel platform. Wilson and McNaughton suggested that this structured re-activations (*replay)* could lead to strengthening of newly formed memory traces, through a Hebbian learning process (Hebb 1949). With replay, memories would eventually become independent of the hippocampus as they are transferred to the neocortex (Marr 1971). Over the past decade, this replay phenomenon has been replicated many times, and the observed physiological and behavioural properties lend some support to its hypothesized role. On a linear track, replay represents place cells sequences both in the forward and reverse order (Foster and Wilson 2006, Diba and Buzsaki 2007), novel trajectories are more likely to be replayed than familiar ones (Cheng and Frank 2008), replay is coordinated with reactivations in the neocortex (Ji and Wilson 2007), and if ripples are suppressed following training on a navigational task, learning is impaired (Girardeau, Benchenane et al. 2009).

More recently, scientists have suggested replay may also play a role in navigational planning. For example, at a choice point place cell sequences representing possible future trajectories are seen to sweep ahead of the animal (Johnson and Redish 2007), and during behavioural immobility the hippocampus has been found to piece together trajectories in a familiar environment (Gupta, van der Meer et al. 2010). In an interesting study by Pfeiffer and Foster (2013) large ensembles of place cells were recorded while an animal carried out a goal-directed navigational task. Pfeiffer and Foster observed awake replay over-represented goals in the environment, and tended to predict the future goal-directed path of the animal (Pfeiffer and Foster 2013).

Planning or *simulations* of experiences have been reported many times from human neuroimaging and neuropsychological studies. For example, Hassabis and colleagues (2007) found that hippocampal patients were impaired at imagining novel experiences (Hassabis, Kumaran et al. 2007). In fact, many human studies imply the hippocampus may play an equal role in recalling and constructing novel experiences (Buckner and Carroll 2007, Schacter, Addis et al. 2012). However, evidence for the cellular mechanism underlying construction of *novel* experiences in the hippocampus has hitherto been mostly lacking.

Dragoi and Tonegawa (2011, 2013) did observe hippocampal place cell sequences being *preplayed* before animals entered a novel environment. In their study they recorded from ensembles of hippocampal place cells in mice while an animal rested on a novel track, before entering an adjacent arm, never before experienced. During periods of rest, they found hippocampal cell sequences for the future novel arm were pre-played. Even more intriguingly they observed preplay, as well, while the animal slept before it ran on either track (Dragoi and Tonegawa 2011). Finally, Dragoi and Tonegawa found, in a follow-up study, that such preplay extended to multiple future novel environments (Dragoi and Tonegawa 2013).

One might thus contend preplay is the candidate mechanism for simulations in the hippocampus, yet from these studies it is not clear whether preplay results from goal-directed influences or passive pre-configurations in the hippocampal circuitry (McNaughton, Barnes et al. 1996); that is, pre-existing synaptic linkages or 'charts'. Here I will describe results that tease apart the two possible influences. We hypothesized preplay reflects the motivational relevance of a novel, unvisited environment, such that a novel behaviourally relevant environment *would* be

preplayed but a behaviourally irrelevant environment would not. Furthermore, we will suggest preplay may represent the neural signature underlying goal-directed planning in novel environments.

4.2. Method

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

4.2.1. Animals

Four male Lister Hooded rats (300-400g in weight and three to six months of age at implantation) were used in this study. Three animals had dual hippocampal implants, each hemisphere containing a 16-channel microdrive. One animal had a single 32-channel microdrive implanted over the right hemisphere[5].

Table 4.1 Total Number of Place Cells Recorded per Animal

Rat	Cell Yield
R1838	37
R505	52
R504	58
R584	66
Total	213

4.2.2. Experimental apparatus and protocol

We used a T-shaped track, raised 43cm off the ground. The track consisted of 10cm wide run-ways covered in black rubber, the stem was 222cm long and the arms 200cm from tip to tip. To prevent animal's forming representations of tracks similar to the one used in the study, animals were not habituated to running on linear tracks. However, all animals had been trained to forage for rice in a square enclosure during screening (see Methods section 2.1.3). Initially, to prevent animals from accessing the arms, a barrier was placed on the stem 22cm from the junction with the arms (see figure 4.1). The barrier consisted of a grill of vertical wooden dowels spaced so the animals could easily see through but could not pass. The apparatus was located in a curtained environment, and two distal landmarks were mounted on the curtain on either side of the maze for orientation. To encourage even distribution of place fields along the track, textured local cues were placed along the stem and arms. These five different cues were sheets of: black

---

[5] A total of 12 rats were implanted for this study. However, we imposed an inclusion criterion whereby they were only run through the experimental procedure if they had at least 35 place cells, held in parallel.

rubber, grey cardboard, black polypropylene, plaster tape, and sand paper each approximately 20cm long and 10cm wide (see Figure 4.1). For sleep sessions, the animal was placed in a cylindrically shaped enclosure (18cm diameter x 61cm high) with a towel placed in the bottom. Prior to the experiment the animal was familiarised with the sleep enclosure for a minimum of 120min. The bin was located beside the track during sleep sessions, but not present during the other sessions.
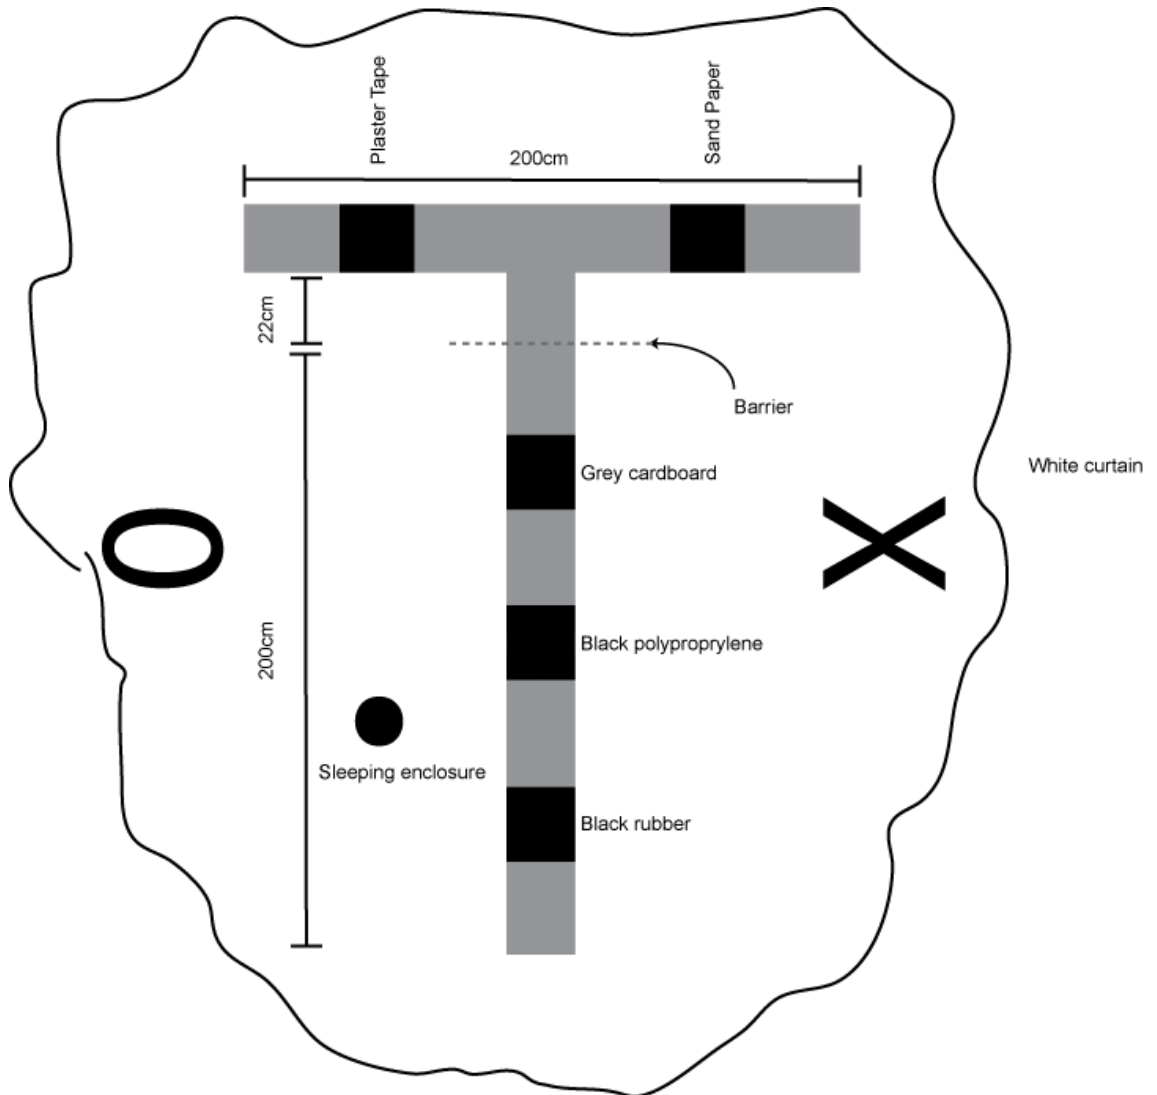
Figure 4.1 Schematic of Experimental Apparatus Black rectangles on track represent texture cues, dotted line transparent barrier, circle sleep enclosure (18cm wide), wiggly border demarcates the white curtains surrounding the environment, with distal landmarks fixed to the curtain with black tape ('X' and '0').

Prior to running or seeing the track, animals were placed in the rest enclosure for at least 60min - SLEEP1. Following SLEEP1, rats were placed at the end of the stem facing away from the track. During RUN1 animals were encouraged to run up and

down the stem by the experimenter rewarding it with a rice grain at each end. During this session the animal was prevented from running onto the arms by the barrier. Once animals had completed 20 laps on the stem, the experimenter stopped baiting the stem and instead baited the end of one of the arms (GOAL-CUE); inaccessible yet visible to the animal (cue located ca. 1m from the animal). The animals had good visual access to the inaccessible arms as they could see between the wooden dowels of the barrier. Moreover, as the barrier was only ~20cm wide the animals could easily observe the ends of the arms, located less than 1m from them. The experimenter stood by the end of the cued arm to elicit further interest. Animals were allowed to remain on the track for at least a further 10 minutes.

Following RUN1, the animals were placed back into the sleep enclosure - SLEEP2. Following SLEEP2, animals were put back on the end of the stem facing away from the track. The barrier at the end of the stem was removed and the animals were allowed to traverse both the stem and arms in alternate L-shaped laps for each arm (figure1) – RUN2. To force the animal to alternate, a barrier (the same as the one used for RUN1) was placed at the junction of the T-maze, and was moved between trials by the experimenter. The animals completed 20 L-shaped laps for each arm. It should be noted, although the sleep of animals was not explicitly monitored, we know from velocity estimates, they were very still during the sleep sessions.
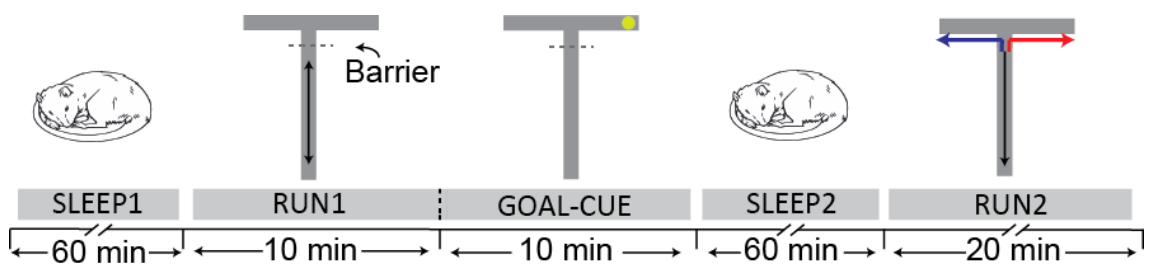
Figure 4.2 Experimental Protocol Prior to running on the T-track, animals were placed in an opaque sleep enclosure for at least an hour (SLEEP1). Following SLEEP1, the animals completed twenty laps on the stem (RUN1), which was then followed by at least 10 minutes of goal-cueing (GOAL-CUE), where the experimenter pretended to bait one inaccessible T-arm (yellow circle). Following cueing, the animals were put back into the sleep enclosure for at least an hour (SLEEP2). Finally, after SLEEP2 the animals traversed the entire extent of the T-track in alternate L-shaped laps (RUN2), as indicated by coloured arrows.

Table 4.2 Duration of Each Experimental Session for All Animals

Rat	Session Duration (min)				
	SLEEP1	RUN1	GOAL-CUE	SLEEP2	RUN2
R1838	60	13	10	60	34
R505	88	10	17	67	19
R504	74	9	11	60	31
R584	75	12	10	72	35

4.2.3 Data Analysis

Template Generation

Ratemaps for the runs on the arms during RUN2 were generated. Prior to ratemap generation, the last 20cm (10%) of each arm were removed to exclude areas where the animal groomed, ate etc. Moreover, position samples where the animal's speed was lower than 3cm/sec were excluded. Animals' paths were linearised and dwell time and spikes binned into 1cm bins, then smoothed with a Gaussian kernel (σ=2cm). Firing rate was calculated for each bin by dividing spike number by dwell time. Templates were generated based on the location of the peak firing rate of each cell (if a cell had more than one place field, the field with the highest firing rate was used). Specifically, four templates were generated, two for each arm – one for runs commencing at the end of the stem and concluding at the end of the arm (TO cued/uncued arm) and vice versa (FROM cued/uncued arm). These templates we refer to as Up Cued Arm (UCA), Down Cued Arm (DCA), Up Uncued Arm (UNA) and Down Uncued Arm (DNA). For example, the Up Cued Arm (UCA) template enumerates the order in which the peaks of firing fields appear on the cued arm for runs towards the end of the arm. So, if cells with ids 3, 20 and 7 were the first three to have peaks on that arm the UCA template would start 3, 20, 7. Cells whose peak firing rate in the linearised ratemap were below 0.5Hz or had less than five contiguous bins with rates below the mean firing rate of the cell were excluded.

For a follow-up analysis we generated templates for the stem based on the RUN1 session. Templates were generated in the same way as for the arms.

R1838
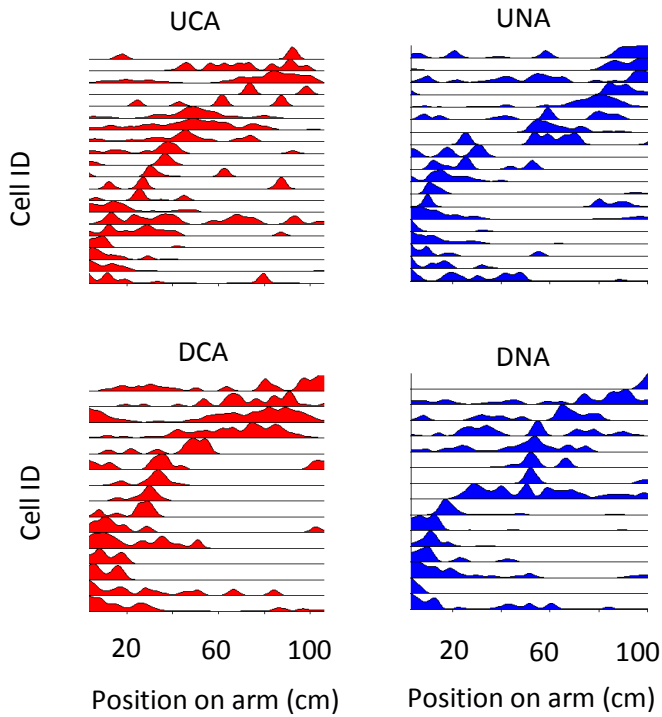
Figure 4.3 Place Cell Sequence Templates Templates generated by ordering the location of the peak firing rate of all cells active on an arm. Red templates represent arms cued during GOAL-CUE, blue templates are for uncued arms. Templates are from R1838. Y-axis contains cell IDs, X-axis is location on arm. UCA = Up Cued Arm, DCA = Down Cued Arm, UNA = Up Uncued Arm, DNA = Down Uncued Arm.

Table 4.3 Number of Cells in Each Template for All Animals

Rat	Template Length (# cells)			
	UCA	**DCA**	**UNA**	**DNA**
R1838	20	15	19	15
R505	36	35	26	32
R504	43	33	40	43
R584	53	45	45	41
Total	152	128	130	131

Preplay Analysis

For each sleep session, periods where at least 15% of cells from a given template fired within 300ms and were bound by at least 50ms of silence were selected as 'spiking events' (for R1838, which had a lower cell yield than the other rats, a minimum of 4 cells were required to be active). If a single cell fired more than one spike within this period, the first spike was counted and other spikes disregarded. The extent to which templates were represented in the spiking events was assessed using template-matching (Lee and Wilson 2002, Foster and Wilson 2006, Diba and Buzsaki 2007). Specifically, the temporal sequence of cells in spiking events and the order of their corresponding peaks in the template was compared. If future template sequence are preplayed in sleep the order in which cells fire during spiking events would be expected to resemble the order in which they fire on the track. In other words, one would expect the spike sequence in events and in future template sequence to be more correlated than predicted by chance.

To determine if each template was being significantly preplayed we compared the proportion of spiking events that individually correlated with that template with the proportion expected by chance (as predicted by a shuffling procedure). Specifically, the Spearman's rank-order correlations between each spiking event and the current template was computed and the number yielding significant ($p <$ 0.05) correlations was recorded; these were labelled 'preplay' events (see figure 4.4a-b). Note, both positive as well as negative significant correlations were categorised as preplay events, as previous studies have reported that replay can reactivate a place cell sequence in the forward or reverse order (see section 1.4.7 above for a detailed explanation of this, and (Foster and Wilson 2006)). Next, the sequence of spikes in each event was randomly permuted (shuffled) 100 times and for each permutation the Spearman's rank-order correlation and matching p-value calculated, again the proportion of significant correlations was recorded; this value represented the proportion of preplay events expected by chance. Finally, the proportion of shuffled events that had a significant correlation with the future template sequence was compared to the number obtained for the unshuffled data using a one sample binomial tests (see figure 4.7d-e).

To assess whether the cued arm was preplayed more than the uncued arm we directly compared the total number of events and number of preplay events (i.e. those with individually significant correlations) for the cued and uncued arm using

a Fisher's Exact test (see table 4.5, and fig4.7d). A similar approach was used for comparing number of preplay events in SLEEP1 and SLEEP2 (see table 4.6 and fig4.7e). However, the Fisher's Exact test assumes that the contingency table marginal are fixed and this was not the case here (i.e. the row and column totals of the condition vs preplay contingency table should be known before the experiment starts(Howell 2007). Because of this we also used a two-sample binomial test to validate the results and confirm they were not method-dependent.

a)   b)

r = -0.98, p < 0.05   r = -0.66, p > 0.05

20cm   100ms   20cm   100ms
Position on arm (cm)   Time (ms)   Position on arm (cm)   Time (ms)
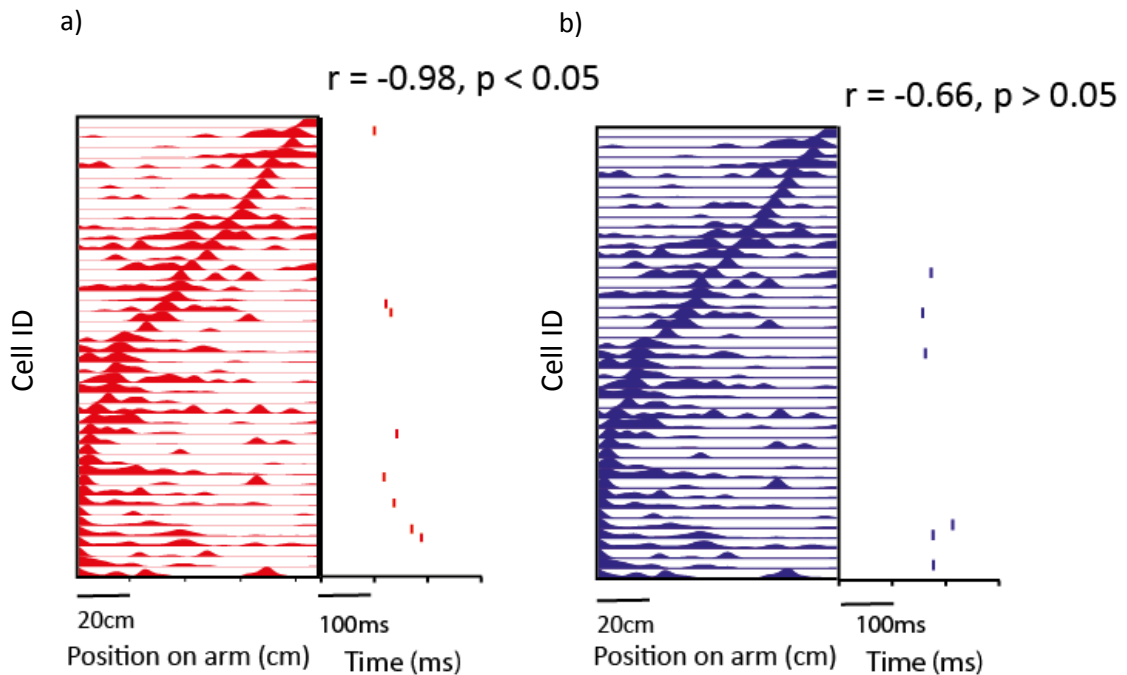
Figure 4.4 Preplay Analysis a) Example of a template sequence (left) for a future run on the cued arm (RUN2), to the right is an example of a preplay event recorded in SLEEP2. The order in which the cells fire during this event are almost in the exact reverse order relative to that seen during RUN2, consequently the correlation between the preplay event and the template sequence is strongly negative – correlation: r = -0.98, p = <0.05. **b)** Example template sequence for a run on the uncued arm in RUN2 (left). To the right is an example of spiking event recorded in SLEEP2. The order in which the cells fire during SLEEP2 does not resemble the order in which they fire in RUN2 – correlation: r = -0.66, p >0.05. Data shown is from rat R584.

A follow-up analysis looked at preplay during GOAL-CUE. To identify spiking events we employed the same criteria as above, but adding an additional speed criterion necessitating spiking events to occur while the animal's speed was less than 10cm/sec.

To control for differences in the number of cells contributing to each template we conducted a further analysis and down-sampled templates for each animal to the length of the shortest template. Specifically, if the shortest template consisted of 20 cells while another template consisted of 30 cells, we randomly removed 10 cells from the longer template. Moreover, we repeated this down-sampling process a 100 times, each time removing a random set of 10 cells. We then ran the preplay analysis as before – estimating the number of preplay events for each down-sampled cued and uncued template and comparing them to each other and chance levels, using a Fisher's Exact Test (and two-sample binomial test) and one-sample binomial test, respectively.

Controlling for Cluster Quality and Proximity to Cued Arm

Two further analyses were carried out. The first to eliminate the possibility that poor cluster isolation may contribute to the observed preplay (such as mistakenly dividing a complex spike into two clusters). The second, to investigate whether proximity to the cued arm during GOAL-CUE could account for the preplay of the cued arm. That is, during GOAL-CUE animals leaned/stretched over to the cued arm more than the uncued arm. Thus, perhaps the enhanced preplay of the cued arm simply reflects replay of the occupied place fields during behavioural cueing.

To address the former, we looked at the relationship between Mahalanobis' distance between pairs of cells active on an arm and the physical distance between their place fields. Mahalanobis' distance (in squared units) between two clusters was estimated using the following formula.

$$D_{(x)} = (x-\mu)^{S-1}(x-\mu)$$

<div align="right">(1)</div>

x is a data point from a cluster, and μ the mean of a different cluster, and S its covariance

Clusters that have been mistakenly assigned to two or more cells are likely to be identified by low Mahalanobis' distance and proximate place fields. We divided cell pairs into two groups – those with short inter-place field distances and those with long inter-place field distances. A short distance was defined as cell pairs with an inter-place field distance below the estimated median distance in the sample and long distances as cell pairs with an inter-place field distance greater than the

median distance. We then looked at the distribution of Mahalanobis' distances for the two groups and compared them to each other using a 2-sample Kolmogorov-Smirnov test.

For the latter analysis, we looked at whether the parts of the arms closer to the animal were more likely to be represented in preplay events than parts of the arm further from the animal during GOAL-CUE. We divided the arm in halves for this purpose, not including the first 20cm as they were common to cued and uncued runs. To estimate positions represented in preplay events we used the peak location of cells active in an event. A preplay event will represent a number of positions, the number and position depending on how many and which cells are involved in an event. Then we simply asked, are positions on the proximal half of the arm more often preplayed than positions on the arm half further away from the animal. To assess this statistically we used a one-sample binomial test.

Local Field Potential Analysis

The local field potential (LFP) from CA1 was recorded at 4.8 kHz throughout the experiment. In order to analyse sharp-wave ripples (O'Keefe and Nadel 1978, Buzsaki 1986) the LFP was band-pass filtered between 80 and 250Hz. An analytic signal was constructed using the Hilbert transform, taking the form:

$$s_a(t_k) = s(t_k) + i\mathsf{H}[s(t_k)]$$

(2)

H specifies the Hilbert transform, $s(t_k)$ is the filtered LFP signal, $t_k = k\Delta$, where $k = 1,...,K$ indexes the time-step and Δ is the inverse of the sampling rate.

An instantaneous measure of power was found by taking the squared complex modulus of the signal at each time point. This measure was then down sampled to 50Hz to match the position sampling rate, and finally was smoothed with a boxcar filter of width 0.1s. Two-tailed independent samples t-tests were used to compare power in the ripple-band during significant (i.e. preplay events) and non-significant spiking events and to compare power during significant spiking events and periods outside of spiking events.

4.2.4. Histology

Coronal sections were made for all brains used in the study, and they were stained with cresyl violet to locate tetrode recording locations. Figure 4.5 below shows examples of tetrode tracts for animal R504 in both hemispheres.
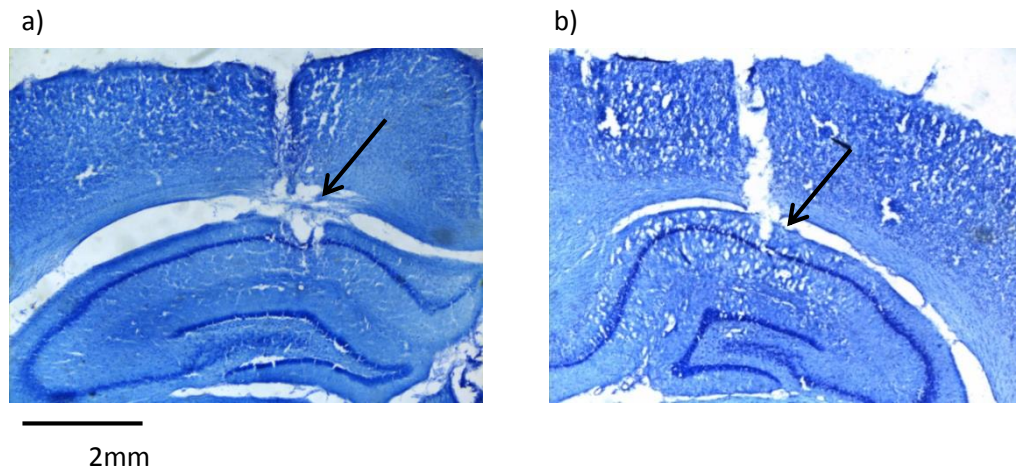
Figure 4.5. Location of Tetrode Recording Example of tetrode tracts from coronal cresyl violet stained sections for animal R504. **a)** Tetrode tract from the left hemisphere, **b)** same as panel a) but for right hemisphere. Arrow heads point to recording location

4.2.5. Behavioral Analysis

To analyse interest in the cued arm during goal-cueing the time the animal spent by the barrier (within 20cm from the barrier) on the cued side was estimated. Specifically, the stem was notionally divided in two width way – cued and uncued side. The time an animal spent on the cued side was divided by the total time spent by the barrier. Thus, a value above 0.5 indicates the animal spent more time on cued side than uncued side. In addition, we also analysed time spent by the barrier, versus another part of the stem equal in size – in this case, the start of the stem. Similar to the cueing-bias measure, we divided the time spent by the barrier by the total time spent by the barrier and start of the stem, a value above 0.5 indicating the animal spent more time by barrier than start of stem. We performed this analysis before and after goal-cueing. We were not able to perform statistical analyses on these performance measures since we only had four data points (four animals).

4.3. Results

4.3.1. Behavioural Observations

The animals had not been trained to run on linear tracks prior to their participation in the experiment. Consequently, when the animals were first placed on the stem during RUN1 they were a bit hesitant to run on the track, yet quickly developed stereotypical running behaviour. Similarly, for RUN2 all animals were hesitant to run on the novel arms of the T-track, and showed much interest in the junction and barrier that controlled which arm they ran on. With time, however, they developed stereotyped behavior, running in alternating L-shaped laps

With regards to the GOAL-CUE period, animals generally increased the time they spent by the barrier versus other parts of the track, such as the start, after goal-cueing. The 'Pre Bias' and 'Post Bias' columns of table 4.4 show this effect On this measure a value of 0.5 indicates equal time spent by the barrier as elsewhere on the stem and values above 0.5 more time spent by the barrier than elsewhere. Table 4.4 shows that before GOAL-CUE animals spent about the same time by the barrier and the start of the stem (i.e. average Pre Bias = 0.48), whereas once goal-cueing had begun they generally spent more time by the barrier (average Post Bias = 0.64). More importantly, all animals spent more time on the cued side of the track by the barrier during GOAL-CUE, leaning over to the cued arm, than on the uncued side (average Cueing Bias = 0.66) – see 'Cueing Bias' column in table 4.4. We quantified this bias in a similar way to the pre/post bias – namely, values above 0.5 indicate more time spent on the cued side of stem. We did not carry out statistical analyses on these behavioural variables due to insufficient power (i.e. only four data points).

Table 4.4. Behavioural Biases Before and After GOAL-CUE

Rat	Pre Bias	Post Bias	Cueing Bias
R1838	0.39	0.61	0.73
R584	0.51	0.81	0.65
R504	0.57	0.77	0.66
R505	0.46	0.38	0.6
mean	**0.48**	**0.64**	**0.66**

4.3.2 Evidence for Preplay and Bias towards Cued Arm

We recorded from a total of 213 place cells across four animals. Place cells active on an arm were ordered according to the (peak) location of their place field, in order to construct place cell template sequences. Activity for runs *to* and *from* the arms was analysed separately. Consequently, we analysed correlations between spiking events in sleep (SLEEP1 and SLEEP2) and the four templates – Up Cued Arm (UCA), Down Cued Arm (DCA), Up Uncued Arm (UNA), and Down Uncued Arm (DNA). However, for most analyses we combined the results for runs in the different directions, to enable a more parsimonious comparison of the cued and uncued arms. Figure 4.6 below shows the template sequences, for the cued and uncued arm for all animals as well as examples of spiking and preplay events for each template.

We recorded a total of 1628 spiking events for the cued arm in SLEEP2 and 1657 for the uncued arm. In SLEEP1 a similar amount of events were recorded (cued arm = 1307, uncued arm = 1670). The number of events recorded per animal varied considerably, see table 4.5-6 below for details of number of spiking and preplay events recorded per animal for both sleep sessions.

Table 4.5 Number of Spiking and Preplay Events in SLEEP2

Rat	Template	Number of spiking events	Number of preplay events	Percentage of preplay events
R1838	Cued	44	9	20.45
	Uncued	26	3	11.53
R504	Cued	516	41	7.95
	Uncued	373	19	5.09
R505	Cued	631	38	6.02
	Uncued	860	29	3.37
R584	Cued	437	32	7.32
	Uncued	398	22	5.53
All rats	Cued	1628	120	7.37
	Uncued	1657	73	4.4

Table 4.6 Number of Spiking and Preplay events in SLEEP1

Rat	Template	Number of spiking events	Number of preplay events	Percentage of preplay events
R1838	Cued	63	3	4.76
	Uncued	35	2	5.71
R504	Cued	311	6	1.93
	Uncued	173	11	6.35
R505	Cued	664	39	5.87
	Uncued	1215	38	3.13
R584	Cued	269	14	5.2
	Uncued	247	21	8.5
All rats	Cued	1307	62	4.74
	Uncued	1670	72	4.31

a)   b)   c)   d)

Cued Arm   **Uncued Arm**

R1838

R504

R505

R584

Cell ID

20cm   100ms   20cm   100ms

Position on   Time (ms)   Position on   Time (ms)
arm (cm)   arm (cm)
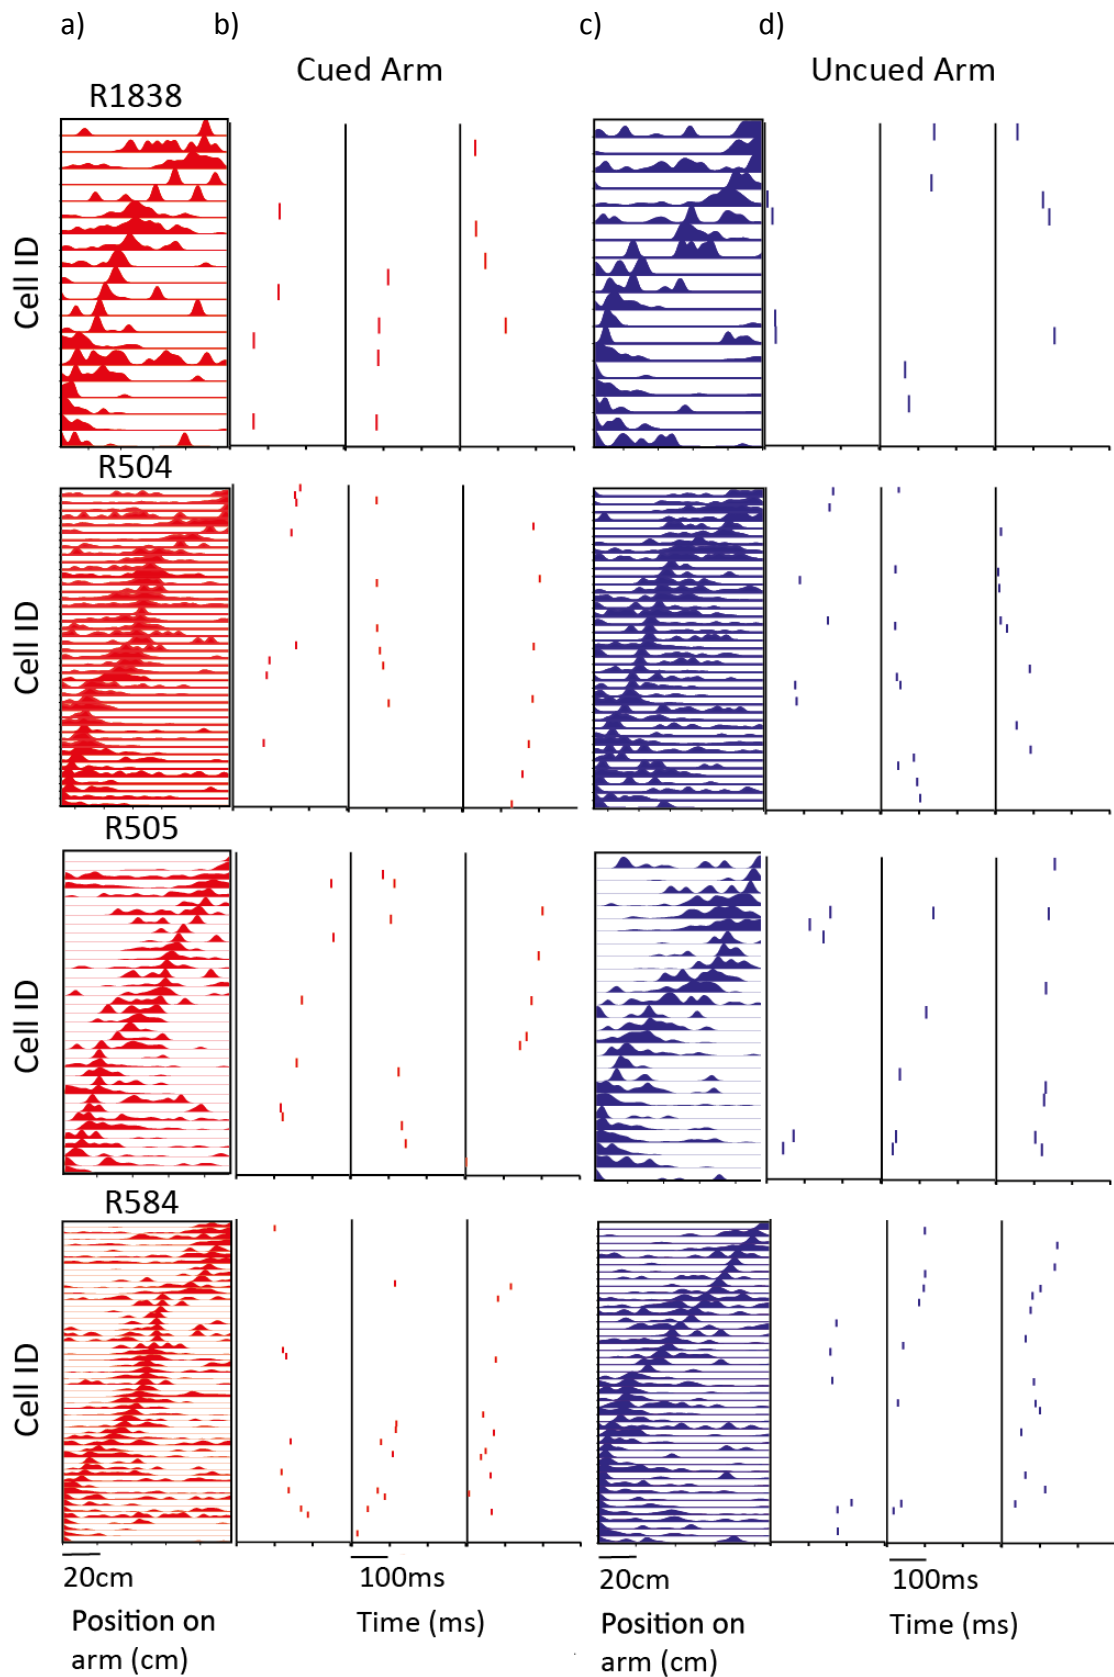
Figure 4.6. Template Sequences and Preplay Events in SLEEP2 a) Template

sequences for cued arm for all four animals. Cells active on cued arm are ordered according

to the location of the peak of their place field. Y-axis indicates cell id, x-axis position on arm

(cm). **b)** Representative preplay events for each of the cued templates. X-axis represents time (ms). **c-d)** Same as a-b but for uncued template. Note not all spiking events in d) are preplay events (i.e. have a significant correlation with the template sequence).

First we assessed whether the number of preplay events observed during SLEEP2 was higher than what would be expected by chance. To address this, we randomly shuffled each spiking event a 100 times and estimated the number of preplay events obtained for this shuffled data. Finally, we compared these chance level estimates to those we obtained in our data using a one-sample binomial test. The cued template was preplayed significantly more than chance (Cued = 7.37%, chance = 4.86%, p = 4.12 x 10^{-6}) (figure 4.7d). On the other hand, the uncued template was not preplayed more than would be expected by chance (uncued = 4.44%, chance = 4.22, p = 0.39). This pattern was present in all four animals (R1838: cued = 20.45%, chance = 6.86%, p = 6.59 x 10^{-4}, uncued = 11.53, chance = 6.77, p = 0.26; R504: cued = 7.95%, chance = 5.14%, p = 0.0044, uncued = 5.09%, chance = 4.4%, p = 0.21; R505: cued = 6.02%, chance = 4.57%, p = 0.038, uncued = 3.37%, chance = 3.82%, p = 0.8229; R584: cued = 7.32%, chance = 4.73%, p = 0.0063, uncued = 5.53%, chance = 4.74, p = 0.19).

Histogram distributions of correlation coefficients for the spiking events in SLEEP2 and the future cued and uncued templates are shown in figure 4.7a-b. The grey bars in each histogram represent the distribution of correlation coefficients obtained from the shuffling procedure and the coloured bars represent the correlation coefficients obtained from the unshuffled data. If spiking events represent a preplay of either template one would expect the bins for the data to be higher than that of the shuffle for highly positive and negative correlation coefficients (i.e. the tails of the distribution); in other words, there would be a higher proportion of (absolute) high correlations in the data compared to the shuffled distribution. Moreover, the reverse pattern would be expected to be present for low correlation coefficient bins. In contrast, if the spiking events do not represent preplay of future place cell templates the distributions of shuffled and data correlation would be expected to be the same. For panel b) in figure 4.7 (which represents the data for the cued template), one can see there are more high (both negative and positive) correlations in the data compared to the shuffle – coloured bars are taller than the grey bars for the tails of the distribution. Moreover, there seems to be fewer low correlations in the data compared to the shuffled

distribution. This is consistent with the hypothesis that spiking events in SLEEP2 preplay the future place cell sequence for the cued arm. A different pattern is evident for the distributions shown in panel a), representing spiking events for the uncued template. The data distribution seems to be rather clustered at low correlations, and the tails of its distribution do not exceed those of the shuffled distribution; indicating the spiking events do not represent significant preplay of the future uncued arm.

To address the main experimental hypothesis of whether preplay can be influenced by motivational variables such as goal-cueing we compared the proportion of significant spiking events (preplay events) for the cued and uncued arms using a Fisher's Exact test. Figure 4.7d below shows the mean proportion of preplay events for all animals combined as well as for individual animals. Overall we found the cued template was preplayed significantly more than the uncued template (cued = 7.37%, uncued = 4.4 %, $p = 3.0 \times 10^{-4}$). Separate analyses for each animal showed a similar trend: all animals exhibited an effect in the same direction though this only reached statistical significance for animal R505 (R1838: cued = 20.45%, uncued = 11.53%, p = 0.26; R504: cued = 7.95%, uncued = 5.09%, p = 0.052; R505: cued = 6.02%, uncued = 3.37%, p = 0.01; R584: cued = 7.32%, uncued = 5.53%, p = 0.18).

To ensure our results were not method specific we re-estimated the statistical significance comparing the proportions obtained for the cued and uncued template using a two sample binomial test. Again we found that the cued arm was significantly preplayed more than the uncued arm ($p = 2.91 \times 10^{-4}$). Difference in the same direction was observed for all four animals, with the difference reaching statistical significance for two out of four animals (R1838: p = 0.13; R504: p = 0.049; R505: p = 0.01; R584: p = 0.11).
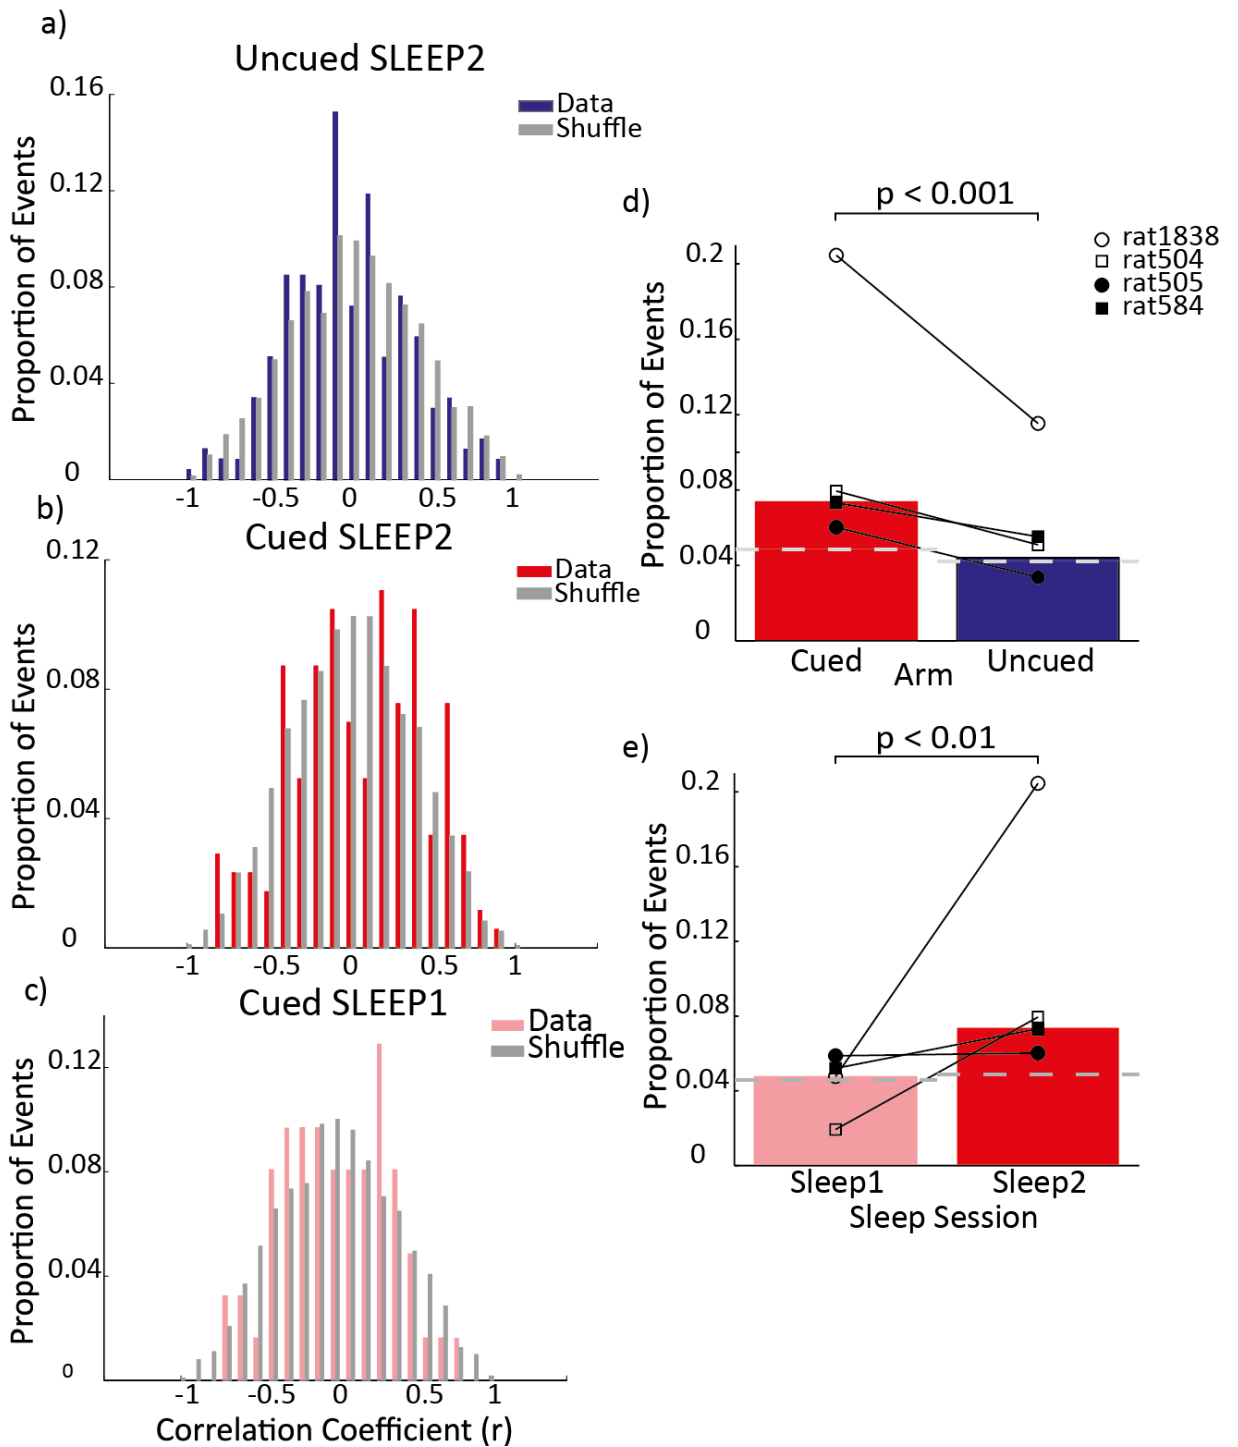
Figure 4.7. Preferential Preplay of Cued Arm in SLEEP2 a) A representative histogram distribution from R584 of correlation coefficients for spiking events for the uncued template against its shuffle distribution in SLEEP2, coloured bars represent correlation coefficients obtained in the data, and the grey bars the correlation coefficients obtained by randomly permuting (shuffling) the order in which the cells fired during spiking events a 100 times. If future place cell sequences are preplayed during SLEEP2 one would expect the coloured bars to be higher than the grey bars for the tails of the distributions; implying there are

more high correlations (both positive and negative) in the data than in the shuffled distribution. Alternatively, if the data distribution does not deviate from its shuffle then that indicates a place cell sequence is not preplayed more than what would be expected by chance. **b-c)** Same as a) but for cued template in SLEEP2 and SLEEP1, respectively. X-axes represent correlation coefficients, y-axis proportion of spiking events **d)** Average proportion of preplayed events, combined across all animals for cued and uncued arm in SLEEP2. **e)** Same as d) but for cued template in SLEEP1 and SLEEP2. Lines represent results for each animal, the gray dashed line the proportion of preplay events one would expect to see by chance, x-axis the arm/sleep session analysed and y-axis proportion of all spiking events.

Thus, it appears, in agreement with the experimental hypothesis, the cued template is preplayed significantly more than the uncued template in SLEEP2, once goal-cueing has occurred. In fact, based on the chance-level analysis, one may conclude that the uncued template is not preplayed at all. Consequently, one might suggest preplay is a result of motivational influences. However, from the foregoing analysis one cannot rule out the effect of pre-configuration in the CA1 circuitry (McNaughton, Barnes et al. 1996); namely, the reason the cued template is preplayed more than the uncued template is because the place cell sequence for that template was already existing in the hippocampus prior to goal cueing. In other words, the cued template was also preplayed in SLEEP1. To rule out this explanation, exactly the same preplay analysis was carried out for spiking events in SLEEP1, and an increase in number of preplay events going from SLEEP1 to SLEEP2 assessed.

4.3.3 Preplay of Cued Arm is not the Result of a Pre-Existing Map in the Hippocampus

To address the question of whether the preferential preplay of the cued template in SLEEP2 was the result of a pre-existing place cells sequence in CA1, spiking events in SLEEP1 were analysed. Firstly, we assessed whether the number of preplay events for the cued template in SLEEP1 significantly exceeded that expected by chance (based on the shuffling procedure described above). Across all animals, the proportion of preplay events for the cued template in SLEEP1 was not significantly higher than chance levels (chance = 4.67%, p = 0.38), based on a one-sample binomial test. For three out of four animals, this pattern was replicated – for rat R505, a weakly significant effect was found. R1838: 4.75% vs. 7.05%, p = 0.66; R504:

1.93% vs. 4.44%, p = 0.99; R505: 5.87% vs. 4.34%, p = 0.025; R584: 5.2% vs. 4.67%, p = 0.38. These results are demonstrated in figure 4.7e. Moreover, figure 4.7c shows the histogram distribution of correlation coefficients for spiking events for the cued template in SLEEP1 for rat R584 against its shuffled distribution. The distribution clearly shows the proportion of high correlations in the data is lower than that of the shuffled data; implying, in SLEEP1, spiking events do not significantly represent the future place cell sequence for the cued arm.

Moreover, to directly assess whether the number of preplay events of the cued arm increased following goal-cueing (as would be consistent with the hypothesis that motivational influences drive preplay) the proportion of preplay events for the cued template observed in SLEEP1 was compared to that observed in SLEEP2. The total proportion of preplay events obtained in SLEEP1 for the cued template was 4.74%, which is significantly lower than that obtained in SLEEP2 (p = 0.0034, see figure 4.7e), based on a Fisher's Exact test. Moreover, this pattern was obtained for all four animals, reaching statistical significance for two. For R1838, SLEEP1 = 4.76%, SLEEP2 = 20.45%, p = 0.013; R504: SLEEP1 = 1.93%, SLEEP2 = 7.95%, p = 1.0×10^{-4}; R505: SLEEP1 = 5.87%, SLEEP2 = 6.02%, p = 0.50; R584: SLEEP1 = 5.2%, SLEEP2 = 7.32%, p = 0.17). A similar pattern of results was obtained when data was analysed using a two-sample binomial test - all data: p = 0.0027; R1838: p = 0.0081; R504: p = 2.85×10^{-4}; R505: p = 0.20; R584: p = 0.11.

However, one might contend that analysing preplay of cued arm in SLEEP1 might not be the appropriate control for addressing the issue of possible pre-existing hippocampal sequences, as SLEEP1 is temporally farther apart from RUN2 (session where animals run on cued and uncued arms) than SLEEP2. To address this concern, we analysed preplay of the template sequence for the stem obtained from RUN1. Similar to before, we estimated the number of preplay events for the stem in SLEEP1 and compared it to chance levels. In SLEEP1 we obtained a total of 1773 spiking events for the stem, however only 4.12% had a significant correlation with the stem sequence in RUN1 (i.e. were preplay events). This proportion did not exceed that expected by chance (chance = 4.08%, p = 0.4408). Moreover, this result was obtained when each animal was analysed separately (R1838: 4.41% vs. 6.01%, p = 0.59; R504: 4.55% vs. 3.5%, p = 0.16; R505 3.57% vs. 4.04%, p = 0.74; R584: 5.17% vs. 4.53%, p = 0.24). Table 4.7 below details the number of spiking events for each animal as well as the number of preplay events. Furthermore, figure 4.8 shows a representative histogram distribution of correlation coefficients for spiking

events from SLEEP1 from rat R584 against its shuffled distribution. The histogram shows the distribution of correlation coefficients for the spiking events for the data is similar to the distribution for the shuffled events. The tails of the data distribution do not consistently exceed those of the shuffled distribution. Moreover the distribution seems clustered at 0; implying the stem was not preplayed in SLEEP1.

Table 4.7 Number of Spiking and Preplay Events for Stem in SLEEP1

Rat	Number of spiking events	Number of Preplay Events	Percentage if Preplay Events
R1838	68	3	4.41%
R504	396	18	4.54%
R505	980	35	3.57%
R584	329	17	5.17%
all rats	**1773**	**73**	**4.12%**
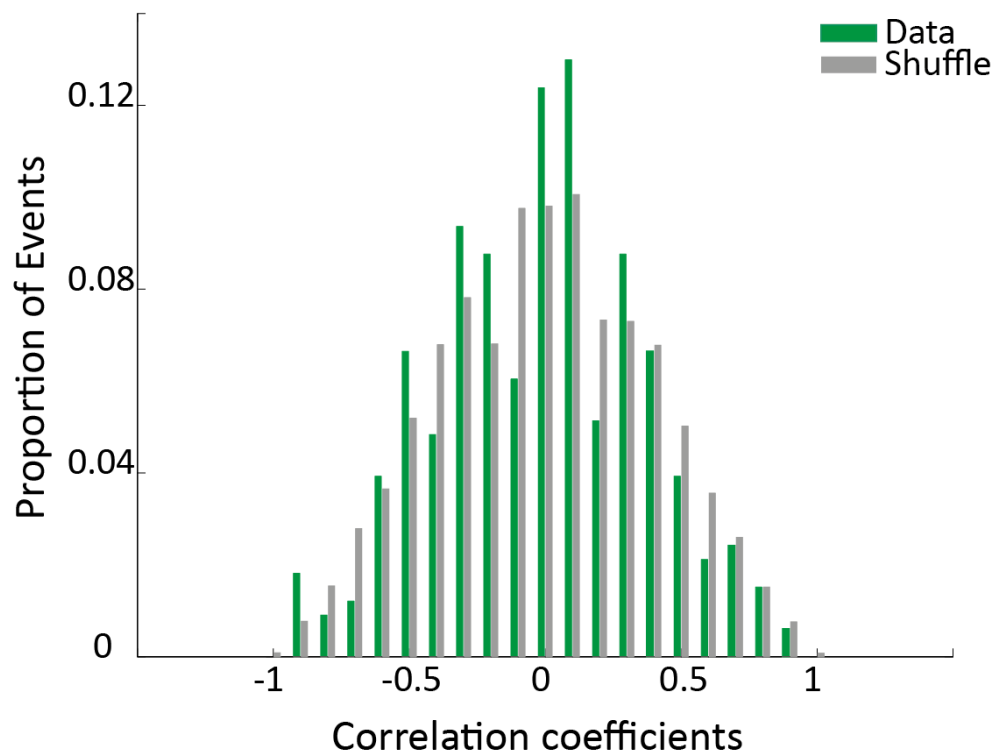
Figure 4.8 Frequency Distribution of Correlations for Stem in SLEEP1 A histogram of correlation coefficients for spiking events obtained for rat R584 for its stem template in SLEEP1. Coloured bars represent correlation coefficients for spiking events in SLEEP1 and the grey bars the correlation coefficients obtained from shuffling the original spiking events.

The tails of the data distribution are not higher than those of the shuffle, moreover the distribution seems peaks at a correlation of 0. This implies the stem template is not preplayed during SLEEP1. X-axis shows correlation coefficients and y-axis proportion of all spiking events.

In sum, considering the cued arm and stem are not preplayed significantly above chance in SLEEP1 it is unlikely that the biased preplay for the cued arm in SLEEP2 is due to pre-existing biases in the CA1 circuitry. Rather, it may be concluded that the bias develops as a consequence of motivational influences, such as prior behavioural cueing.

4.3.4. Preplay of Cued Arm is not confounded by Template Length, Spike Sorting Quality or Proximity to Cued Arm

Upon consideration of the number of cells active on each arm (see Table 4. 3) one might be concerned the preferential preplayed of the cued arm is due to the cued template generally containing more cells than uncued templates. Specifically, for all four animals the template for a run *to* the cued arm (UCA) contains the highest number of cells. On average, cued arm templates contain 35 cells, whereas the uncued arm templates contain 32 cells. A higher number of cells in a template might introduce biases as well as endow the cued templates with higher statistical power. To address this possible confound, we down-sampled templates from each animal so that they were equal in length to the shortest template. The results of this analysis are shown in Figure 4.9a-b. Briefly, down-sampling the templates did not affect the pattern of results described above. In SLEEP2 the cued template was still preplayed significantly more than the uncued template (cued = 7.18%, uncued = 4.37%, $p = 0.0073$), and the cued template was preplayed significantly more in SLEEP2 than in SLEEP1 (SLEEP2 = 7.18%, SLEEP1 = 4.13%, $p = 0.016$), based on a Fisher's Exact test. Moreover, a trend in the right direction was observed when data from individual animals were analysed separately (R1838: cued vs uncued: $p = 0.34$, SLEEP2 vs SLEEP1: $p = 0.12$; R504: cued vs uncued: $p = 0.20$, SLEEP1 vs SLEEP2: $p = 0.003$; R505: cued vs uncued: $p = 0.078$, SLEEP2 vs SLEEP1: $p = 0.49$; R584: cued vs uncued: $p = 0.35$, SLEEP1 vs SLEEP2: $p = 0.17$). Results were found to be very similar when the data was analysed using a two-sample binomial test – all data: cued vs uncued: $p = 0.0045$, SLEEP1 vs SLEEP2: $p = 0.01$; R1838: cued vs uncued: $p = 0.15$, SLEEP1 vs SLEEP2: $p = 0.063$; R504: cued vs uncued: $p = 0.12$, SLEEP1 vs SLEEP2: $p =$

0.0042; R505: cued vs uncued: p = 0.052, SLEEP1 vs SLEEP2: p = 0.19; R584: cued vs uncued: p = 0.17, SLEEP1 vs SLEEP2: p = 0.12. See table 4.8 and 4.9 below for details of number of spiking and preplay events.

Table 4.8 Number of Spiking and Preplay Events in SLEEP2 for Down-Sampled Data

Rat	Template	Number of spiking events	Number of preplay events	Percentage of preplay events
R1838	Cued	24	5	21.52
	Uncued	17	2	9.98
R504	Cued	378	27	7.27
	Uncued	184	9	5.08
R505	Cued	232	13	5.71
	Uncued	629	20	3.2
R584	Cued	286	20	7.08
	Uncued	354	21	5.82
All rats	Cued	920	66	7.18
	Uncued	1184	52	4.37

Table 4.9. Number of Spiking and Preplay Events in SLEEP1 for Down-Sample Data

Rat	Template	Number of spiking events	Number of preplay events	Percentage of preplay events
R1838	Cued	39	3	6.65
	Uncued	22	1	3.61
R504	Cued	215	4	2.05
	Uncued	71	4	5.55
R505	Cued	198	10	5.18
	Uncued	965	29	3.06
R584	Cued	171	8	4.97
	Uncued	213	16	7.46
All rats	Cued	623	26	4.13
	Uncued	1271	50	3.94
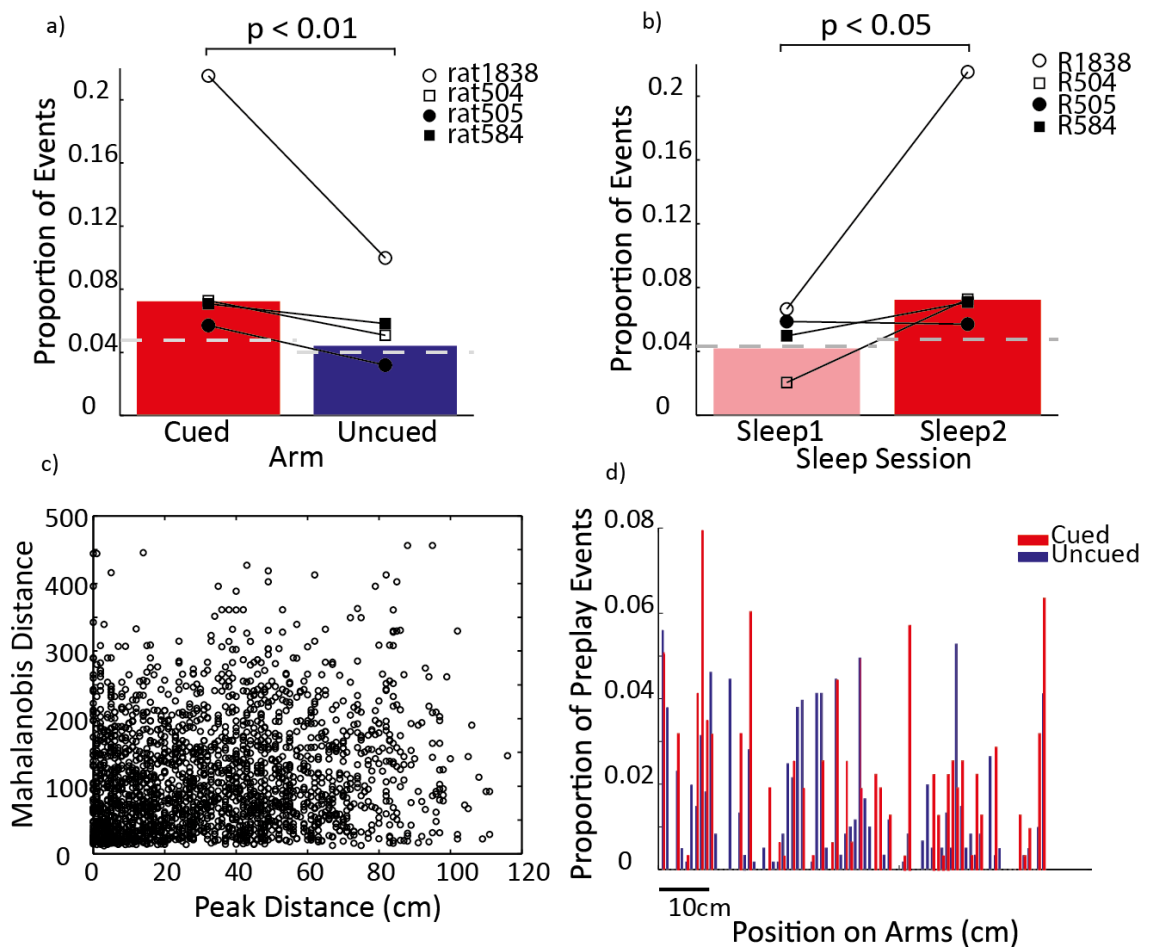
Figure 4.9 Control Analysis for Template Length, Cluster Quality and Proximity to Cued Arm a) Average proportion of preplay events in SLEEP2 for cued and uncued template after down-sampling templates. Each line represents results for individual animals, grey dashed line the proportion of preplay events one would expect by chance, x-axis show arm/sleep session and y-axis the proportion of spiking events **b)** same as a) except showing preplay of cued template in SLEEP2 and SLEEP1. **c)** Mahalanobis'distance (y-axis) for pairs of cells active on an arm plotted against the distance between their place field peaks (cm) (x-axis). **d)** Distribution of positions on the arm represented in preplay events in SLEEP2 for cued and uncued template. X-axis represents location on either arm (cm), and y-axis the frequency with which a particular position on the arm was represented in preplay events. Red bars represent the distribution of preplayed positions for the cued template and blue bars that of the uncued template.

Cluster quality could affect one's estimates of preplay. For example, if complex spikes are mistakenly divided into two clusters the spikes from these clusters would be highly correlated during spiking events, thus artificially inflating correlations. One way to assess this is to look at the relationship between the distance between the place field of pairs of cells and their Mahalanobis' distance. If cells with place

fields close together in space tend to have lower Mahalanobis' distances it would indicate a possible confound. Thus, we analysed this relationship (see Figure 4.9c) and asked whether the distribution of Mahalanobis' distances for short field distances was different from that of long field distances. Comparing the distribution of Mahalanobis' distances for the two groups revealed a non-significant result (2-sample Kolmogorov-Smirnov test: 0.19, p = 0.67). Thus, it seems unlikely that our estimates of preplay are a confound of poor cluster separation.

Finally, one might speculate the enhanced preplay of the cued arm is a result of the animals being closer to the cued arm than the uncued arm during GOAL-CUE. Namely, during goal cueing the animals showed more interest in the cued arm – indexed by time spent on the cued side of the stem, leaning over the arms, etc. Thus, during this period the animals might have occupied the place fields of the cued arm. The enhanced preplay of the cued arm might then merely reflect *replay* of the place fields the animals leaned into during goal cueing. If this was the case, one might expect locations closer to the animal during GOAL-CUE to be preplayed more than locations further away from the animal. Thus, we asked were locations on the first half of the cued arm preplayed more than locations on the second half (the first 20cm of the arm was excluded as it is common to both cued and uncued arms). Combined across the four animals, 54.24% of locations in preplay events belonged to the first half of the arm and 45.76% belonged to the second half of the arm. This difference was not found to be statistically different on a one-sample binomial test (p = 0.22). For the uncued arm, a similar pattern of results was obtained – 47.62% of positions in preplay events belong to first half, 52.38% of positions belong to second half, this difference was not significant (p = 0.56). Moreover, the cued and uncued arms did not differ significantly - first half: 54.24% vs 47.62%, p = 0.13; second half: 45.75% vs 52.38%, p = 0.81. See figure 4.9d for distribution of positions represented in preplay events for cued and uncued arm. Thus, we do not think the preferential preplay of the cued arm reflects *replay* of the place fields occupied during goal cueing.

To summarise, we have here addressed various possible confounds in our data, and found that none of them can account for the enhanced preplay of the cued template in SLEEP2 we observed. Consequently, we do not think the bias in preplay is the result of a higher statistical power for the cued template, poor spike sorting

or physical proximity to the cued arm. Rather we think it is driven by interest in the cued arm.

4.3.5 Physiological and Behavioural Properties of Preplay

Past studies on replay and preplay have found that place cell sequences are activated both in the forward and reverse order. For example, if 5 cells become active on a linear track in the A-B-C-D-E order, in sleep they are either activated in that forward order or the reverse order – E-D-C-B-A (Foster & Wilson, 2006; Diba & Buzsaki, 2007). Some have suggested a functional difference between forward and reverse preplay, namely that reverse replay may play role in reinforcement learning, such as encoding the location of a reward (Barry and Bush 2012, Foster and Knierim 2012), whereas forward replay may predict the future path of the animal (Foster and Wilson, 2006). Considering, both functions are relevant to the current study, one might expect a mixture of forward and reverse preplay events. This is indeed what we found – considering only forward spiking events (events with a positive correlation with a template sequence), 6.93% have a significant correlation with the future place cell sequence, and considering only reverse spiking events (events with a negative correlation with a template sequence), 7.8% can be classified as preplay events. Although there is a subtle difference between the proportion of forward and reverse events, with reverse events being more common, this difference is not statistically significant on a two-sample binomial test ($p = 0.16$). For each animal, reverse preplay events tended to be more frequent than forward preplay events, with the exception of one animal (R1838: forward = 31.81%, reverse = 9.09%, $p = 0.0034$; R504: forward = 6.4%, reverse = 9.4%, $p = 0.09$; R505: forward = 5.57%, reverse = 6.49%, $p = 0.19$; R584: forward = 7.04%, reverse = 7.59%, $p = 0.19$). See figure 4.10a below for a demonstration of these results. Thus, the motivationally-biased preplay we observed in this study is underpinned by preplay events that represent runs on the cued arm in the forward and reverse order.
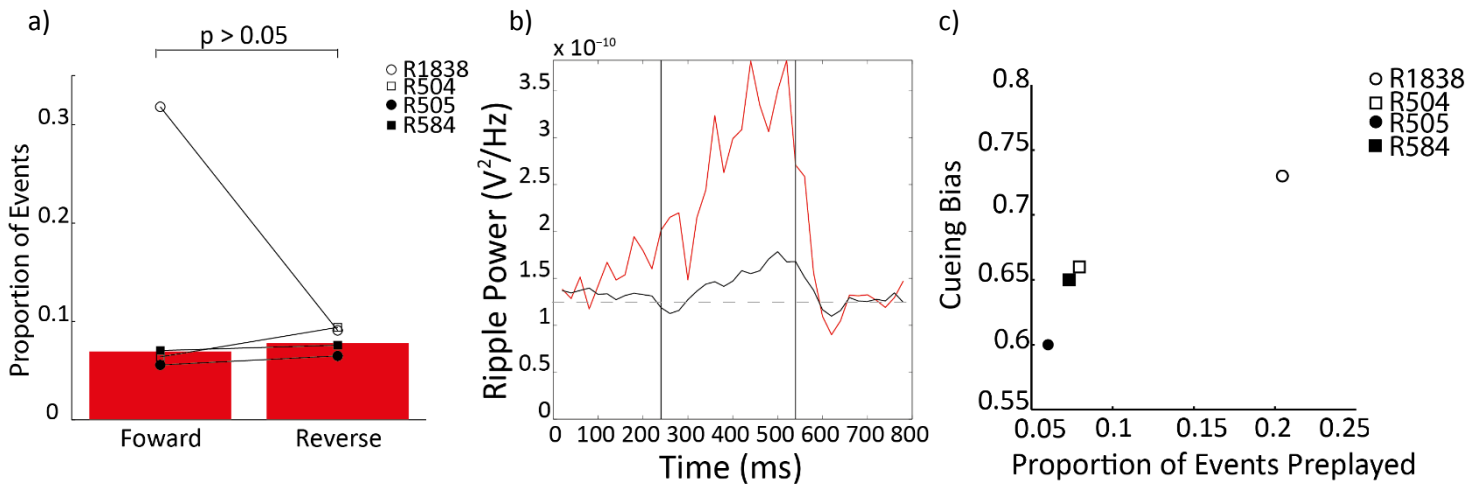
Figure 4.10 Properties of Preplay a) Preplay events activated cued template both in forward and reverse order. X-axis order in which cells are active during preplay events in relation to the place cell sequence of the template, y-axis proportion of spiking events. **b)** Ripple power during preplay (red) and spiking events (black) as well as mean power during non-event (grey dashed) periods. Vertical lines mark start and end of event, respectively. X-axis shows time (ms) and y-axis power in the ripple spectrum (V^2/Hz). **c)** Relationship between cueing bias during GOAL-CUE and amount of preplay in SLEEP2. X-axis – proportion preplay, y-axis – cueing bias.

Previous studies have observed replay and preplay during both awake immobility (such as when a rat is eating food on a track) and during sleep (Karlsson and Frank 2009, Gupta, van der Meer et al. 2010, Dragoi and Tonegawa 2011). Accordingly, we investigated preplay during immobility in GOAL-CUE. Although it is important to bear in mind that the GOAL-CUE period was short (~10min) and the animal was often active during this period; making spiking events, perhaps, relatively infrequent. We identified a total of 695 spiking events during GOAL-CUE, roughly 1000 events fewer than we identified during the sleep sessions. In this small sample, we did not find the proportion of preplay events for either template to significantly exceed chance levels (cued = 4.95%, chance = 4.48%, p = 0.29; uncued = 5.12%, chance = 3.88%, p = 0.14). With a longer recording period and more spiking events we might have seen a significant result. See Table 4.10 below for number of spiking and preplay events for this analysis. In conclusion, we did not observe consistent preplay of cued or uncued arm during goal-cueing as we did in the subsequent sleep session.

Table 4.10. Number of Spiking and Preplay events during GOAL-CUE

Rat	Template	Number of spiking events	Number of preplay events	Percentage of preplay events
R1838	Cued	15	0	0
	Uncued	14	0	0
R504	Cued	36	1	2.78
	Uncued	33	3	9.09
R505	Cued	231	13	5.63
	Uncued	282	9	3.19
R584	Cued	41	2	4.88
	Uncued	43	7	16.28
All rats	Cued	323	16	4.95
	Uncued	372	19	5.11

In the important paper by Wilson & McNaughton (1994), where replay was first reported, its occurrence was tightly correlated with sharp-wave ripples (SWRs). Many studies, since then, have replicated this observation (Karlsson and Frank 2009, Dupret, O'Neill et al. 2010). Thus, we were interested in looking at the relationship between spiking and preplay events and LFP power in the ripple frequency range. Figure 4.10b clearly shows that for preplay events the mean ripple power is low at the start and end of an event and peaks near the middle. Comparing mean ripple power during preplay events to other periods using an independent samples t-test (excluding spiking event periods) we found a large, reliable difference, showing power in the ripple band (80-250Hz) was significantly higher during preplay events than other comparable periods (mean preplay event power = 2.84 x 10^{-10} V^2/Hz (SD = 2.64 x 10^{-10}), mean non-event power = 1.18 x10^{-10} V^2/Hz (SD = 1.33 x10^{-10}), $t_{(653922)}$ = 29.11, p = 3.44 x 10^{-186}). These results were found to be reliable across all four animals (R1838: $t_{(156293)}$ = 5.51, p = 3.63 x 10^{-8}; R504: $t_{(142078)}$ = 23.91, p = 4.66 x 10^{-126}; R505: $t_{(176313)}$ = 24.1, p = 4.55 x 10^{-128}; R584: $t_{(179238)}$ = 6.40, p = 1.60 x 10^{-10}). Interestingly, ripple power during preplay events was also significantly higher than that during spiking events – $t_{(127153)}$ = 18.29, p = 1.48 x 10^{-74} (mean spiking event power = 1.48x 10^{-10} V^2/Hz, SD = 7.59 x 10^{-11}). Again, this result was observed for each animal individually – R1838; $t_{(24298)}$ = 5.78, p =7.71 x 10^{-09}; R504: $t_{(38878)}$ = 14.40, p =6.46 x 10^{-47}; R505: $t_{(28633)}$ = 9.33, p =1.11 x 10^{-20}; R584: $t_{(35338)}$ = 8.68, p = 4.04 x 10^{-18}). See figure 4.10b for a plot of mean ripple power for preplay, spiking and non-events. In sum, in agreement with past studies, we find

that preplay events are associated with a significant increase in the LFP frequency band associated with ripples.

Finally, since we speculate the bias in preplay is driven by the animals' showing preferential interest in the cued arm over the uncued arm during GOAL-CUE, one might hypothesise that the amount of preplay will correspond to how much interest the animal showed in the cued arm. In other words, the cueing-bias measure described earlier, indexing the interest each animal showed in the cued arm, may correlate with the amount the cued arm is then preplayed in SLEEP2. We did indeed find a positive relationship between interest in cued arm and the amount the cued arm was preplayed, see figure 4.10c. As this analysis was only based on four data points, we cannot statistically test this relationship, but this result does lend further support to the notion that the bias in preplay is driven by motivational influences.

4.4 Discussion

The principal finding of this study is that during restfulness, CA1 preplays behaviourally-cued trajectories through unexplored space. The proportion of preplay events we found for the cued environment was significantly higher than that for the uncued environment. In fact, the number of preplay events we observed for the uncued environment did not exceed that which would be expected by chance; possibly indicating the uncued environment was not preplayed at all. Importantly, we did not find evidence for the notion that preplay derives from pre-existing maps in the hippocampus (i.e. pre-configuration (McNaughton, Barnes et al. 1996). The cued environment was not found to be preplayed above chance levels in SLEEP1. Even when we control for the difference in the temporal separation between SLEEP1 and SLEEP2 from RUN2, by assessing preplay for RUN1 in SLEEP1, we still found no evidence of pre-configuration. Finally, we found no evidence for our results to be confounded by differences in statistical power between the cued and uncued arm analyses, spike-sorting quality or physical proximity of the cued arm during GOAL-CUE. Furthermore, preplay events were associated with increased ripple power and the amount of preplay corresponded to the behavioural bias the animals displayed to the cued arm. Together, we believe this provides compelling evidence for the hypothesis that preplay is functional, and driven by goal-directed, motivational influences.

The difference in proportion of preplay events for the cued and uncued template was subtle – namely, 3% on average, corresponding to 47 events. Moreover, a difference of a similar magnitude was found between the proportion of preplay events observed and the proportion expected by chance. Importantly, these effects were replicated for each animal individually. Effects of varying sizes have been reported in the replay literature. For example, Foster & Wilson (2006) reported almost a third of spiking events represented replays of recently run trajectories. While Lee and Wilson (2002), in the first study showing replay of sequential place cell sequences on a linear track, found 10-20% of spiking events represented trajectories (Lee and Wilson 2002). When one examines the replay literature that touches on the behavioral correlate of replay, even smaller effect sizes are reported (Bendor and Wilson 2012, Pfeiffer and Foster 2013). Finally, Dragoi and Tonegawa (2013) reported 6-9% of spiking events represented, future, novel place cell sequences. Thus, the effect size of our results is in accordance with

that of behaviorally relevant replay and preplay of novel trajectories and suggests these effects may be subtle yet reliable.

The motivationally biased preplay we observe here fits with multiple reported findings of an impaired ability to imagine/simulate novel events in hippocampal patients (Hassabis, Kumaran et al. 2007) and findings from functional-MRI studies showing the hippocampus to be just as involved in construction and recollection of events (Buckner and Carroll 2007, Schacter, Addis et al. 2012). Moreover, the reported findings also appear to be in agreement with recent behavioral studies in humans showing the known memory-enhancing effect of sleep is modulated by expectation of future use (Wilhelm, Diekelmann et al. 2011, Born and Wilhelm 2012). On the other hand, these findings do not fit so nicely in to the replay/preplay framework.

Replay has been studied extensively over the past decade, and has been thought to represent the mechanism underlying consolidation (Wilson and McNaughton 1994), although some recent theoretical speculation posit it as well as a mechanism for navigational planning (Johnson and Redish 2007, Erdem and Hasselmo 2012, Pfeiffer and Foster 2013). Preplay has received considerably less attention, although two recent studies have attracted some attention. Dragoi and Tonegawa (2011) found preplay of a novel track, the animal was supposedly not aware of, while the animal was immobile on an adjacent track and during sleep prior to the animal ever running in the environment. Moreover, the same authors reported such preplay could extend to multiple novel, future environments (Dragoi and Tonegawa 2013). The authors interpreted their results as evidence for the existence of pre-configured maps in the hippocampus that become active when an animal enters an environment. Such a finding seems at odds with our results. Not only did we fail to detect preplay in SLEEP1, before the animals had seen the future environment, but also we did not find preplay of the uncued arm in SLEEP2. One might suggest the animals in Dragoi and Tonegawa's studies had cues available to them regarding their future environment, such as when immobile on the familiar track they might have been aware of the presence of another environment. In fact, Dragoi (2014) touches on the role of expectation in driving preplay. However, this explanation does not account for the reported preplay during sleep, prior to the animal even knowing of its future environment. To reconcile this finding with the one described here is difficult. Perhaps, task-demands modulate preplay in such a

manner that only trajectories of future relevance are preplayed. Expectations have been found to modulate the memory-enhancing effects of sleep in humans (Wilhelm, Diekelmann et al. 2011), such that memory for information thought to be relevant in the future was enhanced whereas memory for other information was not. In our study, animals were encouraged to be interested in one environment and not another, whereas in Dragoi and Tonegawa's studies no such biasing occurs. Furthermore, methodological differences between the two studies could also contribute to the diverging results; future research will hopefully shed light on these differences.

As mentioned above, replay has been proposed to underlie memory consolidation in the hippocampus (Wilson and McNaughton 1994), and also navigational planning (Erdem and Hasselmo 2012, Foster and Knierim 2012, Pfeiffer and Foster 2013, Singer, Carr et al. 2013). One indication of a planning function of replay comes from a study by Singer and colleagues (2013) where they observed enhanced coordination during SWRs preceding a correct choice on an alternation task. Moreover, sequences activated in SWRs seemed to represent coherent possible future paths. Another study carried out by Pfeiffer and Foster (2013) found replay events, during a goal-directed, navigational task, over-represented goals and seemed to anticipate the path the animal was about to take. Importantly, such 'prospective replay' only occurred before an animal initiated a goal-directed path The biased preplay of a behavioural-cued environment we observed in the current study may be interpreted as an extension of the navigational planning hypothesis of replay; where preplay allows for navigational planning in novel environments. Although it should be highlighted that future studies would need to be carried out to confirm this role of preplay. Moreover, the contributions of other spatially-modulated cells, such as head-direction and grid cells, in navigational planning should be understood better; evidence has hitherto either been lacking or negative (Brandon, Bogaard et al. 2012).

Before bringing this chapter to conclusion, I would like to discuss some hypotheses regarding the mechanism underlying motivationally–biased preplay. In the current study, a goal cue seems to be the driving factor for preplay. Brain regions coding for goals/rewards are widespread in the brain, and some send direct afferents to the hippocampus. For example, the ventral tegmental area (VTA), a prominent area in the study of reward-encoding in the brain (Schultz 1998), projects to the dentate

gyrus (DG) of the hippocampus (Swanson and Hartman 1975). Specifically, VTA neurons project to the deep polymorphic layer of the DG, which provides excitatory inputs to cells in the granule layer. The granule cells are responsible for the powerful excitatory input to CA3, which then project heavily to CA1. This VTA -> DG -> CA3 -> CA1 loop might provide the means for goal-related information to be relayed to the hippocampus and drive subset of neuronal ensembles to fire. Furthermore, the role of vision plays an obvious role in the current study; the animal sees the experimenter bait an arm the animal cannot access. The visual cortex has been found to interact with the hippocampus during sleep (Ji and Wilson 2007); perhaps, providing the means for recent visual information to be incorporated into spiking events. These suggestions are, at this point, merely speculative. Hopefully, future research will be able to dissect the role of different sets of cues in replay/preplay, and place cell firing more generally, in order for us to gain a deeper understanding of the role of the hippocampus for cognition and behaviour.

To summarise, this study has found preplay during rest to be biased by motivational factors, such as behavioural cueing to a particular unvisited environment. This bias in preplay cannot be simply explained by pre-existing maps in the hippocampal circuit. The function of such preplay remains yet to be elucidated, but we propose it may contribute to the hypothetical function of the hippocampus in navigational planning. Future studies that directly assess the relationship between preplay and navigational behavior might be able to confirm this role.

5. Single Cell Laser-Induced Ablations in the Mouse Cortex

5.1 Introduction

The foregoing chapters have described studies implicating small functional sub-groups of the CA1 pyramidal population in representing, and perhaps planning, future, goal-relevant, trajectories of a rodent. However, with these studies, and most traditional electrophysiological studies, it is important to bear in mind that findings are merely correlational - they do not show that the functional sub-population is *necessary* for successful behaviour. An important goal of systems neuroscience is to draw a causal link between neural activity and behaviour; to meet this end requires selective interference with the neural sub-populations involved in behaviour.

In the past, to understand the role of a particular brain region in behaviour scientists relied on humans incurring damage to an area of the brain. The behaviour or cognitive ability found to be impaired after the damage then allowed scientists to infer a causal link between the activity of a brain region and behaviour. For example, the hippocampectomy patient H.M. suffered as a result of his epilepsy taught us the hippocampus is needed for the formation of new memories and the recall of recent memories (Scoville and Milner 1957). Moreover, lesioning and pharmacological inactivation of single brain regions in rodents was an advancement in this field, allowing scientists to gain enhanced control over the extent of lesions induced and the power to reverse them. The seminal study by Kim and Fanselow (1992), for example, where the hippocampus was lesioned at different delays following training on a novel task revealed the important time-limited role of the hippocampus in memory storage (Kim, Fanselow et al. 1992). However, these techniques are generally crude and result in the permanent loss of a brain region, or, in the case of pharmacological inactivations, they are only slowly reversed.

In recent years, scientists have developed more sophisticated techniques to perturb neural activity. The most prominent of these developments is optogenetics. Optogenetics involves the use of genetically encoded molecules that couple light and neural function –delivered via the injection of a virus or through transgenic lines. The molecule most often used is Channelrhodopsin-2 (ChR2). Channelrhodopsin, when triggered by light, depolarises a cell resulting in an action potential. It can be used to 'turn' neurons on, on a millisecond basis. Importantly,

genetically-defined subpopulations of cells can be selectively targeted with optogenetics; allowing scientists to gain control over neuron sub-types within a brain region. For example, Tsai and colleagues (2009) used ChR2 to drive dopaminergic neurons in the ventral tegmental area (VTA) of the basal ganglia to firing in freely moving rodents, which resulted in behavioural conditioning (Tsai, Zhang et al. 2009); confirming the hypothesised role of these neurons in reward-dependent learning. More recently, optogenetics has been developed to control the activity of neurons in a region based on their projection pattern - by combining the conditional expression of an opsin gene with the presence of retrograde synaptic tracers. For example, Stuber and colleagues (2011) showed that the projection from the basolateral amygdala to the nucleus accumbens plays an important role in reward-related behaviour (Stuber, Sparta et al. 2011).

Although optogenetics represents an obvious step forward in neuroscientists' quest for elucidating the causal links between behaviour and neural activity, the neural sub-unit it enables control over might still be considered too crude. For example, the functional sub-groups of cells described in chapter 3 and 4 are all of the same genetic sub-type and likely display the same projection pattern; thus, probing the causal role of these functional circuits for behaviour could not be achieved with the aforementioned tools. More generally, in many regions of the brain neurons of the same sub-type and with the same projection patterns perform distinct behavioural functions. A small fraction of hippocampal pyramidal cells represent any given environment (O'Keefe and Dostrovsky 1971), and a small, spatially intermingled, population of primary motor cortex axons projecting to layers II/II of the barrel cortex have been found to code for distinct motor behaviours. (Huber, Gutnisky et al. 2012). This emphasises the need to develop methods to interfere with cell groups based on their activity pattern.

Recently, scientists have coupled optogenetics with some of the genetic markers of neural activity. Immediate early genes (e.g. c-Fos, Arc) have been found to be associated with neural activity in a brain region and behavioural learning (Reijmers, Perkins et al. 2007). Liu and colleagues (2012) conditionally expressed ChR2 in activity-induced c-fos labelled neurons in the dentate gyrus of the hippocampus in a contextual fear conditioning paradigm. They found they could induce a fear response in a neutral environment by optogenetically activating the c-Fos labelled neurons from the environment where fear-conditioning had occurred (Liu, Ramirez et al. 2012). However, a problem with immediate early genes is that a) they have a

slow turn-over rate, once a gene has been activated it remains in the cell for days irrespective of activity, thus resulting in an activity-based optogenetic tool with a temporal-resolution spanning days, and b) gene regulation seems to vary depending on cell and gene sub-type, irrespective of activity parameters (Guzowski, Miyashita et al. 2006). Consequently, the use of immediate early gene may not be the method of choice.

Perhaps, a different strategy would be to combine microscopy with optogenetics. The use of two-photon microscopy for imaging neural activity during on-going behaviour has increased dramatically over the past few years due to the development of efficient genetically encoded calcium indicators (GECIs), that allow for chronic, cell-specific targeting (Tian, Hires et al. 2009, Huber, Gutnisky et al. 2012, Zariwala, Borghuis et al. 2012). For example, the GFP-based reporter GCaMP6 that reports neural activity (by increasing its fluorescence) of a cell based on changes in cytoplasmic calcium levels (Chen, Wardill et al. 2013). The principle underlying calcium imaging is derived from the established relationship between intracellular calcium and neural activity (Baker, Hodgkin et al. 1971, Tank, Sugimori et al. 1988); thus, tracking calcium changes in a cell allows inferences to be made about the activity of a cell. Moreover, two-photon imaging of GECIs allows simultaneous study of large ensembles of cells (~2000) (Huber, Gutnisky et al. 2012), with depth limitations up to 1mm below brain surface in the cortex (Denk and Svoboda 1997). An ideal scenario would be to image cells in a brain region during a task, classify cells based on their behavioural correlate and then selectively express a transducer like ChR2 (perhaps using the two-photon laser) in one behavioural class. This would enable selective interference with functionally-defined sub-groups of cells and would allow scientists to infer whether a causal link exists between the behavioural correlate of cells and successful task performance.

Such a tool requires developments on various fronts (although some initial progress has recently been made, see (Rickgauer and Tank 2013). A stepping-stone might be to harness the high spatial resolution of a two-photon microscope and use it to ablate (i.e. lesion) cells belonging to one functional class. It is known that when a cell is exposed to elevated energy doses from a laser of a high-repetition rate the laser can cause lesioning of the exposed tissue. If combined with imaging experiments the two-photon microscope could be used to selectively remove cells of a particular functional class. Functionally-targeted ablations, using the two-photon laser, have been carried out before in worms (Yanik, Cinar et al. 2004) and

zebra fish (Bianco, Ma et al. 2012). Yet, limited work has been done using rodents (see(Canty, Huang et al. 2013). Consequently, we sought out to develop a tool to reliably ablate single-cells in the rodent brain.

To develop a laser-induced ablation technique we chose the barrel cortex – the whisker representation area of the somatosensory cortex (Kleinfeld, Berg et al. 1999, O'Connor, Peron et al. 2010, Kleinfeld and Deschenes 2011). The reasons for this choice are threefold; 1) the two-photon laser has a depth limitation of about ~1mm below brain surface (Denk and Svoboda 1997), thus the technique is limited to the cortex (unless intervening tissue is removed), 2) whiskers of rodents are topologically arranged in the barrel cortex such that one whisker is (almost) only served by a single region ('barrel') of the barrel cortex (Simons and Woolsey 1978, Kerr, de Kock et al. 2007). Such one-to-one mapping allows targeted recording and manipulation of its neural circuitry. 3) Extensive research has been carried out on the barrel cortex using two-photon microscopy (O'Connor, Peron et al. 2010, Huber, Gutnisky et al. 2012, Petreanu, Gutnisky et al. 2012). Furthermore, the barrel cortex, its function and anatomy, have been extensively studied. We will give a brief overview of the literature relating to the barrel cortex here, but will refer the reader to more extensive reviews e.g. (Feldmeyer, Brecht et al. 2013, Petersen and Crochet 2013)

Touch-related information from the whiskers travels from the trigeminal nuclei in the brain stem, where it is transmitted to the thalamus (ventral posterior and posterior thalamus) and finally to the primary somatosensory cortex (Kleinfeld, Berg et al. 1999, Feldmeyer, Brecht et al. 2013), see figure 5.1a). In the cortex, information first arrives in layer IV and from there it is relayed to layers II/III (and layer V). Layers II/III (~100-400μm below brain surface (Hooks, Hires et al. 2011)) are the site of cortico-cortico connections, both within the barrel cortex (Feldmeyer, Lubke et al. 2006) and between other regions of the cortex, such as the primary motor cortex (Hooks, Hires et al. 2011, Petreanu, Gutnisky et al. 2012). In terms of function, the barrel cortex has been found essential for various whisker-dependent tasks (Hutson and Masterton 1986, O'Connor, Clack et al. 2010), however, activity rates are generally low and responses particularly sparse (O'Connor, Peron et al. 2010). Moreover, neural coding for different aspects of sensory behaviour (such as pole-touch and whisker movement) was found to be spatially intermingled (Peron, Iyer et al. 2012). Such sparse functional circuits might be the ideal candidate for the development of a laser-induced ablation technique,

as an entire functional representation could be easily and completely removed. As a result, layers II/III of the barrel cortex was our candidate brain region for the development of this ablation technique.
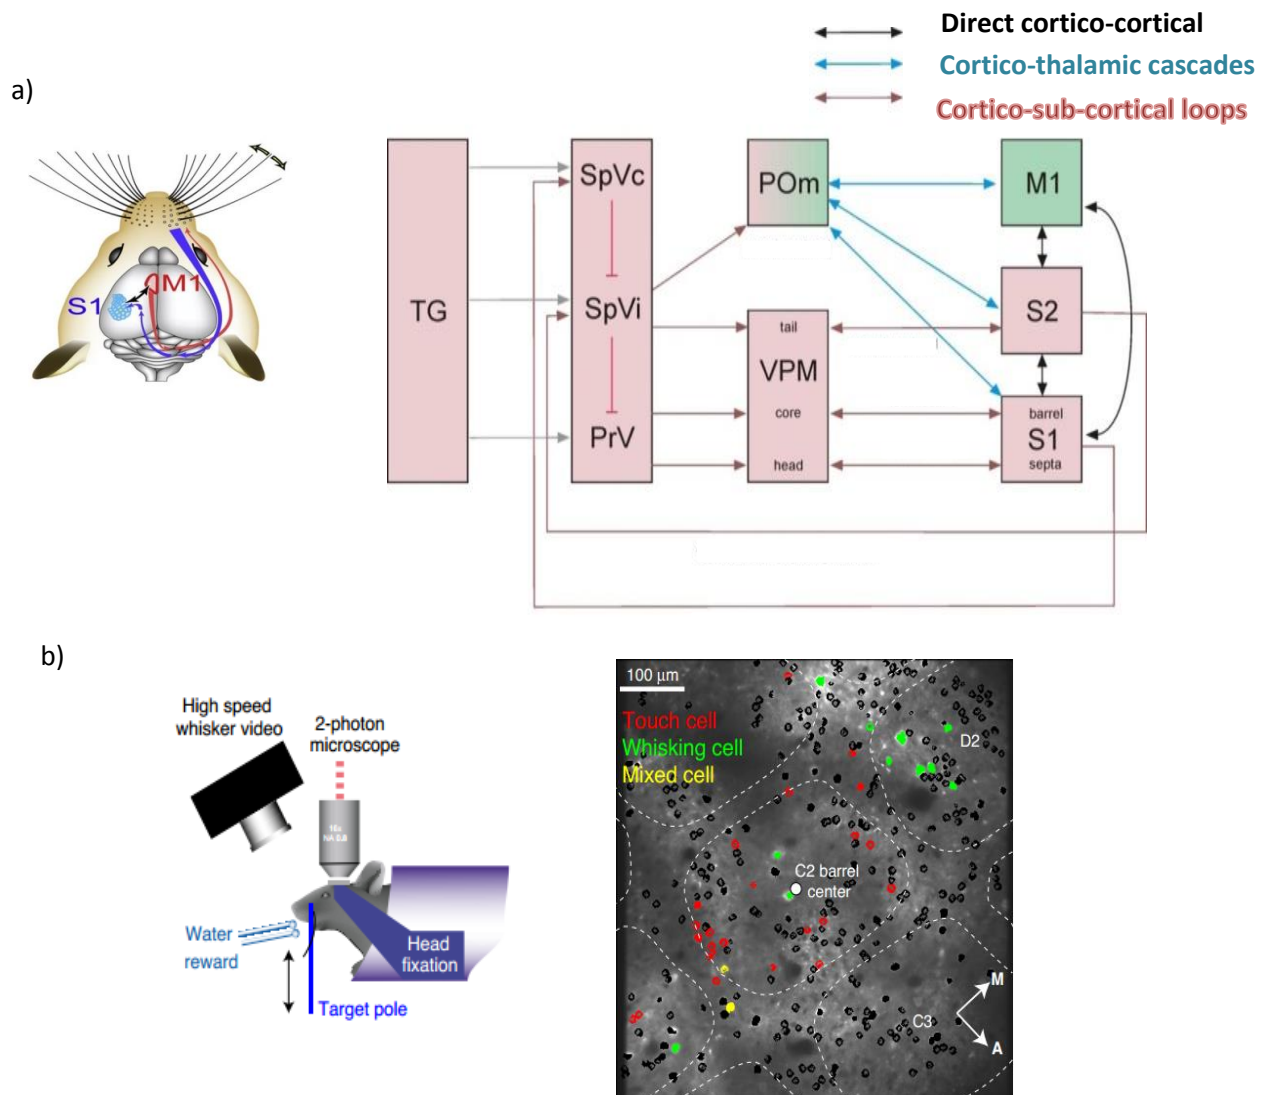
Figure 5.1. The Anatomy and Function of the Barrel Cortex a) Schematic diagram of the ascending pathways connecting the whiskers to the somatosensory cortex. The pathway starts in the trigeminal ganglion (TG) and ascends to different trigeminal nuclei in the brainstem (PrV = principal nucleus, SpVi = the spinal nuclei interpolaris, SpVc = the spinal nuclei caudalis). From the brainstem projection go to the ventral posterior thalamus (VPM) or the posterior thalamus (POm). Finally, from the thalamus projections arrive in the cortex (S1 = primary somatosensory cortex, S2 = secondary somatosensory cortex, M1 = primary motor cortex). Cortico-cortico connections are also shown. **b)** Two-photon imaging carried out in the barrel cortex while head-fixed mice performed a whisker-based object localisation task (left: schematic of experimental setup). The panel on the right shows a two-photon image with regions of interest (ROIS) drawn. Coloured ROIs represent cells showing

sensory-related activity. Panel a) is adapted from Feldmeyer et al. (2013) and panel b) from Peron et al. (2012).

Below, we will describe a method developed that allows one to reliably and reproducibly ablate single cells in the superficial layers of the barrel cortex. Moreover, we will show that the laser-induced ablations are precise. Finally, we will show one of the experimental application of this method; namely, investigating the effect of removing a small proportion of sensory-tuned excitatory cells on a wider network of cells. These results will demonstrate this method has the potential to probe many experimental questions that aim to shed light on the causal links between functional neural circuits and behaviour.

5.2. Method

See general methods for details of equipment, surgical procedures, virus selection, imaging strategy, photomultiplier tube (PMT) settings and fluorescence extraction.

5.2.1 Animals

A total of 11 adult mice (older than postnatal day 60) were used for this study. Some mice contributed to all analyses, while others only to one. Moreover, for some animals, volumetric imaging was carried out where three separate focal planes (20-40μm apart) were imaged simultaneously, whereas for other animals a single plane was imaged. Table 5.1 below shows the experimental details of the animals used.

Table 5.1 Animal Details

Animal	Transgenic Line	Analysis[6]	Imaging Type
208556	LSL-HsB-mCherry X emx1-Cre	AS,AP	Single-plane
209500	Gad2-NLS-mCherry X six3cre	AS,AP,NI,BC	Volumetric
210172	PV-ires-Cre X ai9	AS,AP	Volumetric
210173	PV-ires-Cre X ai9	AS,AP	Volumetric
210408	LSL-HsB-mCherry X emx1-Cre	AS,AP,NI,BC	Volumetric
215063	LSL-HsB-mCherry X emx1-Cre	AS	Volumetric
215064	LSL-HsB-mCherry X emx1-Cre	AS,AP,NI,BC	Volumetric
208576	LSL-HsB-mCherry X emx1-Cre	AS	Single-plane
208560	LSL-HsB-mCherry X emx1-Cre	AS	Single-plane
197526	LSL-HsB-mCherry X emx1-Cre	AS	Single-plane
208576	LSL-HsB-mCherry X emx1-Cre	AS	Single-plane

---

[6] AS = Ablation Success, AP = Ablation Precision, NI = Network Influence, BC = Baseline Control

5.2.2 Ablation Strategy

A scanning beam (excitation wavelength: 880 or 900nm; pulse width ~ 200fs) at elevated powers (170mW – 860mW) irradiated single excitatory cells for 0.2-2 seconds in layers II/III of the barrel cortex, while animals were anaesthetized and head fixed. A 1-10μm of cubic volume was stimulated in each ablation attempt, centred on a cell's cytoplasm. Cell localisation was performed manually. We also ablated in awake animals, however due to the time taken to ablate a handful of cells we considered it less stressful to the animal to keep it anaesthetized during ablations. To irradiate cells we either used a ramping strategy where we progressively increased the laser energy delivered to the cell (e.g. starting from 0.17W and going up to 0.86W, the maximum power delivery for 880nm wavelength). Alternatively, we exposed single cells with a single energy dose. The strategy employed generally varied depending on the depth of cells targeted for ablation – for more superficial cells a ramping strategy was more often employed but for deeper cells a single, high energy dose seemed to work better. The laser beam targeted the cytoplasm of the cells, where GCaMP6 is expressed primarily. Early experiments showed the fluorescence intensity at a targeted site following high exposures to the laser beam was dramatically increased. Moreover, we observed these drastic increases (or 'fluorescence flashes') were associated with tissue removal – defined as the absence of the targeted cell when imaged immediately after ablation and the day after ablation (see section below for a definition of tissue removal). Interestingly, during an ablation attempt which led to fluorescence flashes the measured fluorescence during ablation showed a distinct *drop* in fluorescence, deviating greatly from what would be expected given the energy input from the laser beam at the time. The bottom panel in figure 5.2b shows an example of the measured fluorescence during an ablation attempt that lead to fluorescence flashes. During the attempt, the fluorescence suddenly and transiently drops, this 'fluorescence drop' we used as a signature for ablation. Consequently, we developed an automatic ablation detection algorithm which ran during an ablation and shut off the laser when an ablation occurred. Specifically, we interpolated the expected fluorescence intensity by measuring fluorescence during the first 20ms of laser exposure, whenever the fluorescence intensity deviated more than two standard deviations from the expected intensity the laser beam was automatically cut off. Figure 5.2 below shows examples of measured fluorescence intensity during unsuccessful and successful ablation attempts for a

single cell, and clearly demonstrates how the fluorescence intensity deviates from expectation when an ablation occurs. GCaMP imaging did not take place during an ablation attempt. However, 10-20 frames of a target cell were captured before and just after an ablation attempt. During ablation, the gain of the PMTs was kept lower than during imaging so to prevent PMTs from tripping (see section 2.2. above for details of PMT gain settings).
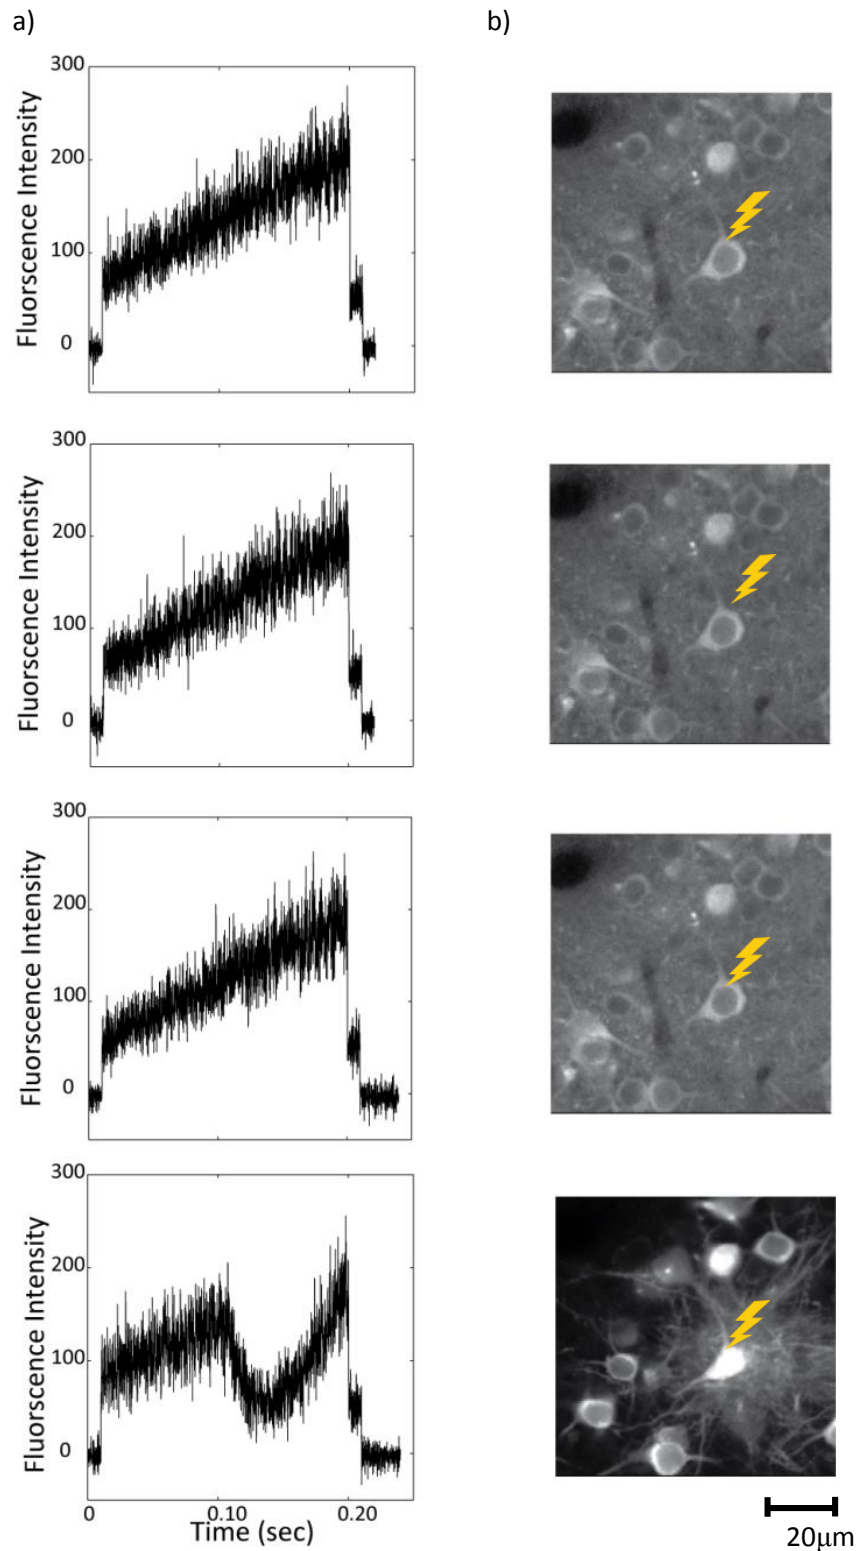
Figure 5.2. Fluorescence Marker of Ablation Examples of four ablation attempts for a single cell – highlighted by the 'lightning' symbol. Each row shows a single attempt. **a)** Measured fluorescence intensity (y-axis) for a single attempt (time on x-axis). Fluorescence intensity is measured from the photons collected by the photon-multiplier tubes (PMTs). For this ablation, a ramping laser power input strategy was used, starting at 170mW and going up to 860mW. As the power increased so did the fluorescence intensity. However, on attempt four (when a cell is successfully ablated), the fluorescence intensity drops in the middle of the attempt. This abrupt deviation from what would be expected given energy

input, was used to develop an automatic ablation detection algorithm. **b)** Calcium images (laser wavelength = 880nm) of GCaMP6S acquired immediately after each ablation attempt. On attempt four, one can see a clear increase in fluorescence emitted by the targeted cell, indicating the cell has been successfully ablated. Images of the cell a day after ablation, where it is clear the cell is no longer present can be viewed in figure 5.5, row 1. The data shown is from mouse 208576.

5.2.3. Definition of Ablation Success

To define whether a cell had been ablated we initially used the approach of checking whether the cell was still present the day after ablation. To assess the presence of a cell we imaged both GCaMP and mCherry at 800nm wavelength (a wavelength that excites mCherry and unbound GCaMP). It was important to image the cell the day after ablation as the absence of the cell immediately after ablation might merely be an indicator of bleaching of fluorescent proteins rather than a reflection of tissue removal.   After initial experiments, we found that fluorescence drop was reliably associated with the absence of the targeted cell, as imaged day after ablation. Thus, for most experiments we used the occurrence of a drop in fluorescence during ablation as an indicator of ablation success.

Another approach to assessing the success of an ablation attempt is to analyse activity changes before and after ablation. To do this, we analysed calcium-based fluorescence changes to an air-puff stimulus before and after ablation (see chapter 2, section 2.2 for a detailed description of the fluorescence extraction process). Specifically, on day 1 of imaging an animal was exposed to an air puff stimulus over 50 trials. Each trial consisted of a 1sec baseline period followed by the onset of a 500ms long air puff stimulus targeted at the animal's right whisker pad. The air puff stimulus was followed by 7.5ms post-stimulus period. We called the first three seconds following the stimulus onset the 'stimulus period'. To assess whether a cell's activity was modulated by the stimulus, we computed the difference between the average fluorescence intensity ($\Delta F/F_0$) during the baseline period and the peak activity during the stimulus period (Stimulus$\Delta F/F_0$ - Baseline$\Delta F/F_0$), we called this difference 'stimulus activity'. If the difference between the two exceeded $0.5\Delta F/F_0$ we categorized the cell as a 'stimulus cell' (see figure 5.3 for an example of a stimulus cell). On day 2 of imaging we ablated a number of stimulus cells (as well as other non-stimulus cells), and ran another session of air puff trials. Finally, on day 3 of imaging we ran a third session of air puff trials. We then looked at the changes in stimulus modulation before and (day) after ablation, using a paired-

samples t-test, a t-statistic with an associated probability below 0.025 was deemed significant. Moreover, we estimated the number of cells that were still categorised as 'stimulus cells' following ablation. To note, the described experimental procedure was also used to assess the precision of ablations and the effects of ablations on network activity. During imaging the voltage gain of the PMTs was kept constant.
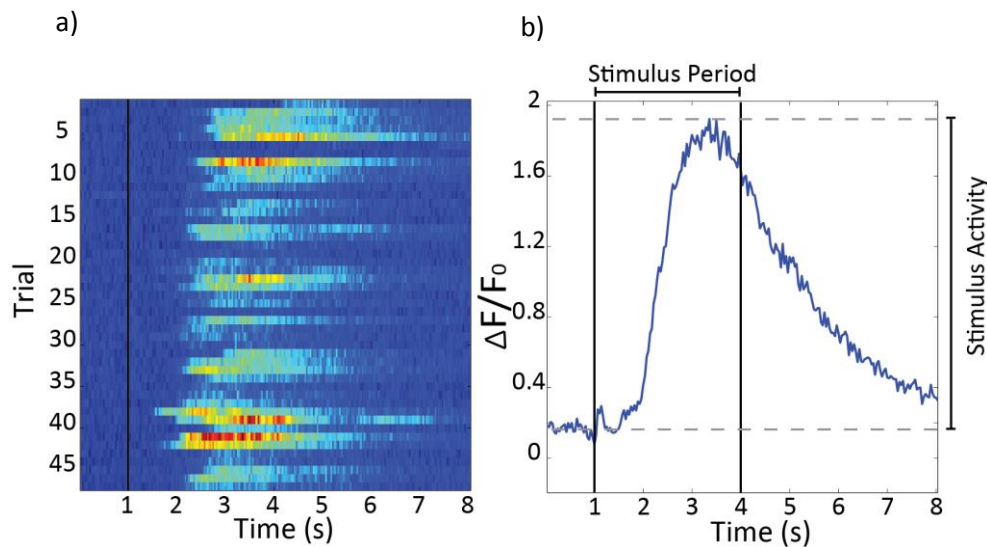
Figure 5.3. Stimulus Activity Definition a) An activity heatmap for a single cell, hotter colours represent greater changes in fluorescence intensity over baseline fluorescence intensity ($\Delta F/F_0$), for one air puff session, black vertical bar represent stimulus onset time. Y-axis shows trial number, x-axis trial time. **b)** The mean activity of cell in a), y-axis shows change in fluorescence over baseline fluorescence ($\Delta F/F_0$), and the x-axis time in a trial. The two black vertical bars mark the start and end of the stimulus period, respectively. If the difference between the mean activity during the baseline period and the peak activity during the stimulus period ('stimulus activity') is higher than $0.5(\Delta F/F_0)$ we labelled the cell a 'stimulus' cell. The difference for this cell was approximately 1.6, thus this is an example of a stimulus cell.

5.2.4. Establishment of Ablation Parameters

To assess energy needed to ablate we calculated the size of the laser pulse that led to successful ablation by taking the average power during successful ablation attempts and dividing it by the laser repetition rate (80 MHz). To assess the total energy required (in joules), we multiplied the average power during ablation by the time taken to ablate. Moreover, to assess how energy and time needed to ablate varied with depth we used Pearson correlations. To analyse relationship between success rate and depth of ablation target we grouped the data into 6 depth groups;

100-150μm, 150-200μm, 200-250μm, 250-300μm, 300-350μm and 350-400μm from brain surface. We then correlated the success rate in each group with depth. For these analyses, we defined success by the occurrence of fluorescent flashes, so to maximize statistical power.

5.2.5. Assessment of the Spatial Precision of Laser-Induced Lesions

To assess whether the laser-induced lesions were limited to the targeted cell body we targeted cells close (< 50μm) to a stimulus cell for ablation and looked at changes in activity before and after ablation for the stimulus cell using the same method as described above. That is, the change in activity before and (day) after ablation was compared using a paired t-test, moreover, we estimated the number of cells still classified as stimulus cells after a neighbouring cell had been ablated. Figure 5.4 below shows an image of two neighbouring cells, one of which was targeted for ablation, as well as the activity of the cell not ablated, before and after its neighbour was ablated.
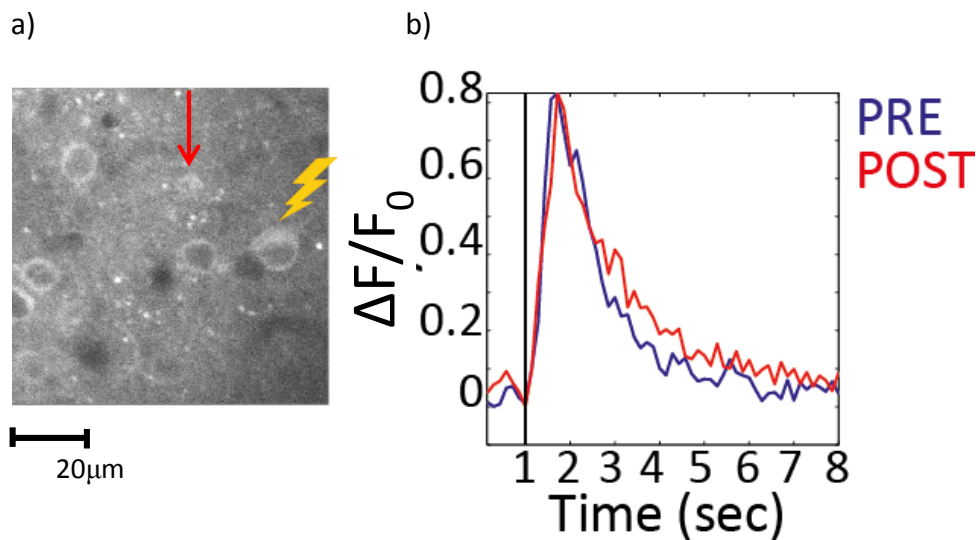
Figure 5.4. Assessing the Precision of Ablations a) A calcium image showing a cell targeted for ablation (lightning) and its neighbouring stimulus cell (red arrow). Image was taken before ablation. **b)** Mean activity of stimulus cell ($\Delta F/F_0$) before and after ablation of its neighbour (blue and red trace respectively). Y-axis shows change in fluorescence over baseline fluorescence ($\Delta F/F_0$) and the x-axis time in second. Although the stimulus activity of the cell seems slightly weaker after ablation of its neighbouring cell it still seems to increase its activity in response to the air puff stimulus.

5.2.6. Assessment of Baseline Changes in Stimulus Activity

To ensure changes in activity we observed in response to ablation did not just reflect the natural fluctuation in network activity between consecutive days we collected data from three mice over three consecutive days, using the air puff procedure described above, and assessed baseline changes in activity between days using a paired-samples t-test. We did this for all cells as well as just stimulus cells. We then compared the baseline absolute changes observed from these control experiments to the changes observed after ablating cells, using an independent samples t-test (using a 2.5% significance level as before). If the activity changes we observed after ablating cells exceeded that found in the control experiments we could be confident the changes were a consequence of the ablation. Moreover, if the changes in stimulus activity for the neighbours of ablated cells did not exceed that found in the control experiments it would provide added evidence that the ablations were focal to the targeted cell.

We also used the data from the control experiment to assess how many cells maintained their functional status between days (i.e. a measure of fidelity). We then used this baseline fidelity measure to assess whether changes we see in functional status of neighbours of ablated cells and the ablated targets differs significantly from what one would expect to see by chance, using a binomial test.

5.2.7 Assessing the Effects of Ablation on a Neuronal Network

To assess what the effects of ablating cells in a small region of layers II/III of the barrel cortex has on the wider network of cells we ablated 30-40 cells from a given barrel cortex volume in three animals and compared changes in stimulus activity for all imaged cells in the volume before and after ablation. The animals used in this experiment were the same as those used for the baseline control experiments described above. Thus, we had three days of baseline data allowing us to get an accurate estimate of fluctuations in network activity between days, which allowed us to compare the changes we saw after ablations to that we saw between days. We compared absolute changes in activity between days using a paired t-test, as before (to note this analysis was also carried out on raw changes (both positive and negative) in activity). Moreover, we ran the analysis including and excluding the ablated cells to ensure any changes we observed when including the targeted cells were not simply due to the targeted cells exerting a strong influence on the estimated network activity. Finally, we correlated changes in stimulus activity of a

cell with the distance to the closest ablated cell using a Spearman correlation, in order to obtain a more comprehensive estimate of the spatial precision of lesions. For the analysis excluding outliers, we defined outliers as cells whose change in stimulus activity exceeded two standard deviations above/below the mean.

5.3 Results

5.3.1 Delivering High Doses of a Two-Photon Laser to a Single Cell Reliably Leads to Ablation

We found irradiating 1-10μm of single cells (at 880/900nm wavelength) in layer II/III of the barrel cortex with a high power laser beam (mean power = 520mW, SD = 195.4) led reliably to the removal of the targeted cells. To initially define tissue removal we imaged GCaMP and mCherry of the targeted cells before, immediately after and a day after attempted ablation. Minutes after ablation the GCaMP images revealed a dramatic increase in the fluorescence emitted from the targeted cell (i.e. fluorescent flashes) and deformation of the nucleus, derived from the mCherry images. The day after ablation, the cell seemed to have been removed from both images. The absence of the cells in both GCaMP and mCherry images, day after ablation, was used as a confirmation the ablation had been successful. Figure 5.5 below shows images of five cells ablated, before, just after and a day after ablation. Calcium images in column c) (DAY 2) sometimes indicate other cells than the stimulated cell (lightning symbol) are absent. The reason for this is that in one ablation session we often ablated multiple cells (not shown) within one field of view. Thus, these 'other' cells were also absent in the DAY 2 images. However, the vasculature between PRE and DAY 2 images is the same, confirming the images are of the same anatomical location.
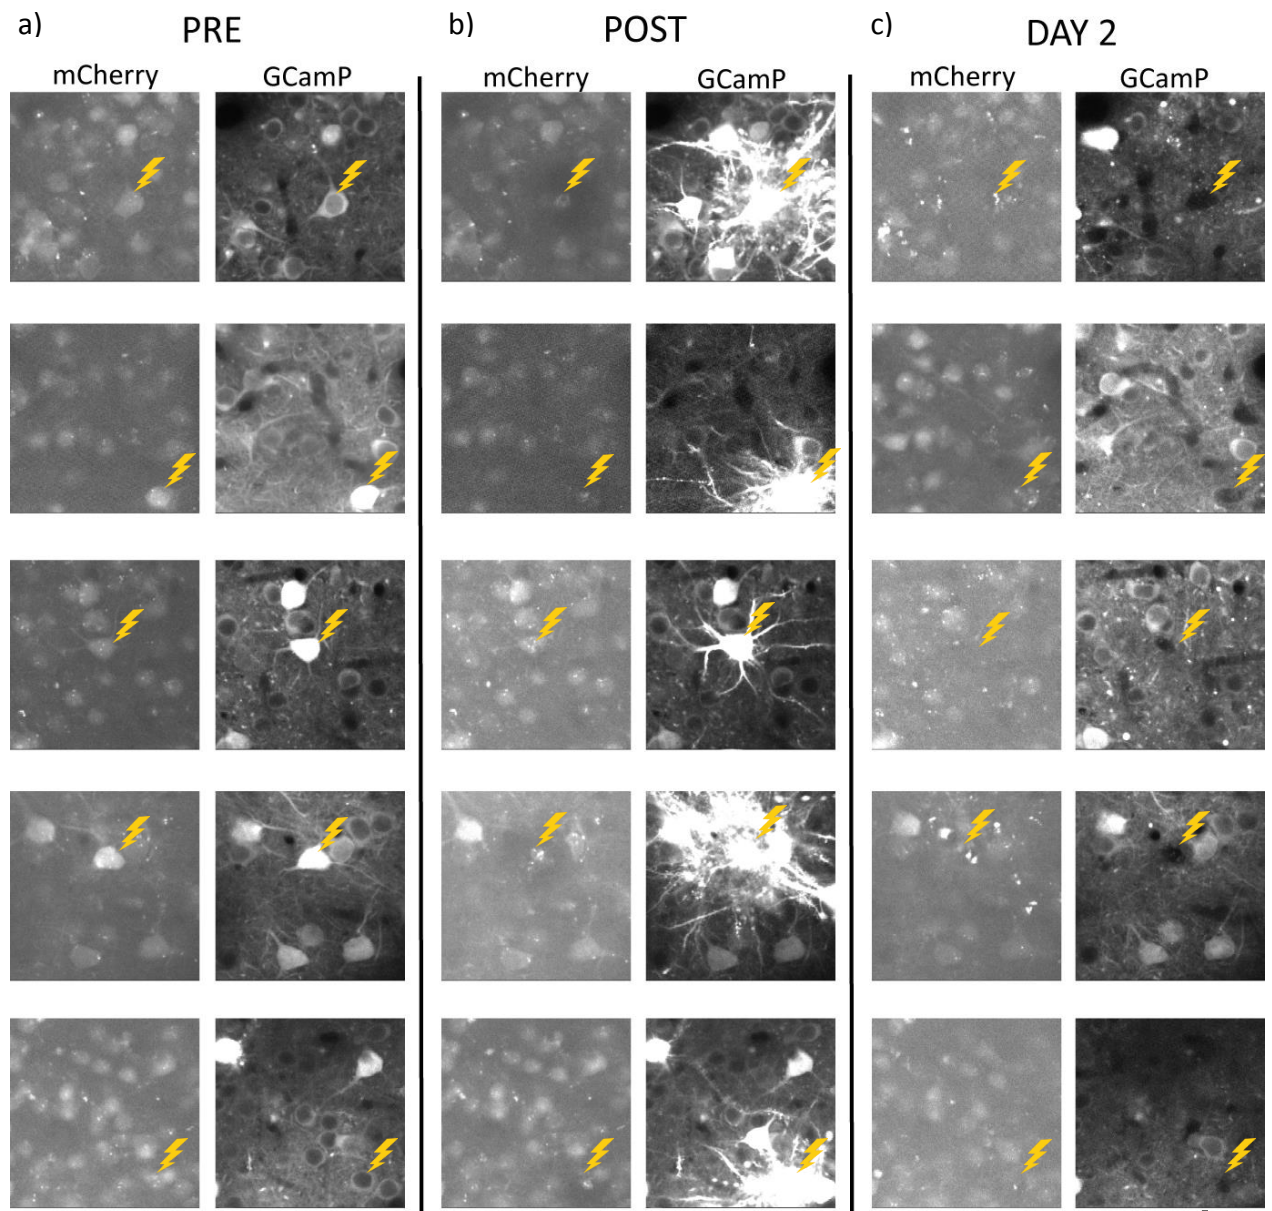
Figure 5.5. Images of Successful Ablation Examples of GCaMP and mCherry images for five cells targeted for ablation. Each row represents one ablation. mCherry images show the nucleus of cells and GCaMP images are calcium images. The lightning icon points to cells targeted for ablation **a)** images taken just before ablation, **b)** images taken minutes after ablation and **c)** images taken a day after ablation. Within a field of view, multiple cells were often ablated (not shown). Hence in panel c), sometimes multiple cells are absent. However, the vasculature between panels a) and c) are the same, confirming that the two images are of the same anatomical location. Rows 1, 2 and 5 are from mouse 208576, row 3 from mouse 208560, and row 4 from mouse 197526.

We observed attempts that lead to the fluorescence flashes were associated with an abrupt and transient drop in fluorescence intensity ('fluorescence drop') during the attempt itself (as demonstrated in figure 5.2 above). We found, the occurrence of a fluorescence drop during ablation was consistently associated with tissue

removal. As such, their occurrence, rather than GCaMP/mCherry images, was used as a marker of success. Using this criterion we found 79% of all attempted ablations (N = 366) were successful. Although the vast majority of ablations were carried out while an animal was anaesthetized (N = 342), we found an identical success rate for the few ablations carried out while the animal was awake (N = 24). The average time taken to ablate was found to be 0.73sec (SD = 0.66), the average size of laser pulse 6.37nJ (SD = 2.41) and the total energy needed 0.36J (SD = 0.37) - total energy was estimated by taking the average laser power during an ablation attempt and multiplying it by the time taken to ablate, whereas to estimate the size of the laser pulse we divided the average power by the laser repetition rate (80mHz).

Another marker for ablation success is the cessation of activity for cells targeted for ablation; that is, a reliable reduction in stimulus modulation of calcium-based fluorescence from cell targets. To analyse ablation-based changes in activity we ran a simple sensory experiment a day before, an hour after and a day after ablations. Specifically, head-fixed mice experienced a 500ms air puff stimulus over their right whisker pad. Each trial was composed of one exposure to the stimulus, which was preceded by a 1sec baseline period and followed by a 7.5sec post stimulus period. The three seconds following the onset of the air puff we referred to as the 'stimulus period', and 'stimulus activity' we defined as the difference in fluorescence intensity ($\Delta F/F_0$) between the baseline and stimulus period (Stimulus$\Delta F/F_0$ - Baseline$\Delta F/F_0$). Each session was composed of 50 trials. We observed a significant reduction, based on a paired-samples t-test, in stimulus activity for sessions a day after ablations compared to those a day before – mean stimulus activity before = 0.36 $\Delta F/F_0$ (SD = 0.34), mean stimulus activity after = 0.23$\Delta F/F_0$ (SD = 0.09), $t_{(107)}$ = 3.86, p = 1.93 x 10^{-4}, see figure 5.7a. Moreover, narrowing the analysis to cells whose difference in baseline and peak stimulus activity exceeded 0.5$\Delta F/F_0$ - i.e. 'stimulus cells', we observed an even more dramatic reduction – mean stimulus activity before = 0.93$\Delta F/F_0$ (SD = 0.52), mean stimulus activity after = 0.24$\Delta F/F_0$ (SD = 0.12), $t_{(17)}$ = 5.04, p = 9.97 x 10^{-5}, see figure 5.7b. In fact, the mean stimulus activity following ablation was below the threshold used for defining whether or not a cell is stimulus tuned. Finally, a day after ablation only one out of 18 stimulus cells still maintained its functional status. Figure 5.6 shows activity plots of ablated cells before and after ablation. These results lend further support to the suggestion that laser-induced ablations are viable and reliable in superficial layers of the barrel cortex.
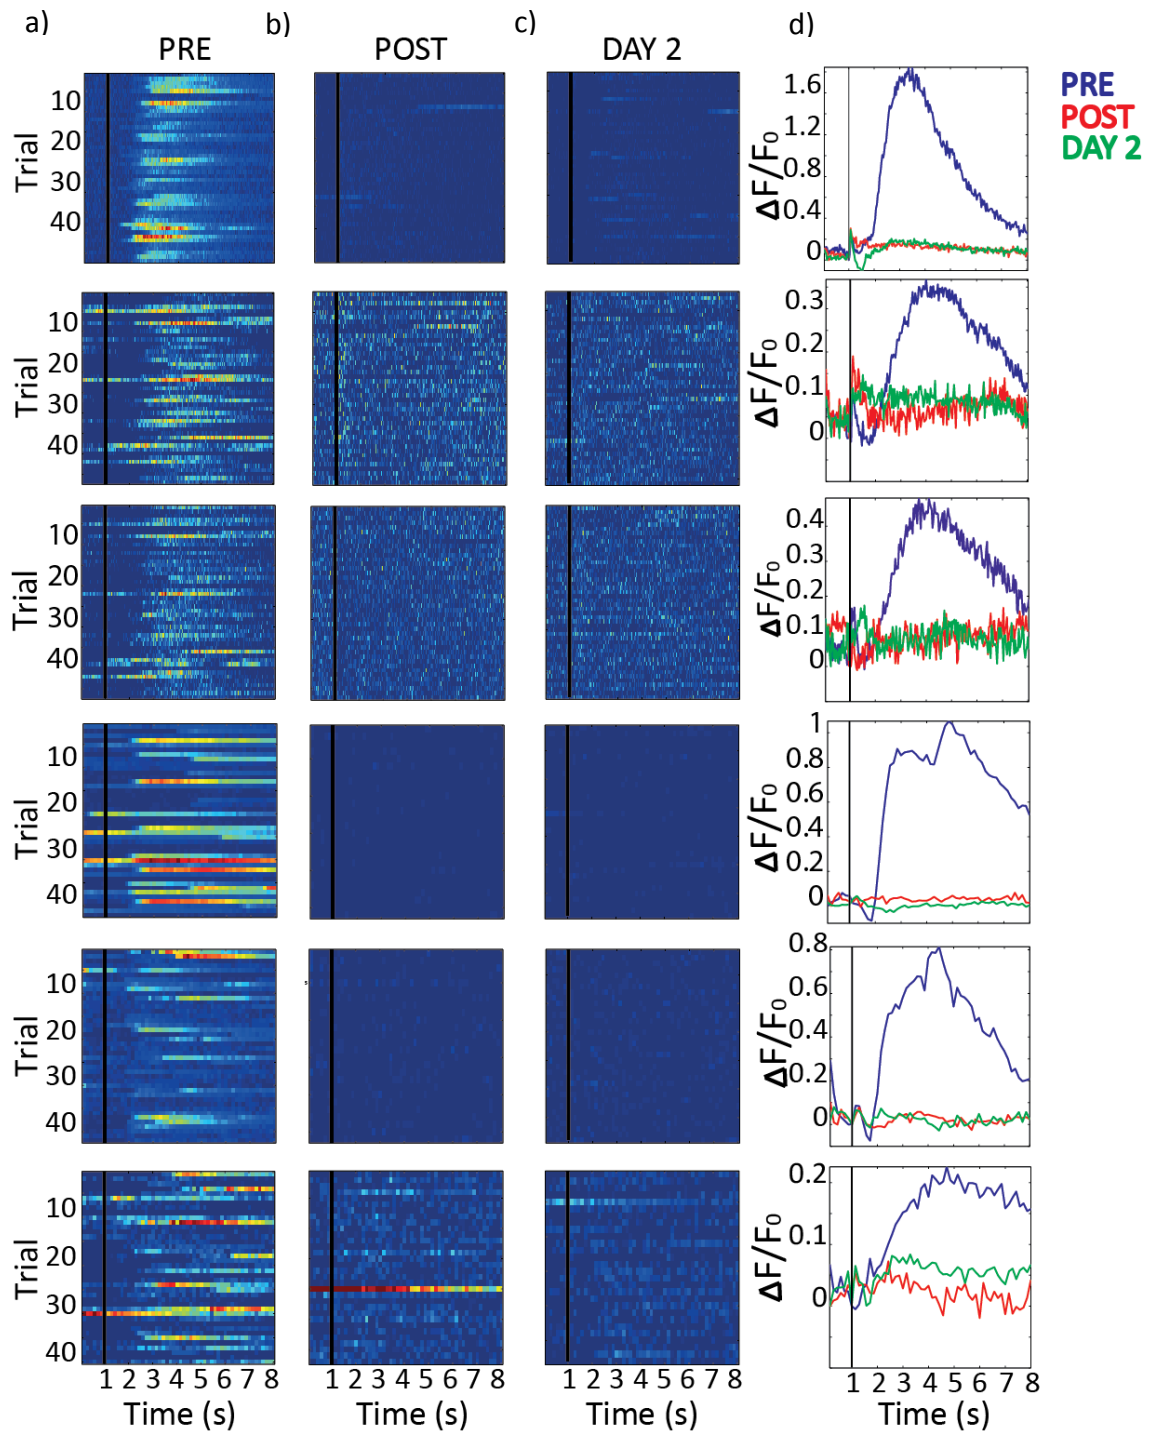
Figure 5.6. Ablation-Based Changes in Stimulus Activity a-c) Heatmaps of fluorescence intensity during an air puff experiment before (PRE), just after (POST) and a day after (DAY 2) ablation. Each row shows one ablation. Y-axis represent trial ID, X-axis time during a trial. Black vertical bar marks onset of air puff stimulus. Hotter colours represent higher fluorescence intensity ($\Delta F/F_0$). **d)** Mean fluorescence intensity over an air puff session before (blue trace), just after (red trace) and a day after (green trace) ablation. Y-axis shows change in fluorescence from baseline fluorescence ($\Delta F/F_0$), and x-axis time in an air puff trial. Rows 1-3 are from mouse 208556, rows 3-6 from mouse 209500.

However, one might expect changes from day-to-day in stimulus activity that are a mere result of inter-session fluctuations or changes in baseline fluorescence levels. That is, cells may change the stimulus tuning between consecutive days that is not a fault of ablations. To estimate expected day-to-day fluctuations in stimulus activity we collected data from three mice on three consecutive days, without ablating any cells. We estimated mean absolute changes in stimulus activity between day1 and 3, and compared these 'baseline' changes to those observed for the ablated cells. Overall, the mean fluorescence intensity change between day1 and day3 was $0.07\Delta F/F_0$ (SD = 0.12), this is lower to the observed $0.13\Delta F/F_0$ (SD = 0.35) we observed for the ablated cells. An independent samples t-test comparing changes in stimulus activity for ablated and control cells was found to be significant $- t_{(2334)} = 8.22$, p = = 3.39×10^{-16}, see figure 5.7c. Moreover, focusing on 'stimulus cells', we found a mean fluorescence intensity change of $0.33\Delta F/F_0$ (SD = 0.54) between days 1 and 3, this was found to be significantly lower to the observed $0.68\Delta F/F_0$ (SD = 0.57) found for the ablated cells $- t_{(132)} = 4.23$, p = 4.29×10^{-5}, see figure 7.6c. Finally, assessing the proportion of cells that maintained their functional status over days on average 49.78% of stimulus cells on one day maintained their functional status when imaged two days later. For the ablated stimulus cells, only 5% (1 cell) maintained their functional status after it had been targeted for ablation, this was found to be significantly lower to the expected fidelity, using a binomial test (p < 0.001).
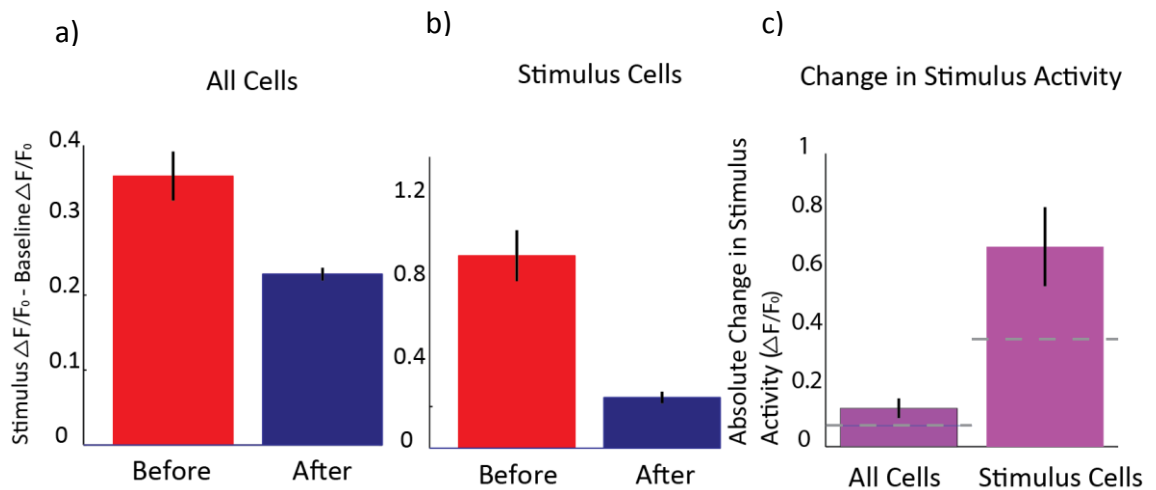
Figure 5.7. Ablation-Based Activity Differences a) Mean stimulus activity before and after ablation. Y-axis represents difference between mean fluorescence intensity during baseline period and the peak fluorescence intensity during the stimulus period (Stimulus$\Delta F/F_0$ - Baseline$\Delta F/F_0$). **b)** Same as a) but for stimulus cells (i.e. cells whose Stimulus$\Delta F/F_0$ - Baseline$\Delta F/F_0$ exceeded 0.5. **c)** Absolute difference between stimulus activity before and after ablation for all cells and only stimulus cells. Grey dashed line represents expected changes from control data. Black vertical line on bars represents standard error of the mean.

Thus, it is clear that delivering high doses of a two-photon laser to single cells in layers II/III of the barrel cortex can cause functional damage and removal of the targeted cell. We found the fluorescence marker for such ablations was a sudden, unexpected, drop in fluorescence from the targeted cell during elevated laser irradiation. However, as light scattering increases with tissue depth (Denk and Svoboda 1997) one would expect the success rate and energy needed to ablate to vary with the depth of a cellular target. Moreover, for the study of functional circuits in layers II/III of the barrel cortex it is important to assess whether the ablation method can be applied to the entire extent of these layers – which stretch to about 400μm below brain surface (Hooks, Hires et al. 2011). As such, we investigated the relationship between success rate and depth of ablations, as well as the energy required to successfully ablate at different depth levels.

5.3.2 Cells Deep in Layers II/III Require more Energy to Ablate

In order to assess whether cells at all depth levels of layers II/III could be ablated with the current method and to establish the depth limitation of this method we analysed the relationship between success rate and depth of ablations, by dividing

the data into six depth groups (see methods) and correlating success rate with depth. We found the success rate of ablations to be high (no less than 75%) up until about 300μm below brain surface. An abrupt decrease in success rate was observed for cells located deeper. Moreover, the relationship between success rate and depth was well described by a linear function – r = -0.82, p = 0.045, see figure 5.8a. See table 5.2 for details of success rate for each depth category.

Table 5.2 Depth Limitations of Laser-Induced Ablations

Depth(μm)	Success Rate (%)	# Ablations
150	84.25	127
200	80.95	147
250	100	10
300	75	56
350	47.62	21
400	20	5

Moreover, the size of the energy pulse required to successfully ablate was higher for deeper targets than shallower targets (r = 0.40, p = 4.6 x 10^{-12}, see figure 5.8b). For ablation targets deeper than 300μm below pia, we found we often had to use the full power of the laser (10.75nJ pulse in this instance) to achieve tissue removal. Although the time taken to ablate did not vary as a linear function with depth (r = -0.10, p > 0.05), the total energy required to ablate did (r = 0.45, p = 4.1 x 10^{-10}, see figure 5.8c).
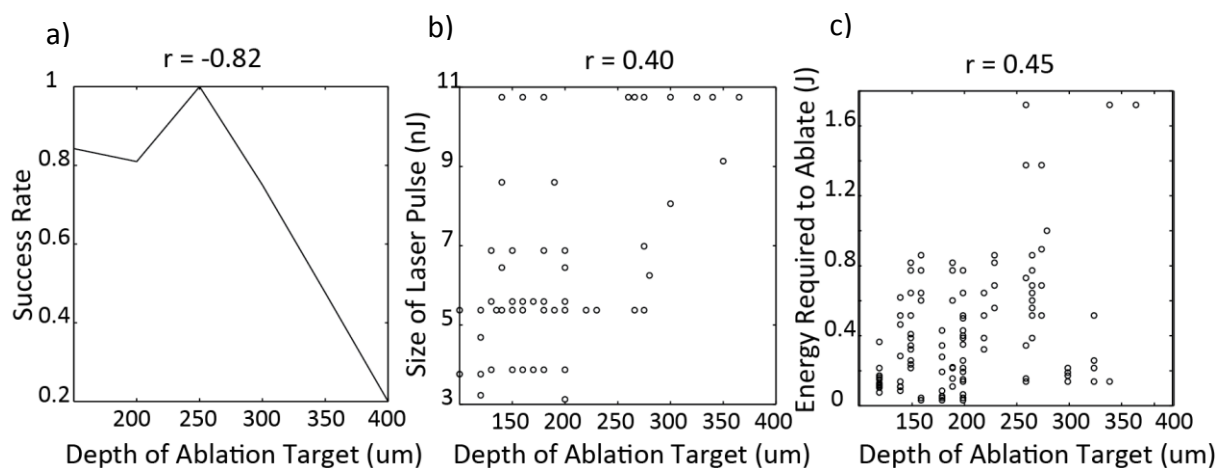
Figure 5.8. The Relationship between Ablation Depth and Success Rate and Energy Required to Ablate a) Success rate plotted against depth relative to brain surface. Y-axis shows success rate, x-axis the maximum depth of the ablation target and the title the correlation between the axes. **b)** Size of laser pulse plotted against depth of ablation target.

Y-axis shows size of pulse in nanojoules, x-axis represents depth of ablations and the title shows the correlation between the two. **c)** Total energy delivered to successfully ablated cells (Joules) plotted against depth of ablations. X-axis shows depth of ablations, as before, and y-axis the total energy delivered. The title shows the correlation between the two axes.

Thus, as expected, it is harder to ablate deeper in the brain with. An absolute depth limit being in the region of 350μm below pia. Furthermore, the size of the energy pulse required to ablate and the total energy needed to be delivered to the tissue increases progressively the deeper the ablation target was.

5.3.3 Assessing the Focality of Ablations: Neighbouring Cells Maintain their Functional Status Following Ablation

Although we have established that one can reliably ablate cells in the superficial layers of the barrel cortex using a two-photon laser, we have not shown lesions are limited to the targeted cell body. Figure 5.5 *does* imply that after targeting a single-cell for ablation only the targeted cell seems to disappear. One way to analyse the spatial precision of ablations is to analyse the activity of cells nearby an ablation site. To address this question we carried out the air puff experiment, described above, ablating cells close (within 50μm) to a stimulus cell. If the activity of the stimulus cell is not reliably affected by the nearby ablation it would suggest the ablations are focal to the ablated cell. Figure 5.9 below shows examples of stimulus cells whose nearby cell was ablated. Although the extent to which the cells are modulated by the air puff stimulus seems somewhat dampened after its neighbouring cell was ablated all cells still show clear stimulus tuning.
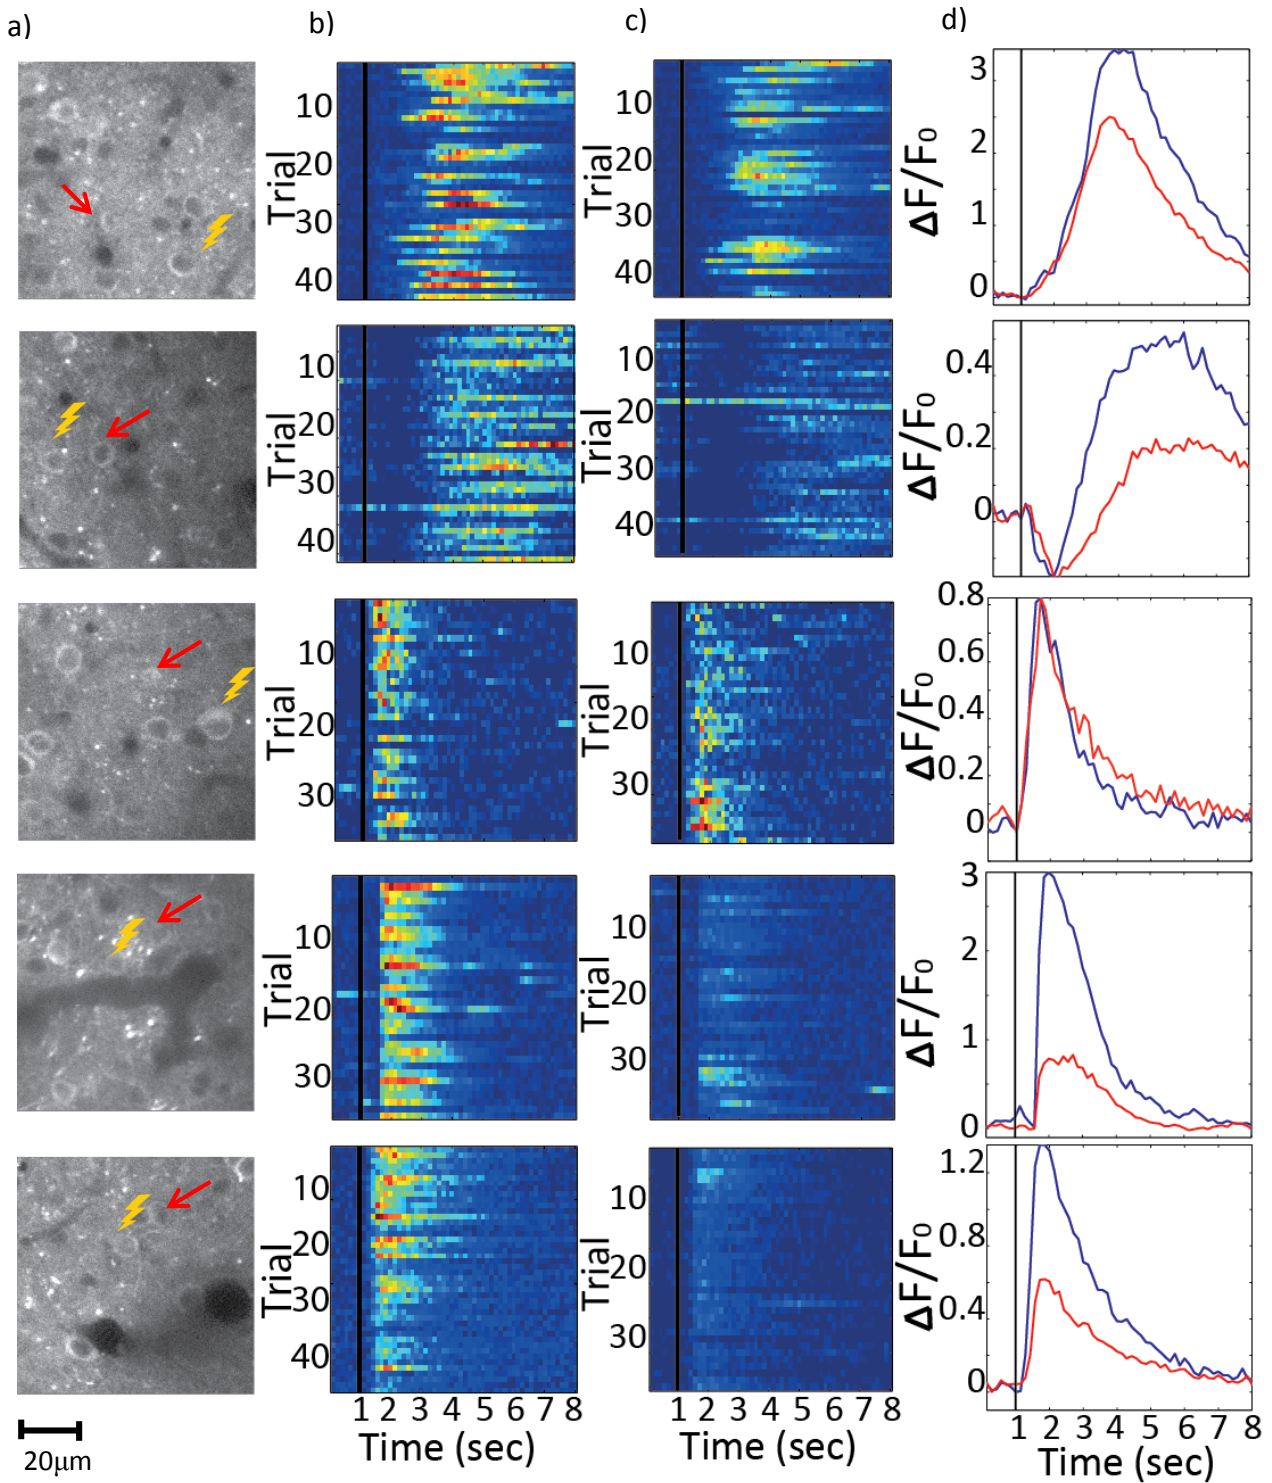
Figure 5.9. The Effect of Ablations on Neighbouring Cells a) GCaMP image of neighbouring cells before one cell (lightning) was ablated, the activity of the neighbouring cell (red arrow) is shown in b-d. **b-c)** Heatmaps of the changes in fluorescence intensity over an air puff session before (b) and after (c) a nearby cell was ablated, hotter colours represent greater changes in fluorescence intensity over baseline fluorescence ($\Delta F/F_0$). X-axis shows time in trial and y-axis trial number. **d)** Mean increase in fluorescence intensity over baseline

fluorescence for a given air puff session. The blue trace represents activity before a nearby cell was ablated and the red trace the activity post ablation. Y-axis shows change in fluorescence intensity over baseline fluorescence ($\Delta F/F_0$), and x-axis time in a trial. Each row represents one ablation. Rows 1, 2 and 5 are from mouse 210408, row 3 and 4 from mouse 209500.

We compared the stimulus activity (stimulus$\Delta F/F_0$ – Baseline$\Delta F/F_0$) of cells neighbouring an ablation site before and a day after ablation. Before ablation, the average fluorescence intensity difference between baseline and the stimulus period was $1.21\Delta F/F_0$ (SD = 0.88), whereas after ablation the fluorescent intensity difference was $0.66\Delta F/F_0$ (SD = 0.71). We found this difference to be significant on an independent-samples t-test – $t_{(25)} = 6.37$, $p = 1.14 \times 10^{-6}$. However, the activity of the stimulus cells neighbouring an ablation site was still significantly higher than that of stimulus cells that were actually targeted for ablation – $t_{(42)} = 2.47$, $p = 0.018$.

These results might indicate that laser-induced ablations cause damage to surrounding tissue in addition to the targeted cell body. However, before coming to such a conclusion one must account for baseline changes in stimulus activity between days. Similar to the analysis above, we looked at the baseline changes in stimulus activity for animals that went through the air puff experiment on three consecutive days, without ablations occurring. We then compared the absolute change in activity of stimulus cells between days from these control experiments to that observed for the cells neighbouring an ablation site. For the cells neighbouring ablations, the mean change in stimulus activity between day was found to be $0.56\Delta F/F_0$ (SD = 0.4535). For the control data, the mean difference observed between day1 and day3 of the experiment was $0.33\Delta F/F_0$ (SD = 0.54). This difference was found to be significant on a two-sample t-test – $t_{(166)} = 2.10$, $p = 0.037$. Moreover, we found 49% of cells from the control experiment maintained their functional status between day1 and day3, this was *not* found to be significantly different to the fidelity we observed for the cells neighbouring an ablation site – namely, 43%, on a binomial test ($p > 0.05$).
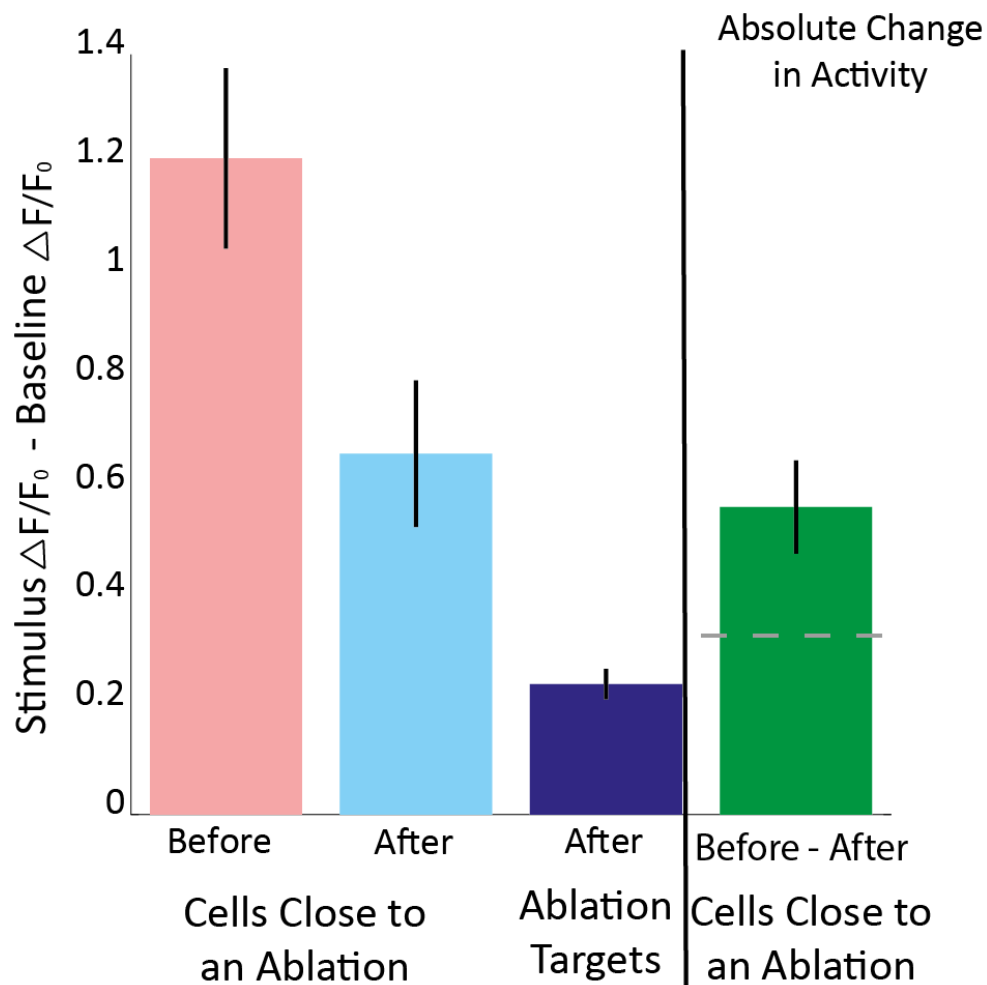
Figure 5.10. Ablation-Based Activity Changes of Cells Neighbouring Ablation Sites

Average activity of stimulus cells neighbouring an ablation site before and after ablation (pink and cyan bars), average activity of ablated cells after ablation (blue bar) and the average absolute difference between the activity of neighbouring cells before and after ablation. Grey dashed line shows the expected change based on control data. Y-axis shows stimulus activity as stimulus$\Delta F/F_0$ – Baseline$\Delta F/F_0$, black vertical lines on bars show standard error of the mean.

In sum, we do observe a significant reduction in stimulus activity for cells that were near to a cell targeted for ablation, and this reduction was significantly greater than what one would expect from day-to-day fluctuations in stimulus activity. However, the stimulus activity after ablation for the neighbouring cells was still significantly higher than that for the ablated cells, and more than half of cells still maintained their functional status after ablation – this fidelity was not found to be significantly different from baseline fidelity estimates. Thus, it does appear cells neighbouring an ablation site survive after a nearby cell was ablated, however their function may be impacted. Perhaps the reduction in activity is simply a result of the loss of

excitatory input from the ablated cells? This issue of collateral damage of ablations will be addressed more thoroughly in the subsequent section.

5.3.4 Ablating a Small Proportion of Cells Results in a Subtle Decrease in Overall Network Activity

We have established laser-induced ablations are possible in the superficial layers of the barrel cortex. Consequently we applied the method to assess what the functional implications are for a network of cells when a small (30-40 cells) sub-set of pyramidal cells are removed from it. That is, although ablations are focal to targeted cell bodies, removing a handful of excitatory cells from a brain region might influence the activity of other cells in a network. To address the network effects of ablations we ablated 30-40 pyramidal cells in a single volume of layers II/III of the barrel cortex for three mice and analysed stimulus activity before and after ablation. Moreover, the animals used in these experiments were the same as the ones used for the control experiments establishing baseline changes in stimulus activity between days. Thus, for this analysis we both had accurate estimates of baseline activity changes between days as well as changes following (day after) ablations.

Over days we saw a gradual drop in overall stimulus activity, activity rates being highest on day 1 and then gradually dropping across days, including the day after ablations (see figure 5.11a). Moreover, this drop was overall consistent across the three mice used (see table 5.3). Stimulus activity was found to be the lowest a day after ablation however it is not clear from this data whether this just reflects the general drop in activity across days. To address this we estimated absolute changes in activity between consecutive days. We found changes in activity between day 1 and 2 to be very similar to the changes seen between days 2 and 3 ($t_{(2227)}$ = 1.51, p = 0.13), however on day 4 (after ablations) the change in activity was significantly greater than that seen between previous days; changes between day 1 and 2 vs changes between days 3 and 4: $t_{(2367)}$ = 5.84, p = 5.87 x 10^{-9}; changes between days 2 and 3 vs those between days 3 and 4: $t_{(2131)}$ = 4.23, p = 2.48 x 10^{-5}, see figure 5.11b, all analyses were carried out using a paired-samples t-test. These results indicate ablating a small proportion of cells in a region of the barrel cortex results in a significant drop in overall network responses to a sensory stimulus. To note, the same pattern of results was observed when we looked at raw changes (both

positive and negative) in stimulus activity between days (changes between day 1 and 2 versus changes between days 2 and 3: $t_{(2227)} = 1.92$, p > 0.05; changes between day 2 and 3 versus changes between day 3 and 4: $t_{(2367)} = 2.73$, p = 0.0064; changes between day1 and 2 versus changes between day3 and 4: $t_{(2131)} = 2,01$, p = 0.0441).

However, one might suggest that the observed changes are a mere result of the removal of the ablated cells from the network. To control for this, we excluded the ablated cells from the analysis, and we found the same pattern of result (see figure 5.11a-b). Namely, across days stimulus activity dropped gradually (although overall activity rates were slightly lower for the first three days), and the absolute change in activity going from day 1 to day 2 was not significantly different from that between day 2 and day 3 ($t_{(2139)} = 0.74$, p > 0.05), while those after ablation were found to be significantly greater than those between the previous days; day 1 and 2 vs day 3 and 4: $t_{(2131)} = 4.23$, p = 2.48 x 10^{-5}; day2 and 3 vs day 3 and 4: $t_{(2367)} = 5.84$, p = 5.87 x 10^{-9}, all results are based on paired-samples t-tests).

Table 5.3. Change in Stimulus Activity across Days for Each Mouse

Animal	Day1	Day2	Day3
an210408	0.32 $\Delta F/F_0$ (SD = 0.40)	0.30 $\Delta F/F_0$ (SD = 0.38)	0.30 $\Delta F/F_0$ (SD = 0.27)
an215064	0.28 $\Delta F/F_0$ (SD = 0.093)	0.23 $\Delta F/F_0$ (SD = 0.10)	0.22 $\Delta F/F_0$ (SD = 0.13)
an209500	0.26 $\Delta F/F_0$ (SD = 0.21)	0.26 $\Delta F/F_0$ (SD = 0.19)	0.23 $\Delta F/F_0$ (SD = 0.18)
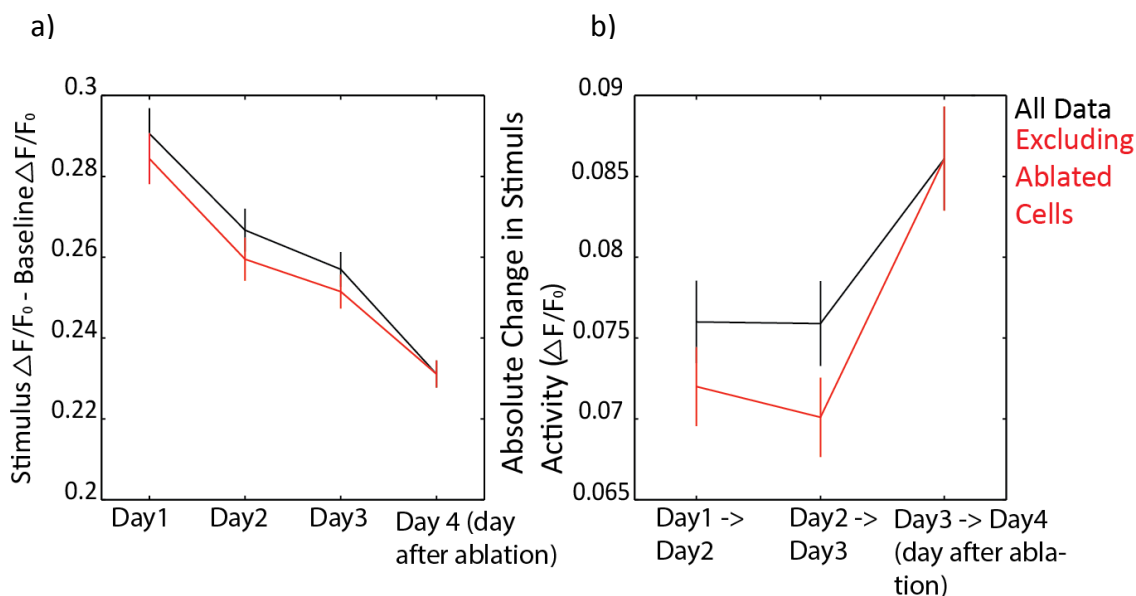
Figure 5.11. Network Activity Before and After Ablations a) Average stimulus activity across days of experiment, y-axis shows stimulus activity (stimulus$\Delta F/F_0$ – Baseline$\Delta F/F_0$), x-axis shows day of experiment – n.b. day 4 was the day following ablations. **b)** Absolute change in stimulus activity between consecutive days, y-axis represents absolute change in

stimulus activity (stimulus$\Delta F/F_0$ – Baseline$\Delta F/F_0$), x-axis the days being compared. Red lines show data after excluding ablated cells. Vertical lines show standard error of the mean.

Perhaps the slight drop in overall stimulus activity is a result of collateral damage. In other words, perhaps cells close to the ablated cells are damaged as well. The analysis described above did not give conclusive evidence for ablations affecting surrounding tissue however the number of cells included in that analysis was quite small (N = 28), thus with the current data one is able to address this problem more comprehensively. Consequently, we estimated the distance between each cell and the nearest ablated cell and correlated absolute change in stimulus activity with these distance estimates. We did not find a correlation between distance to nearest ablation site and change in stimulus activity (r = 0.01, p > 0.05, figure 5.12b). However, this analysis did reveal some outliers that could exert disproportional influence on the results. Thus, we re-ran the analysis excluding outliers (cells that showed a change in stimulus in activity exceeding two standard deviations above the mean, ~0.4). We obtained the same null result – r = -0.02, p > 0.05 (figure 5.12c).

Finally, perhaps the observed null result is a consequence of the inclusion of many cells showing very weak stimulus tuning; this would make detection of a correlation difficult as the activity rates are low before ablation. The average stimulus activity for each day shown in figure 5.11a above indicates activity rates were generally low – below the threshold used to define whether a cell showed stimulus related activity. Thus, to address this possible confound, we limited the analysis to stimulus cells. Figure 5.12d below shows the absolute change in stimulus activity for stimulus cells plotted against distance to nearest ablation site. Again, we do *not* observe a relationship between change in activity and distance to nearest ablation (r = 0.04). Although, it is clear from the plot, a few cells within 100μm from their nearest ablation site show a large change in activity. Thus, perhaps, in some rare instances cells close to an ablation site are affected by the ablation. Yet, *generally*, this does not seem to be the case.
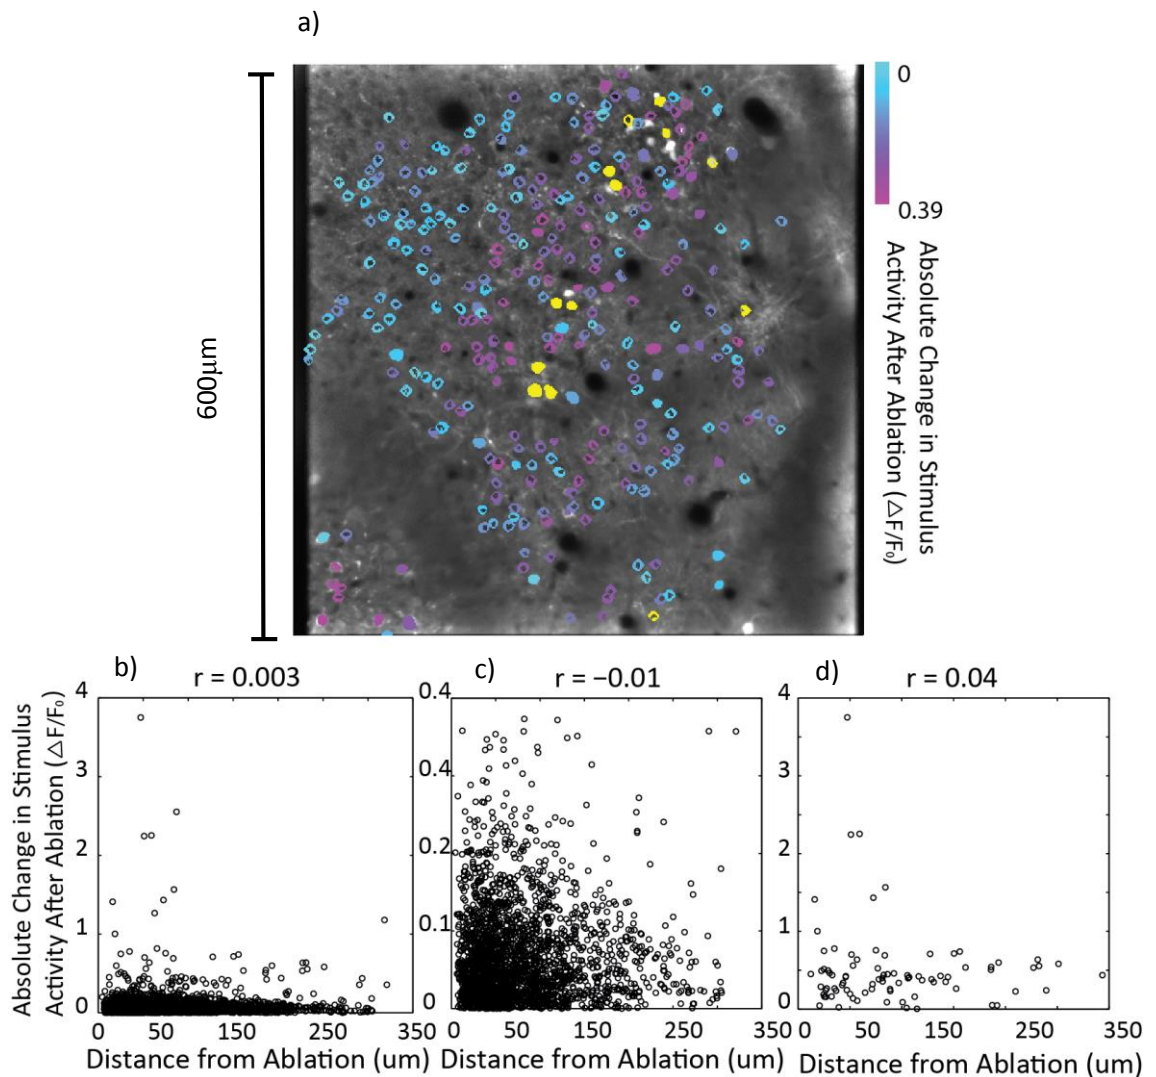
Figure 5.12. The Relationship between Change in Stimulus Activity after Ablation and Distance to Ablation a) A calcium image of one plane imaged in the air puff experiment, cells are colour-coded according to the absolute change in stimulus activity before and after ablation. Cyan represents small changes, while magenta represents larger changes (see colour bar). The ablated cells are colour coded yellow. The image indicates that cells showing a large change in stimulus activity are not clustered around the ablation sites. **b)** Changes in absolute stimulus activity before and after ablation plotted against distance to nearest ablation site. Y-axis shows absolute change in stimulus activity (stimulus$\Delta F/F_0$ – Baseline$\Delta F/F_0$) and x-axis distance to nearest ablation (μm). The title shows the correlation between the two axes. **c)** And **d)** same as b) after outliers have been excluded and when analysis is limited to stimulus cells, respectively.

To summarise, we have found ablating a small proportion of cells in a region of the barrel cortex results in a drop in overall tuning of cells to a sensory stimulus. Moreover, we did not find evidence for a relationship between changes in activity rates and distance to nearest ablation site. Although this analysis should not be

interpreted as giving conclusive evidence for the spatial precision of the ablation, it does suggest that if any collateral damage is incurred it is subtle. Consequently, the observed drop in stimulus tuning *may* rather reflect the functional implication of removing a small portion of excitatory input in a network.

5.4 Discussion

In this study we have described a method that enables consistent ablations of cells in layers II/III of the barrel cortex. These laser-induced lesions appear to be spatially precise. Moreover, we used the method to study the effects of removing a small proportion of excitatory cells from a volume of the barrel cortex and found it results in a subtle reduction in the tuning of the region to a sensory stimulus. These results and their implications will now be discussed in detail. These results and their implications will now be discussed in detail.

We found that cells up to a depth of approximately 300μm below the brain surface could be reliably ablated, with success rates dropping for cells deeper than this. The absolute depth limit of ablations was in the region of 400μm– marking the transition between layers II/III and layer IV (Hooks, Hires et al. 2011). Moreover, we observed a fluorescent signature for ablations that allowed efficient classification of success. Namely, when a cell was successfully ablated we observed an abrupt and transient drop in measured fluorescence intensity at the stimulated site. The mechanism underlying this drop is unclear at this point, but perhaps a sudden and momentary lysis-dependent diffusion of the fluorophore could represent one possible mechanism. Furthermore, cells targeted for ablation showed a significant reduction in their stimulus-related activity that exceeded baseline activity changes, observed across days when ablations did not take place; consistent with tissue lesioning. This effect was accentuated for cells with stronger stimulus-tuning; in fact, the activity of stimulus-tuned cells following ablation did not, on average, exceed the stimulus-activity threshold used (i.e. 0.5. $\Delta F/F_0$).

Although we observed a drop in the stimulus activity of cells neighbouring an ablation site, these cells often maintained their functional status after ablation and were significantly more active than cells targeted for ablation. Moreover, we found no evidence for a relationship between change in stimulus activity of a cell and the distance to the nearest ablation site. This null result was consistently obtained across different analyses addressing confounds such as the inclusion of outliers and large samples of inactive cells. Yet, we did note a small fraction of stimulus cells physically close to an ablation site (< 100μm) exhibited a large difference in their stimulus tuning following ablation. A possible interpretation of this is that *most* of the time ablations are focal to the targeted cell but on some rare occasions the ablations damage nearby tissue. Thus, the current results suggest laser-induced

ablations are likely to be limited to the targeted cell. However, future studies should be devoted to dissecting this issue further. For example, knowing the synaptic connections of ablated cells would assist in predicting what the functional implications of ablations would be. If the network effects of ablations match what would be predicted by connection patterns then this would lend support to the idea that ablations are focal. Moreover, high resolution microscopy, such as electron microscopy, could be employed to visually assess damage – perhaps axons just below or above the targeted cell are damaged in the process. Such assessments are difficult with the current method. Finally, a simple experiment that might shed light on this issue would be to ablate in one volume cells that show stimulus-locked activity and in another volume (or a different animal) inactive cells. If the method is precise then one would expect greater impact on network activity after ablating stimulus-locked cells than after ablating inactive cells. It is important to elucidate these issues in order to appreciate how to best apply this method.

With respect to the effect on network activity, we found ablating a small proportion of pyramidal cells in a volume of the barrel cortex had subtle, yet significant, effect on the brain region as a whole. Specifically, the average stimulus-related activity was significantly lower after ablations compared to that of the preceding days. This reduction was not a mere reflection of the gradual reduction in stimulus-related activity observed across days in baseline experiments, and thus could not reflect changes in baseline fluorescence of cells between days. Although the average stimulus-related activity decreased across baseline sessions the change observed 00was also *not* a result of the ablated cells exerting strong influence on the estimates of network activity prior to their removal as the same pattern of results was observed when they were excluded from the analysis. Therefore, we believe these results demonstrate one of the possible functional consequences of removing a small proportion of pyramidal cells in a network.

The reduction in stimulus-related activity following ablations may reflect the loss of excitatory input to cells the ablated targets were connected to. Pyramidal cells in layer II/III have various extrinsic and intrinsic connections (Feldmeyer 2012). Within layers II/III, pyramidal cells excite other pyramidal cells with a 10-20% probability (Feldmeyer, Lubke et al. 2006). Therefore, removing a number of pyramidal cells would reduce the excitatory drive to other pyramidal resulting in an overall dampening of sensory-related activity. Alternatively, various indirect ways exist for ablations to have the observed effect. One mechanism would be via the primary

motor cortex (M1). A prominent cortico-cortical connection exists between M1 and the barrel cortex (Hooks, Hires et al. 2011). Moreover, M1 has been found to influence sensory tuning in the barrel cortex (Petreanu, Gutnisky et al. 2012). Consequently, ablations could affect other cells in the network via input originating in M1. Ultimately, the effects of ablations reflect the connectivity of the targets – although connectivity in superficial layers of the cortex have been found to be varied (Kwan and Dan 2012), cortical pyramidal cells tend to project to other pyramidal cells (Feldmeyer, Lubke et al. 2006). Thus, removing some of these connections is likely to result in reduction in overall stimulus activity levels.

Before turning to the future directions and experimental considerations of this method I will briefly discuss potential mechanisms underlying ablations. Previous studies describing ablations in biological tissue, using similar methods to ours, have shown the lesions may be the result of plasma formation (Galbraith and Terasaki 2003, Heisterkamp, Maxwell et al. 2005). Plasma – the break-down of a molecule as a result of free-electron densities (caused by strong ionisation), can be induced by high levels of multi-photon absorption in a small focal volume. The plasma formation results in tissue removal, limited to the focal volume – similar to that reported in the current study. However, we have reasons to believe the lesions reported here are not the product of plasma formation. This is mainly due to the superior peak intensities required to induce a plasma. Both Heisterkamp and colleagues (2005) as well as Galbraith and Terasaki (2003) report laser pulse durations shorter than were used in the current study (i.e. 100fs compared to 200-300fs) – the shorter the pulse duration, the higher the peak intensity. Thus, the lesions we induced are likely to be the result of a different mechanism. The lesions could possibly result from thermal damage to the stimulated tissue or induced phototoxicity in the fluorophore. Specifically, the elevated irradiation could generate free radicals or perhaps singlet oxygens in the fluorophore, resulting in damage to the tissue surrounding the fluorophore. To address the possibility that ablations are fluorophore mediated, one could assess whether a cell can be ablated if stimulated at a wavelength outside of its fluorophore excitation spectra.

The applications of this method are manifold. However, to enhance its experimental utility the localisation of cell targets for ablation should be automated. The aforementioned results show it takes, on average, less than a second (0.73sec) to ablate a cell. However, in a given session we ablated only 30-40 cells at most due to the time taken to manually localise cell targets. With

automation one could conceivably ablate hundreds if not thousands of cells within a session. Automation of cell localisation would require some effort on the software development front, yet is viable – in fact, a software of this kind has already been developed by (Hayes, Wang et al. 2012). Once cell localisation is automatic this method could, in theory, be applied to any study looking at the causal role of a sparse neural representation in behaviour. However, keeping in mind the depth limitations, the method might best serve its purpose in the cortex. Moreover, topographically organised brain regions (such as the barrel cortex) of the sensory organ they represent may be the ideal candidate systems. In such systems one could fully remove a neural representation of behaviour and thus robustly assess whether it is required for successful behaviour. For brain regions that do not show such a topological arrangement one could employ the method to assess how a degraded neural representation affects behaviour. For example, one might ablate all cells in an imaged region of the hippocampus that follow a prominent, navigationally-relevant, cue in an environment and observe effects on navigation. Finally, the method could be applied to study the role of particular projections or cell sub-types for behaviour – such as the role of parvalbumin expressing (PV) interneurons, neurons thought to play a crucial role in network inhibition.

However, identifying neural representations underlying behaviour using calcium imaging is limited by the difficulty of inferring neural activity from calcium-based fluorescence. If estimates of the behavioural correlate of a neuron's activity are not accurate one would not know which neurons to ablate or how to interpret the behavioural effects after performing laser-induced ablations. Thus, before bringing this chapter to a close I will give a brief overview of some of the limitations of inferring action potentials from calcium-based fluorescent traces and how one could overcome them.

Calcium imaging is dependent on tracking a *correlate* of neural activity- cytoplasmic calcium concentration (Nakai, Ohkura et al. 2001, Tian, Hires et al. 2009, Chen, Wardill et al. 2013). Although the sensitivity and kinetics of genetically encoded indicators has improved dramatically in recent years, due to improvements to the calcium affinity and dynamic range of the indicators, they are still orders of magnitude slower (i.e. rise time ~0.2sec for GCaMP6s) than the calcium influx observed during neural activity and the duration of a single action potential; resulting in an activity signal of a relatively low temporal resolution (Tian, Hires et

al. 2009, Chen, Wardill et al. 2013). Moreover, the relationship between neural activity and fluorescence intensity (or calcium influx and fluorescence intensity) is non-linear (Nakai, Ohkura et al. 2001), such that changes in fluorescence intensity become progressively smaller the higher the activity rates are. In fact, algorithms that have been used to estimate numbers of action potentials from calcium-based fluorescence traces have found only a moderate relationship between changes in fluorescence intensity and activity rates (Vogelstein, Packer et al. 2010). Finally, long-term imaging of calcium indicators is not without problems. For example, resting calcium levels change between days. Moreover, baseline calcium indicator expression increases across days and sometimes cells develop nuclear expression – although these issues could be ameliorated with the use of transgenic mouse lines. The former could create artefacts that would suggest neurons reduce their activity across days, as fluorescence-based measures of neural activity are normalised to baseline fluorescence levels, and the latter has been found to result in aberrant neural responses (Chen, Wardill et al. 2013). Furthermore, several sources of noise exist that can contaminate the fluorescence signal; such as dark current from the photomultiplier tubes (PMTs), although this is negligible, and shot noise and biological movement. If the voltage gain of the PMTs is kept high enough one can limit considerably the influence of shot noise over the signal-to-noise ratio and as long as biological movement is corrected for during imaging the extent to which this noise source contaminates the signal can be mitigated. Yet, it is essential to keep the PMT gains consistent throughout and between imaging sessions, to enable legitimate comparisons between sessions. Finally, fluorescence signals from the neuropil (areas of the brain devoid of cell bodies, made up of axons, dendrites, dendritic spines, and glia (Kerr, Greenberg et al. 2005)) can further contaminate fluorescence signals from imaged cells. Consequently, it is important to estimate the neuropil signal and subtract it from the cell signal in order to obtain an accurate measure of neural activity.

Thus, when designing an experiment aimed to identify neurons involved in a particular behaviour it is important to bear in mind the aforementioned limitations. For example, it would appear unviable to study the neural coding of behavioural or sensory processes that exceeds the temporal resolution of the rise time of the calcium indicator used. Moreover, experiments that seek to obtain precise estimates of neural activity would not be compatible with calcium imaging. However, many sensory and behavioural experiments would still be viable. For

example, sensory events in an experiment would simply need to be separated in time such that the neural activity underlying each event could be unambiguously segregated. Moreover, many experiments only want to estimate whether a cell is active or not active, estimating *how* active it is, is not always crucial. With regards to long-term imaging, although expression changes over time, Chen and colleagues observed stable responses in V1 neurons infected with GCaMP6 over a period spanning weeks (2-3 weeks). Thus, although long-term imaging may introduce confounds in imagining data, these issues may only start influencing results if an experiment runs for more than two weeks. Consequently, although inferring neural activity from calcium-based fluorescence traces is not free from limitations it still represents a useful tool for studying the neural activity underlying behaviour, and thus to assess the importance of functionally defined cell micro-circuits.

To summarise, we have described a method whereby one can reliably perform highly precise lesions in the superficial layers of the barrel cortex. Moreover, we have demonstrated one of the many experimental applications of this method – namely the functional implications of removing a small sub-set of pyramidal cells on a network. Finally, we have discussed the limitations of using calcium imaging generally to study the behavioural correlate of neural activity and possible future directions for the ablation method. This method, although perhaps only an intermediate step to developing more sophisticated tools, can prove useful in neuroscientists end goal of relating action potentials to behaviour.

6. General Discussion

The aim of this thesis has been twofold: 1) to show a possible functional role of out-of-field activity in the hippocampus, and 2) describe the development and use of a method to interfere with functional circuits in the brain. With respect to the former, I have a) shown ensemble out-of-field place cell activity is higher at the start of goal-directed navigation, b) that it may support accurate navigation and, c) during rest such activity may serve to prepare neural sequences for future motivationally relevant trajectories. For the second objective, I have shown that one can induce cell-specific lesions using a two-photon microscope and described the results of one of its many experimental applications; namely, the functional implications of removing a small proportion of excitatory input on a brain region.

In this chapter, I will, firstly, review the findings from the first two experimental chapters, highlight outstanding issues and discuss their contribution in scientists' quest to understand the mechanisms and functions of the hippocampus. For the last data chapter, I will discuss some methodological advancement that could be made to enhance the usability of the laser ablation method and the potential it has for neuroscientific research in general. Finally, to bring the three chapters together, I will discuss possible applications of the ablation method to elucidating the role of functional micro-circuits in the hippocampus.

6.1 Functional Implications of Out-of-Field Hippocampal Activity

Navigational Out-of-Field Activity as Functional Theta Sequences?

In the first data chapter of this document I described a goal-directed navigational study showing ensemble place cell activity was higher at the start of navigation, and that as the animal approached its navigational goal rates dropped progressively. Moreover, we found the aforementioned relationship to relate to performance on the task; namely, the strength of the relationship positively predicted an animal's performance on the task. Importantly, this relationship could not be accounted for by place field distribution in the environment, thus we assume this effect reflects non-local hippocampal activity; suggesting a functional role for such activity.

In terms, of the hippocampal literature these findings accord with findings from navigational studies carried out on humans using functional MRI (fMRI) (Maguire, Burgess et al. 1998, Hartley, Maguire et al. 2003, Spiers and Maguire 2006, Brown, Hasselmo et al. 2014). These studies have found a) preferential hippocampal activity at the start of navigation (Spiers and Maguire 2006) b) that activity drops linearly in with distance to goal (Howard, Yu et al. 2011), and c) the amount of hippocampal activity during navigation predicts navigational accuracy (Maguire, Burgess et al. 1998, Hartley, Maguire et al. 2003). As a result, the findings described in this study represent an important, and rare, bridge between human imaging and rodent electrophysiological research.

With respect to the electrophysiological literature, the findings may be seen as somewhat novel. Although the data might benefit from more scrutiny in order to reach concrete conclusions about the mechanism underlying the described effect, it could be considered consistent with the proposed role of place cell sequences observed during a theta cycle (Jensen and Lisman 1996, Foster and Wilson 2007, Johnson and Redish 2007). Perhaps, a neural sequence of the future path of the animal is activated in theta cycles. If this were the case, one would expect a transient increase in activity within theta cycles only at the start of navigation (unless the animal re-plans its trajectory in the middle of navigation) and a relationship with performance, as observed in the current study.

As mentioned at various points throughout this thesis, a product of phase precession is place cell sequences of paths just behind and ahead of an animal's current location are represented within a theta cycle (O'Keefe and Recce 1993).

Although 'theta sequences' might be a mere physiological bi-product of phase precession others afford it a functional role (Foster and Wilson 2007, Johnson, Fenton et al. 2009, Foster and Knierim 2012, Gupta, van der Meer et al. 2012). Perhaps, theta sequences represent a kind of 'look-ahead' mechanism (Johnson, Fenton et al. 2009, Foster and Knierim 2012), allowing the animal to anticipate future paths and guide decision making. Moreover, different parts of a theta cycle have been proposed to be propitious for distinct mnemonic tasks (Brankack, Stewart et al. 1993, Wyble, Linster et al. 2000, Hasselmo, Bodelon et al. 2002) – the trough of theta is thought more favourable for encoding whereas the peak of the theta seems better suited for retrieval. With respect to the study described in this document, perhaps the enhanced ensemble activity represents retrieval, within a theta cycle, of the place cell sequence of the goal-directed path ahead, helping the animal plan its navigational trajectory. Yet, as mentioned earlier, this hypothesis needs to be tested directly in a study where a) locations represented at the start of navigation are studied, b) the relationship between start activity and the theta rhythm understood and c) the animals' movement throughout navigation controlled so to rule out behavioural confounds, as discussed in more detail in section 3.4 of chapter 3.

It should be noted that little evidence exists for a functional role of theta sequences. Although they seem to occur with a higher probability at important navigational points, such as choice points (Johnson and Redish 2007), the paths represented tend to *not* correlate with the future behaviour of an animal. Johnson & Redish (2007) found theta sequences at choice points to represent the *possible* future paths of the animal rather than the path the animal *chose* to take. Yet, a recent study implied a structure to theta sequences. Gupta and colleagues (2012) found theta sequences to preferentially represent trajectories between physical landmarks in an environment, rather than trajectories across them (Gupta, van der Meer et al. 2012). Moreover, some studies show place cell sequences in a theta cycle represent places not local to the animal - an apparent 'flickering' between environments, which was found to depend on spatial ambiguity (Jezek, Henriksen et al. 2011). Consequently, if the reported enhanced non-local place cell activity at the outset of navigation reflects a sequence compressed in a theta cycle of the future goal-directed path this would add credence to the suggestion that theta sequences are functional and perhaps afford it a planning function.

In sum, the current study has implied out-of-field hippocampal activity is functional, highlighted necessary follow-up studies, and pointed to the possible mechanism underlying its function. The study of the functional implications of theta sequences is a nascent area of research which will hopefully receive more attention in future years and shed further light on the contribution of this intensively studied brain rhythm to cognition.

Hippocampal Non-Local Activity as Preparation for Goal-Directed Behaviour?

The second data chapter of this thesis related to another type of non-local hippocampal activity but that is, perhaps, better understood than theta sequences and has also been implicated in navigational planning; namely, the synchronized activity of place cells observed during behavioural quiescence (e.g.(Wilson and McNaughton 1994)). To quickly re-iterate the principal finding of this study, we found the hippocampus *pre*plays future place cell sequences for a motivationally relevant, unvisited environment during rest. Although a few previous studies have found hippocampal activity during quiescent behaviours to anticipate the future path of an animal (Pfeiffer and Foster 2013, Singer, Carr et al. 2013), such prospective activity has only been found for familiar environments (with the exception of (Dragoi and Tonegawa 2011, Dragoi and Tonegawa 2013)). Thus, we consider our finding novel and we think it extends the function of the hippocampus to representing unvisited environments. Importantly, we did not find evidence implying the effect was a result of pre-existing maps in the hippocampus (McNaughton, Barnes et al. 1996, Dragoi and Tonegawa 2011, Dragoi and Tonegawa 2013), rather we believe the results are a result of motivational influences that bias which sequences are played out during rest. Such an idea is reinforced by finding the more an animal showed interest in a novel environment the more the future place cell sequence for that environment was preplayed.

The implications of this study are numerous. First of all, the findings imply the highly coordinated activity observed during quiescence does not merely reflect retrieval of recently formed memory traces but also the 'construction' of future traces. This idea might seem quite radical considering the long-standing hypothesis that the hippocampus is important for memory formation and consolidation. However, *replay* of recent place cell sequences, observed by various laboratories, do not account for all spike sequences observed during rest (e.g. (Nadasdy, Hirase et al. 1999, Lee and Wilson 2002). Thus, a part of the spike sequences that do not

correlate with past experience might correlate with future, novel experiences – similar to that found by Dragoi and Tonegawa (2011, 2013). Moreover, an anterograde function of such activity is consistent with developmental findings showing sharp-waves develop before (post-natal days 4-6) spatial cells (Leinekugel, Khazipov et al. 2002), such as place cells and grid cells (Langston, Ainge et al. 2010, Wills, Cacucci et al. 2010). Finally, a form of prospective coding in the activity of place cells is not a novel finding. Various laboratories have found the activity of a place cell is modulated by the future path of the animal. For example on an alternation task, a place cell may not fire at all when the animal is on the central arm of an environment and the future turn will be to the left whereas when the animal is on the same arm but is about to turn right the place cell fires robustly (Frank, Brown et al. 2000, Wood, Dudchenko et al. 2000). Thus, the reported preplay of the animal's future environment may simply extend the prospective coding observed in place cell activity during behaviour to that observed during quiescence.

A second implication of the study is seeing *but not reaching* a novel environment is enough to create a place cell sequence. This finding seems at odds with findings showing that place cells require direct experience with an environment to form a spatial representation of it (Rowland, Yanovich et al. 2011). Rowland and colleagues (2011) blocked the NMDA receptor pharmacologically while animals explored a novel environment they had previously observed but not been into. NMDA blockade should preferentially interfere with newly formed place fields. Consequently, if the pharmacological manipulation fails to de-stabilise the place fields in the novel environment this would indicate the place fields had been formed earlier, while the animal observed the environment. This is *not* what the authors found. Rather, NMDA blockade resulted in de-stabilised place fields of the novel environment; implying place cells need direct experience with a novel environment in order to form stable place fields. To reconcile these divergent findings, one might consider the role motivational relevance of a novel environment plays in place field formation. Perhaps, if a novel environment an animal can only observe seems behaviourally relevant, the mere observation of it can drive the formation of place fields. In the current study, preplay was only observed for an environment an animal had showed an interest in, another environment, not motivationally-relevant but equally visible, was not preplayed. The role of goals and reward in driving place cell activity is ambiguous. In many

foraging and spatial tasks the activity at goal-sites is confounded by animals spending more time at the goal sites. Moreover, when located at goal sites animals are often relatively immobile. Thus, activity associated with these sites might be a result of sharp-wave ripple (SWR) associated population bursts (O'Keefe and Nadel 1978, Buzsaki, Horvath et al. 1992). However, some recent studies have implied place cells may be modulated by goal-locations even after taking into account these behavioural confounds – e.g.(Singer and Frank 2009, Dupret, O'Neill et al. 2010). For example, Dupret and colleagues (2010) found CA1 place cells clustered around navigational goals, whereas CA3 place cells did not. An effect of goals in just one of the CA fields cannot simply be explained by behavioural differences around goal sites. Moreover, Singer and Frank (2009) found CA3 place cells were significantly more active during SWRs following a rewarded trajectory than an unrewarded trajectory (Singer and Frank 2009). Perhaps the current finding represents another example of goal-related modulation of hippocampal place cells. Therefore, the findings form the current study highlight the possible influence of goals in driving place cell activity. Hopefully future studies will be better able to explore the role goals and rewards play in modulating place cell activity – the use of optogenetics methods where one might be able to create a synthetically rewarding experience (by exciting dopaminergic cells with light) might help with this quest.

Lastly, the current study has implication for the role of synchronous activity observed during quiescence. To re-iterate, re-activation of recently formed place cell sequences ('replay') was originally posited as the mechanism underlying memory consolidation in the hippocampus (Wilson and McNaughton 1994). However, activating a place cell sequence not yet experienced does not accord with this function. Recently, theorists have suggested replay may rather serve a prospective planning function in navigation (Erdem and Hasselmo 2012, Foster and Knierim 2012, Pfeiffer and Foster 2013, Erdem and Hasselmo 2014). For example, during a spatial alternation task Singer and colleagues found activity during SWRs to be more synchronized preceding a correct vs an incorrect turn – as indexed by the number of co-active cell pairs (Singer, Carr et al. 2013). More direct evidence for a 'planning' role of replay comes from a study carried out by Pfeiffer and Foster (2013) where they found replay at the start of a goal-directed trajectory predicted the future path of the animal. Importantly, such prospective replay was not found at the initiation of a non-goal-directed trajectory. Although a relationship between preplay and future behaviour was not directly assessed in this study, one might

suggest that preplay is a preparatory mechanism for future novel navigation; this function should be assessed in future studies.

More generally, this and other studies might suggest the mnemonic function of the hippocampus is more constructive and prospective in nature. Firstly, activity presumed to only reflect consolidation may also reflect preparation of place cell sequences for unexperienced yet desired behaviours. Such suggestions agree with neuropsychological findings showing hippocampal patients are impaired at imaging novel experienced (Hassabis, Kumaran et al. 2007). Moreover, reactivations of recent experiences might be selective, such that not just any recent experience is reactivated but rather those relevant to the animals' motivational states and future behaviour. Such functions could represent exciting developments in our understanding of this comprehensively researched brain region and open doors to hosts of studies to test these ideas.

6.2. Synergies between Studies: A Role for Hippocampal Noisy States

The first two data chapters of this thesis have shown hippocampal place cell activity occurring outside a cell's main firing field may be functional; specifically it may prepare the animal for future spatial behaviour. The idea hippocampal activity may be prospective in nature and out-of-field activity may be more than just 'noise' is supported by numerous recent studies (Wood, Dudchenko et al. 2000, Johnson and Redish 2007, Pfeiffer and Foster 2013). Thus, the studies described in this thesis encourage the movement away from traditional views suggesting non-local place cell activity is noise (O'Keefe and Nadel 1978, Muller, Kubie et al. 1987) and rather suggest it may contain a structure, relevant to behaviour.

It has been established that hippocampal place cell activity can be extremely variable (e.g. (Fenton and Muller 1998)). Almost identical passes through a cell's firing field are associated with varying numbers of action potentials. Fenton and Muller (1998) quantified this variability (referred to as 'overdispersion') and found it exceeded that expected by an inhomogeneous Poisson process. This brings one to ask – what is the source of this variability? Johnson and colleagues suggested that such variability might be the result of covert ('cognitive') variables, such as attention, task demands, behavioural state, etc. Thus, by accounting for the covert variable one would be able to account for more of place cell spiking (Johnson, Fenton et al. 2009).

The idea that covert variables can account for neural activity is not a novel idea. Even in the motor and sensory cortices the activity of neurons is often better described by cognitive rather than sensory factors (Georgopoulos, Lurito et al. 1989, Rizzolatti, Fadiga et al. 1996). For example, in a seminal study by Georgopoulos and colleagues (1989) activity of single neurons was recorded from the motor cortex while rhesus monkeys performed a motor task requiring them to move their arm to a position perpendicular to that of a light cue. The neural population vector was computed from the onset of the light cue until the monkey executed an arm movement. Georgopoulos found a gradual shift in the population vector starting at the position of the light rotating to the position perpendicular to the light, where the monkey eventually moved its arm to (Georgopoulos, Lurito et al. 1989). In other words, the motor cortex seemed to be performing a 'mental rotation' of the cue location in order to locate the target position; hence, the

activity of this motor region was better accounted for by a cognitive operation rather than the motor operations of the monkey at the time.

In the hippocampus, findings have accumulated that imply non-spatial factors modulate place cell activity. Some of these are discussed in detail in the first chapter of this document, for example, the influence of navigational goals (Hollup, Molden et al. 2001, Dupret, O'Neill et al. 2010), the prospective/retrospective path of the animal (Frank, Brown et al. 2000, Wood, Dudchenko et al. 2000, Ferbinteanu and Shapiro 2003) and, of course, behavioural state (O'Keefe and Nadel 1978, Buzsaki, Horvath et al. 1992). Moreover, others have suggested attention (Kentros, Agnihotri et al. 2004), task phase (Markus, Qin et al. 1995, Griffin, Eichenbaum et al. 2007) and time (MacDonald, Lepage et al. 2011, Kraus, Robinson et al. 2013) can also account for some of the observed place cell activity. In fact, some of these variables could be interpreted as accounting for 'noisy' out-of-field spikes. To give an example, cells becoming silent on identical passes through the stem of a T-maze can be accounted for by assuming the cells code for the future path of the animal which might differ between identical passes through the stem (Wood, Dudchenko et al. 2000). An alternative explanation of this effect, also assuming the influence of 'covert' variables, is that when the animal is running on the stem it retrieves its past path in order to guide future decision making (Hasselmo and Eichenbaum 2005)– retrieval would result in activation of different sets of place cells during identical runs on the stem, depending on which path the animal had just taken. Moreover, spontaneous re-mapping of place cells in a familiar environment – a phenomenon reported by many studies (Markus, Qin et al. 1995, Knierim, Kudrimoti et al. 1998, Hayman, Chakraborty et al. 2003, Leutgeb, Leutgeb et al. 2005), can partly be explained by task demands (Kentros, Agnihotri et al. 2004) and cues relied on during navigation (Fenton, Lytton et al. 2010). Finally, out-of-field activity often observed at the end and start of a run on a linear track has repeatedly been found to represent a retrieval of a previous experience (Foster and Wilson 2006, Diba and Buzsaki 2007). Thus, an amalgamation of studies has found various 'covert' variables can account for place cell activity traditionally referred to as noise.
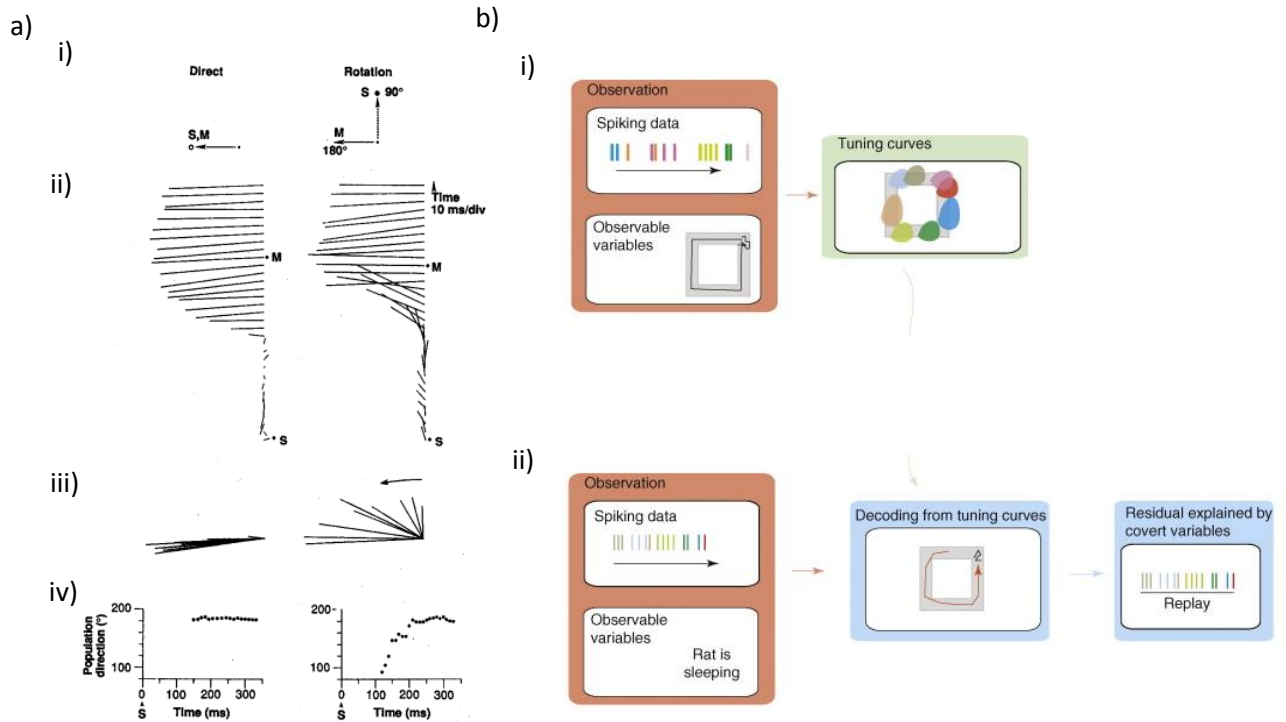
Figure 6.1. The Influence of 'Covert' Variables on Neural Activity a) Example decoded activity of a motor neuron recorded from rhesus monkey while the monkey carries out a task where it had to either reach to a location of a light cue (left) or to a location perpendicular to the cue (right). i) Task, S = light cued, M = movement, ii) decoded population vectors from cue onset until monkey made arm movement. iii) Spatial view of population vectors shown in ii), iv) decoded direction derived from the population vector. **b)** Neural activity accounted for by a covert variable. i) While an animal runs on a square track (bottom left), spikes are recorded (top left) and spatial tuning curves constructed (right). ii) During sleep, spike data is recorded yet no observable variables are available to account for the activity (left). However, if the tuning curves are applied to the spiking data, it is clear the spikes represent the recently traversed square environment (middle). Thus, a covert variable (retrieval of prior trajectory) is able to account for all the spikes in sleep (right).

With respect to the studies described in this thesis, both have described activity occurring outside a place cell's main firing field and in both instances seem functional and perhaps important for future behaviour. Therefore, the reported findings may be interpreted as adding to the body of literature suggesting noisy hippocampal states are functional, and are better accounted for by assuming the influence of covert, perhaps cognitive variables. Finally, we have suggested the cognitive operation influencing the observed out-of-field activity might be preparation for future behaviour – a form of future planning (although this hypothesis remains to be empirically tested).

Before bringing this discussion to a close I will briefly address the possible reason why future planning might be sub-served by both the theta and SWR oscillatory states. Although both could theoretically underpin future planning, place cell sequences in theta cycles might be important for planning of short trajectories – i.e. paths just ahead, whereas SWR-associated place cell sequences might allow for planning over longer distances. Such a hypothesis is consistent with the length of trajectories represented in each state. Namely, phase precession dictates only places just ahead of the animal are represented (O'Keefe and Recce 1993) whereas longer trajectories (up to 10 metres (Davidson, Kloosterman et al. 2009)), have been found to be represented during SWRs. Furthermore, the decay time of the NMDA receptor channel (i.e. 100-150ms (Debanne, Guerineau et al. 1995, Debanne, Shulz et al. 1995) would dictate binding of short paths in theta cycles and possibly longer paths during SWRs. Moreover, findings showing interference with SWR disrupts learning (Girardeau, Benchenane et al. 2009) implies this state is important for learning and memory formation, although evidence for its direct involvement in synaptic modification is lacking. Finally, the lack of sensory input associated with SWRs might also provide propitious conditions for flexible binding of places in order to guide navigation of novel routes.

6.3. A Tool for Assessing the Causal Role of Functional Circuits in Behaviour

In the third data chapter of this thesis, we described the development of a method to reliably and precisely ablate (i.e. lesion) single cells in the superficial layers of the barrel cortex by exposing targeted cells to elevated doses of a femtosecond laser. We showed the laser-induced lesions were spatially precise; most of the time, only incurring damage to the targeted cell and that we could reliably ablate cells along the entire length of layers II/III. Moreover, we described preliminary results from one of the many experimental applications of the method – namely, sensory tuning of a brain region is slightly reduced if a small proportion of sensory-tuned pyramidal cells are ablated. This method holds important potential for elucidating causal links between neural activity and behaviour.

The candidate brain region for the development of this method was the barrel cortex. However, we view this method as a *general* method that could, theoretically be applied to any part of the brain (although perhaps cortical regions would be the best suited areas considering the depth limitations reported). Thus, we believe this method represents a valuable addition to the toolbox of available methods to interfere and manipulate neural activity.

One might ask, how does this method compare to other available methods designed to assess the causal role between neuron spiking and behaviour? If compared to older methods for lesioning tissue – such as aspiration, pharmacological inactivation or electrical stimulation, the advantages of the current method are obvious. The precision of the laser-induced lesions compared to the more traditional methods is far superior; the current method allows for lesions on the micron scale while other methods produce lesions on the millimetre or even centimetre scale, often incurring unwanted damage to other brain regions and fibres of passage in an uncontrolled manner. With respect to more modern techniques such as optogenetics the current method may be considered superior in some aspects but inferior in others. The main disadvantage of the current method compared to optogenetics is that optogenetics allows reversible interference of neural activity, whereas tissue lesioning incurred by laser irradiation is permanent. Moreover, optogenetics can more readily be applied to different parts of the brain. The depth limitations (~400μm) of the current method suggest the cortex may be the best suited brain region for its application. However, such

limitations could be overcome, to an extent, by either using laser sources with a higher repetition rate (to increase the peak intensity and thus depth penetration) or by simply removing intervening tissue for regions deeper in the brain. The important advantage of the current method is that it does *not* rely on genetic markers for targeting neurons – rather neurons can be targeted based on functional (or any other) classification. This advantage is important as within brain regions multiple micro-circuits exist, serving often largely distinct functions (O'Keefe and Dostrovsky 1971, Huber, Gutnisky et al. 2012). Moreover, automating cell detection, as discussed earlier, could allow for ablation of a large numbers of cells (*at leas*t a few hundred cells). Consequently, the current method could play a crucial role in neuroscientists aim to elucidate the necessity of functional neural circuits for successful behaviour.

The experimental applications of this method are many and diverse, especially if cell detection is automated. Although the method is perhaps best suited for sensory brain regions where sensory representations are topographically arranged. For example, one could assess whether sparse neural representations for sensory stimuli often observed in the cortex (Huber, Gutnisky et al. 2012, Peron, Iyer et al. 2012) are required for successful behaviour. Peron and colleagues (2012) found only about 1% of neurons in one barrel column in layers II/III of the barrel cortex are tuned to a particular sensory event (such as touch) in a whisker-based object localization task (Peron, Iyer et al. 2012). One could assess whether this sparse functional representation is necessary for successful object localization by ablating all the stimulus-tuned cells (given the animal can only use the whisker subserved by the targeted barrel column to solve the task). Another experimental application would be to sever axon projections between brain regions, to assess the functional contribution of projections. For example, an important projection from the primary motor cortex (M1) to the barrel cortex exists (Hooks, Hires et al. 2011), known to contribute to sensory coding in the barrel cortex (Petreanu, Gutnisky et al. 2012). One could simply sever axons from M1 to the barrel cortex, and assess the behavioural effects. More computationally-inclined questions could also be addressed with this technique. One could assess how many neurons in a functional circuit are needed to enable an animal to perform a task successfully. Perhaps, a degraded representation is enough? If that is the case, at what point does behaviour break down? Such questions might be also appropriate for non-topographically organized brain regions, such as the hippocampus. Finally, one

could investigate research questions relating to plasticity in the brain. For example, if you remove a functional micro circuit do other neurons take over the function of that circuit? Such experiments would also have important clinical implications.

To summarise, we have described a potentially powerful method for assessing causal links between neural activity and behaviour that is particularly well suited to address research questions where the target cell population cannot be defined genetically. It is clear that its experimental applications are manifold and diverse; thus we believe this method could make an important contribution to neuroscientists' end goal to relate action potentials to behaviour.

6.4. The Application of Laser-Induced Ablations to the Study of Hippocampal Functional Circuits

To conclude this thesis we will briefly describe some of the possible applications of the described ablation method to the study of functional representations in the hippocampus. Firstly, to the study of the proposed navigational planning function and then generally to the study of the role of this brain region in behaviour and cognition.

As mentioned earlier, the method to lesion individual cells by means of laser irradiation is likely best suited for superficial brain regions – within the technique's depth limitation, having a topographical organisation of the sensory receptors it represents. However, the method could still prove useful for sub-cortical brain regions like the hippocampus, if, simply, the cortical tissue above the region is removed - allowing the two-photon microscope to be positioned on the dorsal surface of the hippocampus. In fact, some laboratories have already successfully tried this (Dombeck, Harvey et al. 2010, Lovett-Barron, Kaifosh et al. 2014). Moreover, instead of applying the ablation method to completely remove a functional representation one could assess the behavioural effect of degrading a representation by, for example, ablating a proportion of the population coding for a behaviour.

In relation to the effects described in chapter 3, where enhanced place cell activity was observed at the start of navigation (which was also found to relate to performance), perhaps at the start of navigation a random set of place cells could be ablated, and performance on the task assessed. This approach might prove too crude. If that is the case, a more targeted line of action might be to first identify cells that contribute often to the enhanced activity at navigation initiation and ablate a portion of these. Then the performance of the animals could be compared to that of controls where a random set of place cells are ablated.

With respect to the effects described in chapter 4 – showing the hippocampus preplays motivationally relevant future paths during rest, perhaps the role of such preplay in future navigation could be probed more directly with the ablation method. Namely, an animal navigates in an environment, perhaps a virtual-reality maze, with a part of the environment made inaccessible yet visible. The inaccessible part of the environment might be a corridor, representing a shortcut to a goal. The animal is exposed to the inaccessible corridor for enough time to

record awake preplay data. We would assume sequences active during awake preplay to at least partly represent the future place cell sequence for the inaccessible corridor. Thus, the cell sequence identified during awake preplay would be ablated during subsequent sleep. In a control group we would ablate random cells. Upon re-entry into the maze, with all corridors now accessible, we could assess whether the experimental group is *less* likely to take the previously inaccessible shortcut to the goal than the control group.

More general functions of the hippocampus could also be assessed using the ablation method. For example, is the activity of cells coding for a particular environment needed for successful navigation? Although studies have found the hippocampus as a whole is required for navigation (Morris, Garrud et al. 1982), a causal link between place cells active in an environment and navigation has yet to be established. One could ablate a portion of cells active in a navigational environment and observe effects on performance.

To avoid causing damage to the brain, one could rather target ablations to superficial regions of the brain that might be important for hippocampal processing. The medial pre-frontal cortex, visual cortex or the barrel cortex might be possible candidates. Simultaneously, activity could be recorded from hippocampal place cells, using less invasive methods like electrophysiology. A task that encourages involvement of a the target cortical region would be used, and the effect on place cell activity after ablating cortical cells involved in the task assessed, as well as effects on behaviour. To give an example, animals might carry out a virtual reality navigational task, while the visual cortex is imaged using a two-photon microscope and place cell activity recorded using tetrodes. Cortical cells involved in the task are identified and ablated and the effect on place cells analysed. Such a study might further our understanding of the sensory determinants of place cell activity and the perhaps the role of cortical regions in navigation.

To conclude, we have discussed some of the many applications of the ablation method described in chapter 5 to the study of hippocampal functioning; highlighting the flexibility with which it can be employed and its utility for studying functional neural circuits.

6.5. Conclusions

To conclude, in this thesis I have described two studies showing hippocampal activity traditionally viewed as 'noise' may have a functional role for behaviour. Specifically, we believe such activity can represent future, goal-directed trajectories. Moreover, we have speculated such prospective activity may support planning of future navigational experiences. Finally, we have also described the development of a method that allows for targeted, spatially-precise, cell ablations, which may be ideal for addressing the role of functional circuits, such as those described in the preceding chapters, for behaviour. We think these studies represent important contributions to our ever-advancing understanding of the hippocampus' role in cognition and the means to study the neural basis of behaviour. Hopefully, inspiring many more experiments and further scientific progress.

References

Abrahams, S., R. G. Morris, C. E. Polkey, J. M. Jarosz, T. C. Cox, M. Graves and A. Pickering (1999). "Hippocampal involvement in spatial and working memory: a structural MRI analysis of patients with unilateral mesial temporal lobe sclerosis." Brain and cognition **41**(1): 39-65.

Adey, W. R. (1967). "Hippocwmpal states and functional relations with corticosubcortical systems in attention and learning." Progress in brain research **27**: 228-245.

Aimone, J. B., W. Deng and F. H. Gage (2011). "Resolving new memories: a critical look at the dentate gyrus, adult neurogenesis, and pattern separation." Neuron **70**(4): 589-596.

Allen, G. V. and D. A. Hopkins (1989). "Mamillary body in the rat: topography and synaptology of projections from the subicular complex, prefrontal cortex, and midbrain tegmentum." The Journal of comparative neurology **286**(3): 311-336.

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. and P. Lavenex (2007). Hippocampal Neuroanatomy. The Hippocampus Book. New York, Oxford University Press**:** 37-115.

Amaral, D. G. and M. P. Witter (1995). Hippocampal Formation. The Rat Nervous System, Academic Press. **Morris, R.G.M., Amaral, D.G., Bliss, T.V., & O'Keefe, J.:** 443-486.

Andersen, P., T. Bliss and K. Skrede (1971). "Lamellar organisation of hippocampal excitatory pathways." Experimental brain research **12**: 222-238.

Anderson, M. I. and K. J. Jeffery (2003). "Heterogeneous modulation of place cell firing by changes in context." The Journal of neuroscience : the official journal of the Society for Neuroscience **23**(26): 8827-8835.

Arantius, G. (1587). De humano foetu...Ejusdem anatomicorum observationum liber. Venice.

Baker, P. F., A. L. Hodgkin and E. B. Ridgway (1971). "Depolarization and calcium entry in squid giant axons." J Physiol **218**(3): 709-755.

Baks-Te Bulte, L., F. G. Wouterlood, M. Vinkenoog and M. P. Witter (2005). "Entorhinal projections terminate onto principal neurons and interneurons in the subiculum: a quantitative electron microscopical analysis in the rat." Neuroscience **136**(3): 729-739.

Barry, C. and D. Bush (2012). "From A to Z: a potential role for grid cells in spatial navigation." Neural Systems & Circuits **2**(1): 6.

Barry, C., C. Lever, R. Hayman, T. Hartley, S. Burton, J. O'Keefe, K. Jeffery and N. Burgess (2006). "The boundary vector cell model of place cell firing and spatial memory." Reviews in the neurosciences **17**(1-2): 71-97.

Bassett, J. P. and J. S. Taube (2001). "Neural correlates for angular head velocity in the rat dorsal tegmental nucleus." The Journal of neuroscience : the official journal of the Society for Neuroscience **21**(15): 5740-5751.

Bendor, D. and M. A. Wilson (2012). "Biasing the content of hippocampal replay during sleep." Nature neuroscience **15**(10): 1439-1444.

Bianco, I. H., L. H. Ma, D. Schoppik, D. N. Robson, M. B. Orger, J. C. Beck, J. M. Li, A. F. Schier, F. Engert and R. Baker (2012). "The tangential nucleus controls a gravito-inertial vestibulo-ocular reflex." Current biology : CB **22**(14): 1285-1295.

Blackstad, T. W., K. Brink, J. Hem and B. Jeune (1970). "Distribution of hippocampal mossy fibers in the rat. An experimental study with silver impregnation methods." The Journal of comparative neurology **138**(4): 433-449.

Born, J. and I. Wilhelm (2012). "System consolidation of memory during sleep." Psychological research **76**(2): 192-203.

Borst, J. G., L. W. Leung and D. F. MacFabe (1987). "Electrical activity of the cingulate cortex. II. Cholinergic modulation." Brain research **407**(1): 81-93.

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. M. Andrews and M. E. Hasselmo (2012). "Head direction cells in the postsubiculum do not show replay of prior waking sequences during sleep." Hippocampus **22**(3): 604-618.

Brankack, J., M. Stewart and S. E. Fox (1993). "Current source density analysis of the hippocampal theta rhythm: associated sustained potentials and candidate synaptic generators." Brain research **615**(2): 310-327.

Brodmann, K. (1909). Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues. Leipzig, Germany, Barth.

Brown, T. I., M. E. Hasselmo and C. E. Stern (2014). "A High-resolution study of hippocampal and medial temporal lobe correlates of spatial context and prospective overlapping route memory." Hippocampus.

Brown, T. I., R. S. Ross, J. B. Keller, M. E. Hasselmo and C. E. Stern (2010). "Which way was I going? Contextual retrieval supports the disambiguation of well learned overlapping navigational routes." J Neurosci **30**(21): 7414-7422.

Buckner, R. L. and D. C. Carroll (2007). "Self-projection and the brain." Trends in cognitive sciences **11**(2): 49-57.

Bullock, T. H., G. Buzsaki and M. C. McClune (1990). "Coherence of compound field potentials reveals discontinuities in the CA1-subiculum of the hippocampus in freely-moving rats." Neuroscience **38**(3): 609-619.

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

Burgess, N. and J. O'Keefe (2005). "The theta rhythm." Hippocampus **15**(7): 825-826.

Bush, D., C. Barry and N. Burgess (2014). "What do grid cells contribute to place cell firing?" Trends in neurosciences **37**(3): 136-145.

Buzsaki, G. (1986). "Hippocampal sharp waves: their origin and significance." Brain research **398**(2): 242-252.

Buzsaki, G. (2002). "Theta oscillations in the hippocampus." Neuron **33**(3): 325-340.

Buzsaki, G., Z. Horvath, R. Urioste, J. Hetke and K. Wise (1992). "High-frequency network oscillation in the hippocampus." Science **256**(5059): 1025-1027.

Buzsaki, G., L. W. Leung and C. H. Vanderwolf (1983). "Cellular bases of hippocampal EEG in the behaving rat." Brain research **287**(2): 139-171.

Buzsaki, G. and E. I. Moser (2013). "Memory, navigation and theta rhythm in the hippocampal-entorhinal system." Nat Neurosci **16**(2): 130-138.

Buzsaki, G., P. Rappelsberger and L. Kellenyi (1985). "Depth profiles of hippocampal rhythmic slow activity ('theta rhythm') depend on behaviour." Electroencephalography and clinical neurophysiology **61**(1): 77-88.

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): 115-129.

Canteras, N. S. and L. W. Swanson (1992). "Projections of the ventral subiculum to the amygdala, septum, and hypothalamus: a PHAL anterograde tract-tracing study in the rat." The Journal of comparative neurology **324**(2): 180-194.

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**: 381243.

Canty, A. J., L. Huang, J. S. Jackson, G. E. Little, G. Knott, B. Maco and V. De Paola (2013). "In-vivo single neuron axotomy triggers axon regeneration to restore synaptic density in specific cortical circuits." Nature communications **4**: 2038.

Chen, T. W., T. J. Wardill, Y. Sun, S. R. Pulver, S. L. Renninger, A. Baohan, E. R. Schreiter, R. A. Kerr, M. B. Orger, V. Jayaraman, L. L. Looger, K. Svoboda and D. S. Kim (2013). "Ultrasensitive fluorescent proteins for imaging neuronal activity." Nature **499**(7458): 295-300.

Cheng, S. and L. M. Frank (2008). "New experiences enhance coordinated neural activity in the hippocampus." Neuron **57**(2): 303-313.

Cohen, H. S. (2000). "Vestibular disorders and impaired path integration along a linear trajectory." Journal of vestibular research : equilibrium & orientation **10**(1): 7-15.

Cornwell, B. R., N. Arkin, C. Overstreet, F. W. Carver and C. Grillon (2012). "Distinct contributions of human hippocampal theta to spatial cognition and anxiety." Hippocampus **22**(9): 1848-1859.

Cornwell, B. R., L. L. Johnson, T. Holroyd, F. W. Carver and C. Grillon (2008). "Human hippocampal and parahippocampal theta during goal-directed spatial navigation predicts performance on a virtual Morris water maze." The Journal of neuroscience : the official journal of the Society for Neuroscience **28**(23): 5983-5990.

Cressant, A., R. U. Muller and B. Poucet (1997). "Failure of centrally placed objects to control the firing fields of hippocampal place cells." The Journal of neuroscience : the official journal of the Society for Neuroscience **17**(7): 2531-2542.

Csicsvari, J., H. Hirase, A. Czurko, A. Mamiya and G. Buzsaki (1999). "Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving Rat." The Journal of neuroscience : the official journal of the Society for Neuroscience **19**(1): 274-287.

Csicsvari, J., H. Hirase, A. Mamiya and G. Buzsaki (2000). "Ensemble patterns of hippocampal CA3-CA1 neurons during sharp wave-associated population events." Neuron **28**(2): 585-594.

Daitz, H. M. and T. P. Powell (1954). "Studies of the connexions of the fornix system." Journal of neurology, neurosurgery, and psychiatry **17**(1): 75-82.

Darwin, C. (1873). "Origin of certain instincts." Nature **7**: 417-418.

Davidson, T. J., F. Kloosterman and M. A. Wilson (2009). "Hippocampal replay of extended experience." Neuron **63**(4): 497-507.

Day, M., R. Langston and R. G. Morris (2003). "Glutamate-receptor-mediated encoding and retrieval of paired-associate learning." Nature **424**(6945): 205-209.

Debanne, D., N. C. Guerineau, B. H. Gahwiler and S. M. Thompson (1995). "Physiology and pharmacology of unitary synaptic connections between pairs of cells in areas CA3 and CA1 of rat hippocampal slice cultures." Journal of neurophysiology **73**(3): 1282-1294.

Debanne, D., D. E. Shulz and Y. Fregnac (1995). "Temporal constraints in associative synaptic plasticity in hippocampus and neocortex." Canadian journal of physiology and pharmacology **73**(9): 1295-1311.

Deng, W., J. B. Aimone and F. H. Gage (2010). "New neurons and new memories: how does adult hippocampal neurogenesis affect learning and memory?" Nature reviews. Neuroscience **11**(5): 339-350.

Denk, W. and K. Svoboda (1997). "Photon upmanship: why multiphoton imaging is more than a gimmick." Neuron **18**(3): 351-357.

Diba, K. and G. Buzsaki (2007). "Forward and reverse hippocampal place-cell sequences during ripples." Nature neuroscience **10**(10): 1241-1242.

Dolorfo, C. L. and D. G. Amaral (1998). "Entorhinal cortex of the rat: organization of intrinsic connections." The Journal of comparative neurology **398**(1): 49-82.

Dombeck, D. A., C. D. Harvey, L. Tian, L. L. Looger and D. W. Tank (2010). "Functional imaging of hippocampal place cells at cellular resolution during virtual navigation." Nature neuroscience **13**(11): 1433-1440.

Dragoi, G. and S. Tonegawa (2011). "Preplay of future place cell sequences by hippocampal cellular assemblies." Nature **469**(7330): 397-401.

Dragoi, G. and S. Tonegawa (2013). "Distinct preplay of multiple novel spatial experiences in the rat." Proceedings of the National Academy of Sciences of the United States of America **110**(22): 9100-9105.

Dupret, D., J. O'Neill, B. Pleydell-Bouverie and J. Csicsvari (2010). "The reorganization and reactivation of hippocampal maps predict spatial memory performance." Nature neuroscience **13**(8): 995-1002.

Eichenbaum, H. (2000). "Hippocampus: mapping or memory?" Current biology : CB **10**(21): R785-787.

Eichenbaum, H., P. Dudchenko, E. Wood, M. Shapiro and H. Tanila (1999). "The hippocampus, memory, and place cells: is it spatial memory or a memory space?" Neuron **23**(2): 209-226.

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): 184-188.

Epsztein, J., A. K. Lee, E. Chorev and M. Brecht (2010). "Impact of spikelets on hippocampal CA1 pyramidal cell activity during spatial exploration." Science **327**(5964): 474-477.

Erdem, U. and M. Hasselmo (2012). "A goal-directed spatial navigation model using forward trajectory planning based on grid cells." The European journal of neuroscience **35**: 916 - 931.

Erdem, U. M. and M. E. Hasselmo (2014). "A biologically inspired hierarchical goal directed navigation model." J Physiol Paris **108**(1): 28-37.

Etienne, A. S., R. Maurer and V. Seguinot (1996). "Path integration in mammals and its interaction with visual landmarks." The Journal of experimental biology **199**(Pt 1): 201-209.

Feldmeyer, D. (2012). "Excitatory neuronal connectivity in the barrel cortex." Frontiers in neuroanatomy **6**: 24.

Feldmeyer, D., M. Brecht, F. Helmchen, C. C. Petersen, J. F. Poulet, J. F. Staiger, H. J. Luhmann and C. Schwarz (2013). "Barrel cortex function." Progress in neurobiology **103**: 3-27.

Feldmeyer, D., J. Lubke and B. Sakmann (2006). "Efficacy and connectivity of intracolumnar pairs of layer 2/3 pyramidal cells in the barrel cortex of juvenile rats." The Journal of physiology **575**(Pt 2): 583-602.

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 : the official journal of the Society for Neuroscience **28**(44): 11250-11262.

Fenton, A. A., W. W. Lytton, J. M. Barry, P. P. Lenck-Santini, L. E. Zinyuk, S. Kubik, J. Bures, B. Poucet, R. U. Muller and A. V. Olypher (2010). "Attention-like modulation of hippocampus place cell discharge." The Journal of neuroscience : the official journal of the Society for Neuroscience **30**(13): 4613-4625.

Fenton, A. A. and R. U. Muller (1998). "Place cell discharge is extremely variable during individual passes of the rat through the firing field." Proceedings of the National Academy of Sciences of the United States of America **95**(6): 3182-3187.

Ferbinteanu, J. and M. L. Shapiro (2003). "Prospective and retrospective memory coding in the hippocampus." Neuron **40**(6): 1227-1239.

Foster, D. J. and J. J. Knierim (2012). "Sequence learning and the role of the hippocampus in rodent navigation." Current opinion in neurobiology **22**(2): 294-300.

Foster, D. J. and M. A. Wilson (2006). "Reverse replay of behavioural sequences in hippocampal place cells during the awake state." Nature **440**(7084): 680-683.

Foster, D. J. and M. A. Wilson (2007). "Hippocampal theta sequences." Hippocampus **17**(11): 1093-1099.

Frank, L. M., E. N. Brown and M. Wilson (2000). "Trajectory encoding in the hippocampus and entorhinal cortex." Neuron **27**(1): 169-178.

Freemon, F. R., J. J. McNew and W. R. Adey (1969). "Sleep of unrestrained chimpanzee: cortical and subcortical recordings." Experimental neurology **25**(1): 129-137.

Freemon, F. R. and R. D. Walter (1970). "Electrical activity of human limbic system during sleep." Comprehensive psychiatry **11**(6): 544-551.

Freund, T. F. and G. Buzsaki (1996). "Interneurons of the hippocampus." Hippocampus **6**(4): 347-470.

Fuhs, M. C. and D. S. Touretzky (2006). "A spin glass model of path integration in rat medial entorhinal cortex." The Journal of neuroscience : the official journal of the Society for Neuroscience **26**(16): 4266-4276.

Fyhn, M., S. Molden, S. Hollup, M. B. Moser and E. Moser (2002). "Hippocampal neurons responding to first-time dislocation of a target object." Neuron **35**(3): 555-566.

Gaarskjaer, F. B. (1978). "Organization of the mossy fiber system of the rat studied in extended hippocampi. I. Terminal area related to number of granule and pyramidal cells." The Journal of comparative neurology **178**(1): 49-72.

Gaarskjaer, F. B. (1978). "Organization of the mossy fiber system of the rat studied in extended hippocampi. II. Experimental analysis of fiber distribution with silver impregnation methods." The Journal of comparative neurology **178**(1): 73-88.

Galbraith, J. A. and M. Terasaki (2003). "Controlled damage in thick specimens by multiphoton excitation." Molecular biology of the cell **14**(5): 1808-1817.

Gallagher, M. and M. T. Koh (2011). "Episodic memory on the path to Alzheimer's disease." Current opinion in neurobiology **21**(6): 929-934.

Georgopoulos, A. P., J. T. Lurito, M. Petrides, A. B. Schwartz and J. T. Massey (1989). "Mental rotation of the neuronal population vector." Science **243**(4888): 234-236.

Gilbert, P. E., R. P. Kesner and W. E. DeCoteau (1998). "Memory for spatial location: role of the hippocampus in mediating spatial pattern separation." The Journal of neuroscience : the official journal of the Society for Neuroscience **18**(2): 804-810.

Girardeau, G., K. Benchenane, S. I. Wiener, G. Buzsaki and M. B. Zugaro (2009). "Selective suppression of hippocampal ripples impairs spatial memory." Nature neuroscience **12**(10): 1222-1223.

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 : the official journal of the Society for Neuroscience **16**(24): 8027-8040.

Gothard, K. M., W. E. Skaggs, K. M. Moore and B. L. McNaughton (1996). "Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-

based navigation task." The Journal of neuroscience : the official journal of the Society for Neuroscience **16**(2): 823-835.

Grastyan, E., G. Karmos, L. Vereczkey and L. Kellenyi (1966). "The hippocampal electrical correlates of the homeostatic regulation of motivation." Electroencephalography and clinical neurophysiology **21**(1): 34-53.

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.

Griffin, A. L., H. Eichenbaum and M. E. Hasselmo (2007). "Spatial representations of hippocampal CA1 neurons are modulated by behavioral context in a hippocampus-dependent memory task." The Journal of neuroscience : the official journal of the Society for Neuroscience **27**(9): 2416-2423.

Gupta, A. S., M. A. van der Meer, D. S. Touretzky and A. D. Redish (2010). "Hippocampal replay is not a simple function of experience." Neuron **65**(5): 695-705.

Gupta, A. S., M. A. van der Meer, D. S. Touretzky and A. D. Redish (2012). "Segmentation of spatial experience by hippocampal theta sequences." Nature neuroscience **15**(7): 1032-1039.

Guzowski, J. F., T. Miyashita, M. K. Chawla, J. Sanderson, L. I. Maes, F. P. Houston, P. Lipa, B. L. McNaughton, P. F. Worley and C. A. Barnes (2006). "Recent behavioral history modifies coupling between cell activity and Arc gene transcription in hippocampal CA1 neurons." Proceedings of the National Academy of Sciences of the United States of America **103**(4): 1077-1082.

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): 801-806.

Hamam, B. N., T. E. Kennedy, A. Alonso and D. G. Amaral (2000). "Morphological and electrophysiological characteristics of layer V neurons of the rat medial entorhinal cortex." The Journal of comparative neurology **418**(4): 457-472.

Harris, E., M. P. Witter, G. Weinstein and M. Stewart (2001). "Intrinsic connectivity of the rat subiculum: I. Dendritic morphology and patterns of axonal arborization by pyramidal neurons." The Journal of comparative neurology **435**(4): 490-505.

Hartley, T., C. Lever, N. Burgess and J. O'Keefe (2014). "Space in the brain: how the hippocampal formation supports spatial cognition." Philosophical transactions of the Royal Society of London. Series B, Biological sciences **369**(1635): 20120510.

Hartley, T., E. A. Maguire, H. J. Spiers and N. Burgess (2003). "The well-worn route and the path less traveled: distinct neural bases of route following and wayfinding in humans." Neuron **37**(5): 877-888.

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): 941-946.

Hassabis, D., D. Kumaran and E. A. Maguire (2007). "Using imagination to understand the neural basis of episodic memory." The Journal of neuroscience : the official journal of the Society for Neuroscience **27**(52): 14365-14374.

Hassabis, D., D. Kumaran, S. D. Vann and E. A. Maguire (2007). "Patients with hippocampal amnesia cannot imagine new experiences." Proceedings of the National Academy of Sciences of the United States of America **104**(5): 1726-1731.

Hasselmo, M. E., C. Bodelon 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 H. Eichenbaum (2005). "Hippocampal mechanisms for the context-dependent retrieval of episodes." Neural Netw **18**(9): 1172-1190.

Hayes, J. A., X. Wang and C. A. Del Negro (2012). "Cumulative lesioning of respiratory interneurons disrupts and precludes motor rhythms in vitro." Proc Natl Acad Sci U S A **109**(21): 8286-8291.

Hayman, R. M., S. Chakraborty, M. I. Anderson and K. J. Jeffery (2003). "Context-specific acquisition of location discrimination by hippocampal place cells." The European journal of neuroscience **18**(10): 2825-2834.

Hebb, D. O. (1949). The Organisation of Behaviour: A Neuropsychological Theory. New York, Wiley.

Heisterkamp, A., I. Z. Maxwell, E. Mazur, J. M. Underwood, J. A. Nickerson, S. Kumar and D. E. Ingber (2005). "Pulse energy dependence of subcellular dissection by femtosecond laser pulses." Optics express **13**(10): 3690-3696.

Hill, A. J. (1978). "First occurrence of hippocampal spatial firing in a new environment." Experimental neurology **62**(2): 282-297.

Hok, V., P. P. Lenck-Santini, S. Roux, E. Save, R. U. Muller and B. Poucet (2007). "Goal-related activity in hippocampal place cells." The Journal of neuroscience : the official journal of the Society for Neuroscience **27**(3): 472-482.

Hollup, S. A., S. Molden, J. G. Donnett, M. B. Moser and E. I. Moser (2001). "Accumulation of hippocampal place fields at the goal location in an annular watermaze task." The Journal of neuroscience : the official journal of the Society for Neuroscience **21**(5): 1635-1644.

Hooks, B. M., S. A. Hires, Y. X. Zhang, D. Huber, L. Petreanu, K. Svoboda and G. M. Shepherd (2011). "Laminar analysis of excitatory local circuits in vibrissal motor and sensory cortical areas." PLoS biology **9**(1): e1000572.

Howard, L. R., Y. Yu, R. Mill, L. C. Morrison, R. Knight, M. Loftus, L. Staskute, L. Healy and H. J. Spiers (2011). Human hippocampus encodes Euclidean distance and future path to goals during real-world navigation. Society for Neuroscience. Washington, D.C. **Neuroscience Meeting Planner:** Program no. 288.219/UU283.

Howell, D. C. (2007). Statistical Methods for Psychology. United States, Thomson Wadsworth.

Huber, D., D. A. Gutnisky, S. Peron, D. H. O'Connor, J. S. Wiegert, L. Tian, T. G. Oertner, L. L. Looger and K. Svoboda (2012). "Multiple dynamic representations in the motor cortex during sensorimotor learning." Nature **484**(7395): 473-478.

Hutson, K. A. and R. B. Masterton (1986). "The sensory contribution of a single vibrissa's cortical barrel." Journal of neurophysiology **56**(4): 1196-1223.

Insausti, R., M. T. Herrero and M. P. Witter (1997). "Entorhinal cortex of the rat: cytoarchitectonic subdivisions and the origin and distribution of cortical efferents." Hippocampus **7**(2): 146-183.

Jeffery, K. J., J. G. Donnett, N. Burgess and J. M. O'Keefe (1997). "Directional control of hippocampal place fields." Experimental brain research **117**(1): 131-142.

Jeffery, K. J. and J. M. O'Keefe (1999). "Learned interaction of visual and idiothetic cues in the control of place field orientation." Experimental brain research **127**(2): 151-161.

Jensen, O., M. A. Idiart and J. E. Lisman (1996). "Physiologically realistic formation of autoassociative memory in networks with theta/gamma oscillations: role of fast NMDA channels." Learning & memory **3**(2-3): 243-256.

Jensen, O. and J. E. Lisman (1996). "Hippocampal CA3 region predicts memory sequences: accounting for the phase precession of place cells." Learning & memory **3**(2-3): 279-287.

Jezek, K., E. J. Henriksen, A. Treves, E. I. Moser and M. B. Moser (2011). "Theta-paced flickering between place-cell maps in the hippocampus." Nature **478**(7368): 246-249.

Ji, D. and M. A. Wilson (2007). "Coordinated memory replay in the visual cortex and hippocampus during sleep." Nature neuroscience **10**(1): 100-107.

Johnson, A., A. A. Fenton, C. Kentros and A. D. Redish (2009). "Looking for cognition in the structure within the noise." Trends in cognitive sciences **13**(2): 55-64.

Johnson, A. and A. D. Redish (2007). "Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point." The Journal of neuroscience : the official journal of the Society for Neuroscience **27**(45): 12176-12189.

Jouvet, M., F. Michel and J. Courjon (1959). "[Electric activity of the rhinencephalon during sleep in cats]." Comptes rendus des seances de la Societe de biologie et de ses filiales **153**(1): 101-105.

Kaitz, S. S. and R. T. Robertson (1981). "Thalamic connections with limbic cortex. II. Corticothalamic projections." The Journal of comparative neurology **195**(3): 527-545.

Karlsson, M. P. and L. M. Frank (2009). "Awake replay of remote experiences in the hippocampus." Nature neuroscience **12**(7): 913-918.

Kelsch, W., S. Sim and C. Lois (2010). "Watching synaptogenesis in the adult brain." Annual review of neuroscience **33**: 131-149.

Kentros, C. G., N. T. Agnihotri, S. Streater, R. D. Hawkins and E. R. Kandel (2004). "Increased attention to spatial context increases both place field stability and spatial memory." Neuron **42**(2): 283-295.

Kerr, J. N., C. P. de Kock, D. S. Greenberg, R. M. Bruno, B. Sakmann and F. Helmchen (2007). "Spatial organization of neuronal population responses in layer 2/3 of rat barrel cortex." The Journal of neuroscience : the official journal of the Society for Neuroscience **27**(48): 13316-13328.

Kerr, J. N., D. Greenberg and F. Helmchen (2005). "Imaging input and output of neocortical networks in vivo." Proc Natl Acad Sci U S A **102**(39): 14063-14068.

Kim, J. J., M. S. Fanselow, J. P. DeCola and J. Landeira-Fernandez (1992). "Selective impairment of long-term but not short-term conditional fear by the N-methyl-D-aspartate antagonist APV." Behavioral neuroscience **106**(4): 591-596.

Kleinfeld, D., R. W. Berg and S. M. O'Connor (1999). "Anatomical loops and their electrical dynamics in relation to whisking by rat." Somatosensory & motor research **16**(2): 69-88.

Kleinfeld, D. and M. Deschenes (2011). "Neuronal basis for object location in the vibrissa scanning sensorimotor system." Neuron **72**(3): 455-468.

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

Kohara, K., M. Pignatelli, A. J. Rivest, H. Y. Jung, T. Kitamura, J. Suh, D. Frank, K. Kajikawa, N. Mise, Y. Obata, I. R. Wickersham and S. Tonegawa (2014). "Cell type-specific genetic and optogenetic tools reveal hippocampal CA2 circuits." Nature neuroscience **17**(2): 269-279.

Kohler, C. (1985). "Intrinsic projections of the retrohippocampal region in the rat brain. I. The subicular complex." The Journal of comparative neurology **236**(4): 504-522.

Kolb, B. and I. Q. Whishaw (1996). Fundamentals of Human Neuropshychology. New York, W.H. Freeman and Company

Kraus, B. J., R. J. Robinson, 2nd, J. A. White, H. Eichenbaum and M. E. Hasselmo (2013). "Hippocampal "time cells": time versus path integration." Neuron **78**(6): 1090-1101.

Kwan, A. C. and Y. Dan (2012). "Dissection of cortical microcircuits by single-neuron stimulation in vivo." Current biology : CB **22**(16): 1459-1467.

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): 1576-1580.

Lee, A. K. and M. A. Wilson (2002). "Memory of sequential experience in the hippocampus during slow wave sleep." Neuron **36**(6): 1183-1194.

Lein, E. S., E. M. Callaway, T. D. Albright and F. H. Gage (2005). "Redefining the boundaries of the hippocampal CA2 subfield in the mouse using gene expression and 3-dimensional reconstruction." The Journal of comparative neurology **485**(1): 1-10.

Leinekugel, X., R. Khazipov, R. Cannon, H. Hirase, Y. Ben-Ari and G. Buzsaki (2002). "Correlated bursts of activity in the neonatal hippocampus in vivo." Science **296**(5575): 2049-2052.

Leung, L. W. and J. G. Borst (1987). "Electrical activity of the cingulate cortex. I. Generating mechanisms and relations to behavior." Brain research **407**(1): 68-80.

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 **309**(5734): 619-623.

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 : the official journal of the Society for Neuroscience **29**(31): 9771-9777.

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.

Lisman, J. E. (1999). "Relating hippocampal circuitry to function: recall of memory sequences by reciprocal dentate-CA3 interactions." Neuron **22**(2): 233-242.

Lisman, J. E. and M. A. Idiart (1995). "Storage of 7 +/- 2 short-term memories in oscillatory subcycles." Science **267**(5203): 1512-1515.

Liu, X., S. Ramirez, P. T. Pang, C. B. Puryear, A. Govindarajan, K. Deisseroth and S. Tonegawa (2012). "Optogenetic stimulation of a hippocampal engram activates fear memory recall." Nature **484**(7394): 381-385.

Long, L. L., J. R. Hinman, C. M. Chen, M. A. Escabi and J. J. Chrobak (2014). "Theta Dynamics in Rat: Speed and Acceleration across the Septotemporal Axis." PLoS One **9**(5): e97987.

Losonczy, A., B. V. Zemelman, A. Vaziri and J. C. Magee (2010). "Network mechanisms of theta related neuronal activity in hippocampal CA1 pyramidal neurons." Nature neuroscience **13**(8): 967-972.

Lovett-Barron, M., P. Kaifosh, M. A. Kheirbek, N. Danielson, J. D. Zaremba, T. R. Reardon, G. F. Turi, R. Hen, B. V. Zemelman and A. Losonczy (2014). "Dendritic inhibition in the hippocampus supports fear learning." Science **343**(6173): 857-863.

Luo, A. H., P. Tahsili-Fahadan, R. A. Wise, C. R. Lupica and G. Aston-Jones (2011). "Linking context with reward: a functional circuit from hippocampal CA3 to ventral tegmental area." Science **333**(6040): 353-357.

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): 143-152.

MacDonald, C. J., K. Q. Lepage, U. T. Eden and H. Eichenbaum (2011). "Hippocampal "time cells" bridge the gap in memory for discontiguous events." Neuron **71**(4): 737-749.

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): 921-924.

Maguire, E. A. and D. Hassabis (2011). "Role of the hippocampus in imagination and future thinking." Proceedings of the National Academy of Sciences of the United States of America **108**(11): E39.

Maguire, E. A., R. Nannery and H. J. Spiers (2006). "Navigation around London by a taxi driver with bilateral hippocampal lesions." Brain : a journal of neurology **129**(Pt 11): 2894-2907.

Markus, E. J., Y. L. Qin, B. Leonard, W. E. Skaggs, B. L. McNaughton and C. A. Barnes (1995). "Interactions between location and task affect the spatial and directional firing of hippocampal neurons." The Journal of neuroscience : the official journal of the Society for Neuroscience **15**(11): 7079-7094.

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.

McNaughton, B. L., C. A. Barnes, J. L. Gerrard, K. Gothard, M. W. Jung, J. J. Knierim, H. Kudrimoti, Y. Qin, W. E. Skaggs, M. Suster and K. L. Weaver (1996). "Deciphering the hippocampal polyglot: the hippocampus as a path integration system." The Journal of experimental biology **199**(Pt 1): 173-185.

McNaughton, B. L., C. A. Barnes and J. O'Keefe (1983). "The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats." Experimental brain research **52**(1): 41-49.

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): 663-678.

Mitchell, S. J. and J. B. Ranck, Jr. (1980). "Generation of theta rhythm in medial entorhinal cortex of freely moving rats." Brain research **189**(1): 49-66.

Mittelstaedt, M. L. and H. Mittelstaedt (1980). "Homing by path integration in a mammal." Naturwissenschaften **67**: 566–567.

Moita, M. A., S. Rosis, Y. Zhou, J. E. LeDoux and H. T. Blair (2004). "Putting fear in its place: remapping of hippocampal place cells during fear conditioning." The Journal of neuroscience : the official journal of the Society for Neuroscience **24**(31): 7015-7023.

Morris, R. G., P. Garrud, J. N. Rawlins and J. O'Keefe (1982). "Place navigation impaired in rats with hippocampal lesions." Nature **297**(5868): 681-683.

Mosko, S., G. Lynch and C. W. Cotman (1973). "The distribution of septal projections to the hippocampus of the rat." The Journal of comparative neurology **152**(2): 163-174.

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 : the official journal of the Society for Neuroscience **7**(7): 1951-1968.

Muller, R. U., J. L. Kubie and J. B. Ranck, Jr. (1987). "Spatial firing patterns of hippocampal complex-spike cells in a fixed environment." The Journal of neuroscience : the official journal of the Society for Neuroscience **7**(7): 1935-1950.

Nadasdy, Z., H. Hirase, A. Czurko, J. Csicsvari and G. Buzsaki (1999). "Replay and time compression of recurring spike sequences in the hippocampus." The Journal of neuroscience : the official journal of the Society for Neuroscience **19**(21): 9497-9507.

Nakai, J., M. Ohkura and K. Imoto (2001). "A high signal-to-noise Ca(2+) probe composed of a single green fluorescent protein." Nat Biotechnol **19**(2): 137-141.

O'Connor, D. H., N. G. Clack, D. Huber, T. Komiyama, E. W. Myers and K. Svoboda (2010). "Vibrissa-based object localization in head-fixed mice." The Journal of

neuroscience : the official journal of the Society for Neuroscience **30**(5): 1947-1967.

O'Connor, D. H., S. P. Peron, D. Huber and K. Svoboda (2010). "Neural activity in barrel cortex underlying vibrissa-based object localization in mice." Neuron **67**(6): 1048-1061.

O'Keefe, J. (1976). "Place units in the hippocampus of the freely moving rat." Experimental neurology **51**(1): 78-109.

O'Keefe, J. (1979). "A review of the hippocampal place cells." Progress in neurobiology **13**(4): 419-439.

O'Keefe, J. and N. Burgess (1996). "Geometric determinants of the place fields of hippocampal neurons." Nature **381**(6581): 425-428.

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): 573-590.

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): 171-175.

O'Keefe, J. and L. Nadel (1978). The Hippocampus as a Cognitive Map, Claredon, Oxford.

O'Keefe, J. and L. Nadel (1979). "Précis of O'Keefe & Nadel's The hippocampus as a cognitive map." Behavioral and Brain Sciences **2**(04): 487-494.

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

O'Keefe, J. and A. Speakman (1987). "Single unit activity in the rat hippocampus during a spatial memory task." Experimental brain research **68**(1): 1-27.

Olton, D. S., M. Branch and P. J. Best (1978). "Spatial correlates of hippocampal unit activity." Experimental neurology **58**(3): 387-409.

Pare, D. and D. R. Collins (2000). "Neuronal correlates of fear in the lateral amygdala: multiple extracellular recordings in conscious cats." The Journal of neuroscience : the official journal of the Society for Neuroscience **20**(7): 2701-2710.

Peron, S., V. Iyer, Z. Guo, T.-S. Chen, D. Kim, D. Huber and K. Svoboda (2012). Towards imaging complete representations of whisker touch in the mouse barrel cortex. Society for Neuroscience**:** KK12/677.612.

Petersen, C. C. and S. Crochet (2013). "Synaptic computation and sensory processing in neocortical layer 2/3." Neuron **78**(1): 28-48.

Petreanu, L., D. A. Gutnisky, D. Huber, N. L. Xu, D. H. O'Connor, L. Tian, L. Looger and K. Svoboda (2012). "Activity in motor-sensory projections reveals distributed coding in somatosensation." Nature **489**(7415): 299-303.

Petsche, H. and C. Stumpf (1962). "[The origin of theta-rhytm in the rabbit hippocampus]." Wiener klinische Wochenschrift **74**: 696-700.

Pfeiffer, B. E. and D. J. Foster (2013). "Hippocampal place-cell sequences depict future paths to remembered goals." Nature **497**(7447): 74-79.

Pickel, V. M., M. Segal and F. E. Bloom (1974). "A radioautographic study of the efferent pathways of the nucleus locus coeruleus." The Journal of comparative neurology **155**(1): 15-42.

Pickenhain, L. and F. Klingberg (1967). "Changes of electrophysiological and behavioural indicators during the delay period of conditioned reflexes in rats." Activitas nervosa superior **9**(3): 330.

Pikkarainen, M., S. Ronkko, V. Savander, R. Insausti and A. Pitkanen (1999). "Projections from the lateral, basal, and accessory basal nuclei of the amygdala to the hippocampal formation in rat." The Journal of comparative neurology **403**(2): 229-260.

Pitkanen, A., M. Pikkarainen, N. Nurminen and A. Ylinen (2000). "Reciprocal connections between the amygdala and the hippocampal formation, perirhinal cortex, and postrhinal cortex in rat. A review." Annals of the New York Academy of Sciences **911**: 369-391.

Powell, T. P., R. W. Guillery and W. M. Cowan (1957). "A quantitative study of the fornixmamillo-thalamic system." Journal of anatomy **91**(4): 419-437.

Ramon y Cajal, S. (1893). "Estructura del asta de Ammon y fascia dentata. ." Ann Soc Esp Hist Nat(22).

Ranck, J. B., Jr. (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.

Rapp, P. R. and M. Gallagher (1996). "Preserved neuron number in the hippocampus of aged rats with spatial learning deficits." Proceedings of the National Academy of Sciences of the United States of America **93**(18): 9926-9930.

Reijmers, L. G., B. L. Perkins, N. Matsuo and M. Mayford (2007). "Localization of a stable neural correlate of associative memory." Science **317**(5842): 1230-1233.

Ribak, C. E., J. E. Vaughn and K. Saito (1978). "Immunocytochemical localization of glutamic acid decarboxylase in neuronal somata following colchicine inhibition of axonal transport." Brain research **140**(2): 315-332.

Rickgauer, J. P. and D. W. Tank (2013). Optical imaging and stimulation of neural activity at cellular resolution in awake, mobile

mice Society for Neuroscience. San Diego**:** 871.807/MMM873-DP810.

Rizzolatti, G., L. Fadiga, V. Gallese and L. Fogassi (1996). "Premotor cortex and the recognition of motor actions." Brain research. Cognitive brain research **3**(2): 131-141.

Robertson, R. T. and S. S. Kaitz (1981). "Thalamic connections with limbic cortex. I. Thalamocortical projections." The Journal of comparative neurology **195**(3): 501-525.

Rowland, D. C., Y. Yanovich and C. G. Kentros (2011). "A stable hippocampal representation of a space requires its direct experience." Proceedings of the National Academy of Sciences of the United States of America **108**(35): 14654-14658.

Sainsbury, R. S. (1998). "Hippocampal theta: a sensory-inhibition theory of function." Neuroscience and biobehavioral reviews **22**(2): 237-241.

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): 758-762.

Schacter, D. L., D. R. Addis, D. Hassabis, V. C. Martin, R. N. Spreng and K. K. Szpunar (2012). "The future of memory: remembering, imagining, and the brain." Neuron **76**(4): 677-694.

Schultz, W. (1998). "Predictive reward signal of dopamine neurons." Journal of neurophysiology **80**(1): 1-27.

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.

Seress, L. and C. E. Ribak (1983). "GABAergic cells in the dentate gyrus appear to be local circuit and projection neurons." Experimental brain research **50**(2-3): 173-182.

Sherrill, K. R., U. M. Erdem, R. S. Ross, T. I. Brown, M. E. Hasselmo and C. E. Stern (2013). "Hippocampus and retrosplenial cortex combine path integration signals for successful navigation." J Neurosci **33**(49): 19304-19313.

Shipley, M. T. (1975). "The topographical and laminar organization of the presubiculum's projection to the ipsi- and contralateral entorhinal cortex in the guinea pig." The Journal of comparative neurology **160**(1): 127-145.

Simons, D. J. and T. A. Woolsey (1978). "Response properties of vibrissa units in rat SI somatosensory neocortex." Journal of neurophysiology **41**: 798-820.

Singer, A. C., M. F. Carr, M. P. Karlsson and L. M. Frank (2013). "Hippocampal SWR activity predicts correct decisions during the initial learning of an alternation task." Neuron **77**(6): 1163-1173.

Singer, A. C. and L. M. Frank (2009). "Rewarded outcomes enhance reactivation of experience in the hippocampus." Neuron **64**(6): 910-921.

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): 1865-1868.

Solstad, T., E. I. Moser and G. T. Einevoll (2006). "From grid cells to place cells: a mathematical model." Hippocampus **16**(12): 1026-1031.

Soriano, E. and M. Frotscher (1989). "A GABAergic axo-axonic cell in the fascia dentata controls the main excitatory hippocampal pathway." Brain research **503**(1): 170-174.

Spiers, H. J., N. Burgess, T. Hartley, F. Vargha-Khadem and J. O'Keefe (2001). "Bilateral hippocampal pathology impairs topographical and episodic memory but not visual pattern matching." Hippocampus **11**(6): 715-725.

Spiers, H. J., N. Burgess, E. A. Maguire, S. A. Baxendale, T. Hartley, P. J. Thompson and J. O'Keefe (2001). "Unilateral temporal lobectomy patients show lateralized topographical and episodic memory deficits in a virtual town." Brain : a journal of neurology **124**(Pt 12): 2476-2489.

Spiers, H. J., R. M. Hayman, A. Jovalekic, E. Marozzi and K. J. Jeffery (2013). "Place Field Repetition and Purely Local Remapping in a Multicompartment Environment." Cerebral cortex.

Spiers, H. J. and E. A. Maguire (2006). "Thoughts, behaviour, and brain dynamics during navigation in the real world." NeuroImage **31**(4): 1826-1840.

Spiers, H. J. and E. A. Maguire (2007). "A navigational guidance system in the human brain." Hippocampus **17**(8): 618-626.

Spiers, H. J., E. A. Maguire and N. Burgess (2001). "Hippocampal amnesia." Neurocase **7**(5): 357-382.

Stackman, R. W., A. S. Clark and J. S. Taube (2002). "Hippocampal spatial representations require vestibular input." Hippocampus **12**(3): 291-303.

Stuber, G. D., D. R. Sparta, A. M. Stamatakis, W. A. van Leeuwen, J. E. Hardjoprajitno, S. Cho, K. M. Tye, K. A. Kempadoo, F. Zhang, K. Deisseroth and A. Bonci (2011). "Excitatory transmission from the amygdala to nucleus accumbens facilitates reward seeking." Nature **475**(7356): 377-380.

Sugar, J., M. P. Witter, N. M. van Strien and N. L. Cappaert (2011). "The retrosplenial cortex: intrinsic connectivity and connections with the (para)hippocampal region in the rat. An interactive connectome." Frontiers in neuroinformatics **5**: 7.

Suzuki, S. S. and G. K. Smith (1985). "Single-cell activity and synchronous bursting in the rat hippocampus during waking behavior and sleep." Experimental neurology **89**(1): 71-89.

Swanson, L. W. and W. M. Cowan (1975). "Hippocampo-hypothalamic connections: origin in subicular cortex, not ammon's horn." Science **189**(4199): 303-304.

Swanson, L. W. and B. K. Hartman (1975). "The central adrenergic system. An immunofluorescence study of the location of cell bodies and their efferent

connections in the rat utilizing dopamine-beta-hydroxylase as a marker." The Journal of comparative neurology **163**(4): 467-505.

Swanson, L. W. and C. Kohler (1986). "Anatomical evidence for direct projections from the entorhinal area to the entire cortical mantle in the rat." Journal of Neuroscience **6**: 3010-3023.

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.

Tank, D. W., M. Sugimori, J. A. Connor and R. R. Llinas (1988). "Spatially resolved calcium dynamics of mammalian Purkinje cells in cerebellar slice." Science **242**(4879): 773-777.

Taube, J. S. (1995). "Place cells recorded in the parasubiculum of freely moving rats." Hippocampus **5**(6): 569-583.

Taube, J. S. (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, Jr. (1990). "Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis." The Journal of neuroscience : the official journal of the Society for Neuroscience **10**(2): 420-435.

Taube, J. S., R. U. Muller and J. B. Ranck, Jr. (1990). "Head-direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations." The Journal of neuroscience : the official journal of the Society for Neuroscience **10**(2): 436-447.

Thompson, S. M. and R. T. Robertson (1987). "Organization of subcortical pathways for sensory projections to the limbic cortex. I. Subcortical projections to the medial limbic cortex in the rat." The Journal of comparative neurology **265**(2): 175-188.

Tian, L., S. A. Hires, T. Mao, D. Huber, M. E. Chiappe, S. H. Chalasani, L. Petreanu, J. Akerboom, S. A. McKinney, E. R. Schreiter, C. I. Bargmann, V. Jayaraman, K. Svoboda and L. L. Looger (2009). "Imaging neural activity in worms, flies and mice with improved GCaMP calcium indicators." Nat Methods **6**(12): 875-881.

Tolman, E. C. (1948). "Cognitive maps in rats and men." Psychological review **55**(4): 189-208.

Tsai, H. C., F. Zhang, A. Adamantidis, G. D. Stuber, A. Bonci, L. de Lecea and K. Deisseroth (2009). "Phasic firing in dopaminergic neurons is sufficient for behavioral conditioning." Science **324**(5930): 1080-1084.

Tse, D., R. F. Langston, M. Kakeyama, I. Bethus, P. A. Spooner, E. R. Wood, M. P. Witter and R. G. Morris (2007). "Schemas and memory consolidation." Science **316**(5821): 76-82.

Tsodyks, M. V., W. E. Skaggs, T. J. Sejnowski and B. L. McNaughton (1996). "Population dynamics and theta rhythm phase precession of hippocampal place cell firing: a spiking neuron model." Hippocampus **6**(3): 271-280.

Van der Werf, Y. D., M. P. Witter and H. J. Groenewegen (2002). "The intralaminar and midline nuclei of the thalamus. Anatomical and functional evidence for participation in processes of arousal and awareness." Brain research. Brain research reviews **39**(2-3): 107-140.

van Groen, T. and J. M. Wyss (1992). "Connections of the retrosplenial dysgranular cortex in the rat." The Journal of comparative neurology **315**(2): 200-216.

van Haeften, T., L. Baks-te-Bulte, P. H. Goede, F. G. Wouterlood and M. P. Witter (2003). "Morphological and numerical analysis of synaptic interactions between neurons in deep and superficial layers of the entorhinal cortex of the rat." Hippocampus **13**(8): 943-952.

Vanderwolf, C. H. (1969). "Hippocampal electrical activity and voluntary movement in the rat." Electroencephalography and clinical neurophysiology **26**(4): 407-418.

Verwer, R. W., R. J. Meijer, H. F. Van Uum and M. P. Witter (1997). "Collateral projections from the rat hippocampal formation to the lateral and medial prefrontal cortex." Hippocampus **7**(4): 397-402.

Vogelstein, J. T., A. M. Packer, T. A. Machado, T. Sippy, B. Babadi, R. Yuste and L. Paninski (2010). "Fast nonnegative deconvolution for spike train inference from population calcium imaging." J Neurophysiol **104**(6): 3691-3704.

Wallace, D. G., D. J. Hines, S. M. Pellis and I. Q. Whishaw (2002). "Vestibular information is required for dead reckoning in the rat." The Journal of neuroscience : the official journal of the Society for Neuroscience **22**(22): 10009-10017.

Wang, M. E., E. G. Wann, R. K. Yuan, M. M. Ramos Alvarez, S. M. Stead and I. A. Muzzio (2012). "Long-term stabilization of place cell remapping produced by a fearful experience." The Journal of neuroscience : the official journal of the Society for Neuroscience **32**(45): 15802-15814.

West, M. J., L. Slomianka and H. J. Gundersen (1991). "Unbiased stereological estimation of the total number of neurons in thesubdivisions of the rat hippocampus using the optical fractionator." The Anatomical record **231**(4): 482-497.

Whishaw, I. Q. and H. Maaswinkel (1998). "Rats with fimbria-fornix lesions are impaired in path integration: a role for the hippocampus in "sense of direction"." The Journal of neuroscience : the official journal of the Society for Neuroscience **18**(8): 3050-3058.

Whishaw, I. Q. and C. H. Vanderwolf (1973). "Hippocampal EEG and behavior: changes in amplitude and frequency of RSA (theta rhythm) associated with spontaneous and learned movement patterns in rats and cats." Behavioral biology **8**(4): 461-484.

Wilhelm, I., S. Diekelmann, I. Molzow, A. Ayoub, M. Molle and J. Born (2011). "Sleep selectively enhances memory expected to be of future relevance." The Journal of neuroscience : the official journal of the Society for Neuroscience **31**(5): 1563-1569.

Wills, T. J., F. Cacucci, N. Burgess and J. O'Keefe (2010). "Development of the hippocampal cognitive map in preweanling rats." Science **328**(5985): 1573-1576.

Wills, T. J., C. Lever, F. Cacucci, N. Burgess and J. O'Keefe (2005). "Attractor dynamics in the hippocampal representation of the local environment." Science **308**(5723): 873-876.

Wilson, M. A. and B. L. McNaughton (1993). "Dynamics of the hippocampal ensemble code for space." Science **261**(5124): 1055-1058.

Wilson, M. A. and B. L. McNaughton (1994). "Reactivation of hippocampal ensemble memories during sleep." Science **265**(5172): 676-679.

Winner, B., Z. Kohl and F. H. Gage (2011). "Neurodegenerative disease and adult neurogenesis." The European journal of neuroscience **33**(6): 1139-1151.

Wood, E. R., P. A. Dudchenko, R. J. Robitsek and H. Eichenbaum (2000). "Hippocampal neurons encode information about different types of memory episodes occurring in the same location." Neuron **27**(3): 623-633.

Wyble, B. P., C. Linster and M. E. Hasselmo (2000). "Size of CA1-evoked synaptic potentials is related to theta rhythm phase in rat hippocampus." Journal of neurophysiology **83**(4): 2138-2144.

Wyss, J. M. and T. Van Groen (1992). "Connections between the retrosplenial cortex and the hippocampal formation in the rat: a review." Hippocampus **2**(1): 1-11.

Yanik, M. F., H. Cinar, H. N. Cinar, A. D. Chisholm, Y. Jin and A. Ben-Yakar (2004). "Neurosurgery: functional regeneration after laser axotomy." Nature **432**(7019): 822.

Yartsev, M. M. and N. Ulanovsky (2013). "Representation of three-dimensional space in the hippocampus of flying bats." Science **340**(6130): 367-372.

Zariwala, H. A., B. G. Borghuis, T. M. Hoogland, L. Madisen, L. Tian, C. I. De Zeeuw, H. Zeng, L. L. Looger, K. Svoboda and T. W. Chen (2012). "A Cre-dependent GCaMP3 reporter mouse for neuronal imaging in vivo." J Neurosci **32**(9): 3131-3141.

Zipser, D. (1986). "A model of hippocampal learning during classical conditioning." Behavioral neuroscience **100**(5): 764-776.