Rebecca Nutbrown

Thesis submitted for the degree of

Doctor of Philosophy

in

Neuroscience

Wolfson Institute for Biomedical Research

University College London

2021

Summary of contributions

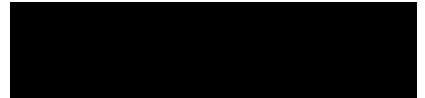Throughout thesis: behavioural training was carried out by Theshika Jeyaratnam and Beckie Nutbrown.

Chapter 3: surgeries, and experiments were carried out by Beckie Nutbrown. Software was developed by Lloyd Russell. Analysis was conduced by Lloyd Russell and Beckie Nutbrown.

Chapter 4: Software was developed by Lloyd Russell and Beckie Nutbrown

Chapters 5: surgeries were carried out by Beckie Nutbrown and Theshika Jeyaratnam. Experiments were carried out by Beckie Nutbrown. Software was written by Lloyd Russell and analyses were conducted by Lloyd Russell and Beckie Nutbrown.

Declaration

'I, Rebecca Nutbrown 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.'

Rebecca Nutbrown

June 30th, 2021

Two photon interrogation of hippocampal subregions CA1 and CA3 during spatial behaviour

Abstract

The hippocampus is crucial for spatial navigation and episodic memory formation. Hippocampal place cells exhibit spatially selective activity within an environment and form the neural basis of a cognitive map of space which supports these mnemonic functions. Hebb's (1949) postulate regarding the creation of cell assemblies is seen as the pre-eminent model of learning in neural systems. Investigating changes to the hippocampal representation of space during an animal's exploration of its environment provides an opportunity to observe Hebbian learning at the population and single cell level. When exploring new environments animals form spatial memories that are updated with experience and retrieved upon re-exposure to the same environment, but how this is achieved by different subnetworks in hippocampal CA1 and CA3, and how these circuits encode distinct memories of similar objects and events remains unclear. To test these ideas, we developed an experimental strategy and detailed protocols for simultaneously recording from CA1 and CA3 populations with 2P imaging. We also developed a novel all-optical protocol to simultaneously activate and record from ensembles of CA3 neurons. We used these approaches to show that targeted activation of CA3 neurons results in an increasing excitatory amplification seen only in CA3 cells when stimulating other CA3 cells, and not in CA1, perhaps reflecting the greater number of recurrent connections in CA3. To probe hippocampal spatial representations, we titrated input to the network by morphing VR environments during spatial navigation to assess the local CA3 as well as downstream CA1 responses. To this end, we found CA1 and CA3 neural population responses behave nonlinearly, consistent with attractor dynamics associated with the two stored representations. We interpret our findings as supporting classic theories of Hebbian learning and as the beginning of uncovering the relationship between hippocampal neural circuit activity and the computations implemented by their dynamics. Establishing this relationship is paramount to demystifying the neural underpinnings of cognition.

Publications

Work and skills developed during this thesis have contributed to the following publications:

Robinson, T. M., Descamps, A. L., Russell, L. E,. Buchholz, O. M., Antonov, K. G., Lau, J., **Nutbrown, R.,** Schmidt-Hieber, C., Häusser, M. (2020). 'Targeted activation of hippocampal place cells drives memory-guided spatial behaviour. *Cell* 183(7):2041-2042.

Russell, L. E., Dalgleish, H. W. P., **Nutbrown, R.,** Gauld, O. G., Herrmann, D., Fisek., M., Packer., A., & Häusser, M. (2021, in prep). 'All optical interrogation of neural circuits during behaviour'. *Nature Protocols.*

Acknowledgements

Impact Statement

How do we remember? Why do we have problems accessing memories in conditions like dementia? How can we teach machines to learn? These are some of the most important questions of our time and have been the focus of intense study since the dawn of experimental cognitive science. By understanding how memory functions under normal conditions, we can better understand the mechanisms for dysfunction in disease. Through identifying the rules by which information is processed, stored, and recalled, we increase the possibility of harnessing this knowledge to solve memory access problems, and optimize learning. A distant possibility will be the augmentation of normal human abilities and restoration of lost function with synthetic assistance. Artificial intelligence (AI) and neuroscience research are synergistic, since understanding our own brain improves our ability to build better artificial ones, and vice versa. Hence the intersection between theory and experiments holds the key to unlocking a more complete understanding of how human brains learn and remember.

To be able to achieve such advances, we need tools to read and write specific and precise patterns of activity into real biological neural circuits underpinning memory processing. To do this we can leverage cutting edge tools in optical biology, to investigate the hippocampus, a brain area involved in memory processing. We can build on the 2014 Nobel prize winning discovery of "place cells", neurons whose activity patterns contribute to our brain's positional system - our inner GPS. These cells provide distinct patterns of neural activity with a behavioural correlate and therefore make spatial navigation a tractable entry point into understanding cognitive functions such as memory. For the first time we can read neural activity and 'know' what the animal is 'thinking' about where it is in space.

Hippocampal subregion CA3 has been proposed to underpin memory recall (remembering), given its biological connectivity architecture resembles that of artificial network models with memory recall functions. Our new optical tools, applied to interrogating place cells, offers the opportunity to re-ignite collaborations between protagonists of artificial and biological intelligence. As such,

we are now in a position to probe these brain regions and test long standing theories of their memory functions. In this thesis I describe a novel surgical method to optically access CA1 and CA3 simultaneously and developed a miniaturized 2D virtual reality system to precisely control the virtual environment experienced by behaving animals on our all-optical set up. I then applied these developments to probe functional properties in different two hippocampal subregions simultaneously during spatial navigation. Together with recent work in the field, my work spearheads progress towards spatial memory becoming the first higher order cognitive function to be understood in mechanistic detail.

Abbreviations

2P – Two-photon

3D - Three dimensional

AP – Action potential

AMPAR - Alpha-Amino-3-Hydroxy-5-Methyl-

4 -Isoxazole Propionic Acid Receptor

A-P - Anterior-posterior

C – Caudal

C1V1 – Two-photon activatable opsin

CA (1-4) - Cornu ammonis 1-4

D - Dorsal

ΔF/F – Measure of fluorescence change

DG - Dentate gyrus

FOV – Field of view

Galvo – Galvanometer mirror

NAPARM – Near-Automatic Photoactivation

Response Mapping

NMDAR - N-methyl-D-aspartate receptors

M - Medial

PaS - parasubiculum

PMT – Photon multiplier tube

POR - Postrhinal cortex

PrS - Presubiculum

PV – Parvalbumin positive interneuron

ROI(s) – region(s) of interest

S – Stimulus

SLM – Spatial light modulator

SNR – Signal-to-noise ratio

ST – Soma targeted

Sub – Subiculum

ERC - Entorhinal cortex

EPSP - Excitatory postsynaptic potential

HD Head direction

HF Hippocampal formation

L - Lateral

Le – Left

LFP - Local field potential

LTD - Long term depression

LTP - Long term potentiation

MEC - Medial entorhinal cortex

V - Ventral

VR - Virtual reality

R – Rostral

Ri – Right

SL - Stratum lucidum

CA3 pyr - CA3 pyramidal cell layer

IML - Inner molecular layer

DG-IB - DG inner blade

DG-OB - Outer blade dentate gyrus

SO - Stratum oriens

CA1 pyr - CA1 pyramidal cell layer

SR - the stratum radiatum

Table of Contents

List of Figures

15

List of Tables

17

1

Introduction

1.1 The cognitive revolution

1.1.1 Bringing matters of the mind into the laboratory

Psychologists have long been fascinated by spatial memory, but it was not until the middle of last century when a series of findings propelled its study into the domain of neuroscience and brought the cognitive revolution to psychology (Jeffery, 2018). In the 1800's Herman Ebbinghaus pioneered the scientific study higher mental functions. He was the first to systematically study human learning and memory and determine early principles of memory encoding and recall (Ebbinghaus, 1885). The work of Ebbinghaus created the foundation for the quantitative study of memory in the laboratory.

In parallel with this scientific study of human learning and memory, a new generation of experimental psychologists described many fundamental relationships between the environment and behaviour underpinning associative learning in animals. These behaviourists including Pavlov, Watson, Thorndike and Skinner - generally posit that complex behaviours are elicited by a cue functioning as a stimulus (S), that becomes progressively associated with a precise behaviour, the response (R), through obtaining rewards. Such associative explanations of learning and memory suggest spatial behaviour, i.e. navigation, is generated as a chain of rewarded responses to environmental stimuli, selected in a trial and error learning process (Thorndike, 1898).

While behaviourism revolutionized psychological schools of thought, quantifying the observable and providing a logical explanation of behaviour, it avoided physiological descriptions given the intangible nature of neural activity at that time (Jeffery, 2018). Bridging memory theory, and a mechanistic explanation of this cognitive function, has not been technologically possible until recent years. In this section I introduce the idea of a cognitive map and research on the neurobiological substrate underpinning it. These key concepts have inspired our current

understanding of learning and memory in modern systems neuroscience and laid the foundations for memory to be the first higher-level cognitive function understood at the mechanistic level.

1.1.2 Tolman's cognitive map: flexible organization of knowledge

While powerful in systematically quantifying the seemingly unquantifiable, behaviourist theory fails to explain the capacity of animals to perform flexible behaviour. The ability for animals to update choices rapidly and the influence of motivational factors clashes with the fixed S-R chain of events posited by the behaviourist paradigm. In his seminal work Tolman used a 'sunburst maze' to demonstrate that rats would preferentially choose the shortest path to the goal despite having never explored that path before (Tolman, 1948). Tolman concluded that navigation behaviour could be flexible. Tolman also noted that the information represented during a spatial experience can be retrieved in a subsequent exposure to the same environment; a process he thought to be dependent on recalling a previous memory, and part of a broad learning mechanism he defined as 'latent learning' (Tolman, 1948). Together, this classic series of behavioural experiments cemented the possibility of an internally stored representation of space that was allocentric or world-centered. The idea that animals possess a 'cognitive map', integrating multiple cues, that can be used flexibly, allowing adaptive navigation, subsequently became the foundational theory of spatial memory (O'Keefe & Nadel, 1978). Tolman's cognitive map has been invaluable in enabling us to develop a framework to understand later seminal discoveries behind how animals navigate in the world (O'Keefe & Nadel, 1978).

Tolman's cognitive map was based on observation of an animal's actions and makes no provision as to where such a complex system could exist in the brain. However, other theorists and experimentalists were beginning to make a deeper connection between the brain and the underlying architecture and computations driving outward displays of behaviour. Broadly, Karl Lashley showed that lesions in the cerebral cortex had predictable effects on behaviour in animals (Lashley, 1929, 1950). On a different level of understanding, Donald Hebb conceptualised ideas that account for complex brain functions at the level of neural circuits (Hebb 1949). Hebb's notion of the cell assembly, the idea that information propagates through groups of neurons, connected

together (Hebb, 1949), propelled the investigation of behaviour from psychology into the domain of neuroscience. A concept which we will come back too later in this chapter. Both Lashley and Hebb insisted that studies of the brain were critical to understand behaviour, this now sounds obvious, but was originally a debatable question (Jeffrey, 2018). Lashley and Hebb both searched for a specific locus in the brain where information is stored but could not find one- instead, tentatively suggesting memory might be a more distributed system.

1.2 The locus of memory processing

1.2.1 Memory deficits and hippocampal damage

In the following years, the search for such a brain structure largely responsible for memory processing continued. A pivotal turning point occurred when Scoville and Milner (1957) reported severe loss of memory in the tragic case of Henry Gustav Molaison, an epileptic patient who became known as patient H.M. H.M. suffered bilateral damage to two thirds of his anterior hippocampi during surgery to cure his epilepsy (Scoville & Milner, 1957). While partially successful, the surgery resulted in profound anterograde amnesia (an inability to form new memories; Scoville & Milner, 1957). H.M. could no longer recall recent memories, yet his long-term memory and intelligence remained intact (Scoville & Milner, 1957). Other amnesiac patients with hippocampal lesions have shown a similar pattern of deficits; while their intelligence was spared, they become unable to encode episodic memories and their recollections of past events or spatial sceneries were schematic, and lacking contextual details (Rosenbaum et al., 2000). For example, patients with bilateral damage to hippocampal subregion CA3 had problems retrieving episodic memories (Miller et al., 2020). In addition to loss of function studies, Maguire and colleagues have shown London taxi drivers (who navigate all day) have an enlarged posterior hippocampus (Maguire et al., 2000). Collectively, these studies highlight the different types of memory function, and the crucial role the hippocampus plays in supporting declarative memory, localising the search for a neural substrate to the hippocampal formation, and inspiring a new era of hippocampal physiology (Eichenbaum, 2013).

21

Conscious or 'declarative' memories can be classified as semantic memories comprising of factual knowledge about the world (for example, Kensington is a borough in London), and episodic memories, that depict unique experiences (for example, your first time on the tube in London). Episodic memories associate items and events with the spatial and temporal context in which they were experienced (Eichenbaum, 2017). Memories must be first encoded as a permanent trace or 'engram', be maintained and consolidated over time and, finally become accessible to recall (Eichenbaum, 2017). Most studies discussed in this chapter focus on memories whose elements are organized in space. However, there is evidence to suggest that the neuronal mechanisms underlying spatial memories also apply to conscious memories organized in other, non-spatial for example, social (Tavares, Mendelsohn & Grossman & 2015) or conceptual domains (Aronov, Nevers, & Tank, 2017; Behrens et al., 2018; Constantinescu, O'Reilly, & Behrens, 2016). Therefore, spatial memory may hold the key to understanding the neural processes that organize knowledge in the complex, high-dimensional, non-spatial cognitive map that Tolman originally envisaged (Behrens et al., 2018).
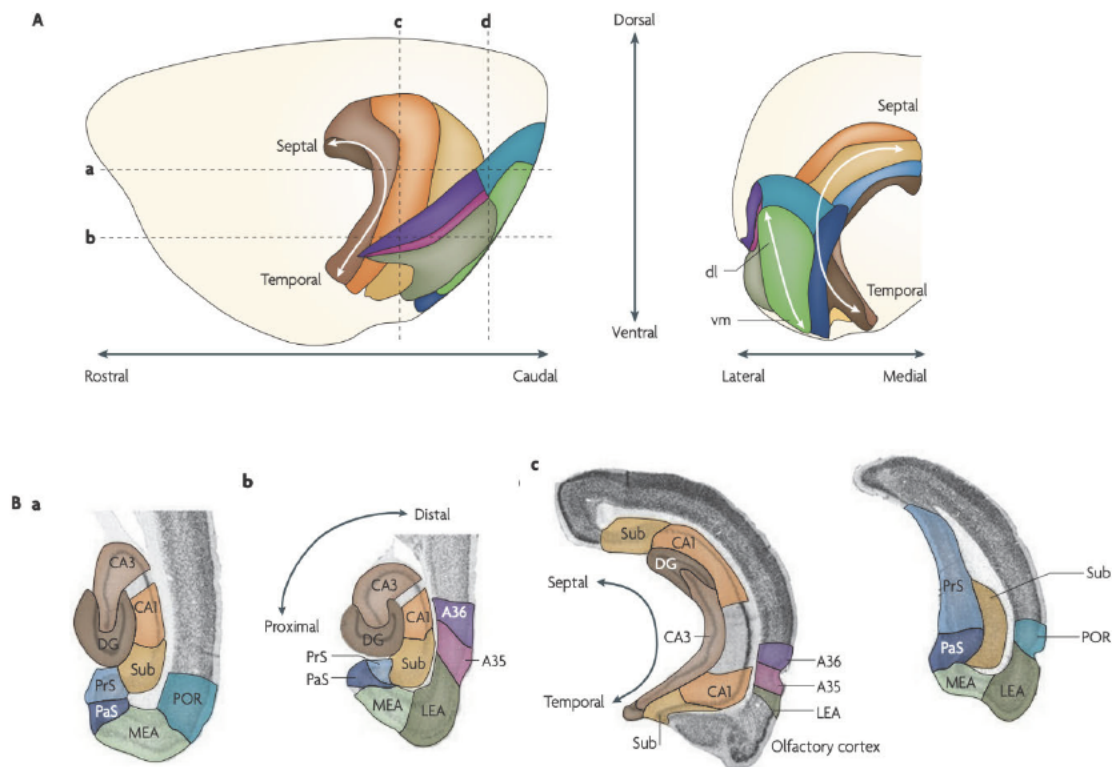
Figure 1-1 Hippocampal anatomy and circuitry.

Anatomical overview of the mouse spatial-cognitive system showing lateral (A; left) rostro-caudal view and (A; right) coronal view showing medio-lateral extent. B) The hippocampal formation shown in a more dorsal (left) and ventral (right view) and comprises CA1 (orange), CA3 (light brown), DG (dark brown) and the subiculum (Sub; light yellow; note (CA2 not shown due to the small anatomical size). Surrounding areas form the parahippocampal region including ERC (medial portion in light green and lateral in dark green), presubiculum (PrS; medium blue) and parasubiculum (PaS; dark blue and the perirhinal cortex (consisting of Brodmann areas (A) 35 (pink) and 36 (purple)) and the postrhinal cortex (POR; blue-green). Three axes are indicated are: the long or septotemporal axis (also referred to as the dorsoventral axis, the transverse or proximodistal axis, which runs parallel to the cell layer and starts at the DG; and the radial or superficial-to-deep axis, which is defined as being perpendicular to the transverse axis. From (Adapted from: Van Strien, Cappaert, & Witter, 2009).

23

1.2.2 Hippocampal anatomy and circuitry

The hippocampal formation is a bilateral C-shaped structure with a complex 3D architecture, situated in the caudal part of the forebrain, over the diencephalon (Figure 1.1). The hippocampus is one of several brain regions that together comprise the hippocampal formation. The hippocampal formation is a prominent C-shaped structure bulging in the floor of the temporal horn of the lateral ventricle. The hippocampus proper consists of three major subfields Cornu Ammonis area 1 (CA1), CA2 and CA3, dentate gyrus (DG) and the subiculum. All of these regions have a trilaminar organisation, for which the nomenclature differs by region but the roles are fairly consistent: a deep layer, comprising a mixture of afferent and efferent fibres and interneurons; a more superficial cell layer, where densely-packed principal neurons are found; and the most superficial molecular layer which is predominantly cell-free (Figure 1.2; (Andersen, Bliss, & Skrede, 1971). In the DG, the deep layer (named the hilus) is enclosed in a 'U' shape by the principal granule cell layer and stratum moleculare. In the CA and subiculum regions, the deep layer is the stratum pyramidale, and then the molecular layer(s). The molecular layer of the CA regions are subdivided into additional layers. In CA3, the unique stratum lucidum, receives mossy fibre axon inputs from the DG; the stratum radiatum, contains the apical dendrites of neurons located in the CA3 pyramidal layer; and the most superficial stratum lacunosum moleculare, contains the apical tufts of the apical dendrites and perforant path fibres (Andersen, et al., 1971). The lamination in CA2 and CA1 is similar, except there is no stratum lucidum (Van Strien et al., 2009).
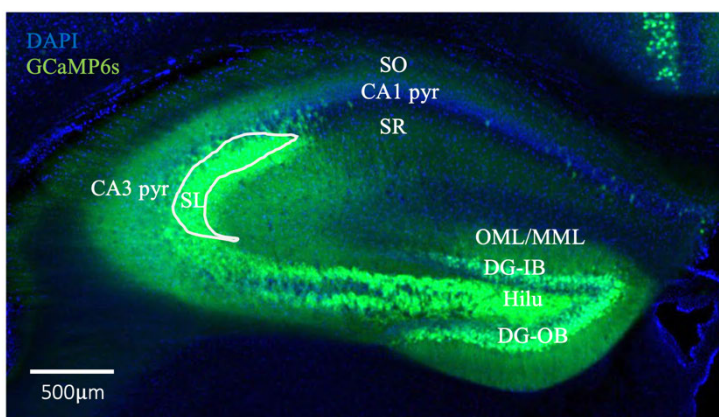
Figure 1-2 Trilaminar organisation of the rodent hippocampus.

A) Mouse coronal histological hippocampal section with DAPI (blue) labelling cell bodies and GCaMP6s (green): stratum lucidum (SL), CA3 pyramidal cell layer (CA3 pyr); (C) DG: outer and middle molecular layer (OML/MML), inner (DG-IB) and outer blade (DG-OB) of the dentate gyrus, hilus; (D) CA1:

stratum oriens (SO), CA1 pyramidal cell layer (CA1 pyr) and the stratum radiatum (SR). Scale bar: 500 μm. Authors own image.

1.2.3 Hippocampal projection pathways

Intensive investigation of circuitry in the hippocampal formation has enabled mapping of the connections inside (intrinsic connections) as well as brain areas outside of the hippocampal formation (extrinsic connections). The DG, CA3 and CA1 are interlinked by strong predominantly forward connections forming the canonical 'tri-synaptic loop'. While an oversimplification of a much more elaborate and parallel input system within the hippocampus (Knierim et al., 2014), it provides a useful overview of information flow though this structure (Figure 1.3). The main input to the hippocampal formation originates in the superficial layers of the entorhinal cortex (ERC; layers II and III), with the majority projecting onto DG granule cells via the perforant path and a small input directly to CA1 (Witter, 1993). DG projects to CA3 pyramidal neurons via mossy fibres which provide a sparse but powerful connection, hence this synapse operates as a 'detonator' (typically producing depolarization and firing the post-synaptic neuron; Henze, Wittner, & Buzsáki, 2002; Vyleta, Borges-Merjane, & Jonas, 2016). Each CA3 cell receives ~46 mossy fiber inputs, so that the sparseness of this mossy fiber-CA3 connectivity is ~0.005% (Treves & Rolls, 1992a; Rolls, Treves, Robertson, Georges-François, & Panzeri, 1998). CA3 also receives two further inputs, including one from ERC via the perforant pathway. By contrast, these inputs are more numerous, but weaker than mossy fibre input (Ishizuka et al., 1990). However, the largest number of synapses on the dendrites of CA3 pyramidal cells is provided by the recurrent collaterals of CA3 cells themselves (Amaral & Witter, 1989, 1995; Amaral et al., 1990; Ishizuka, Cowan, & Amaral, 1995; Amaral, 1993; Ishizuka et al., 1990; Witter, 2007), such that effectively the CA3 system provides a single network, with a connectivity estimate of 2% between CA3 neurons (given that the connections are bilateral; Guzman et al., 2016; Rolls, 2013). Interestingly, the CA3-CA3 recurrent collateral system is more extensive in macaques than in rats (Kondo, Lavenex, & Amaral, 2009), suggesting that in primates, the CA3 network operates even more as a single network than in rodents, equipping them with a larger storage capacity (Rolls, 2013). Finally, the pyramidal cells of CA3 project on to pyramidal cells of CA1 via the Schaffer collateral

system. Schaffer collateral axons are topographically organized along the septotemporal axis; CA3 cells closer to the DG project to CA1 cells septal to their location (Ishizuka et al., 1990). CA3 cells closer to CA1 project to cells more temporally located (Ishizuka et al., 1990). To close the loop, information from CA1 is projected back to the deep layers of the ERC (especially layer V) both directly and indirectly through the subiculum and parahippocampal structures (Delatour & Witter, 2002) and prefrontal cortex (Verwer, Meijer, Van Uum, & Witter, 1997).
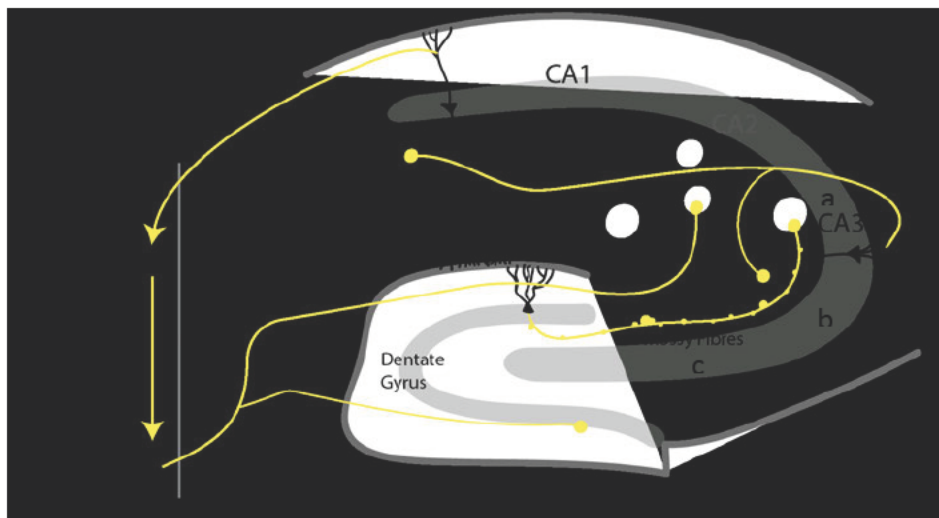
Figure 1-3. Coronal section and wiring in the hippocampal circuit underpinning the 'tri-synaptic' loop.

Schematic illustrating the circuitry within the hippocampus. CA3 pyramidal cells receive two parallel inputs from the ERC, a direct input via the perforant path and an indirect input from the DG via powerful, but sparse mossy fiber detonator synapses (Henze et al., 2000; Vyleta et al., 2016). Direct ERC connections to CA3 might deliver cues for retrieving information from CA3, even when input is incomplete (Rolls, 2013). CA3 also receives input from excitatory recurrent connections between CA3 pyramidal neurons, known as the associative/commissural loop (Witter, 2003). CA3 projects to CA1 via Schaffer collaterals and CA1 transfers information out of the hippocampus to the ERC and other regions. CA3 can be further subdivided along the transverse axis into CA3a, most distal from DG, CA3c, proximal to DG, and CA3b, intermediate (Lorente De Nó, 1934). Image authors own.

It is important to note that several back projections have been reported highlighting limitations to the overly simplified 'tri-synaptic' circuitry. For example infrequent projections from CA3 to

temporal DG (Li, Somogyi, Ylinen, & Buzsáki, 1994), from CA1 back to CA3 and direct inputs from superficial EC to the understudied CA2 region (Chevaleyre & Siegelbaum, 2010), from which CA2 axons then project to the superficial layers of CA1 and proximal CA3 (Mercer, Trigg, & Thomson, 2007). In addition to these back projections, CA3 recurrent collateral show a particular organization along the septotemporal axis in that each neuron targets cells nearby (Sun et al., 2017). The complexity of these connections has raised several questions about their putative functions which will be revisited later in this chapter. Furthermore, CA3 pyramidal neurons respond to perforant pathway stimulation prior to granule cell stimulation, suggesting that mossy fiber input may not initiate CA3 neuron activation (Do et al., 2002; Yeckel & Berger, 1990). These deviations to the propagation of intrinsic activity as described by the tri-synaptic loop suggests the hippocampus may be more likely arranged as a set of parallel microcircuit loops, where each loop is formed by a hippocampal subfield (Hainmueller & Bartos, 2020).

1.2.4 Hippocampal subfield CA3

As with other areas in the brain, cellular populations initially considered homogeneous have been subsequently shown to be comprised of distinct cell types with specific circuit structure and function, when investigated at sufficient resolution. This is certainly true for the hippocampus (interesting examples include mature vs immature DG granule cells, and thorny vs athorny CA3 pyramidal neurons; Hunt et al., 2018; Toni & Schinder, 2016), but this should not be surprising given Lorente de Noˊ (1934) originally divided CA3 into three further subregions (CA3a, b, c; Fig. 1A) along the proximodistal axis (Lorente De Nó, 1934). Proximal CA3 (CA3c) lies closest to DG, whereas distal CA3 (CA3a) lies closest to CA2, with mid-CA3 (CA3b) interposed (Lorente De Nó, 1934). This parcellation has been forgotten until recently. Despite the majority of theoretical literature treating CA3 as one homogenous entity, there is marked proximo-distal heterogeneity in CA3 pyramidal neuron dendritic morphology (Sun et al., 2017), axonal projections (Sun et al., 2017; Tamamaki & Nojyo, 1995) and recurrent connectivity (Ishizuka, Weber, & Amaral, 1990). Given that neural architecture determines the computations possible, these intricacies suggest a functional gradient, potentially with intra-subregions serving different computational roles in

hippocampal processing (Lee & Kesner, 2004) and will be discussed in more detail later in this chapter.

1.3 The cellular basis of learning and memory

1.3.1 Associative learning

Learning is the acquisition of skill or knowledge. Subsequently, memory is the storage and expression of what has been acquired. To understand how neural networks might implement learning changes on a cellular level, Hebb and other theorists working on mathematical models of the brain, created the concept of a 'neural network' (Amit, 1989; Hebb, 1949; Hopfield, 1982; Marr, 1971). This is a term synonymous with models of distributed connectivity in which neurons are abstracted to nodes and associated by changing modifiable connections (plastic synapses) based on learning rules (Hebb, 1949). Hebb postulated that a synaptic modification for learning and memory occurs as a consequence of coincident pre- and postsynaptic activity (Hebb, 1949). Experimental evidence for plastic synapses, a physiological correlate of Hebb's rule, lagged almost 20 years, until the pioneering work of Bliss and Lømo and their discovery of long-term potentiation (LTP; Bliss & Lømo 1959). They found that repetitive high frequency stimulation trains (15Hz for 15 seconds) of the perforant pathway in rabbit hippocampal slices, resulted in a rapid and long-lasting increase in the strength of these synapses that could persist for many days (Bliss & Gardner-Medwin, 1971, Bliss & Lømo 1959; Figure 2). LTP involves both an increase in the synaptic response and an increase in neuronal excitability (EPSP-to-spike potentiation; Bliss et al., 2014). In contrast to LTP, long term depression (LTD) refers to an activity-dependent reduction in the efficacy of neuronal synapses (Bear & Abraham, 1996). LTD is induced by low frequency stimulation either alone or paired with a small postsynaptic depolarization (Dudek & Bear, 1993; Kirkwood & Bear, 1994; Linden & Connor, 1995; Mulkey, et al., 1992). Combined LTP and LTD allow activity-dependent bidirectional modification of synaptic strength and serve as promising candidates for the synaptic basis of learning and memory (Bliss & Collingridge, 1993; McCormack, Stornetta, & Zhu, 2006; Siegelbaum & Kandel, 1991).
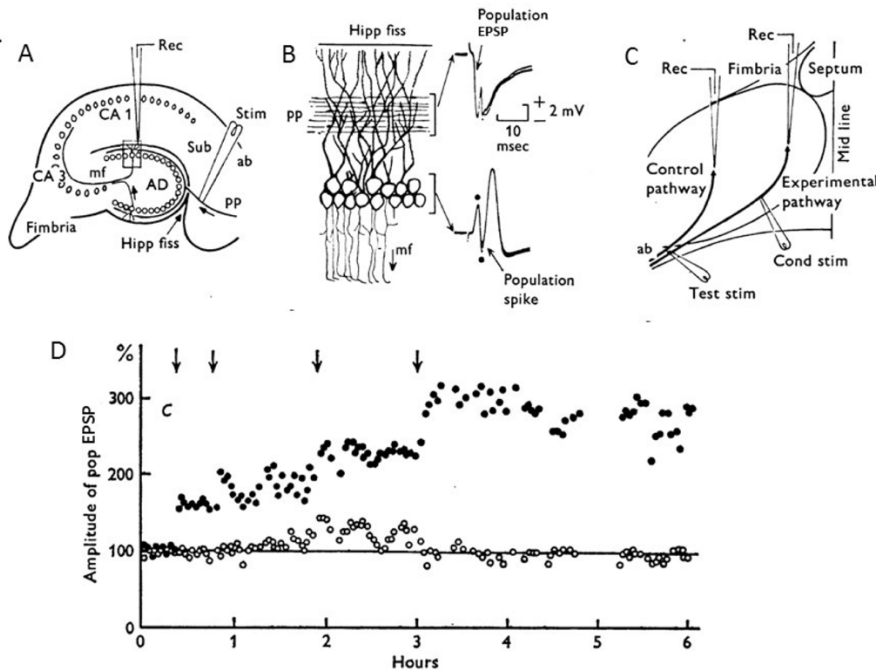
Figure 1-4 The first example of LTP in the DG of the anaesthetized rabbit.

A-C. Anatomy of the hippocampus; A) population potentials from synaptic and granule cell body layers (B), and placement of stimulating and recording electrodes (C). The arrangement of the two stimulating electrodes in (B) allowed the rostral electrode (Test stim) to activate the perforant path while the second conditioning electrode (Cond stim) was placed to activate only fibres projecting to granule cells nearer the midline (experimental pathway). (B). High-frequency trains (15Hz for 15 seconds) were delivered at intervals to the experimental pathway (arrows in D) via the conditioning stimulating electrode (labelled Cond stim). LTP of the population (field) EPSP in the experimental pathway (filled circles) but not the control pathway (open circles) following multiple episodes of high-frequency stimulation. Adapted from Bliss & Lømo, (1973).

To further support the role of activity induced synaptic plasticity in learning and memory, later work delineated the underpinning molecular mechanisms. Seminal research demonstrated that LTP in CA1 requires glutamate and sufficient post synaptic depolarisation to remove the voltage dependent magnesium block in the N-methyl-D-aspartate receptors (NMDARs; Collingridge, Kehl, & McLennan, 1983). This magnesium block means the NMDARs themselves are not sensitive to low levels of glutamate but their neighbouring AMPARs are, subsequently resulting depolarisation of the postsynaptic neurons (via AMPARs) removing the magnesium block, allowing calcium ions to pass into neurons via the NMDAR (Song et al., 1999). Three key

29

supporting expression mechanisms were found to implement these longer term structural synaptic changes. At potentiated synapses there is an increase in the probability of presynaptic neurotransmitter release, an increase in single channel conductance of AMPARs and an increase in the number of AMPARs (Song et al., 1999; Shi et al., 1999; Zhu, Qin, Zhao, Van Aelst, & Malinow, 2002). Collectively, these changes make these synapses more sensitive to glutamate (Bliss et al., 2014). Hence providing the molecular adaptations to support an activity plastic 'Hebbian' synapse (Bliss, Collingridge, & Morris, 2014).

To characterize the temporal requirements for the induction of LTP and LTD, Levy and Steward, (1983) varied the relative timing of a strong and a weak input from the ERC to the DG and found that synaptic modification depended on the temporal order (Levy & Steward, 1983). Potentiation occurred when the weak input preceded the strong input by < 20ms, and reversing the order led to LTD (Levy & Steward, 1983). Subsequent studies further demonstrated the importance of the temporal order of pre- and postsynaptic spiking in synaptic modification, delineating the critical window on the order of tens of milliseconds (Bi & Poo, 1998; Debanne, Gähwiler, & Thompson, 1998; Magee, & Johnston, 1997; Markram, Lübke & Frotscher, 1997) Such spike-timing-dependent plasticity (STDP; Abbott, & Nelson, 2000) has now been observed at excitatory synapses in a wide variety of neural circuits (Boettiger, et al., 2001; Feldman, 2000; Sjöström, et al., 2001). Compared with the correlational forms of synaptic plasticity, STDP captures the importance of causality in determining the direction of synaptic modification, which was originally implied in Hebb's postulate.

To clarify this mechanistic understanding of neural circuit changes in response to experience, Goddard and colleagues formalized the concepts of co-operativity and associativity as defining characteristics of LTP:

- Co-operativity refers to the need to activate a threshold number of inputs, which may serve to filter out non-salient information (McNaughton et al., 1996).

- Associativity refers to the property whereby a strong stimulus can enable a weak stimulus (that by itself is below threshold for LTP). To elicit LTP when the two independent pathways are co-activated in close temporal and spatial proximity (Barrionuevo & Brown, 1983)

Research supports these processes during behaviour, for example, spatial memory task performance was disrupted by administration of pharmacologic agents that block NMDA receptors, preventing synaptic plasticity (Morris et al., 1986). This body of work supports a synaptic-plasticity theory of associative learning and memory formation (Bliss et al., 2014; Bliss & Collingridge, 2013, Martin, Grimwood & Morris, 2000; Morris et al., 1986).

1.3.2 Encoding memories

Hippocampal subregion CA3 has been implicated in the initial encoding and storage of memory, in part due to the theoretical notion that CA3 is thought to effectively operate as a single network given its recurrent collateral connections (Kesner & Rolls, 2015; Rolls, 2013). Therefore, it follows that CA3 can allow arbitrary associations between inputs originating from very different parts of the cerebral cortex (Kenser & Rolls, 2005). For example, object-place mnemonic function involves associations between information originating in the temporal visual cortex about the presence of an object, and information originating in the parietal cortex regarding its location (Kenser & Rolls, 2005; Rolls, 2013). Although CA3 recurrent connectivity is not fully uniform (Ishizuka et al., 1990; Witter, 2007), the connectivity in CA3a and b is considered to have the properties of an interconnected auto-association network (Sun et al., 2017) and therefore associate diverse afferent inputs to CA3 as described above (Kesner & Rolls, 2015; Rolls, 2013).

A simplified hypothesis of how the DG, CA3 and ERC circuitry works together to encode and store memories is as follows: DG granule cells acts as a competitive network (employing unsupervised learning) to produce the sparse yet efficient representation in CA3 neurons (Rolls & Kesner, 2016). This non redundant representation is necessary for the autoassociation properties, namely the ability to re-activate or recall an existing stored memory pattern (also known as

pattern completion and discussed later in this chapter; Kesner & Rolls, 2015; Lee & Kesner, 2004; Rolls, 1990b, Treves & Rolls, 1992, 1994). To achieve this DG granule cells perform orthogonalization (or pattern separation), preparing the DG representations for the CA3 network (Kesner & Rolls, 2015; Lee & Kesner, 2004; McHugh et al., 2007; Rolls, 1989, 1990b, Rolls, Treves, Robertson, Georges-François, & Panzeri, 1998, Treves & Rolls, 1992, 1994). Given the small number of DG mossy fiber-CA3 connections, there is a randomizing effect on the representations created in CA3, such that representations are maximally different to each other, enabling storage of a large number of similar events (Kesner & Rolls, 2015; Lee & Kesner, 2004; Rolls, 1990b, Treves & Rolls, 1992, 1994) and minimizing interference between distinct memories (Hopfield, 1982; Treves and Rolls, 1991; Rolls, 2008; Rolls, et al., 1998). After DG granule cells have recruited unique CA3 pyramidal cell groups during memory encoding, they become associated with the ERC input patterns (hetero-association) representing the new information and other CA3 inputs (auto-associations; Rolls, et al., 1998).

Hence a single active DG granule cell could send a mossy fiber axon to CA3 targeting 10–15 pyramidal cells with 'giant' mossy-fibre synapses and subsequently determine whether the set of connected CA3 cells become co-active, by combining the detonator properties with heterosynaptic and associative plasticity (Rebola et al., 2017). Strong mossy-fibre synapses at proximal pyramidal cell dendrites can elicit large dendritic events that support instantaneous reorganization of pyramidal cell place fields in CA1 and CA3 (Bittner, Health, & Grienberger, 2017; Diamantaki, et al., 2018; Diamantaki, et al., 2016). Hence, mossy fibre-mediated dendritic events promote heterosynaptic plasticity at perforant pathway and CA3 recurrent inputs (Buzsaki 1989; Rebola et al., 2017). Once synaptic plasticity has strengthened the intrinsic and extrinsic connections of these ensembles (Do, Martinez, Martinez, & Derrick, 2002; Tsukamoto et al., 2003), perforant pathway inputs may serve to reactivate them, promoting recall of the associated memory without the need for mossy-fibres (Rolls, 2018).

Within this framework for memory encoding, memories are stored at perforant pathway and recurrent CA3-CA3 connections (Rolls, 2018). Indeed, lesion studies suggest a functional

dissociation between the perforant pathway–CA3 and the perforant pathway–DG–CA3 paths (Lee & Kesner, 2004). Mossy-fibre afferents were required for memory acquisition but not for retrieval, whereas, perforant pathway afferents initiated memory retrieval from the CA3 without mossy-fibre input (Lee & Kesner, 2004). Moreover, acute DG inactivation impaired memory encoding but not recall in most (Kheirbek, Drew, & Burghardt, 2013; Lassalle, Bataille, & Halley, 2000; Madroñal, et al., 2016), but not all studies (Denny, Kheirbek, & Alba, 2014; Park, et al., 2016). Notably, granule cell ensembles are reactivated during the retrieval of recent, but not remote, memories (Tayler, Tanaka, & Reijmers, 2013). Thus, DG granule cell activity may be required during initial encoding and early, intrahippocampal memory consolidation but can be dispensable for recall (Hainmueller & Bartos, 2020). Additional support for memory storage in perforant pathway-CA3 connections comes from the requirement of NMDARs in CA3 pyramidal cells neurons to enable single trial learning (Nakazawa et al., 2002, Nakazawa et al., 2003; Wagatsuma et al., 2017) and the ability to recall memories in the in the absence of mossy-fibre signaling (Tsukamoto et al., 2003). These findings support the hypothesis that memories in the CA3 are ultimately stored at perforant pathway and recurrent CA3 connections (Nakazawa et al., 2002).

1.3.3 Consolidating memories

Stabilising CA3 ensembles, essentially consolidating memories is thought to occur through the reactivation of these ensembles during sharp wave ripples (SWRs; Buzsáki, 1989; Sasaki, Piatti & Hwaun, 2018). Once CA3 ensembles are stabilized, recall may become independent of mossy-fibre inputs (Tsukamoto et al., 2003). SWR complexes are a prominent endogenous hippocampal activity pattern that occur during quiet behavioural states (Buzsaki, Horvath, & Urioste, 1992; Csicsvari & Hirase, 2000, Ólafsdóttir, Carpenter, & Barry, 2017) Massive synchronized excitation from CA3 pyramidal neurons discharges distributed subsets of CA1 pyramidal neurons and interneurons (Buzsaki et al., 1992). This collective activity results in high-frequency oscillations (150~250 Hz ripples) in the CA1 pyramidal layer (Schomburg, Anastassiou, Buzsáki, & Koch, 2012; Ylinen et al., 1995). Ripple frequency is conducive for LTP induction, and selective disruption of

33

SWRs has been shown to result in spatial memory deficits (Buzsáki, 1989; Girardeau, Benchenane, Wiener, Buzsáki, & Zugaro, 2009). Hence, during SWRs, hippocampal cell ensembles reinstate previously encoded network activity patterns, a process that is key for memory consolidation (van de Ven, Trouche, McNamara, Allen, & Dupret, 2016).

While the morphological hallmark of CA3 neurons is large thorny excrescences (complex spines) characteristic of giant postsynaptic mossy-fiber boutons, a new class of CA3 pyramidal neurons has been discovered. Whole-cell patch-clamp recordings from neurons in stratum pyramidale in acute mouse hippocampal slices, revealed a bimodal distribution of firing frequencies and subsequently two groups of CA3 pyramidal neurons with distinct morphological and physiological features (Hunt, Linaro, Si, Romani, & Spruston, 2018). Namely, regular-spiking neurons decorated with complex pines, and athorny neurons, lacking spines (Hunt et al., 2018). Athorny CA3 (10-20% CA3 population) occupied a distinct cytoarchitectural position (deep-layer SP in CA3a and CA3b; Hunt et al., 2018), with a distinct physiological phenotype, in that their intrinsic busting initiated SWRs (Hunt et al., 2018). Thus, our current understanding follows that memory consolidation involves reactivation or replay of previous activity patterns during specific network states (SWRs) to consolidate those memories, and athorny CA3 pyramidal neurons likely initiating this process (Hunt et al., 2018).

1.3.4 Recalling memories

Recalling or remembering involves accessing and reactivating that stored activity pattern. Theories posit that such reactivation occurs through amplified activation in the recurrent CA3-CA3 network (McNaughton et al., 1996; McNaughton & Morris, 1987; Rolls & Treves, 1990; Treves & Rolls, 1992a). While this has been difficult to test experimentally, there is evidence that CA3 plays a crucial role in memory recall. Nakazawa and colleagues (2002) ablated NMDA receptors in CA3 pyramidal neurons in mice trained in a Morris water maze task. The mutant mice acquired and retrieved the spatial reference memory but were impaired when performing with familiar cues removed (Nakazawa et al., 2002). These results suggest NMDA receptor-

dependent synaptic plasticity mechanisms in CA3 are critical for auto associative memory recall (Nakazawa et al., 2002). However, all synaptic inputs onto CA3 pyramidal neurons are affected by inactivation of NMDARs, therefore it is not possible to predict which pathway the removal of plasticity is causal to behavioural deficits. Furthermore, loss of NMDA receptors likely also impacts spike transmission at mossy fibre–CA3 pyramidal cells (Hunt, Puente, Grandes, & Castillo, 2013; Rebola, Carta, Lanore, Blanchet, & Mulle, 2011). Contradictory findings have shown that genetic removal of the NMDA receptor subunit GRIN1 from CA3 pyramidal neurons did not affect performance of mice when navigating to previously experienced platform locations (Nakazawa et al., 2003). However, their ability to rapidly acquire memory of a novel hidden platform locations was impaired (Nakazawa et al., 2003). Hence, CA3 NMDA receptors play a crucial role in rapid hippocampal encoding of novel information (Nakazawa et al., 2003). The locus coeruleus in the brainstem is thought to provide CA3 with the novelty signal to enable this single-trial learning (Wagatsuma et al., 2017).

While the synaptic mechanisms underpinning recall remain enigmatic, successful memory recall requires sufficient synaptic efficacy and network connectivity (Guzman, Schlogl, Frotscher, & Jonas, 2016). Recent work suggests initial estimates of CA3-CA3 macro connectivity are much sparser than previously thought. Simultaneous recordings of up to 8 CA3 pyramidal neurons, revealed that the connection probability was 0.92% (Guzman et al., 2016). However, Guzman and colleagues found the efficacy of unitary connections was high (Guzman et al., 2016). Therefore, coincident firing of a small number of presynaptic cells was sufficient to initiate action potentials in a postsynaptic cell (Guzman et al., 2016). These properties of recurrent CA3–CA3 synapses allow efficient information encoding by small neuronal ensembles (Guzman et al., 2016). Guzman et al., (2016) also demonstrated that real-size network models (with 330,000 neurons, representing CA3 in one hemisphere of a rat; Amaral, Ishizuka, & Claiborne, 1990) with a realistic connection probability of ~1% can recall stored activity patterns, when disynaptic connectivity motifs were included in the model architecture (Guzman et al., 2016). However, whether CA3 network properties in vivo are consistent with these assumptions remains unclear and the mechanisms generating the connectivity motifs are unknown (Guzman et al., 2016). It is possible

these motifs arise from structural plasticity in synchronously active ensembles as discussed in Section 1.3. Ultimately, advanced neural circuit mapping during behaviour, combined with single-cell electrophysiology is required to define the set of plasticity rules at play in identified CA3 cells to enable episodic memory encoding in CA3 circuits.

1.4 Spatially tuned neurons in the hippocampal formation

1.4.1 Place cells

In 1971, O'Keefe and Dostrovsky made their landmark discovery that neurons in the rat hippocampus had spatial receptive fields. Using single unit electrophysiological recording in freely moving rats, they found single neurons (place cells) from the CA1 region are selectively active when the animal visited a particular location in the environment (O'Keefe & Dostrovsky, 1971). Different cells had different place fields, such that collectively they formed a representation which tiles the environment (O'Keefe & Dostrovsky, 1971). In principle this enables the animal's location to be inferred from the joint activity of a fairly small population of neurons. Given these observations and inspired by Tolman's proposal that navigation is guided by internal cognitive maps, O'Keefe and Nadel (1978) suggested that place cells are the basic element of a distributed allocentric cognitive map of the animal's environment, and form the neural substrate supporting spatial cognition (O'Keefe and Nadel, 1978).

This notion that the hippocampus supports spatial cognition was becoming increasingly apparent through different behavioural measures making cognitive processes explicit and measurable. Evidence from hippocampally lesioned animals demonstrated that they could not solve a radial maze task (Olton & Samuelson, 1976) or other spatial tasks (reviewed in O'Keefe & Nadel, 1978). Nor could lesioned animals learn to find a hidden platform the Morris water task (Morris 1981) lacking local cues to location and therefore requiring integration of environmental allocentric cues. In a further line of support, since their discovery place cells have been found in all

36

hippocampal subregions (CA1, 2 and 3) and the DG (Lu et al., 2015; Mankin, Diehl, Sparks, Leutgeb, & Leutgeb, 2015; Muller & Kubie, 1987; O'Keefe & Dostrovsky, 1971), and have been identified across multiple species including, mice (Tonegawa et al., 1996) monkeys (Rolls et al., 1989), bats (Ulanovsky & Moss, 2007; Yartsev & Ulanovsky, 2013), pigeons (Bingman, et al., 1996) and humans (Ekstrom et al., 2003). Such widespread, cross species findings suggest these cells are crucial for a universal function shared across animals.

How place cells support spatial cognition is still an intense topic of debate. However it has been rapidly determined that the representation they form is not a pixel, but rather a population code, as a given place cell may have place firing fields in multiple environments (Alme et al., 2014) and place cells are non-topologically organised. Hence, for a particular location in space, only a subset of place cells will be active, and anatomically proximal cells do not necessarily have adjacent place fields in physical space, forming a distributed representation of the environment (Wilson & McNaughton, 1993; Fenton et al., 2008, Redish et al., 2001). Collectively, place cells represent spatially tuned neurons encoding the current position of an animal, such that sequences of place cells provide a distributed population code for past (replay) and future (preplay) trajectories though space (Diba et al., 2007; Foster & Wilson, 2006; Karlsson et al, & 2009; Pfeiffer & Foster, 2013).

1.4.2 Other spatially tuned cells

A whole body of literature has since revealed a number of additional neurons encoding unique (or the conjunction of) information relevant for spatial navigation such as position, distances, objects, direction, geometry and speed (Grieves & Jeffery, 2017; Hardcastle, et al., 2017). Ranck and colleagues, reported the discovery of spatially sensitive neurons in the dorsal presubiculum (Ranck, 1984; Taube et al., 1990a, 1990b). These head direction cells did not have place fields but fired when the animal faced in a particular direction (Ranck, 1984). A key insight emerging from the head direction system was that activity across the neural population is coherent; the relative firing directions of the neurons was the same in every environment (Ranck, 1984). Directional

signals likely form the core of the spatial representation; however, one outstanding question is whether the place cell map is metric, that is, has explicit information about distance and direction, or whether it is just associative. Distance representation was not confirmed until the discovery of grid cells in the medial entorhinal cortex (MEC; Hafting et al., 2005). Grid cells have multiple regularly spaced place fields, in a triangular or hexagonal pattern, which repeats across the entire available space (Hafting et al., 2005). The scale of the grid increases from dorsal to ventral MEC (Brun et al., 2008). Collectively the discovery of place and gird cells won the 2014 Nobel Prize in Physiology or Medicine for 'discoveries of cells that constitute a positioning system in the brain'.

It is clear that place and grid cells are part of a diverse and intermingled network of cell types, with distinct functions contributing to a cognitive map. Other spatially tuned cells include border cells which were found in the subiculum (Barry, Lever, Hayman, et al., 2006; Lever, Burton, Jeewajee, O'Keefe, & Burgess, 2009), social place cells, cells that encode the locations of other agents on the map (Danjo, Toyoizumi, & Fujisawa, 2018; Omer, Maimon, Las, & Ulanovsky, 2018) band cells (Krupic, Burgess, & O'Keefe, 2012) and cells that encode the vector relationships to borders (Solstad, Boccara, Kropff, 2008) objects (Høydal, Skytøen & Andersson, 2019), rewards and goals (Sarel, Finkelstein & Las, 2017; Figure 1.5). The discovery of these functionally specialized neurons in hippocampal–ERC circuit, with clear behavioural correlates, provides the possibility to uncover how cognitive processes, like memory arise (Fyhn, Hafting, Treves, Moser, & Moser, 2007)

It is unclear how these differently specialised cell types interact on a local and population level. It is likely that individual place cells receive inputs from both grid and border cells, possibly with grid cells providing self-motion-based distance information and border cells providing position in relation to geometric boundaries (Bush, Barry, & Burgess, 2014; Zhang et al. 2014). It also seems plausible that the strongest input may originate from grid cells, which, in the superficial layers of the medial entorhinal cortex, are several times more abundant than border cells (Sargolini, et al., 2006; Solstad et al., 2008; Boccara et al., 2010). In the majority of circumstances, the two input

classes are likely to be coherent and redundant, however if one is absent then the other might be sufficient to generate localized firing in the hippocampus. This perhaps explains why place cells are still found with disruption to the grid system or how animals can navigate when grid patterns become distorted (reviewed in Carpenter & Barry, 2016; Krupic, Bauza, Burton, & O'Keefe, 2016). The evidence collected in the last few decades has moved towards developing a mechanistic understanding of how the hippocampal system, forms, updates and retrieves a map-like representation of space supporting navigation (Moser, Rowland, & Moser, 2015). Deployment of new and sophisticated tools including genetic manipulation and optophysiological methods for recording and manipulating neurons, will allow us to go further than just observing the distributed neural code underpinning Tolman's cognitive map, to more completely understand how such a map might function (Jeffery, 2018).
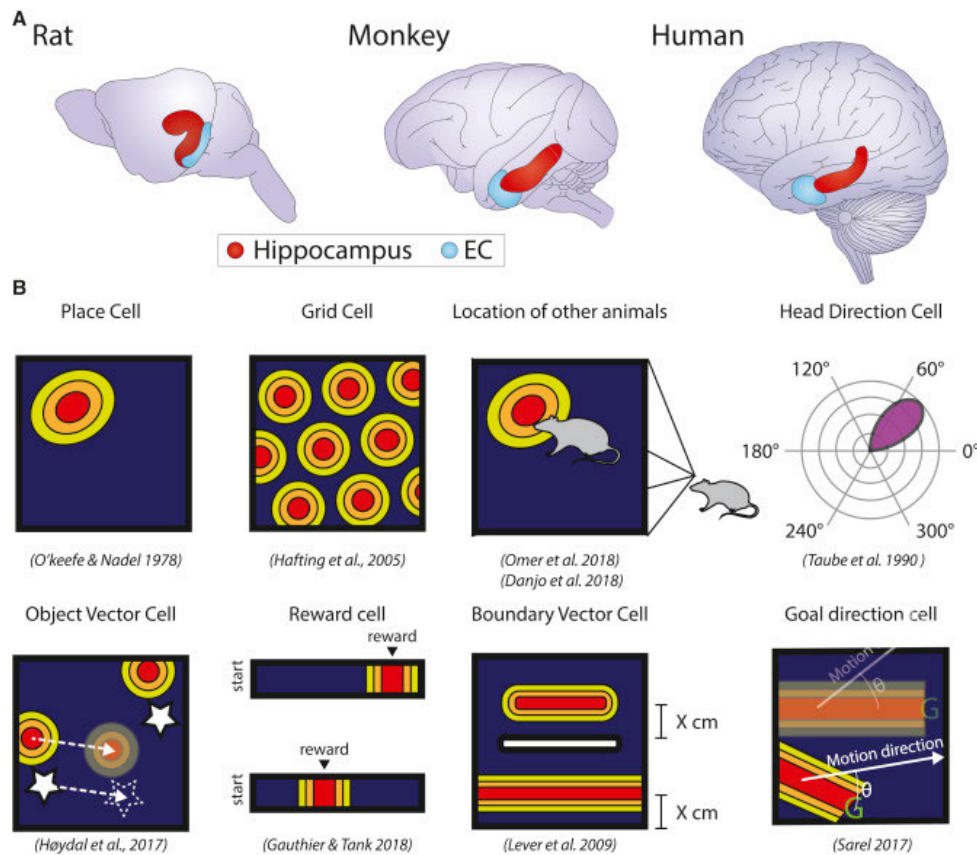
Figure 1-5 Different cell types in the hippocampal formation represent different spatial variables.

A) Anatomical location of the hippocampus (red) and entorhinal cortex (blue) across species. (B) Place cells are active when an animal is in a specific location (O'Keefe & Dostrovsky, 1971). Grid cells are active when an animal is in one of multiple locations on a triangular lattice (Hafting, Fyhn, Molden, Moser, & Moser, 2005). Social place cells (Danjo et al., 2018; Omer et al., 2018) are active in one animal when it observes that another animal is in a particular location. Head-direction cells (Jeffrey S Taube, Muller, & Ranck, 1990) are active when an animal's head is facing a particular direction. Object-vector cells (Høydal et al., 2018) are active when an animal is in a particular direction and distance from any object. Reward cells (Gauthier & Tank, 2018) are active when an animal is in the proximity of a reward. Boundary vector cells (Lever et al., 2009) are active at a given distance away from a boundary in allocentric orientation. Goal direction cells (Sarel et al., 2017) are active when the goal (shown as a green G) of an animal is in a particular direction relative to its current direction. From Behrens et al., 2018.

1.4.3 The paradox of place cell remapping

This thesis will focus more specifically on place cells given the relative ease of optically accessing the hippocampus and existing frameworks for identifying these cells. Years of electrophysiological studies have revealed descriptive place cell features. In most open field electrophysiological recordings only a subset (approximately 33%) of anatomically proximal pyramidal cells have place fields in a given environment, even though all cells are similarly active during short-wave sleep and antidromic stimulation (Thompson & Best, 1989, Henze et al., 2000). However, if place cells are active in a given environment, the stability of their representation is still debated, and dependent on multiple factors (Kentros et al., 1998; Cacucci, Wills, Lever, Giese, & O'keefe, 2007; Mankin et al., 2012). Early studies found that when an animal enters a new environment, new place fields form within a matter of minutes (Kentros et al., 1998, Thompson & Best, 1990) and can be stable for months (Thompson & Best, 1990). This apparent long-term stability of place fields implies that the spatial representation is recalled and not created de novo each time the animal enters a familiar environment, as expected for neurons of a spatial memory system (McNaughton et al., 1996). However, more recent work finding long-term instability in place fields (Kentros, Agnihotri, & Streater, 2004) has challenged this notion, raising the question: how can a changing spatial neural representation reliably retrieve the appropriate memory of a previously encountered environment?

An interesting property of place cells is their propensity to remap. Remapping refers to any manipulation-induced changes in the firing of place cell. Hence, an individual place cell is part of many independent representations. Muller and Kubie investigated the effects of changing the most salient visual cues in a cylindrical environment and introducing various local cues (Muller & Kubie, 1987). Whether place cells remapped completely or only partially depended on the extent of the manipulations for example, changing the colour of a cue card to removal of the cue card. These manipulations induced changes to the firing of place cells and have since been defined as:

1. **Local remapping**: Changes in rate coding with no change in the location of place fields, preserving the population code for space (Leutgeb et al., 2005).

2. **Partial remapping:** Changes in location of some but not all place fields, typically in fields near a specific manipulated cue (Muller & Kubie 1987; Bostock, Muller, & Kubie, 1991).

3. **Global remapping**: Re-arrangement of place fields throughout the environment, resulting in changes in location of place fields, usually observed when comparing environments that differ in several features including shape, colour, texture, and location (Leutgeb et al., 2004; Fyhn et al., 2007).

Local or global changes to the environments reflects a dual rate and position coding, such that each pyramidal neuron is active when an animal traverses a distinct location in space and its mean firing rate when in this location varies with contextual features (Leutgeb et al., 2005).

Global remapping suggests different environments appear to be encoded by separate spatial maps (O'Keefe & Nadel, 1978), providing a unique population-wide neural code for context specificity (Plitt & Giocomo, 2021). In contrast, changes to contextual features like the colour or shape of surrounding walls (Leutgeb et al., 2005), odors (Anderson & Jeffery, 2003), or task (Allen, Nick, Rawlins, Bannerman, & Csicsvari, 2012) within the same space can elicit substantial changes in firing rates while the location of place fields is unaffected. Such rate remapping can affect behavioural decisions (Allen et al., 2012) and appears to reflect the presence of multiple context-dependent memories stored within a single spatial map. Furthermore, the degree of familiarity of the environments to the animal might also account for an unstable cognitive map. While O'Keefe and Burgess (1996) found that the place fields stretched or compressed along with the aspect ratio of the environment, other findings suggest distortion of hippocampal maps might only occur when the manipulated environment is novel, as place maps might be more plastic during the initial exposures to an environment (Barry, Hayman, & Burgess, 2007). Collectively, results highlight that hippocampal maps are not rigid, and the environmental conditions underpinning their different forms of plastic modification are not fully understood. The majority of the research on place coding comes from relatively small samples of CA1 neurons. To better understand the seemingly paradoxical nature of dynamic and stable representations we need to investigate larger populations in and beyond the CA1 subregion.

42

1.4.4 Orthogonalization of place maps

The orthogonal nature of place-cell maps is not shared by any of the known spatial cell types upstream of the hippocampus, as any overlap between active subsets in two environments is not larger than expected by chance (Leutgeb et al., 2004). The apparent independence of the place-cell maps contrasts with the functional rigidity of the grid-cell population (Hafting et al., 2005). For example, changes in the environment which lead to global remapping in CA1, induced changes in the firing locations of simultaneously recorded grid cells, except these changes were coherent among grid cells (Fyhn et al., 2007). Such that, for grid cells with similar spacing, the firing locations shifted in the X Y plane between environments, but the distance and direction of grid displacements were similar across the population (Fyhn et al., 2007). Similarly, internal and spatial coherence has been observed in head direction and border cells, respectively (Taube et al., 1990; Solstad et al., 2008; Taube & Burton, 1995; Yoganarasimha, Yu, & Knierim, 2006). The discovery of remapping and the uncorrelated nature of place maps, suggests multiple orthogonal representations, and therefore strongly implicates the hippocampus in the accurate storage and retrieval of declarative memories (Moser et al., 2015). The capacity of place maps stored in the hippocampus is not known, but if place maps are expressions of individual memories, this should be high, making remapping a necessity (Rolls, 2013).

Furthermore, these observations suggest that remapping might not be generated in the ERC, but in the hippocampus itself. Given this line of thought, upstream hippocampal areas have been suggested as likely candidates (Latuske, Kornienko, Kohler, & Allen, 2018). Hippocampal subregion CA3 has received relatively little investigation, with the majority of experimental studies involving electrode recordings from CA1. Practically, CA1 is the first cell layer encountered in the rodent hippocampus during surgery advancing from dorsal surface of the brain. Furthermore, a wealth of knowledge exists for spatial memory coding in CA1, but CA3 remains relatively unexplored (Delatour & Witter, 2002; Gonzales, Galvan, Rangel, & Claiborne, 2001).

Despite this paucity of experimental knowledge, the connectivity in CA3 circuitry has inspired much theoretical work on its function.

1.5 Neural network models memory function

1.5.1 Emergent neural functions

The notion that neural circuits are built for an emergent function is not novel and originates in the ideas put forward by Hebb and Marr (Hebb, 1949; Marr, 1971; Yuste, 2015). We now know most neural circuits in the brain have distributed connectivity, whereby each neuron receives inputs from many others and outputs to a large population. Therefore, the contribution of an individual neuron is overshadowed, as function is an emergent property of collective activity (Carrillo-Reid, Yang, Bando, Peterka, & Yuste, 2016a; Marr, 1971). As mentioned, Hebb proposed that information propagates by sequentially activating groups of neurons, or 'Hebbian cell assemblies' (Hebb, 1949). These neuronal ensembles display synaptic plasticity when they fire simultaneously (Hebb's rule), enabling groups of coactive neurons to functionally bind together (Hebbian ensemble; Hebb, 1949). Hebbian assemblies (or groups of coactive neurons) have many properties considered to be signatures of an attractor network (Knierim & Zhang, 2012; McNaughton et al., 1996; Rolls, 2013).

Attractor networks are a popular computational construct used to model different brain systems (Knierim & Zhang, 2012; Hopfield, 1984). An attractor is a convenient general concept for describing the stability of a dynamical system, whose state evolves in time (Amit, 1989). In an intuitive sense, an attractor refers to a collection of states that attract neighbouring states toward that collection (Knierim & Zhang, 2012). Conceptually this is illustrated by the point attractor, which is a single, stable, equilibrium state of the system (Figure. 1.7 top left). The basin of attraction refers this stable state which 'attracts' the activity of the circuit as it evolves over time (Figure 1.7; Knierim & Zhang, 2012). The simplicity and explanatory power of these models have been useful in conceptualising how distributed memory functions might work (Knierim &

Zhang, 2012). Moreover, there is similarity between attractor network architectures and hippocampal circuitry (McNaughton et al., 1996; Rolls, 2013; 1992; Rolls & Treves, 1990). Given the structure of network connectivity underpins processing and consequently shapes the functional properties of the network, it follows that these networks might share similar computational properties (Amit, 1989; McNaughton & Morris, 1987; Rolls, 2013).

1.5.2 Auto-associative memory networks

Auto-associative models allow arbitrary associations between often distributed inputs. They employ a recurrent neural architecture that implements a Hebbian learning rule whereby synapses are strengthened when there is correlated presynaptic and postsynaptic activity, and vice versa (Hopfield, 1982; Marr, 1971; Rolls, 2007). A well-known example is a Hopfield network, which allows multiple stable states or 'attractors', each corresponding to a stored memory pattern (Hopfield, 1982). Hippocampal mnemonic function has been modelled both as a point attractor neural network and as a continuous attractor network (Rolls, 2007; Rolls et al., 2013; Treves & Rolls, 1992; Tsodyks, 1999). The development of formalisms to describe 'Hopfield nets' have provided tools for analysing and understanding the properties of hippocampal memory networks (De Almeida, Idiart, & Lisman, 2007). The biological correlate of these models is generally considered to be the CA3 subregion, with its extensive recurrent connectivity (Amaral et al., 1990; Rolls & Kesner, 2006; Rolls, 2007) and Hebbian plasticity at CA3-CA3 synapses (Bains, Longacher & Staley, 1999). Hence, this plastic circuitry fits assumptions underlying the Hopfield formalism (discussed in detail below), providing reason to speculate that memory functions in CA3 might operate according to similar principles (Rolls, 2013; Samsonovich 2007)

Theoretical analysis of Hopfield nets has revealed multiple properties of these memory networks. Marr (1971) formulated the idea that such a network is capable of a memory recall process known as pattern completion, the notion that a complete memory can be recalled using only a part of the memory as a cue. That is, if a memory is stored in a recurrent network as a distributed pattern of synaptic weights, the network can reinstate the full activity pattern even if the original input pattern is incomplete, degraded, or corrupted (McNaughton et al., 1996;

45

McNaughton & Morris, 1987; Rolls & Treves, 1990; Treves & Rolls, 1992a). Hence a memory cue is a subset of the cells that encode a memory. When this cue is presented to the network, it excites the corresponding cells, thereby reproducing a part of the memory activity pattern (Rolls, 2013). Then, through the recurrent excitatory connections, these cells excite all the cells that represent the memory, notably including those that did not receive direct input from the cue (McNaughton et al., 1996; McNaughton & Morris, 1987; Rolls & Treves, 1990; Treves & Rolls, 1992a). Such a process seems anecdotally important for human memory; for instance, the sight of a backpack (the cue) may evoke the view from the last hike. A second important property of Hopfield nets is that a memory, when activated, persists in its activity (Kesner, 2013). This occurs simply because the excited cells continue to excite each other (reverberation; Kesner, 2013). This persistence might be a mechanism of short-term memory (or working memory). Recordings from various brain regions revealed that reactivation of a memory can cause persistent firing even after the stimulus is removed (Wang, 2001). A third and critical aspect of Hopfield nets is their attractor property (De Almeida, et al., 2007). Once a memory becomes active, it can be perturbed by external or internal factors that cause errors in firing (Rolls, 2013). In some dynamical systems such a perturbation worsens over time (Deco & Rolls, 2003). However, if this were to occur in a memory system, the memory would eventually become unrecognizable and lost (De Almeida, et al., 2007; Deco & Rolls, 2003). In the case of attractor networks, however, a perturbation does not cause progressive worsening; rather, the firing pattern is restored (attracted) to its original state (Deco & Rolls, 2003). This repair process is possible because the memory is redundantly encoded, such that the cells that remain active can reactivate the cells that stopped firing (Deco & Rolls, 2003; Rolls, 2013; Rolls & Treves, 1990). Conversely, if the perturbation is strong enough, the firing pattern could potentially change to another memory pattern (Rolls, 2013). This abrupt, nonlinear response is a feature of attractor networks (De Almeida, et al., 2007; Deco & Rolls, 2003). Hence, a complementary idea to pattern completion is pattern separation, i.e., two similar input patterns are stored as more dissimilar patterns to reduce the probability of errors in recall (Rolls & Kesner, 2016). We can use these key functional properties of Hopfield networks, to address the potential of CA3 to function as an auto-associational recurrent network (De Almeida, et al., 2007).

$$\delta \, w_{ij} = k \cdot r_i \cdot$$
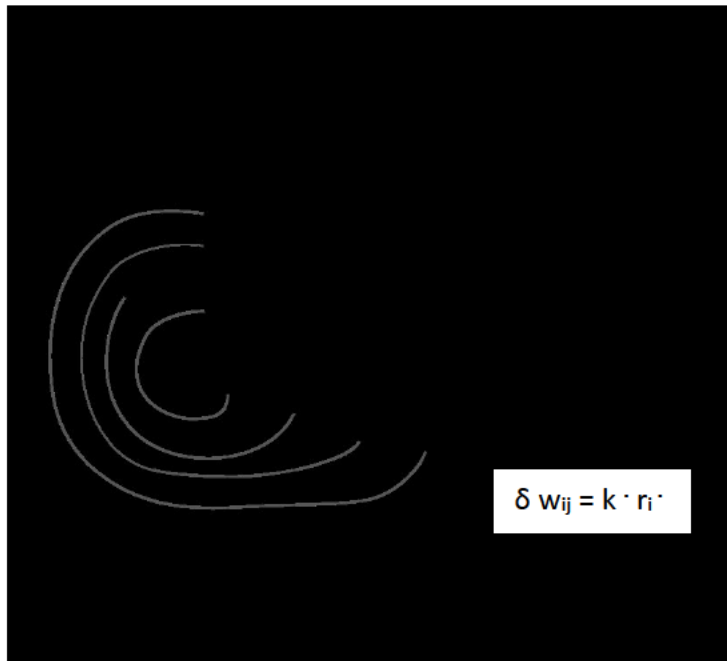
Figure 1-6 The Hopfield auto-associative network architecture

The connectivity architecture of an auto-association or attractor neural network consisting of 4 'neurons' (thick black lines) with feedback (recurrent) connections, such that each neuron outputs to all other neurons. In this illustration, k is a constant, r_i is the activation of the dendrite (the postsynaptic term), r'_j is the presynaptic firing rate, and $\delta \, w_{ij}$ is the change in the synaptic weight w_{ij}., where w_{ij} refers to the j'th synapse onto the i'th neuron. Hence, every pair of units *i* and *j* in a Hopfield network has a connection described by the connectivity weight. The learning rule for change in the synaptic weights is shown in the grey box (Rolls, 2013). The constraint that weights are symmetric ensures that the energy function decreases monotonically while following the learning rules (Mackay, 2003). Image authors own.

1.5.3 Resemblance between CA3 and autoassociative network architecture

If the recurrent CA3-CA3 synaptic connectivity forms the subcellular substrate enabling symmetric memory recall; that is, the whole of the memory can be retrieved from any part (Guzman et al., 2016; McNaughton & Morris, 1987; Rolls, 2013), then multiple functional predictions exist to test this auto-association hypothesis. For example, we can assess the response of the CA3 network and animal behavioural in conditions of uncertainty. Training animals in different environments along a spectrum e.g. a square and a circular environment, tests the prediction that when presented with an environment ambiguous between these end points, hippocampal neurons will fall into an attractor state that represents one of the two previously learned environments, but not a mixture of these environments. Hence, potential attractor dynamics of CA3 will promote a sigmoidal relationship between changes in the inputs to CA3 and changes in the CA3 output representation (O'Reilly, & McCelland, 1994). These separate attractor states refer to independent spatial maps of the environments (O'Reilly, & McCelland, 1994). Such behaviour could reflect the ability of place cells to remap, changing their representation to that of a different environment (Foster & Wilson, 2006; Neunuebel & Knierim, 2014; O'Reilly, & McCelland, 1994).

Remapping has been interpreted as an example of attractor dynamics, supporting the notion that the hippocampus functions as an autoassociative memory system (Neunuebel & Knierim, 2014; Wills et al., 2005). To experimentally probe this, Wills et al. attempted to assess whether the boundaries between attractor states, survive manipulations to the external correlates (Wills et al., 2005). Parametric changes made to the geometry of the animals' environment, as it was morphed between two familiar configurations, did not affect CA1 place cell activity until a threshold was reached, after which the entire network moved state (remapped), to represent the alternative configuration (Figure 1.7 top right; Neunuebel & Knierim, 2014; Wills et al.,2005). Similar results were also found in humans, where fMRI was combined with virtual navigation (Steemers et al., 2016). Linear morphing between familiar environments resulted in nonlinear

changes in human hippocampal activity patterns (Steemers et al., 2016). These findings suggest hippocampal neurons fall into an attractor state, representing one or the other previously learned environments (Neunuebel & Knierim, 2014; Steemers et al., 2016; Wills, 2005).
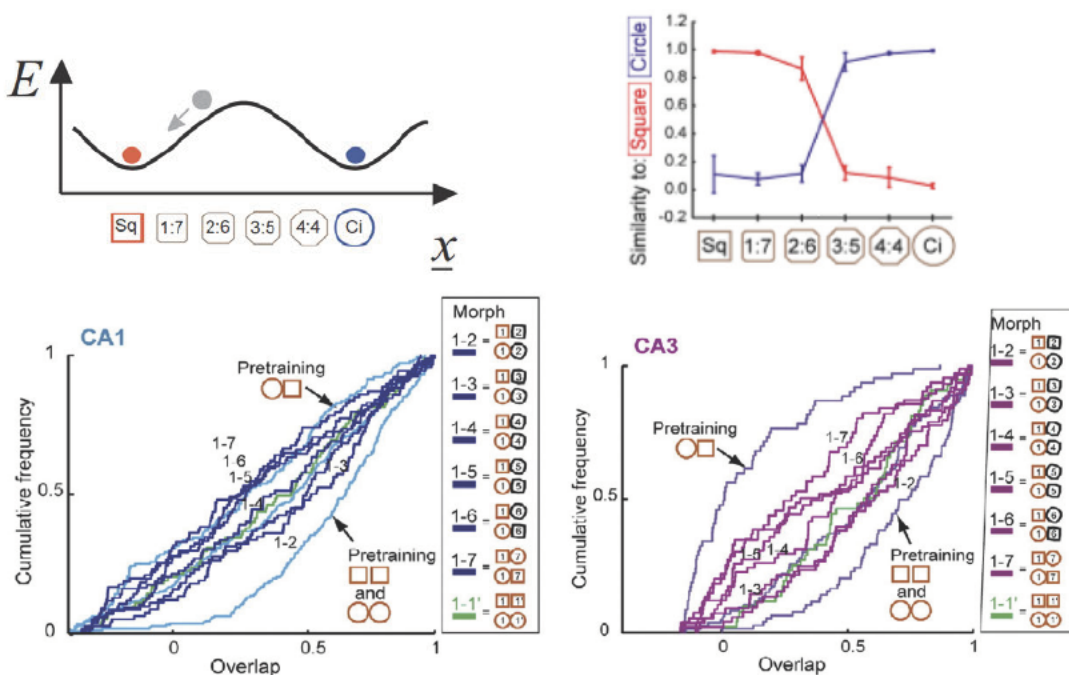
Figure 1-7 Progressive and step changes in ensemble spatial representations in the hippocampus.

Top left) Schematic illustration of the landscape of 'energy space' of an attractor network. Different hippocampal representations of the square and circle environment (below plot) are due to separate attractors for each shape Representations lie at the 'basin of attraction' (see red and blue balls). Representations of intermediate shapes (grey) evolve into the attractor basin of either red or blue shape (indicated by the grey ball 'rolling' down). Top right) Parametric shape changes to the animals' environment did not affect place cell activity until a threshold, after which entire network moved state to represent the alternative configuration. Place cell remapping was assessed using Pearson's r correlation over corresponding pairs of visited bins. A cell was considered stably remapped if the mean Square-Circle correlation was 0.5 < mean Square-Square and Circle-Circle correlation. Adapted from Wills et al., (2005). Bottom left) CA1 and CA3 (bottom right) place cell spatial representations vary progressively along a continuum to represent intermediate environments. Cumulative distribution functions for overlap scores in CA1 (bottom left) and CA3 (bottom right) do not show a sigmoidal function in response to morphing environments (0-1 square to circle overlap). Adapted from Leutgeb et al., (2005).

On contrast, Leutgeb et al. found that network inputs from environmental changes that varied progressively, did not result in sharp and coherent transition between pre-established (learned) representations (Leutgeb et al., 2005). Rather, correlations between encoded patterns led to linearized responses along the morphing sequence (Figure 1.7 bottom; Leutgeb et al., 2005). These linear and non-linear hippocampal responses to graded inputs can be reconciled by crucial differences between the two experiments. Namely, the degree of orthogonalization of the spatial representations originally established (whether complete remapping took place or not; Leutgeb et al., 2005). Leutgeb et al.'s morphing procedure was incremental, whereas Will's used a semi-scrambled sequence (Leutgeb et al., 2005; Wills, 2005). Pre-established representations might be more prone to plastic modification when the environment is changed gradually (Barry et al., 2006), such that scrambled rapid discrimination and testing performed by Wills et al. could have countered the cumulative effect of gradual synaptic changes (Leutgeb et al., 2005). Similarly, the differences in the training procedure likely also resulted in different prior beliefs about the frequency of 'intermediate' world presentation, and subsequently distinct CA1 remapping patterns and is discussed later in Section 4 (Plitt & Giocomo, 2021).

1.6 Recent methodological advances

1.6.1 Optophysiology in CA1

Unambiguously identifying the determinants of the different remapping types will provide new insights into how the hippocampus generates, stores and recalls multiple internal representations. A further restriction of earlier electrophysiological studies is their limited cell population size. To better understand emergent dynamics, we need to study the behaviour of more neurons (Yuste, 2015). Advances in optophysiology utilises the correlation between neural activity and changes in intracellular free calcium (Smetters, Majewska, & Yuste, 1999; Yuste & Katz, 1991; Baker, Hodgkin, & Ridgway, 1971). Two photon calcium imaging (2P imaging), combined with genetically encoded calcium indicators (GECIs) such as GCaMP, a single

fluorophore sensor used as a calcium indicator protein, can track the activity of neuronal populations of specific cell types (Chen et al., 2013).

To understand the conditions underlying neural remapping behaviour, we need to carefully titrate inputs to the hippocampal subregions though better control of the animal's external environments. Dombeck and colleges developed a virtual reality behavioural paradigm to allow head-fixed navigation along a 2D virtual corridor and were the first to optically identify place cells (see Chapter 2 for more details; Dombeck et al., 2010; Harvey et al., 2009).

1.6.2 Precise virtual environment manipulation

Another experimental advantage to navigating in a virtual world, is the ability to precisely control the environment of the animal. Consequently, VR opens the possibility to make parametric world changes and observe the neural responses. Recently, Plitt & Giocomo (2019) designed virtual worlds such that the dominant features of the world were morphed between two end points and recorded from thousands of CA1 pyramidal cells. They altered the frequency in which animals experienced the intermediate worlds, to assess how prior beliefs about stimulus frequency affects remapping of CA1 representations (Plitt & Giocomo, 2021). They found CA1 cells recorded in the infrequent condition maintained their representation until a threshold morph value and then coherently switched to a different representation, forming two discrete representations (Plitt & Giocomo, 2021). Whereas, CA1 cells recorded in the frequent condition formed a spectrum of representations for different morph values (Plitt & Giocomo, 2021). These results indicate that CA1 remapping represents an optimal estimate of the animal's location in a multidimensional stimulus space, where environment representations are activated proportional to their probability (Plitt & Giocomo, 2021). Hence, hippocampal circuits represent location in physical space but also in abstract space where remapping might be a generalisation of this phenomenon (Behrens et al., 2018; Plitt & Giocomo, 2021). These results also offer a unifying explanation for the work mentioned earlier using morphed real environments (Leutgeb et al., 2005; Wills et al., 2005). Whether dynamics emerge in CA1 or are inherited from an earlier stage

of processing remains unknown (Plitt & Giocomo, 2021). To fully understand hippocampal processing and where such dynamics originate, we need to access to large populations of neurons from other hippocampal regions.

Recently a similar surgical procedure to the one developed in this thesis was published, enabling 2P optical access for the first time to subregion CA3, along with the DG (Hainmueller & Bartos, 2018). During virtual navigation of familiar and novel worlds, Hainmueller and Bartos (2018) used 2P imaging to record from CA1, CA2/3 and DG populations. They found, surprisingly, that CA1-3 showed context-specific representations and stronger remapping than DG granules cells. While this contrasts with the notion discussed earlier, in which DG neurons mediate pattern separation, amplifying small context differences (Allera, Posani, & Schmidt-hieber, 2019), there have been reports of a dissociation between mature and young DG granule cells, arbitrating pattern completion and separation, respectively (Nakashiba et al., 2012). These studies pave the way to for investigation of previously understudied hippocampal regions. Hence, combined VR and 2P imaging offers the opportunity to precisely assess whether the CA3 network is governed by nonlinear dynamic responses to input patterns. Therefore, the one of the aims of this thesis is to address this, by imaging genetically defined CA3 neurons during virtual navigation in familiar and novel environments, morphed along a 5-environment continuum. We also image adjacent processing areas simultaneously to directly compare CA1 and CA3 processing within the same animal.

1.6.3 Simultaneous optical recording and manipulation in vivo

As mentioned, a methodological goal in neuroscience to provide conclusive evidence of a link between distributed neuronal ensemble activity and behaviour. The need to selectively manipulate specific groups of neurons was noted decades ago (Crick 1979). However, with the advent of optogenetic tools, this has become a reality and today light is widely used to perturb neuronal activity by targeted expression of opsins in genetically identified neurons (Yizhar *et al.* 2011). Microbial opsins confer light sensitivity to some species of bacteria or algae (Boyden *et al.*

2005). By harnessing the tools of molecular biology, we can deliver these genes to specific subclasses of cells allowing the experimenter to activate, or silence, populations of excitatory, inhibitory, or neuromodulatory neurons in specific layers or regions of the brain, a method now known as optogenetics (Boyden et al., 2005; Luo *et al.* 2018).

While Hebb's postulate has been difficult to directly test at single cell resolution in awake behaving animals, 2P microscopy and 2P optogenetic advancements have begun to open up these experiments (Boyden et al., 2005; Packer, Russell, Dalgleish, & Häusser, 2014; Rickgauer, Deisseroth, & Tank, 2014, Yuste 2015). The ability to perturb the circuit and record the local network response is necessary for causal conclusions (Packer et al., 2015). An 'all optical' strategy would support such conclusions, by combining simultaneous 2P imaging (readout) and 2P optogenetics (manipulation; Carrillo-Reid et al., 2016a; Emiliani, Cohen, Deisseroth, & Hausser, 2012; Packer et al., 2015; Rickgauer et al., 2014; Zhang, Russell, Packer, Gauld, & Häusser, 2018). It requires two laser sources, paired with two probes with minimal cross-talk, to allow targeted manipulation of functionally defined cell populations (Emiliani et al., 2012; Packer et al., 2015) Combined, these developments facilitate the appropriate level of spatiotemporal control over circuit manipulations to provide detailed insights into the precise neural computations that support network phenomena such as pattern completion and associated behaviours (Emiliani et al., 2012; Packer et al., 2015; Yuste, 2015).
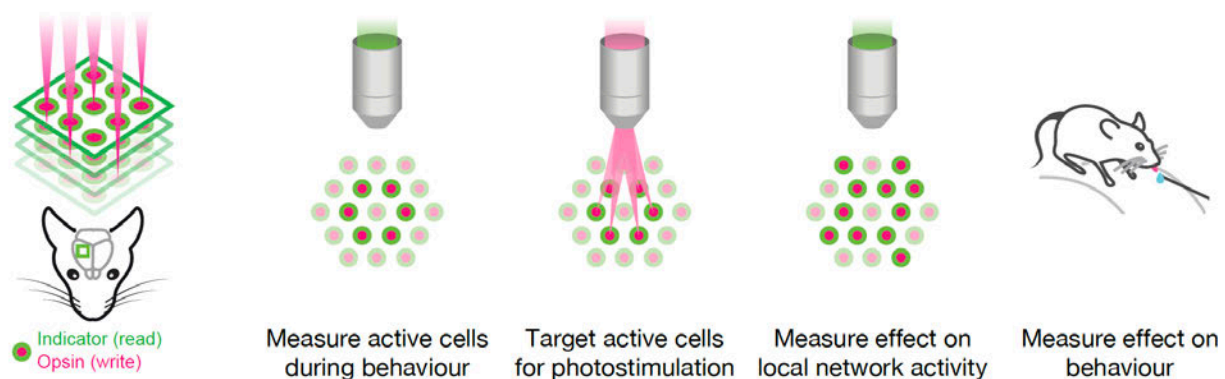
Figure 1-8 Targeted optogenetics closes the loop between neural activity and behaviour.
A schematic illustrating how targeting optogenetics can be used to probe the neural code in brain circuits. The figure highlights the close interplay between behavioural experiments, optical readout of patterns of activity, and replay of the same patterns in behaviourally relevant neurons using optogenetics. Hence, targeted optogenetics allows the precise membership of the neuronal ensemble, to be tested directly to investigate their importance for the neural code driving the behaviour.

An all optical strategy has been used to probe neural circuits in vivo. Rickgauer and colleagues simultaneously measured and manipulated local hippocampal CA1 circuits in vivo, during virtual spatial behaviour (Rickgauer Diesseroth & Tank, 2014). Spatially defined activity, mirroring natural place field activity, was 'imposed' by stimulating single place cells or silent cells (cells with no spatially tuned firing) when the mouse traversed a virtual location, in the case of place cells, a location that preceded the cells natural firing field (Rickgauer et al., 2014). Furthermore, a similar all optical approach has also been used in layer 2/3 visual cortex to assess functional connectivity. Carilio-Reid et al. demonstrated single-cell photostimulation, after population photostimulation training (photostimulating the entire field of view; FOV; ~ 60 neurons), could reliably recall a specific group of neurons not coactive previously, and that these 'imprinted' ensembles reoccurred spontaneously over consecutive days (Carrillo-Reid, Yang, Bando, Peterka, & Yuste, 2016b). Therefore, photostimulation can drive neuronal activity to create artificial ensembles with similar properties to those occurring naturally (Carrillo-Reid et al., 2016b). While remarkable examples of manipulating neural dynamics, these studies are limited by the light delivery being restricted to single cells or large populations (Packer at al., 2015). To circumvent

such limitations, a spatial light modulator (SLM) can be incorporated, to split the photostimulation laser into multiple beamlets (Packer et al., 2015). These beamlets can be individually directed to target multiple neurons for simultaneous photostimulation (Emiliani et al., 2012; Packer et al., 2015; Zhang et al., 2018).

1.6.4 Functional manipulations at single cell resolution in vivo

Recent investigations have demonstrated a link between neural ensembles and behaviour in visual cortex, providing the first causal evidence that ensembles implement internal representations of perceptual states (Carrillo-Reid, Han, Yang, Akrouh, & Yuste, 2019; Carrillo-Reid & Yuste, 2020, Marshel, et al., 2019; Russell et al., 2019). Targeted activation of behaviourally relevant ensembles in layer 2/3 of mouse primary visual cortex improved performance on a go-nogo task and activating nonselective neurons decreased behavioural performance, demonstrating that neuronal ensembles can control a learned behaviour (Carrillo-Reid et al., 2019). Similar results were found in the hippocampus in our own lab: Robinson et al., (2020), used targeted optogenetics and 2P imaging to selectively activate place cells encoding behaviourally relevant locations in a VR environment. Robinson et al. found that some place cells remap or shift their fields, and this was sufficient to bias licking behaviour during a spatial memory task (Robinson et al, 2020). This provides the first evidence of direct influence of place cell activity in supporting spatial navigation behaviour (Robinson et al, 2020).

These 'all optical' methods have recently been developed further to enable deeper imaging over multiple planes, making it possible to perform volume photo-stimulation and volume imaging across cortical layers (Marshel, Kim, Machado, Quirin, Benson, et al., 2019; Russell et al., 2019). In part these advancements were made possible by new more sensitive, somatically targeted opsins, that have reduced concerns over heating/damage and off-target effects when accessing deeper areas (Podgorski & Ranganathan 2016, Picot *et al.* 2018). Russell et al. demonstrated that the behavioural effect of targeted activation of cortical ensembles depends on their functional identity, such that ensembles that normally represent the stimulus have the most potent effect,

either boosting or inhibiting detection behaviour depending on task difficulty (Russell et al., 2019). Marshall et al. employed a similar method and showed that optogenetic targeting of orientation-selective ensembles elicited correct behavioural discrimination (Marshel et al., 2019). They also found cortical layer–specific population dynamics emerging after optogenetic stimulation correctly predicted behaviour (Marshel et al., 2019). Collectively, these studies demonstrate that it is now possible map cell numbers, layers, and network dynamics, via the precise control of naturally occurring, widely distributed neuronal ensemble elements corresponding to neural representations (Marshal et al., 2019; Russell et al., 2019). These recent all-optical experiments signal that we are moving closer to the prospect of recording thousands of cells with single cell and single spike resolution while being able to stimulate with the same spatial and temporal resolution.

Such technological advancements have made dream experiments of 'playing in' patterns of recorded activity in arbitrary spatiotemporal patterns while observing the behaviour a reality. Hence it is logical to consider extending these new methods to investigate deep structures like hippocampal subregion CA3. One could imagine an experiment where volume imaging and targeted stimulation enables access to neurons in CA3 ensembles, making it possible to directly probe the pattern completion properties of CA3 and simultaneously 'read out' downstream CA1 activity. In addition, these methods would enable precise manipulation of spatial representations in CA3, such that eliciting an associated learned spatial behaviour could enable circuit dissection of conditions required to move from one attractor state to another. Work presented in this thesis has taken the first few steps towards this greater goal.

1.7 Aims of this thesis

This thesis will describe the development of a surgical protocol to optically access hippocampal subregion CA3, with CA1 remaining intact (Chapter 2). We then optimise a genetic approach to selectively express an activity indicator in a genetically defined population of CA3 pyramidal neurons (Chapter 2). Following this, we combine the latest optical manipulation advancements

to develop the first all optical strategy to functionally activate neurons in vivo in CA3 (Chapter 2) and use this method to conduct a proof-of-principle all optical experiment to assess functional connectivity in CA1 and CA3 networks in vivo (Chapter 2). We also developed a new miniaturized VR hardware and parametrically morphed VR world designs to enable concurrent 2P imaging and virtual navigation (Chapter 3). We then explore attractor dynamics in subregions CA1 and CA3 by performing simultaneous 3D imaging in both areas while parametrically manipulating the virtual environments during navigation (Chapter 4). Finally, we adapt the VR task to include a behavioural decision conditional upon different worlds to assess the extent to which neural activity matches spatial behaviour (Chapter 5).

2

Materials and Methods

2.1 Animals

All surgical procedures were carried out under license from the UK Home Office in accordance with the Animal (Scientific Procedures) Act 1986.

For imaging experiments:

We used transgenic GCaMP6s mice (Emx1-Cre;CaMKIIa-tTA;Ai94 (Madisen *et al.* 2015) of both sexes aged between P40 and P73. Doxycycline treatment in drinking water from birth to P49 prevented interictal activity in the Ai94 mouse line (Steinmetz *et al.* 2017). These mice expressed GCaMP6s in all excitatory neurons and therefore we used them as our CA1 and 'dual area' mice. We also used transgenic GCaMP6s (GRIK4-Cre;CaMKIIa-tTa;Ai94) of both sexes (P48 - P65) to selectively express GCaMP6s in CA3 pyramidal neurons (Nakawaza et al., 2002) as our 'CA3 mice' (see Chapter 2 for more details).

For 'all-optical iterrogation' experiments:

We used transgenic GCaMP6s mice (Emx1-Cre;CaMKIIa-tTA;Ai94 (Madisen *et al.* 2015) as described above and injected them with ChrMine (Table 1; see Section 5.2; Marshel et al., 2019) to co express opsin in GCaMP positive neurons.

Mice were housed on a 12-h light–dark cycle in groups of 2–4 mice. No statistical methods were used to predetermine sample size. Experiments were not randomized, and investigators were not blinded to allocation during experiments and data collection.

Table 1 Calcium activity indicators and red shifted opsin dilutions.

Various calcium activity indicator and opsins used in the current experiments and their dilutions with sterile buffer. Precise dilutions are listed.

Activity indicators	Dilution	Red Shifted Opsins	Dilution
AAV1-Syn-GCaMP6f-WPRE-SV40	1:10	AAVdj-CaMKIIa-C1V1-TS-P2A-mCherry-WPRE	1:10 1:8
AAV1-CAG.Flex-GCaMP6s WPRE-SV40	1:20 1:12	AAV9-CAG-DIO-ChroME-mRuby3-WPRE-SV40	1:30 1:40
AAV1-Syn-GCaMP6s-WPRE-SV40	1:20 1:16	AAV9-CaMKII-ChrimsonR-Kv2.1-mScarlet	1:10 1:20
pGP-AAV1-Syn-GCaMP7f-WPRE	1:20	AAV2/9-CaMKII-C1V1-Kv2.1-mScarlet	1:10 1:8
Tg mouse: Emx1-Cre;CaMKIIa-tTA;Ai94-GCaMP	-	AAV8.CAMKIIa.ChRmine.mScarlet.kv2.1.WPRE	1:100
AAV8-CaMKIIa-GCaMP6m-p2a-ChRmine-TS-Kv2.1			1:10

2.2 Viral injections and head plate installation

All surgical procedures were performed in a stereotactic apparatus under anaesthesia with ~2% isoflurane and analgesia using 0.1 mg kg^{-1} buprenorphine. A small (~0.5–1mm diameter) craniotomy was made over the hippocampus and 1μl of mixed GCaMP and opsin (see Table 1 for dilutions, exact GCaMP/opsin combinations used are listed in figures); were injected into CA1 (bregma; A/P −2.0mm; M/L 1.4mm; D/V −1.2mm) or CA3 (bregma; A/P −1.8mm; M/L 1.9mm; D/V −1.8mm) and a rate of 100nl per minute. In the same surgery session, the animal's scalp was removed bilaterally from the midline to the temporalis muscles and they were implanted with a stainless-steel head plate (with a 5mm circular hole), fixed to the skull with dental cement (Super-Bond C&B, Sun-Medical) and centered on the injection site. The head plate was oriented horizontally for CA1 imaging implantations and with ~5-15° lateral angle for CA3 implantations.

Mice recovered from surgery for at least 7 days before the next surgery. Postoperative analgesic treatment was continued with administration of 0.7µl ml^{-1} Rimadyl for 3 days.

2.3 Imaging window implantation

Cortical excavation and imaging window implantation were performed more than 7 days after the initial virus injection, according to published protocols (Dombeck et al., 2010). A craniotomy (diameter 3mm) was made cantered on the previous injection site. Parts of the somatosensory cortex and posterior parietal association cortex for CA3 access, and parietal cortex and parts of visual and hindlimb sensory cortex for CA1 access, were gently aspirated, while being irrigated with chilled saline solution. Aspiration continued until the external capsule was exposed. The outer part of the external capsule was gently peeled away, leaving the inner capsule and the hippocampus undamaged. The imaging window implant consisted of a 3-mm diameter coverslip (Coopers Needle Works) glued to the bottom of a stainless-steel cannula (1.2–1.9-mm height). The window was lowered into the craniotomy until the glass was in contact with the external capsule and then fixed to the skull using dental cement (Super-Bond C&B, Sun-Medical). Mice recovered from window implantation for 14 days with post-operative administration of 0.7µl ml^{-1} Rimadyl for 3 days. During imaging, the objective was tilted a similar amount (~0-20°), so the long-axis was perpendicular to the imaging window.

2.4 Building VR worlds for use with 2P microscopy

We used a virtual reality MATLAB engine (ViRMEN; Aronov & Tank, 2014) to design and run virtual worlds (Figure 3.1). The VR world was projected onto a mirrored arc which reflects this projection into a plastic dome, surrounding the mouse and providing an immersive 270° view of the virtual world. The mouse was head-fixed on a cylindrical treadmill (20cm diameter Styrofoam ball) floating on a column of air flowing through a custom designed, treadmill base to reduce friction during movement (Figure 3.1 A). Mouse movement was tracked using an incremental rotary encoder (Kuebler) to read forward ball speed. ViRMEN was used to render a closed loop visual

61

virtual reality environment based on movement input from the mouse running on the ball. Rewards were controlled by a TTL output from the virtual reality engine and were given at the end of the VR track. Each reward consisted of 1μl of water delivered via a custom made lickometer with a piezo-electric sensor. This behavioural set up enables head-fixed navigation on a virtual track enabling concurrent 2P microscopy (Figure 3.1; videos).

2.5 Morphed 2D environments

We designed 5 virtual 2D environments or worlds that included visual stimuli chosen so that the extremes of the stimulus distributions could be gradually and convincingly morphed together (Figure 4.2). The tracks did not change in length or the presence of salient landmarks. The aspects of the stimulus that did change were, i) the frequency and orientation of object textures (i.e., low frequency, oriented stripes to high frequency vertical stripes), ii) the non-monochrome colours of the objects (RGB values from green to blue), and iii) the colour of the background of the visual scene (green to blue), iv) x-y positioning of objects outside the main track corridor (Figure 4.2). Reward location was indicated with a grey and white star on Animals received a reward if they licked 3 or more times within this 30cm reward zone. However, for discrimination tasks, animals must lick 6 or more times in the 'correct' reward zone for that trial to initiate a reward (described below).

2.6 Behavioural training & non-discrimination task

Fourteen days post aspiration, we placed mice in the virtual environment for 10–30 minutes daily, with gradually increasing timespans. After 7-10 days of habituation, mice were placed on water restriction, maintained at 80% initial weight. Animal health and bodyweight was continually monitored to check for signs of distress/dehydration. Water rewards not obtained during a given testing session were given ad libitum post-session. During water restriction, behavioural training continued, and conditional rewards were obtained upon reward-related licking at the end of the track within the reward zone (500-520cm) marked by a white star cue. A

trial is defined as a single traversal of the virtual track from 0cm to >500cm. Correct trials resulted in a water reward conditional upon stopping in the reward zone. After reward collection, mice experienced a brief tone, 2s of darkness followed by virtual teleportation back to the beginning of the track for the next trial. Incorrect trials occurred when the mouse was stationary for >200 seconds or failed to reduce running speed <5cm/s in the reward zone. Incorrect trials were punished with a 2s 'whiteout', followed by virtual teleportation to the beginning of the track for a new trial. Mice were imaged whilst navigating once they completed 10 training session and a minimum of 20 correct trials in 30 minutes. Most mice trained in the green world (world 1) first and a control group (n = 5) trained in the blue world (world 5) first.

The experiment consisted of 6 consecutive days of imaging. Day 1: Imaging whilst navigating in world 1 with conditional rewards as described above. Days 2 and 3: included the presence of a novel, previously unseen environment (world 5) interleaved randomly for 2-6 trials, after approximately 10 'warm up' trials in the original training world. Days 4-6; included the addition of 3 unseen intermediate environments. The trial structure was the same as previously stated, except the world presented was chosen at random for 2-4 trials. If performance deteriorated, defined as 5 or more consecutive incorrect trials then a manual reward was given in the reward location, to maintain motivation. Well-trained mice ran until they received around 0.8 - 1 mL of liquid (~200 rewards).
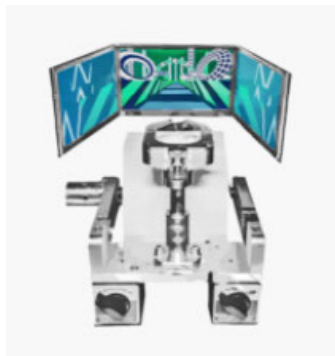
Figure 2-1 VR behavioural set up.

Photograph showing the VR set up including the Styrofoam running linear wheel, the 3-screen arrangement, and head mount. See Chapter 3 for further details.

2.7 Behavioural training & discrimination task

For discrimination experiments mice followed the same behavioural training protocol as described above, except they also experienced an additional world (world 5) during the training phase, such that world 5 and world 1 were randomly interleaved for 2-4 trials. Moreover, for the discrimination task, all worlds had 2 reward zones (see Figure 4.12 A & B). During training mice in the R1 condition learnt that the reward was present in world 1 in the first reward zone and present in reward zone 2 in world 5. For the animals in the R2 condition, they experienced the flipped contingency to counterbalance rewarded location across animals, such that, in world 1, the reward was in reward zone 2 and in world 5 the reward was in reward zone 1. After approximately 10 sessions with at least 3 sessions of consistently achieving ~70% success, mice were imaged while completing the discrimination task. The following day they were imaged while completing this discrimination task and additionally presented with 3 intermediate world 'probe' trials interleaved.

Probe trials involved the presentation of world 2, 3, and 4. In order to increase the number of trials in a session and avoid confounds, probe trials were unrewarded. Trial decisions were initiated by licking 6 or more times in one or the other reward locations. In the end point worlds (worlds 1 or 5), licks were either incorrect or correct, depending on the condition the mouse was trained in. Correct trials resulted in reward delivery, a blackout and teleportation to the start of the track for the next trial. Incorrect trials involved a whiteout (as described above) and teleportation for the next trial. For probe trials (worlds 2, 3, 4), there was no incorrect or correct responses per se, rather a decision was assumed when 6 or more licks were given in one or the other reward zones. This decision was not rewarded but did result in a 0.5s blackout and teleportation to the beginning of the track for the next trial. No decision was made if the animals ran off the end of the track, this resulted in a whiteout and teleportation to the beginning for the next trial. Additionally, if animals stopped in or out of the rewarded zones for more than 5

seconds, this too resulted in a whiteout and teleportation to the beginning of the track for the next trial. Probe trials comprised 16.5% of the total trial types. We pre-calculated the trial orders to ensure this trial frequency was presented.

2.8 Two photon imaging

Neural activity correlates with changes in intracellular free calcium (Helmchen, Imoto & Sakmann, 1996; Smetters, Majewska, & Yuste, 1999; Yuste & Katz, 1991; 'Baker, Hodgkin, & Ridgway, 1971). Calcium imaging utilises this property to track the activity of neuronal populations, combined with genetically encoded calcium indicators (GECIs). We performed two photon imaging of GCaMP6s dynamics with a resonant scanning microscope (Ultima II, Bruker Corporation) using a Chameleon Ultra II laser (Coherent) driven by PrairieView. A 16×/0.8-NA water-immersion objective (Nikon) was used for all experiments. GCaMP was imaged at an excitation wavelength of 930 nm (70-100mW average power at the sample depending on plane depth), and mRuby3, or mScarlett (conjugated to the opsin) was imaged at 765nm (~50mW average power at the sample). For some experiments, the imaging path was coupled with an electrically tunable lens (ETL; Optotune EL-10-30-TC, Gardasoft driver) to allow high speed volumetric imaging, spanning a 100µm range with 33.3 µm spacing between planes. Images (512 × 512 pixels,~800 × 800 µm FOV) were acquired at 29.9Hz for a single plane, and ~ 7Hz for 4 planes (~ 600×600µm FOV). An orbital nose piece was used to maximise imaging quality (Aliñanes et al., 2018) by calculating the tilt of the sample, relative to the microscope, and rotating the objective to be perpendicular to the implanted imaging window (Figure 6.1 A).
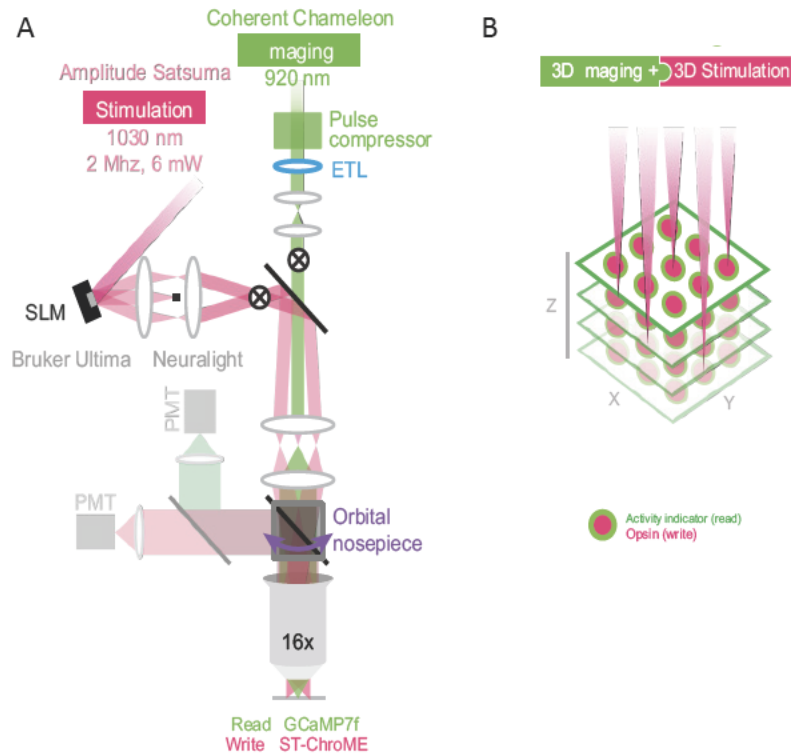
Figure 2-2 Three dimensional (3D) all optical microscope design and functionality.

A) Simplified optical layout of the SLM-based two-photon patterned photostimulation (See Figure 1 B in Packer et al., 2015 for more details), two-photon resonant scanning, in vivo microscope. For a full detailed diagram (see Packer et al., 2015; Russell et al 2019). PMT1 (pink), PMT2 (green), photomultiplier tubes. ETL, electrically tunable lens. B) Schematic of the methodological goal. A calcium sensor generates an optical readout of neuronal activity (green; read), and a somatically restricted opsin enables targeted photostimulation (pink; write). Hence, single cell resolution photostimulation of user-selected neurons on different Z planes enables simultaneous imaging and stimulation in a 3D volume of brain in vivo.

2.9 Two photon optogenetic stimulation

Two-photon photostimulation was carried out using a fibre laser at 1030nm (Satsuma, Amplitude Systemes, 2 MHz rep rate). The laser beam was split using a reflective SLM (7.68 x 7.68mm active area, 512 x 512 pixels, OverDrive Plus SLM, Meadowlark Optics/Boulder Nonlinear Systems) which was installed in-line of the photostimulation path (Neuralight, Bruker Corporation; Fig. 3). The SLM enabled flexible 3D beam splitting to target user defined positions in the sample, across a FOV volume of ~600 × 600 × 100μm (Figure 6.1 B). Phase masks used to generate focused

66

beamlet patterns in the sample were calculated via the weighted Gerchberg-Saxton algorithm (Zhang, Russell, Packer, Gauld, & Häusser, 2018). The targets were weighted according to their location relative to the center of the SLM's FOV to compensate for the decrease in diffraction efficiency when directing beamlets to peripheral positions. We calibrated the targeting of SLM spots in imaging space by burning arbitrary patterns using the photostimulation laser in a fluorescent plastic slide, before taking a volumetric stack of the sample with the imaging laser. We manually located the burnt spots and the corresponding affine transformation from SLM space to imaging space was computed. For 3D stimulation patterns, we interpolated the transformation required from the nearest calibrated planes. Spiral photostimulation patterns (three rotations, 20μm diameter) were generated by moving all beamlets simultaneously with a pair of galvanometer mirrors conjugate to the SLM plane, to spiral scan the focus over the cell body of each neuron targeted (Carrillo-Reid et al., 2016; Packer et al., 2015; Zhang et al., 2018). Power on sample for photostimulation was kept at 6mW per cell and evenly distributed across beamlets generated by the SLM (Packer et al., 2015). Imaging and photostimulation parameters were chosen to give sufficient imaging quality, as imaging at higher power or lower scan speeds and/or smaller FOVs could lead to photoactivation (Packer et al., 2015; Zhang et al., 2018).

2.10 Subregion boundary mapping

To image CA1 and CA3 in the same FOV, we manually drew masks around what we judged to be subregion CA3 (see Figure 4.1; ROIs coloured according to ROI size) using FIJI and custom scripts. Masked ROIs were considered CA3 and we discarded ROIs 100μm from the mask boundary to determine the remaining ROIs as CA1

2.11 Photoresponsivity mapping

To find photostimulation-responsive cells, we semi-automatically detected cell locations from expression images and stimulus-triggered average or pixel-correlation images. We then used a

novel high throughput photomapping protocol we devised to arranged in a 13x 13 pixel grid of target centroids over the image (Figure 4.2). We then grouped target sites into patterns of 10-30 sites per patterns (depending on the total number of target sites making up the grid). Each pattern was numbered 1:max pattern number and randomly targeted with 6mW of spiralled photostimulation simultaneously at a stimulation rate of 20Hz, for 500ms for each pattern, and performed 8 trials (Russell et al., 2019; Figure 4.2). The phase mask, galvanometer positioning and Pockels cell control protocol were generated using custom MATLAB software (NAPARM, available: https://github.com/llerussell/Naparm) and performed by the photostimulation modules of the microscope software (PrairieView, Bruker Corporation) and SLM control software (Blink, Meadowlark). Photostimulation trials for the photoresponsivity mapping were triggered at a fixed rate from an output line on a DAQ card (National Instruments). For experiments, photostimulation trials were triggered manually through the stimulation software.

2.12 Photoresponse metric

To increase speed of data analysis immediately prior to the experiment we streamed the raw acquisition samples to custom software, that performed online registration and motion correction. To measure neuronal responses we extracted the mean fluorescence in a ~500ms window (4 frames) starting immediately after the photostimulation ended (due to the associated light artefact contaminating activity in those traces) and subtracted the mean fluorescence in the ~ 1s baseline (7 frames) before photostimulation (Russell et al., 2019). We then divided this response by the std in that 1s of baseline to give a signal to noise metric. If on a given trial, for a given cell, this value was > 1, the response was scored as a success (Russell et al., 2019). Photostim cells were those that showed successful activation on > 50% trials.

2.13 Photostim ROI exclusion zones

To reduce off-target photostimulation artefacts we excluded cells within a 30μm diameter cylinder (extending through the 4 axial planes) when analysing the network response to

68

photostimulation (Russell et al., 2019). We updated the target stimulation patterns based on the ROIs segmented by Suite2P within this 30µm lateral disk around each SLM target locations (Russell et al., 2019).

2.14 All optical input-output experiment

To select cells for photostimulation we first took a volumetric (4 plane) stack of our FOV which encompassed CA1 and CA3 and then manually selected all ROIs. We then performed a stimulation mapping experiment to filter the list of all ROIs to just the photoresponsive neurons (See Section 5.11). We then programmatically designed groups of random 1, 5, 10 and 20 neurons spatially clustered within 200um of each other. We selected cells on the bottom 2 planes of our volumetric stack to bias towards CA3 only patterns and the top 2 planes to bias towards CA1 only patterns. We selected a total of 6 patterns for each number of cells group in each FOV. During the experiment we stimulated each of the patterns with 6mW per cell at a rate of 20Hz for 500ms. We stimulated the next pattern after 5 seconds. The stimulation patterns were delivered in randomised order. In each recording block we repeated each pattern 10 times and repeated the recording blocks 3-5 times to collect 30-50 trials per each stimulation pattern. We post hoc filtered the stimulation by double-checking that the target cells responded (see Section 5.12) and additionally confirmed that the stimulated cells resided in the intended region (either CA1, or CA3, but not both) by referring to a post-hoc labelled area boundary image (see Section 5.10). If a stimulation pattern was made up of less than 80% of cells in one subregion, it was discarded as a contaminated pattern. Pattens were also discarded if the intended targets did not show significant responses to photostimulation.

2.15 Pre-processing: Imaging frame registration, ROI segmentation and neuropil correction

Analysis was performed using ImageJ and custom scripts written in MATLAB (2019). Raw imaging data was pre-processed using Suite2P (available at https://github.com/cortex-lab/Suite2P).

69

Briefly, Suite2P spatially registers all frames in the movie to each other, registration was sped up 4-fold using a NVIDIA Quadro M5000 Graphical User Interface. Next, a clustering algorithm detects regions of interest (ROI) within the image. ROIs were manually curated through an interactive GUI to discard spurious ROIs. Suite2P also provides a neuropil estimation, as the fluorescence signals recorded can be contaminated by signals from another cell, or neurites (axons and dendrites). Given each pixel's fluorescence trace is a sum of the ROI fluorescence plus the neuropil signal at that location, neuropil was subtracted from the fluorescence trace of each ROI. Change in fluorescence over average baseline fluorescence (ΔF/F) was calculated for each ROI. Baseline fluorescence activity was calculated by averaging activity across the trace and removing the lowest 20%. These processed calcium traces were used for subsequent analysis.

2.16   Spatial analysis: place field identification

Place cells were identified according to previously published methods (Plitt & Giocomo et al., 2021). Briefly, only periods where mouse running speed was >5cm/s were included. Sessions during which mice exceeded 20 trials in 30 minutes with 80% correct trials were used for spatial analysis. To determine place cell identity for a given cell, we created a null distribution of S.I. scores for each cell from their circularly permuted time series (we shifted the position relative to the activity trace within each trial). We shuffled 1000 times per ROI and ROIs exceeding 99% of permutations were considered to be significant 'place cells' (see Giocomo & Plitt, 2021 for details). Following the SI shuffle procedure we applied an additional criteria in that place cells must have atleast one place field. These fields were identified as follows: for each ROI, we detect candidate fields by constructing an average place map, smoothing that map with a gaussian filter of width 20cm and identifying regions of the track on which the ROI has activity >20% of its maximum response on that track (Dombeck et al., 2010). These initial candidate fields must be at least 50 cm wide. Note the criteria of 20% max response and using gaussian smoothed DF/F traces (not deconvolved) means this is a lenient criterion. We then refine the definition of the fields. We ask if those fields have events on individual trials by applying a threshold of 2 SD (computed during baseline periods). If the ROI activity level is above 2 SD in the field on a given

trial that trial is considered responsive. Place field reliability is defined as the proportion of trials showing significant responses. The magnitude of the field response is defined as the maximum response on those significant trials. We then compute the same reliability score on circularly permuted traces to build a null distribution using the activity statistics from the real data. We repeat this shuffle procedure 1000 times. We keep the place fields for each ROI that are associated with a reliability score greater than the 99th percentile of the individual shuffle distributions.

2.17   K-means analysis

In order to unbiasedly select the optimal population activity maps that represented the two familiar endpoint worlds (W1 and W5) we used K-means clustering to cluster all single trial place maps, across all worlds, into two clusters (see Low et al., 2020 for details). Briefly, we constructed single trial place maps for all cells by averaging $\Delta F/F$ values per each position bin. We then selected only those cells that displayed place modulated activity (place cells, in any world of the recorded session) and selected the middle portion of the track where the representation was most divergent across worlds (see Figure 4.11). This resulted in a 3D array of place maps where the dimensions of the array were cells by position by trial. All trials, in all 5 worlds were included. We min-max normalised all activity for each cell to range between 0 and 1. We then used the k-means clustering algorithm to cluster this array of trial maps into 2 clusters, along the 'trials' dimension. This resulted in two centroids in this 3D space that represented two average place maps (two arrays of cells by position) that by definition were maximally different from one another. These two cluster centroids allowed us to classify single trials as belonging to one of the two clusters, and thus specify which map a given trial was using. We renumbered clusters so that the most frequently used cluster map for W1 was map 1, allowing comparison across recordings. We then split the trials according to the real word identity (W1 through W5) and averaged the number of trials in a given world which were classified as using the optimal most frequent W1 map. For Figure 4.19, we repeated this same analysis but at the single cell level instead of the population, allowing each cell to have different 'optimal' W1 and W5 maps.

2.18   Spatial information score calculation

A spatial information (SI) score was calculated by adapting this measure from a traditional method of SI assessment typically used in electrophysiology work (Skaggs, McNaughton & Markus, 1993) and imaging (Hainmueller & Bartos, 2018). Here the mean ΔF/F for each bin was calculated and used to approximate spikes in that position. SI was then calculated using: $\sum_{i=1}^{N} \lambda_i ln \frac{\lambda_i}{\lambda} p_i$ where λ_i and p_i are the average ΔF/F and fraction of time spent in the i^{th} bin respectively. λ is defined the overall ΔF/F and N is the number of bins on the track. Hence, the amount of SI is inferred from differences between the calcium activity (reported as bits per event).

2.19   Data presentation & statistical analysis

No statistical methods were used to predetermine sample size. The experiments were not randomised and investigators were not blinded to allocation during experiments and outcome assessment. Data are presented as mean ± SD or S.E.M (see Figure legend) unless otherwise indicated. Statistical analyses were performed using independent two-sample t-tests or two-way ANOVAs (stated in text).

2.20   Histology and confocal microscopy

During the characterization of our transgenic mouse models and following the end of experiments, mice were deeply anaesthetized with ketamine and perfused transcardially with 4% paraformaldehyde (PFA) in phosphate-buffered saline (PBS). Collected brains were fixed by 4% PFA in 0.1M PBS and 24 hours later, processed using a vibratome in 100μm sections, mounted on slides, and coverslipped with Vectorshield mounting medium with DAPI (Vectorlabs). Sections were imaged using a Zeiss LSM 700 confocal microscope with a 20x

objective, using Zeiss Zen Black software. Processing of confocal images was performed with ImageJ (https://imagej.net/Fiji/)

3

CA3 'all optical' approach

3.1 Introduction

Memory encoding results from durable changes in the activity of synaptic circuits that lead to the storage of neural activity patterns in sparsely distributed neural ensembles (Hebb, 1949; Rebola et al., 2017). Early hippocampal anatomical studies and connectivity mapping revealed the unique rich internal architecture of the CA3 subfield (Amaral & Witter, 1989, 1995; Amaral et al., 1990; Ishizuka, Cowan, & Amaral, 1995; Amaral, 1993; Ishizuka et al., 1990; Witter, 2007). This subsequently motivated theorists to draw parallels between the artificial network architecture of recurrent memory models and connectivity in the CA3 circuit, resulting in the proposition that the hippocampal CA3 circuit architecture might function similarly (Knierim & Zhang, 2012; Treves & Rolls, 1992, 1994; Rolls, 2007, 2013; Rolls & Kesner, 2006 Samsonovich 2007).

3.1.1 Methodological limitations and limited exploration of CA3

Despite heavy theoretical interest in CA3, most experimental work has focused on hippocampal subregion CA1 (Rolls, 2007, 2013; Rolls & Kesner, 2006). Consequently, the functional properties of CA3 remain enigmatic (Delatour & Witter, 2002; Guzman et al., 2016). Much of the neural behaviour observed in CA1 is assumed to be inherited from earlier stages in the hippocampal processing loop, in particular CA3 inputs via the Schaffer collaterals (For more details see Chapter 1; Plitt & Giocomo, 2018; Leutgeb et al., 2004). If we are to make progress in understanding hippocampal function at a systems level, it is imperative that we are able to record from and ideally manipulate ensembles in other hippocampal areas beyond CA1 (Hainmueller & Bartos, 2020).

The sparsity in experimental investigations of CA3 is due to the relative difficulty in accessing this subregion and dissociating the subregion neuronal identity (Leutgeb et al., 2004; Leutgeb et al., 2005; O'Keefe & Dostrovsky, 1971). Electrophysiological studies recording from CA3 do exist, but are limited and involve significant intra hippocampal damage due to the insertion of a recording

electrode(s) from the dorsal surface of the brain, accessing a relatively small number of CA3 neurons, typically <40 (Gold & Kesner, 2005; Lee & Kesner, 2003; Leutgeb et al., 2004; Leutgeb et al., 2005; Leutgeb et al., 2007). Moreover, electrophysically recorded CA3 neurons are mostly from CA3c, given the curved hippocampal structure and difficulty separating neurons from neighbouring CA2. While genetic selection of CA3 cells exists, it has typically been used in selective loss of function studies (Nakazawa et al., 2002, Nakazawa et al., 2003). To capture dynamics in sparsely distributed ensembles, we need access to larger populations of genetically identified CA3 neurons.

3.1.2 Advancements in optics and genetics to target CA3

In addition to observing the CA3 neural population, it is favourable to manipulate defined neurons in CA3 and observe the response of the local network. Such 'reading and writing' of neural activity would permit causal conclusions and conclusive statements on CA3 function in vivo (Carrillo-Reid et al., 2019; Marshal et al., 2019; Packer et al., 2015; Russell et al., 2019; Zhang et al., 2018). Advancements in all-optical technology, particularly the advent of new red-shifted, soma-targeted opsins (Marshal et al., 2019), multiplane simultaneous photostimulation and recording enabling deep imaging (Marshal et al., 2019; Russell et al., 2019), suggest that all-optical control of CA3 neural populations is within reach (see Section 1.6.3).

At the time of our methodological developments, to the best of knowledge, there were no published 2P in vivo CA3 imaging studies investigating CA3 function. Therefore, we aimed to:
develop a transgenic CA3 mouse model, similar to the mouse models used in Nakazawa and colleagues (2002, 2003) and combine this with the surgical "excavation" techniques developed initially for CA1 optophysiology (Dombeck et al., 2010; Harvey et al., 2009). This combination enables 2P optical access to genetically defined CA3 pyramidal neuron populations in vivo.
We also aimed to characterise our 2P CA3 FOVs, by identifying consistent morphological and anatomical markers of CA3 to allow dual hippocampal area recordings, from subregion CA1 and CA3 within the same animal. Finally, we aimed to extend the recent optical manipulation

76

developments to CA3 and create the first 'all optical' CA3 mouse, enabling simultaneous 'reading and writing' of neural activity in CA3, opening the possibility to directly probe CA3 dynamics during behaviour.

3.2 Results

3.2.1 Developing a CA3 mouse model

To target expression of calcium indicator and opsin in CA3, we developed a mouse line based on the GRIK4 gene. GRIK4 is a gene coding for a glutamate receptor subunit of the kainate type. GRIK4 is expressed in specific populations of neuronal cells in the cerebellum and all layers of the frontal and parahippocampal cortices (Knight et al., 2012). In the hippocampus, strong GRIK4 expression is found in the stratum pyramidale and stratum lucidum of CA3, in cell processes in CA1 in polymorphic cells including mossy fiber neurons in the hilus, and DG granule cells (Knight et al., 2012; and demonstrated in our axial section in Figure 1.2). While high levels GRIK4 may play a role in behavioural phenotypes such as autism, it has been used to specifically identify subregion CA3 (Azahara et al., 2020; Knight et al., 2012; Nakazawa et al., 2002, Nakazawa et al., 2003).

To develop a CA3 mouse model in our laboratory, we restricted expression of GCaMP6s to CA3 pyramidal neurons by crossing a GRIK4-Cre driver line expressing cre-recombinase to CA3 under the GRIK4 promoter, with a cre- dependent GCaMP6s reporter line (Nakazawa et al., 2002). We characterised our mouse model and found GCaMP6s to be selectively expressed in CA3 (Figure 2.1, A) and sometimes CA3 and DG (Figure 2.1 B). Interestingly, we found expressing GCaMP6s under a CAMKII protomer yielded dense expression in CA3 pyramidal neurons, whereas our initial attempts expressing GCaMP6s under a Rose26 promoter yielded sparse CA3 expression. Given our experimental questions rely on recording from a large population of confirmed CA3 neurons, we sought it necessary to have our CA3 population densely labelled.

3.2.2 Optically accessing CA3

Next, we developed a novel surgical procedure to optically access CA3. We took inspiration from CA1 excavation surgeries and removed the cortex overlaying the hippocampus (Dombeck et al., 2010; Harvey et al., 2009). We chose a rostral cannula placement in dorsal hippocampus where CA3 cell population are closer to the dorsal surface of the brain (Figure 2.1 A). Our cortical excavation and cannula positioning enabled optical access to CA3 in 84% of animals (n = 270 attempted experimental CA3 mice). Our procedure leaves CA1 intact and also optically accessible (Figure 2.1 C). No intra-CA3 subregion specific genetic labelling was attempted, but given microscopy limitations (focal length), imaging was likely restricted to CA3a/b (Figure 2.1). We characterised our CA3 pyramidal neurons in vivo by two-photon volumetric imaging. The principal constituents of the CA3 recurrent network are typically defined by their anatomical signature, the presence of thorny excrescences (mossy fiber synapses; Sun et al., 2017; Gonzales, Galvan, Rangel, & Claiborne, 2001; Lauer & Senitz, 2006; Lorente De Nó, 1934), although not all CA3 pyramidal neurons share this morphology (Hunt et al., 2018). Closer inspection of our CA3 pyramidal neurons revealed these classic subcellular components (Figure 2.1 D, inset, white arrows), providing clear landmark features of CA3 (video 5). To increase the size of the population we had optical access to, we used an electronically tuneable lens (ETL) to change the axial focus of the imaging laser, allowing 3D volumes of tissue to be imaged (Russell et al., 2019). This microscopy modification enabled an additional 3 axial planes though CA3, extending the FOV up to 400μm deeper without needing to lower the objective (Figure 2.1 E). Using the ETL, we were able to observe the unique CA3 anatomical feature- stratum lucidum (shown in Figure 2.1 E, Plane 4, by the absence of cells).

79

A GCaMP6s DAPI 500μm

B CA1 ori. pyr. rad. CA2 luc. pyr. ori. m.l g.c.l. hil. CA3 DG DAPI Flex GCaMP6s L A P M

C CA1 GCaMP6s DAPI DG CA3 500μm

D GCaMP6s

E x z y Plane 1 GCaMP6s 200μm Plane 2 Plane 3 Plane 4

Plane 4 GCaMP6s Luc.
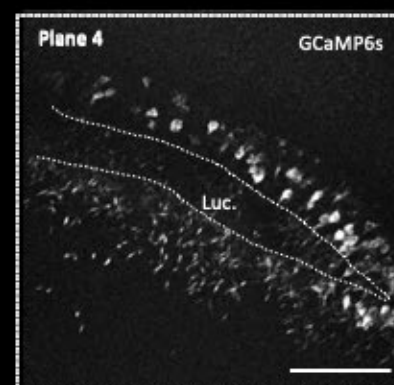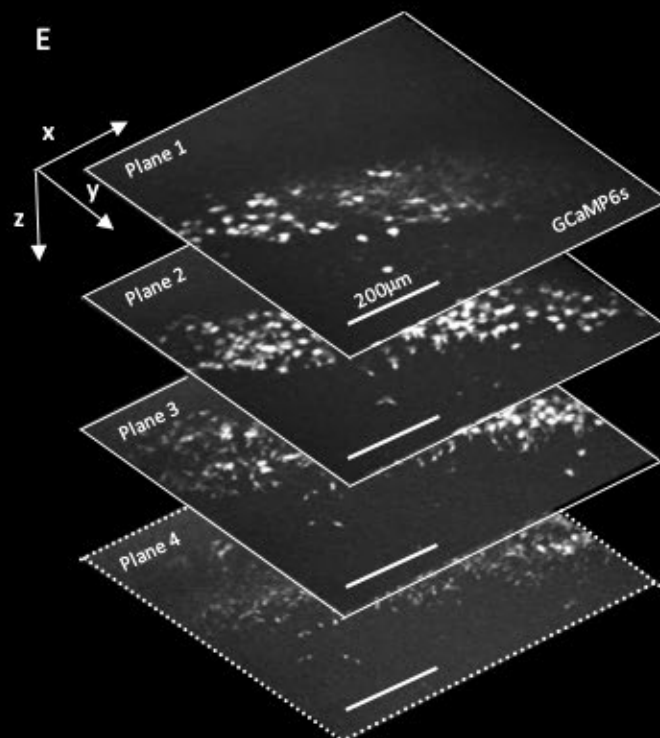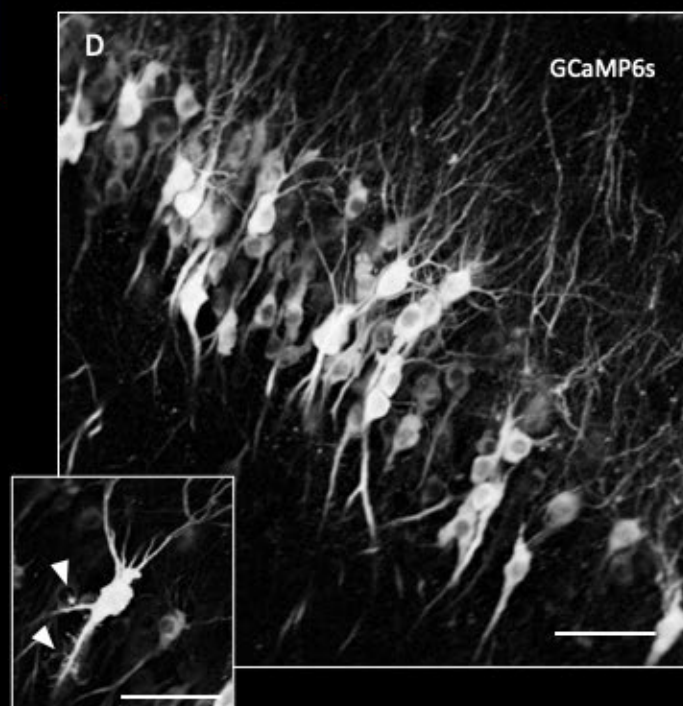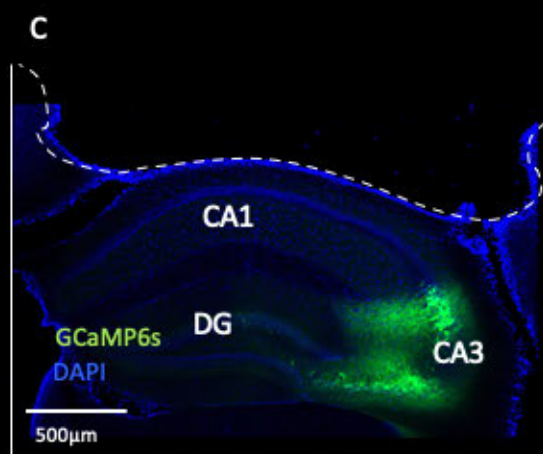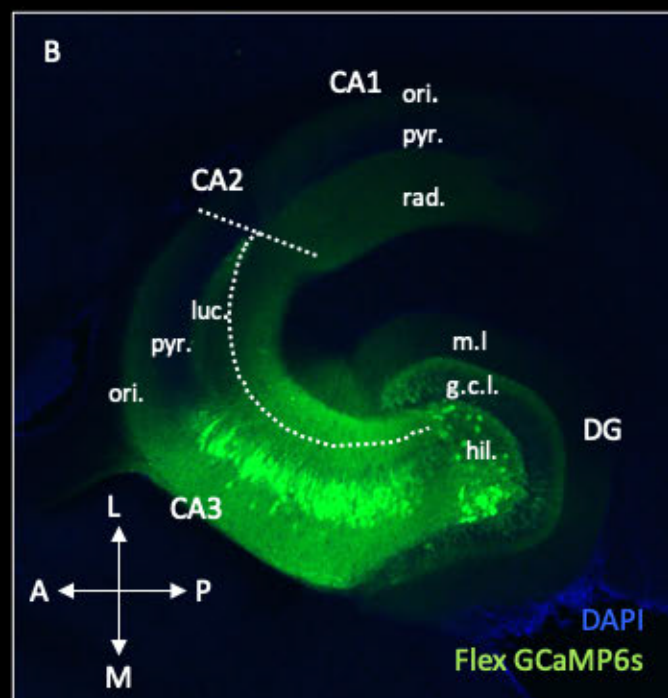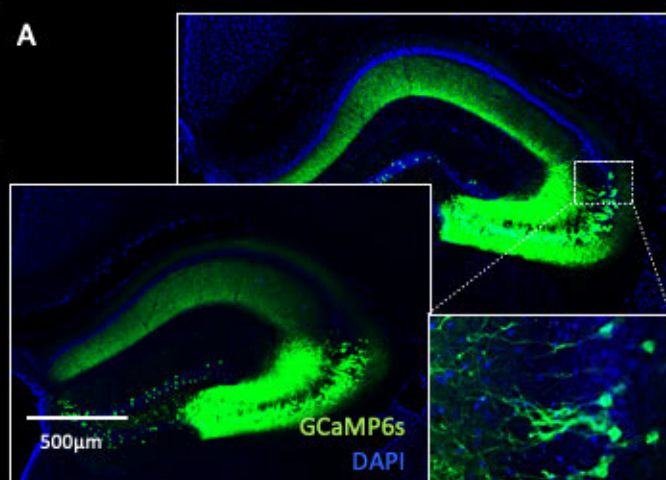
Figure 3-1 Accessing and visualising genetically identified pyramidal neurons in CA3.

A) Coronal histological sections of dorsal hippocampal proceeding rostrally (lower image). GCaMP (green) expression is restricted to CA3, with some low levels of expression in DG and processes projecting to CA1. White track line inset, zoom in of CA3 neurons in the transgenic (GRIK4-Cre;CaMKIIa-tTa;Ai94) mouse. DAPI (blue) labels cell bodies.

B) Axial confocal section of dorsal hippocampus from a GRIK4-Cre mouse expressing Cre-dependent GCaMP6s. GCaMP expression is restricted to CA3 resulting in a 'band' of neurons visualised axially. DAPI (blue) labels cell bodies. The trilaminar hippocampal structure is labelled, indicating stratum oriens (ori.), stratum pyramidale (pyr.), stratum radiatum (rad.), stratum lucidum (in CA3 only), and in the DG, the hilus (hil.), granule cell layer (g.c.l.) and molecular layer (m.l.) are labelled.

C) Coronal histological hippocampal section with CA1, CA3, and DG labelled, white track lines outline cortex excavated for canula implantation, GCaMP6s (green) and DAPI (blue).

D) Representative in vivo FOV in a GRIK4-Cre mouse expressing cre dependent GCaMP6s and showing angled pyramidal neurons. White inset, subcellular resolution visualising, apical and basal dendrites, with large boutons, a CA3 marker (scale bars 50µm & 25µm, respectively).

E) Representative 2P in vivo 800µm x 800µm FOV of 4 planes spaced 33µm apart axially, track line border outlines striatum lucidum on plane 4, a unique feature of CA3 (Witter, 2007).

3.2.3 Co-expression of GCaMP and opsin in CA3

Next, we tried to express optogenetic actuators (opsins) with GCaMP in the same CA3 neurons. Co-expression of opsin and GCaMP has never been attempted in CA3 to our knowledge. We therefore tested various combinations of red-shifted opsins and GCaMPs to optimise the expression level (Table 1). We used red-shifted opsins to minimise optical cross talk by minimising overlap in absorption spectra (Zhang et al., 2011). All opsins employed the Kv2.1 sequence tag to ensure soma-targeting, and reduce off target opsin expression in cellular processes, enhancing the ability to map connectivity with single-cell resolution (Russell et al., 2019). We tried multiple dilution combinations of different animal genotypes (Table 1).

81

Table 2. Co expression strategies combining actuator and activity indicators.

Summary of the total animal genotype and the opsin- GCaMP combinations attempted, with their co expression and photoresponsivity listed. Parenthesis indicates transgenically expressed GCaMP for 6s and bicistronic GCaMP for 6m (see _here_ for more expression strategy details)

Genotype	ST- Opsin	GCaMP	Dilutions tried Opsin	Dilutions tried GCaMP	Total animals	Co- Expression	Photo responsive
GRIK-Cre	ChroME	GCaMP7f	1:30, 1:40, 1:50, 1:60	1:10, 1:15, 1:20	47	13	2
GRIK-Cre	ChrimsonR	GCaMP6s	1:10, 1:40, 1:80	1:10, 1:20	50	0	-
GRIK-Cre	C1V1	GCaMP6s	1:10	1:10, 1:15, 1:20	64	0	-
GRIK-Cre	ChRmine	GCaMP6s	1:1, 1:5, 1:10	1:15	4	3	0
GRIK-Cre	ChRmine	(GCaMP6m)	1:10	-	31	13	0
GRIK-Cre X CaMKII-G6s	ChroME	(GCaMP6s)	1:30, 1:40	-	22	18	6
GRIK-Cre X CaMKII-G6s	C1V1	(GCaMP6s)	1:10	-	4	0	-
Wild-type	ChrimsonR	GCaMP6s	1:40, 1:80	1:10, 1:20	21	0	-
Wild-type	ChroME	GCaMP7f	1:30, 1:40,	1:10, 1:20	24	7	3
Wild-type	ChRmine	(GCaMP6m)		-	17	8	5
Wild-type	C1V1	GCaMP6s	1:10, 1:30, 1:50	1:5, 1:10, 1:20	25	0	-
EMX-Cre	Chrmine	(GCaMP6s)	1:1, 1:5, 1:10, 1:100	-	12	8	3

To confirm our surgery coordinates, we injected 20nl of DiI into 16 mice (Figure 2.2). We leveraged viral spread of up to 0.8mm to achieve expression in CA3 without cell damage from a

bolus injection directly into CA3. Among the co-expression strategies we tested, we found co-expression of ChrMine and ChroMe with GCaMP in 28.14% of animals in which CA3 was successfully surgically accessed (n = 76/270 animals), whereas we failed to co-express the Chrimson and C1V1 with GCaMP in all of our attempts. We experienced two main issues affecting the reliability of opsin/GCaMP co expression, namely, expression of one probe or the other probe (Figure 2.3 bottom two rows) or expression of both probes but not co expressed in the same cells (Appendix Figure 7.2). To circumvent these co-expression reliability issues, we tried a bicistronic construct, providing a single integrated ST-ChRmine/GCaMP6m virus (Table 2 in Methods & Materials). While this dual opsin/GECI construct solved our issue, GCaMP6m was difficult to visualise in most of our animals (Figure 2.4). This construct was only tagged with GFP, and the low fluorescence rendered dynamics in our FOVs difficult to detect, consequently making motion correction during behavioural experiments error prone. We subsequently tried multiple dilution combinations of the bicistronic construct (Table 1) and injecting ST-ChrMine with trangenically expressed GCaMP6s (Table 1). Independently injecting ST-ChrMine and GCaMP6s improved GCaMP visibility issues, but to fully assess the success of our co-expression strategy, we tested opsin functionality.
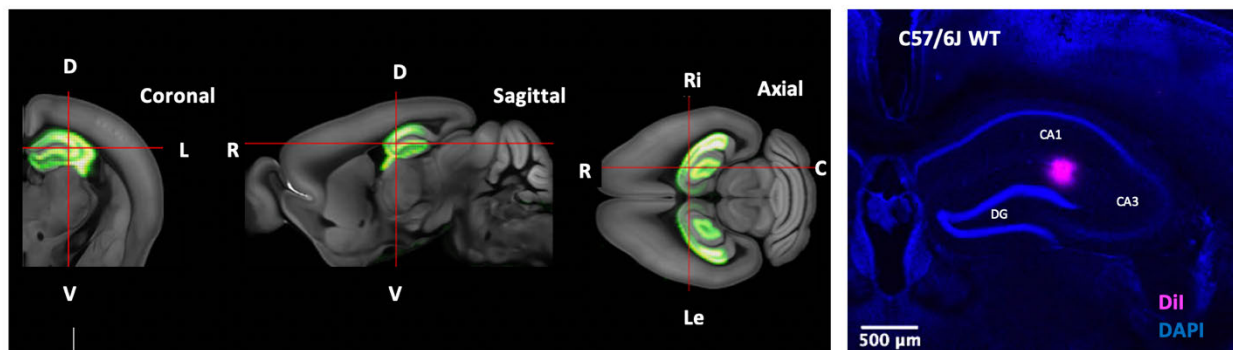
Figure 3-2 Surgically targeting hippocampal subregion CA3.

Top left) Computer generated surgery targettting in sagittal, coronal and axial hippocampal sections, doral (d)-ventral (v), medial (M)-lateral (L) and left (Le)- right (Ri) oreintations labelled (generated using www.connectivity.brain-map.org. Top right) Coronal histological section showing injection target in pink (DiI) in a wild-type (WT) mouse. DAPI labels cell bodies. Scale bar 500µm.

GRIK4- GCaMP6s

Flex-ChroME

200μm

GCaMP6s

Merge

EMX-GCaMP6s

ChrMine

200μm

GCaMP6s

Merge

GRIK4-cre

C1V1

200μm

GCaMP6s

Merge

C57/6J WT

Chrimson

200μm
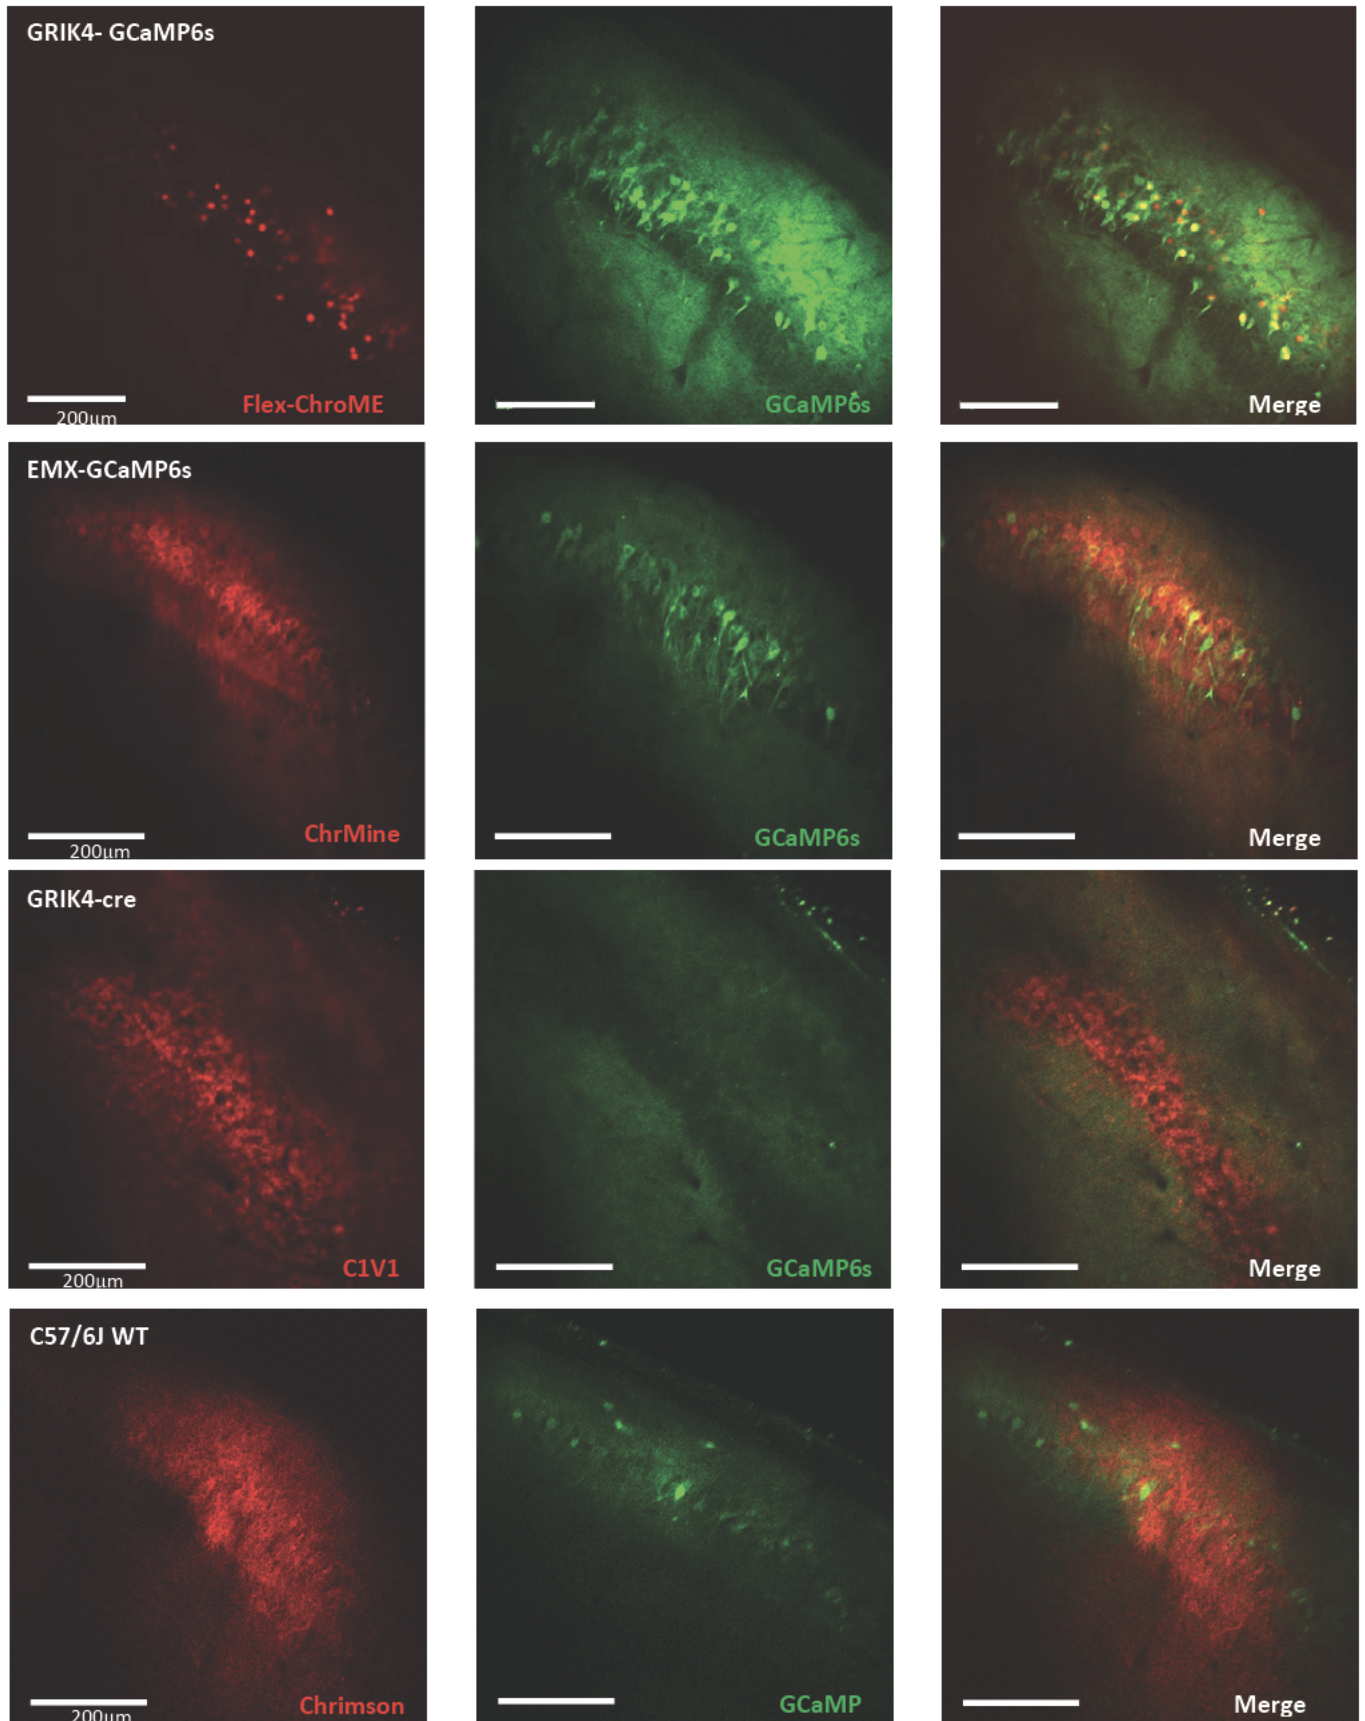
GCaMP

Merge

Figure 3-3 Co-expression of activity indicators and optogenetic actuators in CA3 in vivo.

Two photon FOVs showing expression of opsin (red), GCaMP6s (green) and co-expression (yellow; merge) in CA3 of multiple mice (genotype listed in top left corner, first panel). Top left inset) ST-ChroME is a nuclear localised opsin and therefore not expressed in cellular processes and the fluorophore is not expressed in the somatic membrane. Upper two rows, opsins we successfully co expressed (ChroMe & ChrMine) and bottom two rows opsins we failed to co express with GCaMP (C1V1) and opsin (Chrimson). All scale bars 200μm.

3.2.4 All-optical evaluation of co-expression strategies

Next, we assesed co-expression strategies functionality by photoactivating groups of CA3 neurons with two-photon holographic stimulation while recoding their response with two-photon calcium imaging in awake mice. (Packer et al., 2015). To do this we photoactivated our FOV using a novel NAPARM protocol descibed in Methods & Materials (Figure 2.4 top). Briefly, this high throughput photomapping protocol involved sequentially simultating groups of target sites arranged in a 13x13 pixel grid (Figure 2.4). We quantified the responsivity of target sites to our NAPARM photostimulation protocol using our response metric (Russell et al., 20019; in Materials & Methods), We found 28.36 % (n = 19/67 animals) of co-expressing animals that were photostimulated showed neural responses that met our criteria (neural response metric in 2.21). We found ST-ChrMine and ST-ChroME responded to the NAPARM protocols, activating targeted cells above our threshold for responsivity. We also found these opsin/GCaMP combinations to be unreliable, in that co-expression often failed (Figure 2.3) and in the case of successful co expression, photoresponsivity often failed (27% success n = 19/70; see Appendix). Taken together, most strategies yield a small number of all-optically addressable neurons, but we had the most functional success activating neurons infected with the bicistronic ChrMine construct. However, dim GCaMP dynamics during endogenous activity meant we preferred our second most functionally successful strategy of independently injecting AAV's GCaMP and ChrMine.
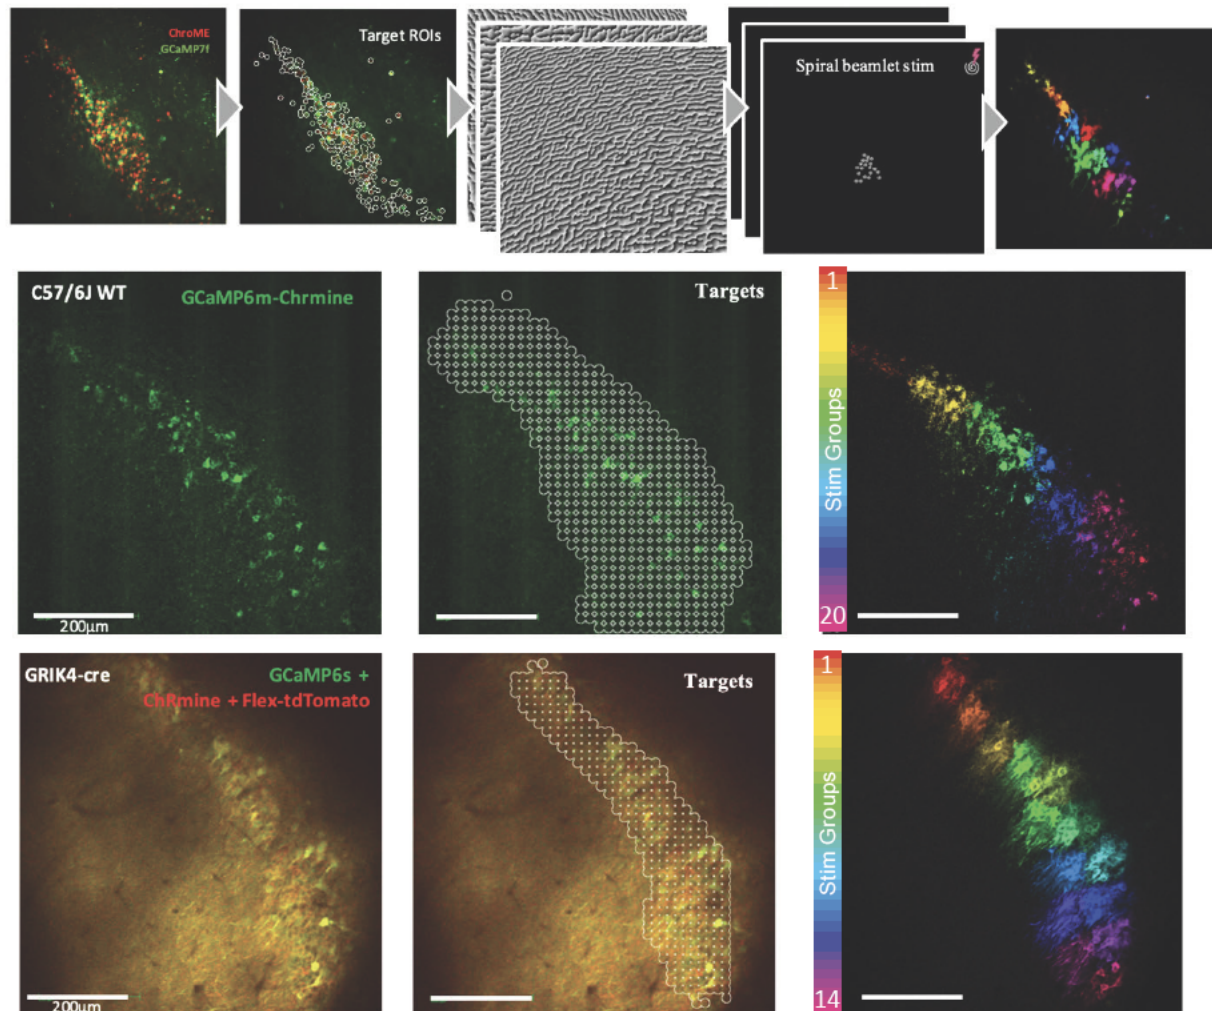
Figure 3-4 An all optical tool kit for control of neural activity in CA3 in vivo.

Top: Example NAPARM protocol for assessing photoresponsivity of multiple neurons in vivo. The leftmost panel shows an example FOV of hippocampal CA3 neurons co-expressing viral GCaMP7f and ST-ChroME (left to right), all pixel targets, centred on the identified ROI, 3 example SLM phase masks, calculated separately for each cluster of ROIs to be photostimulated simultaneously, galvanometer mirror locations are also calculated but not shown, finally, example spiral scan of SLM generated beamlets for 1 cluster of ROIs. Middle row: CA3 FOV in a wild type mouse, co expressing GCaMP6 (green) and ChrMine (no colour, no fluorophore), Next, ROIs selected for photostimulation (white rings). Next, target sites were grouped into 20 different patterns of 20 cells each and targeted for sequential photostimulation to confirm responsivity. Pixel intensity indicates the change in florescence caused by photostimulation. Colour corresponds to the photostimulation pattern which caused the largest change in activity.

Bottom row, same as in middle row but for a GRIK-cre mouse expressing GCaMP6s (green) and ChRmine (red; all scale bars, 200µm, FOVs 512 x 512 pixels).
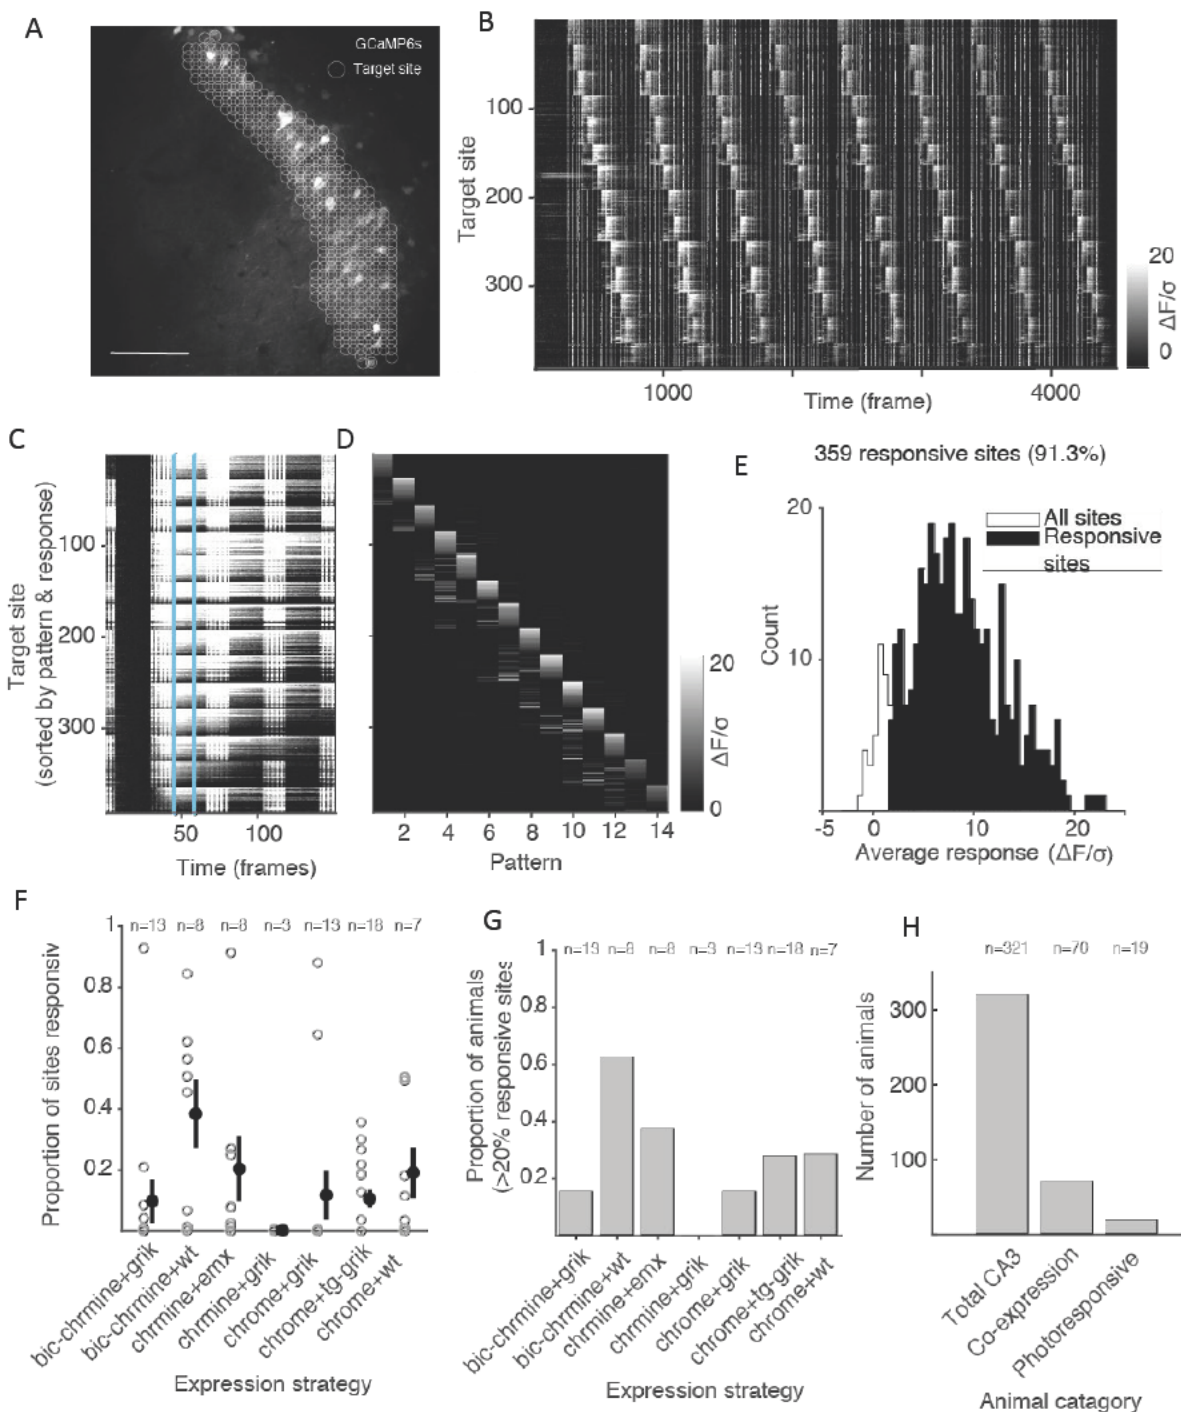
Figure 3-5 An unreliable all optical strategy for controlling CA3 neurons with light.

A) Example FOV expressing GCaMP6s (white) with target sites overlaid (white; scale bars, 200µm).

B) Activity throughout the mapping protocol showing 8 trials for each stimulation pattern for all target sites.

C) Photo stimulation triggered average trace for target sites. Blue line outlines the 50ms response window used for subsequent analysis. Stimulation laser light artefact shown.

D) Target sites are coloured by average response to photo stimulation and sorted by the pattern number and response size.

E) Histogram of the average response sizes of all sites photoactivated. Unfilled bars = all sites, filled bars = photoresponsive sites.

F) Summary plot of the proportion of responsive sites as a function of expression strategy. Mean Std (black) individual experiments (grey rings). Mean ± Std responsive sites: bic-ChrMine + GRIK 0.098 ± 0.26, bic-ChrMine + WT 0.38 ± 0.32, ChrMine + EMX 0.20 ± 0.30, ChrMine + GRIK 0.002 ± -0.003, ChroME + GRIK 0.12 ± 0.29, ChroME + tg-GRIK 0.11 ± 0.13,. ChroMe+ WT 0.19 ± 0.219 (total n = 70).

G) Summary plot of the expression strategies where co-expression of an opsin/GECI combination was achieved. Proportion of animals defined as having responsive ROIs in CA3 (i.e. > 20% target sites respond and meet our neural response metric) as a function of expression strategy. Mean: bic-ChrMine + GRIK 0.15, bic-ChrMine + WT 0.63, ChrMine + EMX 0.38, ChrMine + GRIK 0, ChroME + GRIK 0.15, ChroME + tg-GRIK 0.28. ChroMe+ WT 0.29 (total n = 70).

H) Summery plot showing the total number of CA3 animals attempted (n = 321) and the number of animals with (n = 70) and the number with photoresponsive neurons in CA3 (n = 19).

3.2.5 All-optical experiment

Finally, we conducted a proof of principle experiment, probing the network input-output function in hippocampal subregion CA1 and CA3 within the same animals (n = 6; see Material & Methods for details on subregion segmentation within animal, Figure 2.6 C). Briefly, for this experiment we selectively photoactivated groups of 1, 2, 5, 10 or 20 cells in one subregion and simultaneously recorded activity in that that region as well as in the other subregion (see Figure 2.6 A & B). We quantified the response in all neurons in a 500ms window after the photostimulation offset and sorted cells by subregion and into directly photostimulated target (Figure 2.6 D targets) and the non stimulated background (bg; Figure 2.6 C & D). We defined the network response as the average response across all non stimulated neurons either within or across the stimulated subregion (274 ± 133 nonstimulated CA1 neurons, 197 ± 50 non stimulated CA3 neurons, n = 6 animals).

88

Following CA1 photostimulation, we found no response on average in the local CA1 network (-0.005 ± 0.050 ΔF/F, P = 0.349, Wilcoxon signed rank test; Figure 2.6 G). However, surprisingly activity in CA3 networks were suppressed by CA1 stimulation (-0.018 ± 0.042 ΔF/F, P = 0.002. CA1 vs. CA3, P = 0.0001, Wilcoxon signed rank test, Figure 2.6 G). As we increased the number of stimulated cells in CA1, we saw a trend whereby the CA1 network became less suppressed while the CA3 network became more suppressed (Figure 2.6 G). Conversely, as a result of stimulation delivered to CA3, we found that the CA1 network responses were suppressed across all numbers of stimulated cells (CA1 network response: -0.017 ± 0.033 (ΔF/F), P = 0.012. CA3 network response: 0.004 ± 0.058 (ΔF/F), P=0.439. CA1 vs CA3 P = 0.006, Wilcoxon signed rank test; Figure 2.6 H). As we stimulated more CA3 cells, the local CA3 network response switched from suppressed to facilitated with the response increasing with increasing number of photostimulated CA3 cells (Figure 2.6 H). This increasing excitatory recruitment was seen only in CA3 cells when stimulating other CA3 cells, and not in CA1, perhaps reflecting differences in network connectivity, particularly the greater number of recurrent connections in CA3 serving to amplify excitation.
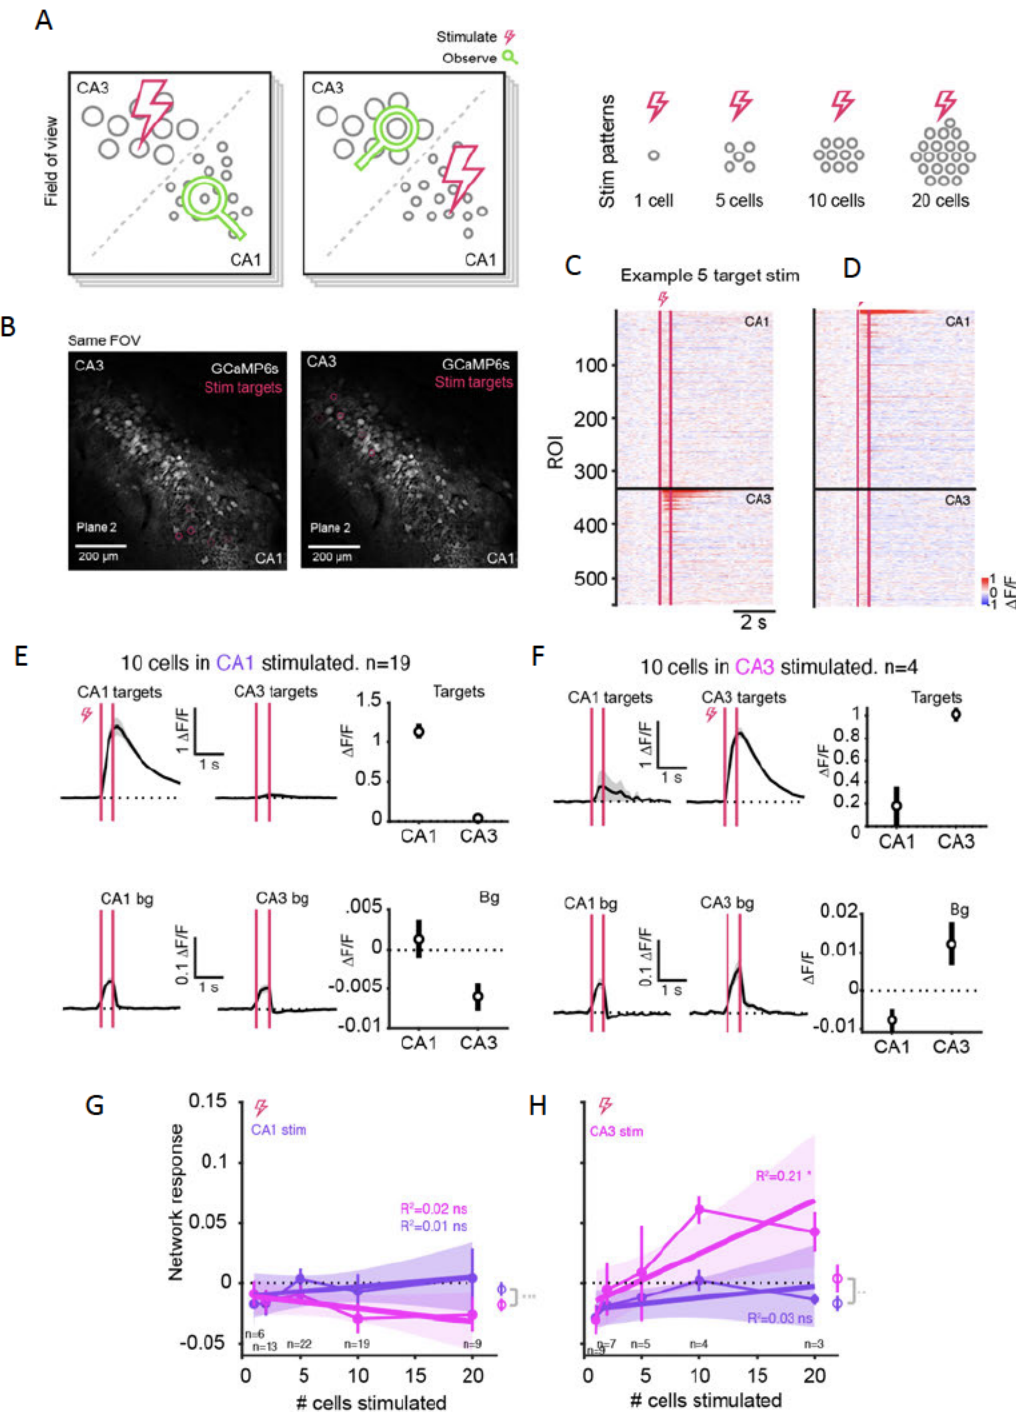
Figure 3-6 An all-optical investigation of CA1 and CA3 network connectivity within animals.

A) Left) Schematic to demonstrate the experimental set up of photostimulating (pink) one subregion while observing (green) neural responses in other non photostimulated subregion. Right) Example photostimulation patterns, including simultaneous stimulation of 1, 2, 5, 10, 20 cell patterns during the subregion stim phase.

B) Example FOV showing 1 planes from a 4 -plane volume stack of CA1 and CA3 neurons expressing GCaMP6s (white). Photostim targets are ringed (pink) and brighter rings indicate targets on that particular imaging plane.

C) Photostim heatmaps showing neural activity (ΔF/F) across ROIs in CA1 (above the black line) and CA3 (below the black line), before, during and after photostimulation (pink lines) of 5 ROIs. Warmer colours = higher responses).

D) Same as in C) but for 5 CA3 ROI photostim targets within that example animal in C).

E) Photostimulus triggered average traces during stimulation of 10 CA1 cells, averaged across all experiments. Top row: target cells. Bottom row: non stimulated background cells. Right: average responses (ΔF/F) across all cells in that population.

F) Same as in E) but for CA3 stimulation.

G) Input output function for CA1 stimulation. Average response (ΔF/F) in the non stimulated network (purple = CA1, pink = CA3) following increasing number of cells stimulated in CA1.

H) Same as in G) but for CA3 photostim.

Grey shaded areas and error bars = S.E.M. Coloured shaded areas = confidence interval of linear fit. Red-pink lines indicate the photostim window. Thicker solid lines = means. All *s indicate Wilcoxon rank sum tests (n.s. non significant, $*P < 0.05$ **, $P < 0.01$, ***, $P < 0.001$).

3.3 Discussion

Understanding the functional connectivity between neural circuits in the hippocampus has been well studied, but not in vivo at the systems level, given limited experimental tools. With the advent of hippocampal imaging and advances in optical manipulation techniques, we were able to develop and combine our CA3 transgenic mouse model, new surgical methods and optimised optics to optically access CA3 while leaving CA1 intact. We achieved 3D 2P imaging of genetically and anatomically defined CA3 populations in vivo. Moreover, our 2P characterisation defined clear subregion unique features for identifying CA3 neurons in non transgenic animals and opened up the possibility for us to simultaneously image both CA1 and CA3 subregions. Our proof of principle all optical, dual subregion experiment demonstrates the potential of the all optical approach beyond CA1 and outlines the feasibility of probing the input-output function of multiple hippocampal subregions at single neuron resolution.

3.3.1 Optical access to defined CA3 neural populations

We were able to develop a novel surgical approach to optically access CA3 with CA1 intact. However, recently two other groups have published similar surgical methods after the development of our protocol (Hainmueller & Bartos, 2018; Rashid et al., 2020; Schoenfeld et al., 2021). However, Hainmueller and Bartos, did not target genetically defined CA3 neurons, instead defining their FOV based on the angle of the neurons and hence inclusion of CA2 neurons is likely. Rashid and colleagues used a similar GRIK-4 mouse model to ours but with sparse GCaMP expression in CA3 pyramidal neurons (Hainmueller & Bartos, 2018; Rashid et al., 2020). These studies provide support for our surgical method, confidence in our 2P CA3 FOV and also addressed procedural concerns regarding disrupted neurophysiology and behaviours. For example, Bartos and Hainmueller assessed the impact of their similar CA3 surgical procedure and found no difference in the physiology of CA3 neurons *in vitro*, and *in vivo*, and behavioural measurements before and after cannula placement (Extended Data Figure. 3 in Hainmueller &

Bartos, 2018). Moreover, we were able to define our CA3 neurons morphologically, providing confidence in non-genetic identification of CA3 using visible mossy fiber boutons and large pyramidal neurons soma diameters. This enabled us to identify CA3 'online' and subsequently imaging multiple subregions in the same animals for the first time.

3.3.2 All optical access to defined CA3 neural populations

We successfully co-expressed 2 opsins (ChrMine and ChroMe) with GCaMP combinations and activated targeted neurons in CA3. However, much work is required to optimize expression strategies in terms of co expression reliability and functional activation in deep brain structures. Our preparation had a number of uncontrolled variables that might have affected the success of the all optical strategy. Cre expression in CA3 is variable in transgenic animals, and low copies of the gene could have resulted in dim GCaMP expression. CA3 is a small deep substructure, difficult to target and therefore our targeting success could have been improved with an electronic stereotaxic frame. Moreover, we relied on viral spread to of our opsin/GCaMP combinations to infect the target area, however, it is possible the range of spread is affected by the titres of opsins, which are batch dependent and changed regularly given these molecular biology tools were also under development. Furthermore, the proportion of successful targets reported in this chapter may be a conservative quantification of the photoresponsive network in each mouse. Our high throughput pixel grid was wider than expression in our FOV (even with the exclusion criteria) and therefore a large number of photstim targets were included in the photoresponsivity analysis but did not target cells.

Other groups within our lab and outside our lab have had similar issues with opsin/GECI co expression to those discussed here. There are a number of future steps to take to begin to control some of these variables. The most optimal long-term solution would be to create an 'all optical mouse' where the opsin and GCaMP are transgenically expressed. While this would require significant investment in molecular techniques and mouse breeding, an alternative long-term strategy would be to design a new bicsistronic virial construct in which the GCaMP is brighter,

93

this would reduce the main problems with co expression and GCaMP visualisation. Despite these drawbacks, we provided proof of principle evidence extending the all-optical tool kit into hippocampal subregion CA3.

3.3.3 All optical exploration of CA1 and CA3 network function

To demonstrate the utility of our CA3 all-optical approach, we probed the local and long-range functional connectivity of two hippocampal subregions in vivo. Our results suggest CA3 recurrent collaterals can be activated to increase excitation in the network, but also pose a challenge to the long-standing view of information flow through the trisynaptic pathway. This notion assumes a unidirectionally excited chain (see Chapter 1.2.3), regulated by local inhibitory circuits and overlooks the complexities of intrahippocampal circuitry.

It is possible that reduced CA1 suppression and increased CA3 inhibition in response to CA1 photostimulation, could result from increasing activation of long-range, cross-regional interneurons. Studies have shown that the spread of excitation is not only controlled by local circuit inhibitory interneurons, but CA1 inhibitory neurons exist and project back to CA3 and hilar regions (described in Sik, Ylinen, Penttonen, Buzsáki, 1994; Harris et al., 2018). These inhibitory neurons with widespread longitudinal and inter-regional projection fields might play a role in precise population activity synchronization by timing the occurrence of action potentials in the principal cells and hence contribute to the coordinated induction of synaptic plasticity in distributed networks (Sik et al., 1994).

Moreover, our all-optical results also suggested that photostim in CA3 increased the CA3 network response, which as mentioned is likely a function of the known recurrent CA3- CA3 network connectivity amplifying excitation. Interestingly, CA3b, which we assume our 2P FOV includes, is known to receive the strongest net excitation from recurrent collaterals (Sun et al., 2017). It is less clear why the CA1 network response was suppressed during CA3 photoactivation, given the traditional view that excitatory CA3 Schaffer collaterals innervate CA1 targets and propagate

excitation into CA1 circuits (Ishizuka et al., 1990). Recent evidence suggests excitatory (CA1 pyramidal) and inhibitory (CA1 parvalbumin-positive) neurons are innervated by the same presynaptic inputs (CA3 Schaffer collaterals; Kwon, et al., 2018). It is likely the particular connection motifs and connection strength differs, possibly resulting in a net inhibitory response. Although quantitative descriptions of connectivity motifs and network architectures, especially of excitatory and inhibitory neurons innervated by the same type of cell remain elusive and much work is needed to clarify our understanding of the precise excitatory- inhibitory balance within and between hippocampal subregions (Kwon, et al., 2018).

However, it is known that electrical or optogenetic stimulation of the Schaffer collaterals in vitro causes rapid monosynaptic excitation of CA1 pyramidal cells followed by delayed, disynaptic inhibition. This 'canonical EPSP-IPSP sequence' is emblematic of feedforward inhibition (Alger & Nicoll, 1982; Lovett-Barron et al., 2012; Miles, 1990; Pouille & Scanziani, 2001; Schwartzkroin, 1975) and it is possible that the relatively slow GCaMP kinetics and imaging speed miss this initial excitation, resulting in our recording CA1 network inhibition. Collectively, our efforts represent the first steps towards all optical investigation into hippocampal subregion CA3 and the first evidence of in vivo inter hippocampal subregion functional manipulation in an attempt to probe functional connectivity.

4

Mini virtual world development

4.1 Introduction

The location-specific firing of hippocampal place cells during navigation represents a salient neural correlate of spatial information in the mammalian brain (O'Keefe & Dostrovsky, 1971). A wealth of subsequent work has defined and refined our understanding of place cell properties (see Chapter 1). To better understand this neural code, we need access to large numbers of cells to observe their emergent dynamics and collective neural behaviour. Two-photon microscopy enables us to achieve this goal by using calcium indicators as a proxy for neural activity in populations of defined cell types (Dombeck et al., 2010; Yuste 2015).

4.1.1 2P Imaging in the hippocampus

Initially, optically accessing the hippocampus for 2P microscopy was challenging given the limited working distance of the objective and subcortical position of the hippocampus (Dombeck et al., 2010). To solve this problem Dombeck and colleagues developed a novel surgical approach to unilaterally removing part of the cortex (including parietal cortex and parts of visual and hindlimb sensory cortex) overlying the hippocampal external capsule (Dombeck et al., 2010). By implanting a 'sunken' imaging window they were able to optically access subregion CA1. This cortical lesion initially raised concerns that cellular and network properties within the hippocampus may be altered. However, spatial properties of CA1 place cells measured electrophysiologically in lesioned and un-lesioned mice were similar, possibly because the excavated cortical regions do not provide strong direct projections to the hippocampus (Dombeck et al., 2010).

There have also been concerns regarding differences between optophysiologically and electrophysiologically recorded place cells. While similar to classically identified place cells, optically identified place cells do show some differences. Notably they have larger mean place field widths, lower spatial information scores and represent a smaller proportion of the overall population of cells when compared to electrophysiologically identified place cells (Dombeck et

al., 2010). These properties are perhaps not surprising given the reduced vestibular input received by a head-fixed mouse and given the non-linearity of calcium indicators thresholds activity signals. When the vestibular input of rats was temporally knocked out through injection of tetrodotoxin, tetrode recordings showed a temporary decrease disruption in location-specific firing of CA1 place cells (Stackman, Clark, & Taube, 2002), hence, reducing the ability of the hippocampus to function as a spatial map. Overall, these drawbacks are outweighed by the ability to study larger subsamples of genetically defined hippocampal populations during navigation with 2P imaging (Chen et al., 2013; Dombeck et al., 2010).

4.1.2 Imaging during behaviour in VR

The use of 2P microscopy to assess spatially tuned cells presents a paradox because high resolution 2P microscopy requires no sample motion while simultaneous requiring navigation-based behaviour. Harvey and colleagues developed a solution to this problem by creating a method for virtual navigation in head fixed animals (Harvey et al., 2009). They projected a computer-generated virtual environment into a dome and positioned a head fixed but freely moving mouse on a polystyrene ball, essentially allowing navigation thorough virtual space (Harvey et al., 2009). Virtual navigation is increasingly becoming a widely adopted method in conjunction with opto- and electro-physiology to investigate spatial behaviour (Aronov & Tank, 2014). However, it has been typically conducted with a 'dome' (toroidal screen; Dombeck et al., 2010; Harvey et al., 2009; Robinson et al., 2020; Schmidt-Hieber & Hausser, 2013) or large screen set up (Hainmueller & Bartos, 2018; Plitt & Giocomo, 2021), to provide an 'immersive experience'. Ceiling texture i.e. the surface above the mouse, is considered particularly important given the positioning of rodent eyes and their large field of view (Harvey et al., 2009; Hölscher, et al., 2005). However, due to space limitations on our microscope configuration, the positioning of essential hardware (ETL; to enable volume imaging), prevents the use of a dome, top ceiling screen or large screen set up. Therefore, we sought to miniaturise a custom VR and hardware to fit onto our microscope. We aimed to use mini-VR to flexibly manipulate different world

parameters for our various experiments and interface with our microscopy software to assess rodent spatial behaviour

4.2 Results

4.2.1 Hardware

Our mini VR system was constructed using the following hardware to connect 3 screens arranged in an arc (picture 1 and 2, Figure 3.1):

3 x 1024 x 600 HD Raspberry Pi screens

1 x UBS hub for power with power supply

1 x HDMI/DisplayPort hub

3 x HDMI cables

Custom removable screen holder on a magnetic base (see picture insert in Figure 3.1 A)

4.2.2 Software

We used ViRMeN, a free opensource MATLAB-based software package for designing and running rodent VR experiments (Aronov & Tank, 2014). ViRMeN provides an inbuilt graphics engine that performs 3D rendering of virtual environments on a computer monitor and a Matlab toolbox for programming experimental logic and manipulating environments in real time (Aronov & Tank, 2014). We used NIDAQ tools to interface the digital signals controlling the delivery of rewards through a custom build lick port, trigger 'stim' TTL pulses in stimulation experiments, and detect lick events.

4.2.3 World design

To design our virtual environments we took inspiration from previous VR worlds and in-house experience. We also wanted to maximise optic flow and contrast, therefore decided to use a

narrow-2D textured corridor (Figure 3.1 B). We used monochrome, blue and green colours given mouse bi-chromatic vision, making it difficult for them to perceive red hues (Goymer, 2007). Our 2D textures and simple colours enable easy parametric morphing (Figure 3.1). The rewarded area is marked on the floor and walls with a monochrome 'X' (Figure 3.1 B after 'END').

4.2.4 Experiment logic

We used the ViRMeN MATLAB toolbox to implement our experimental logic. That is, teleportation upon ending a trial, blackout and whiteout for reward and punishment, respectively, and delivery of water reward upon reaching a specific marked location on the track. This is similar to Robinson et al., (2020) which is consistent with other previously reported rodent VR systems (Alegra et al., 2019; Dombeck et al., 2010; Harvey et al., 2009 Hainmueller & Bartos, 2018; Plitt & Giocomo, 2021; Schmidt-Heiber & Hausser, 2013). The main difference with our VR system is that we do not punish licking outside the reward zone as this was not an issue for our expert mice (Figure 3.2 'Expert'), possibly due to our longer track length and delivery of unsweetened water (Robinson et al., 230cm; our track 480cm). We also made our VR software capable of initiating stimulation and delivering predefined trial structures. We built a GUI that enables the experimenter to change the upcoming trial, deliver manual rewards and switch experiment conditions therefore allowing 'in experiment' changes if necessary.
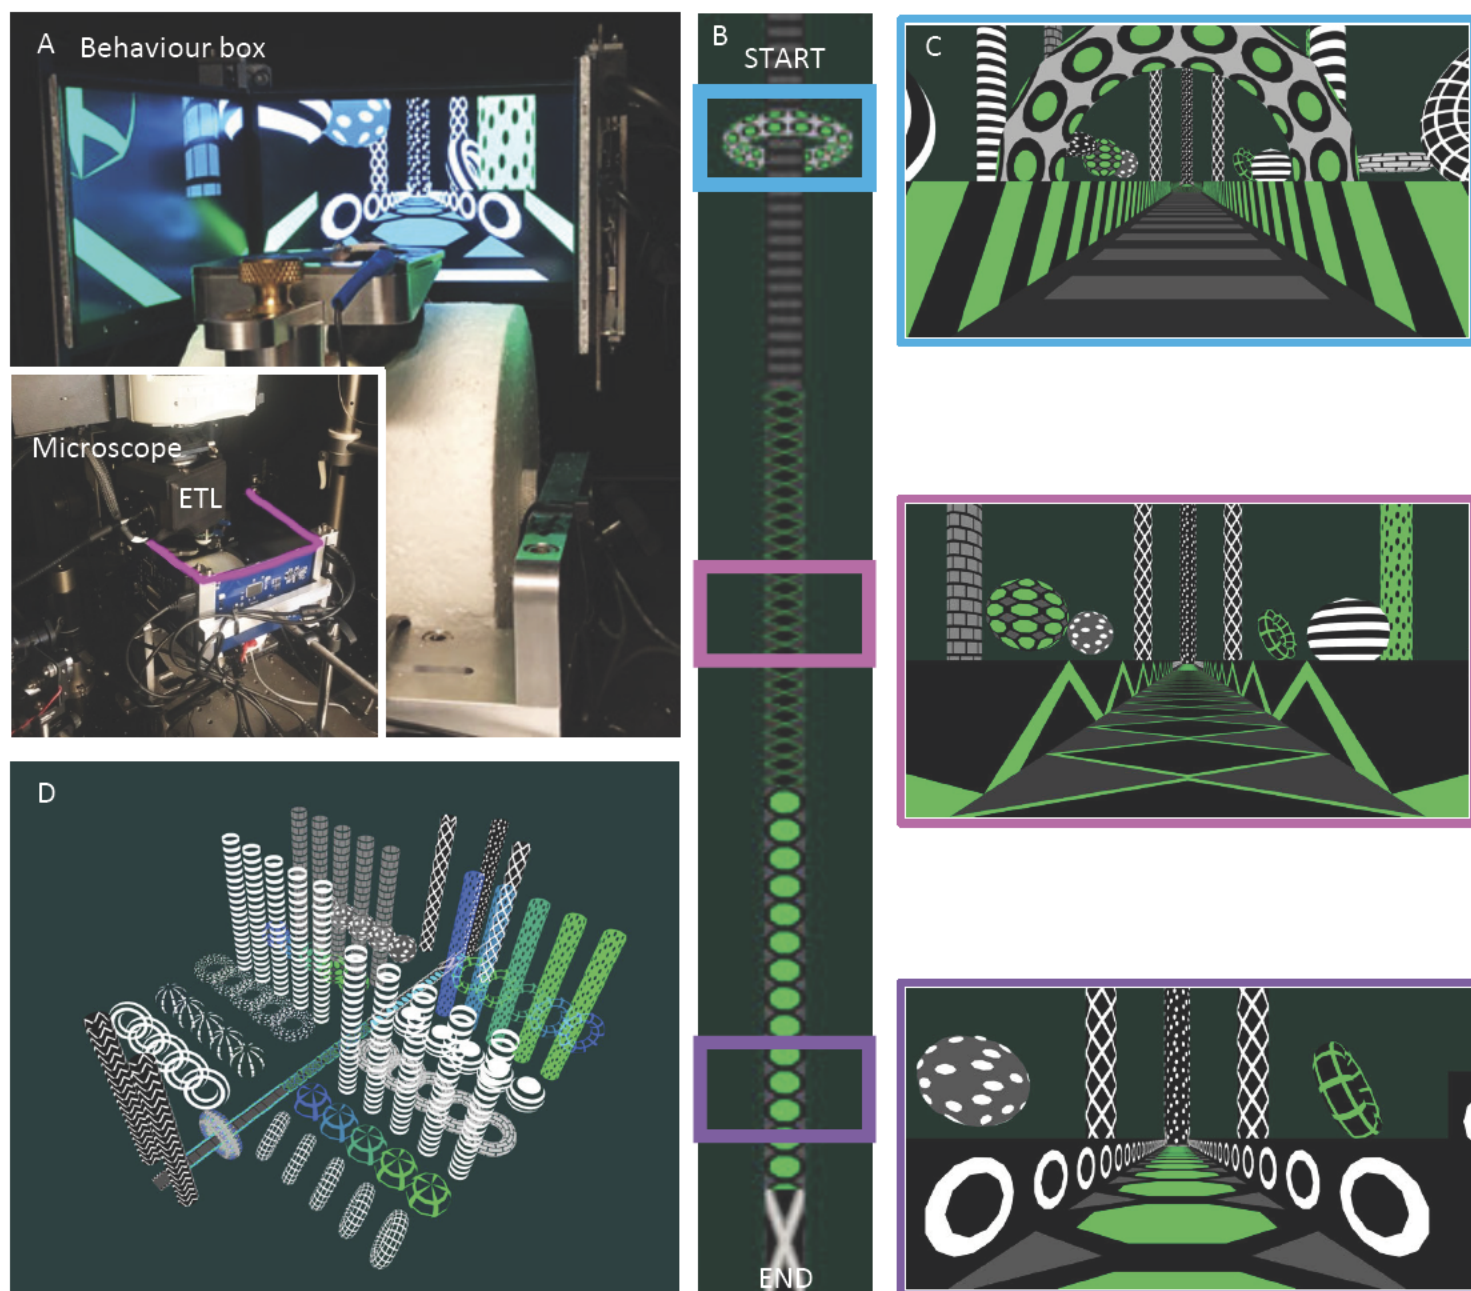
Figure 4-1 Modular mini-VR world design & layout for mouse navigation behaviour. A) Three mini screen configurations, surrounding a freely moving polystyrene cylinder in our custom built behaviour box and insert, on our 2P microscope rig. B) Top-down view of the 480cm virtual track, divided into 3 parts, shown by different floor textures. The start is ringed with a textured torus and the end marked by the reward zone, a grey, white and black cross. C) Coloured boxes correspond to coloured edged cross section, showing the track as a corridor formed from low walls, and spaced objects beyond those walls. Each section of the track is visually distinct, given floor and wall

textures and different objects, to improve optic flow. D) Our modular library of objects and textures enables flexible construction of novel worlds and controlled changes to specific features.

4.2.5 Single mouse

To assess the animal's ability to navigate in our VR system, we investigate their run and lick behaviour as a measure of learning the navigation task. We defined correct behaviour as reducing speed to < 5cm/s and licking 3 or more times in the rewarded area to initiate a water reward. All mice learned the task defined above, and 70.4% (n = 19/27) met our experimental behaviour standards of running a trial per minute, in addition to the 80% correct trials per session.

4.2.5.1 Naïve vs expert mouse behaviour

Initially mice typically lick at most positions along the virtual track (black dots Figure 3.2 A) and run slowly, often < 5cm/s completing few trials in a 30-minute session. Over time their running speed increases and 'expert mice' run directly to the reward zone, indicated by a smoother grey trajectory in Figure 3.5 B > 35cm/s. Subsequently, mice complete more trials, with shorter durations in longer training sessions (session 1: 58.58 ± 22.34s, n = 31, session 10: 58.58 ± 22.34s average trial duration n = 19; Figure 3.3). If the mice make mistakes, they often self-correct after only a few incorrect trials (Figure 3.2 B). Expert mice also show preparatory run behaviour, in that running slows before entering the reward zone having learn that this location is rewarded (Figure 3.2 velocity decreased upon entry into the shaded reward area and increases after mice have received the reward before the next trial). Lick behaviour also shows marked changed in expert mice. After learning the reward location mice lick almost exclusively in this zone, except a small number of pre-emptive licks prior to entering the reward area (Figure 3.2).
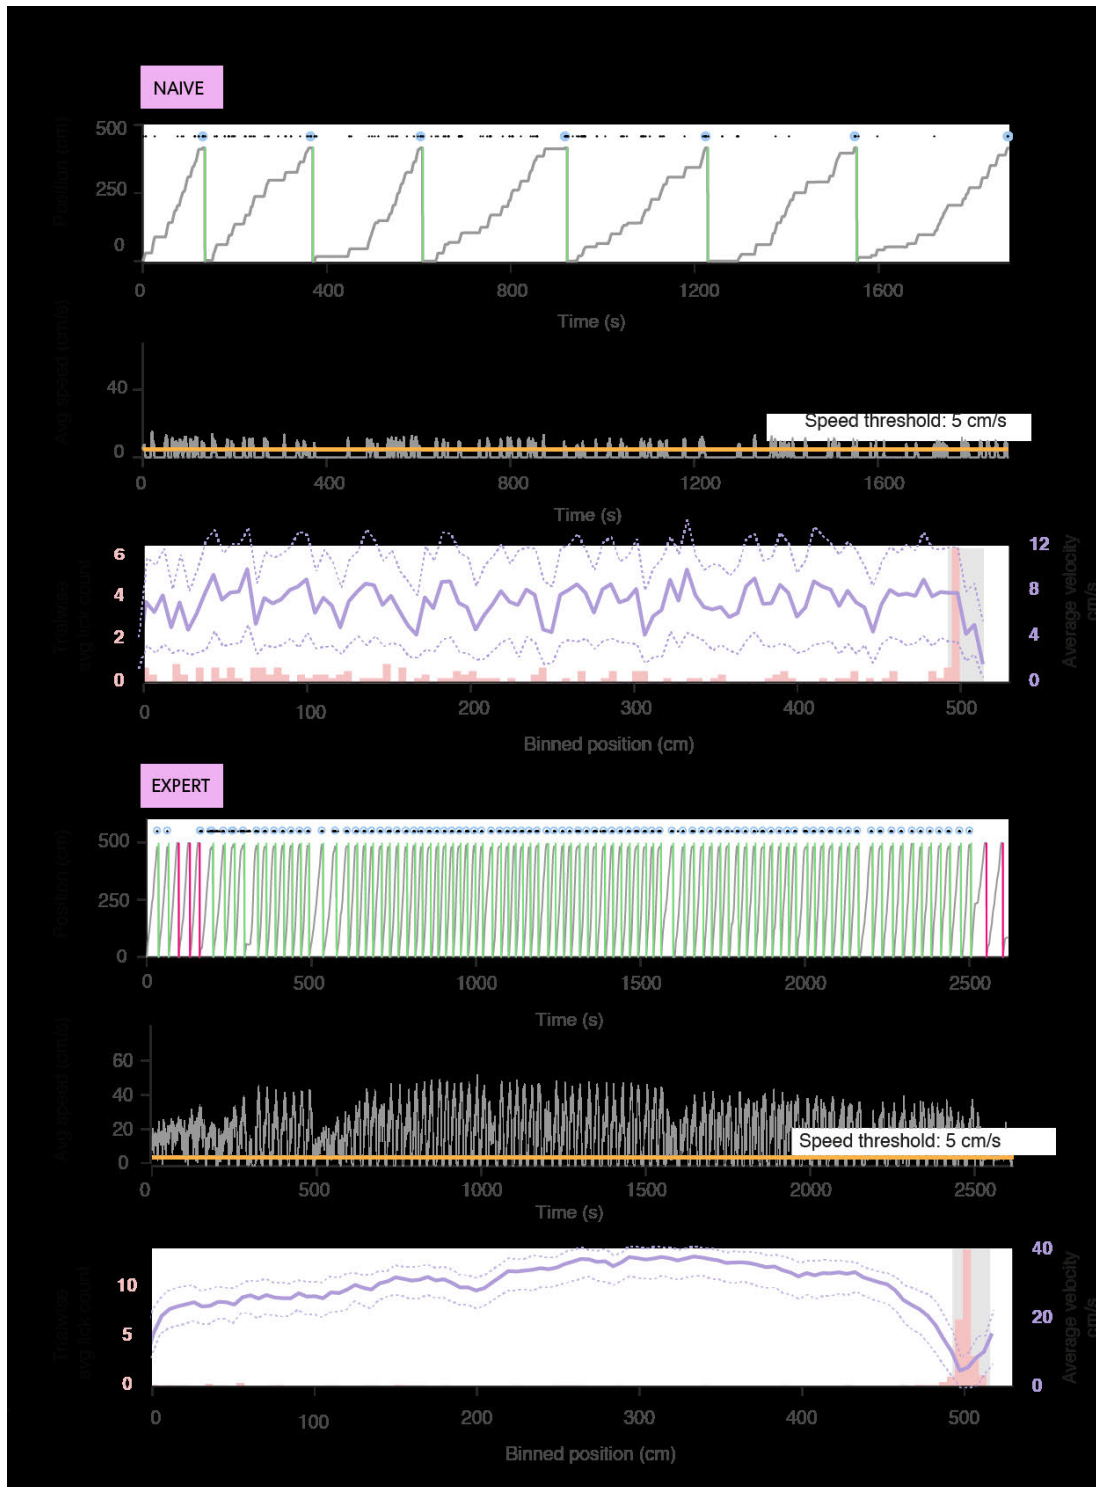
Figure 4-2 Lick and run behaviour in the naïve and expert mouse in mini VR.

A) Naïve mouse (top). Run trajectory (grey line) through virtual space as a function of time (seconds). Green lines indicate correct trials and teleportation to the beginning of the track. The above black dots about the trace represent licks and blue rings correspond to rewards. Initially the mouse is auto rewarded irrespective of lick

103

initiation. Naïve mouse middle trace) Grey line represents average running speed (cm/s) as a function of time over the duration of the session. Yellow threshold is draw at 5cm/s for later place cell analysis, where < 5cm/s is excluded. Naive bottom) Histogram of trial wise mean lick count (pale pink) along binned positions down the virtual track (5cm position bins). Grey shaded area marks the reward zone. Licks increase in the reward zone upon delivery of the reward. Overlaid is trial wise mean velocity (lilac; cm/s) along virtual position (lilac track lines = ± 1SD). Expert (below).

B) Expert mouse. Same as for naïve mouse but more trials and longer overall sessions. Naive top) red lines here indicate incorrect trials.

4.2.6 All mice

4.2.6.1 Run behaviour

To quantify the changes from naïve to expert mice across animals, we trained 31 mice of both sexes in our VR system. We found that the changes described in Figure 3.2 were robust across most animals, with all animals showing some degree of these learning changes over time. Run behaviour increases in speed (session 1; mean ± std 8.90 ± 2.7 cm/s, n = 31) up to session 10 (mean 12.63 ± 3.20 cm/s, n = 19), after which there is a plateau for subsequent sessions (Figure 3.3). We quantified run behaviour changes as speed change into the reward zone. This metric enables us to detect 'correct' run behaviour, and not animals who increase their speed over sessions, but do not slow down in the reward zone and therefore do not display 'correct' run behaviour. We found that with training there is a decrease in speed change into the reward zone (mean 2.70 ± 1.98, n = 31, session 1, vs, mean -0.10 ± 0.69, session 10, n = 19), reflecting that as the mice learn, they slowdown in anticipation of receiving the reward (Figure 3.3). Again, this change plateaus after 10 sessions.

4.2.6.2 Lick behaviour

We quantified licking specificity across sessions for all animals. We found a general decrease in the number of licks outside the reward zone as the number of sessions increases, until session 10 when this change in licking plateaus (Figure 3.3; session 1: 31.59 ± 22.27, n = 31 vs. session 10: 10.14 ± 8.39, n = 19; Figure 3.3 B). Although the data is noisy and confounded by the possible delivery of manual rewards during early training sessions (see peaks of > 100 licks). We also

104

assessed lick behaviour by quantifying mean licks in the reward zone. Mice showed a general increase up to session 10 in mean licks (session 1: 9.33 ± 6.32 licks, n = 31; session 10: 19.32 ± 9.43, n = 19; Figure 3.3 B top). To measure the specificity of licking, i.e. to create a metric to detect 'correct' licking, defined as licking in the reward zone and withholding licking outside the reward zone, we measured lick fidelity. We compute lick fidelity as licks in reward zone minus licks outside reward zone, as a proportion of the total licks. Hence, on average, initially mice licked more outside the reward zone (indicated by -vr lick fidelity; -0.47 ± 0.37, n = 31; Figure 3.3. bottom), gradually decreasing, until session 10, where lick fidelity plateaus (+ve lick fidelity indicates more licks in the reward area; Session 10; 0.38, ± 0.36, n = 19; Figure 3.3. bottom).

4.2.7 Task engagement

To ensure our animals are actively navigating in the VR, we increased the task difficulty by moving the rewarded locations on each trial (Figure 4.4 A right & left). We structured trials such that the reward zone was randomly positioned in one of three possible locations for a randomly selected number of trials (2-6 trials; arrows in Figure 4.4 A). The navigation task was the same as described previously, and we used two versions of this task; one version where the reward zone was moved by larger amounts (at positions 225, 347, 430 cm along the track; pink in Figure 4.4) and another version where it was moved by smaller amounts (390, 430, 470 cm; yellow in Figure 4.4 left). We found that all animals tested (n = 14 and n = 6) were able to complete the task with few differences between run and lick behaviours on trials in each of the reward locations (Figure 4.4 C & D). We found no significant difference in the trial outcomes (% correct) for each of the rewarded locations in each version of the task (Figure 4.4 C & D bottom: Wilcoxon rank sum test with Bonferroni correction, $P > 0.05/3$). This suggests that our mice are actively engaged in the navigation task with flexible behaviour, rather than internally tracking distance.
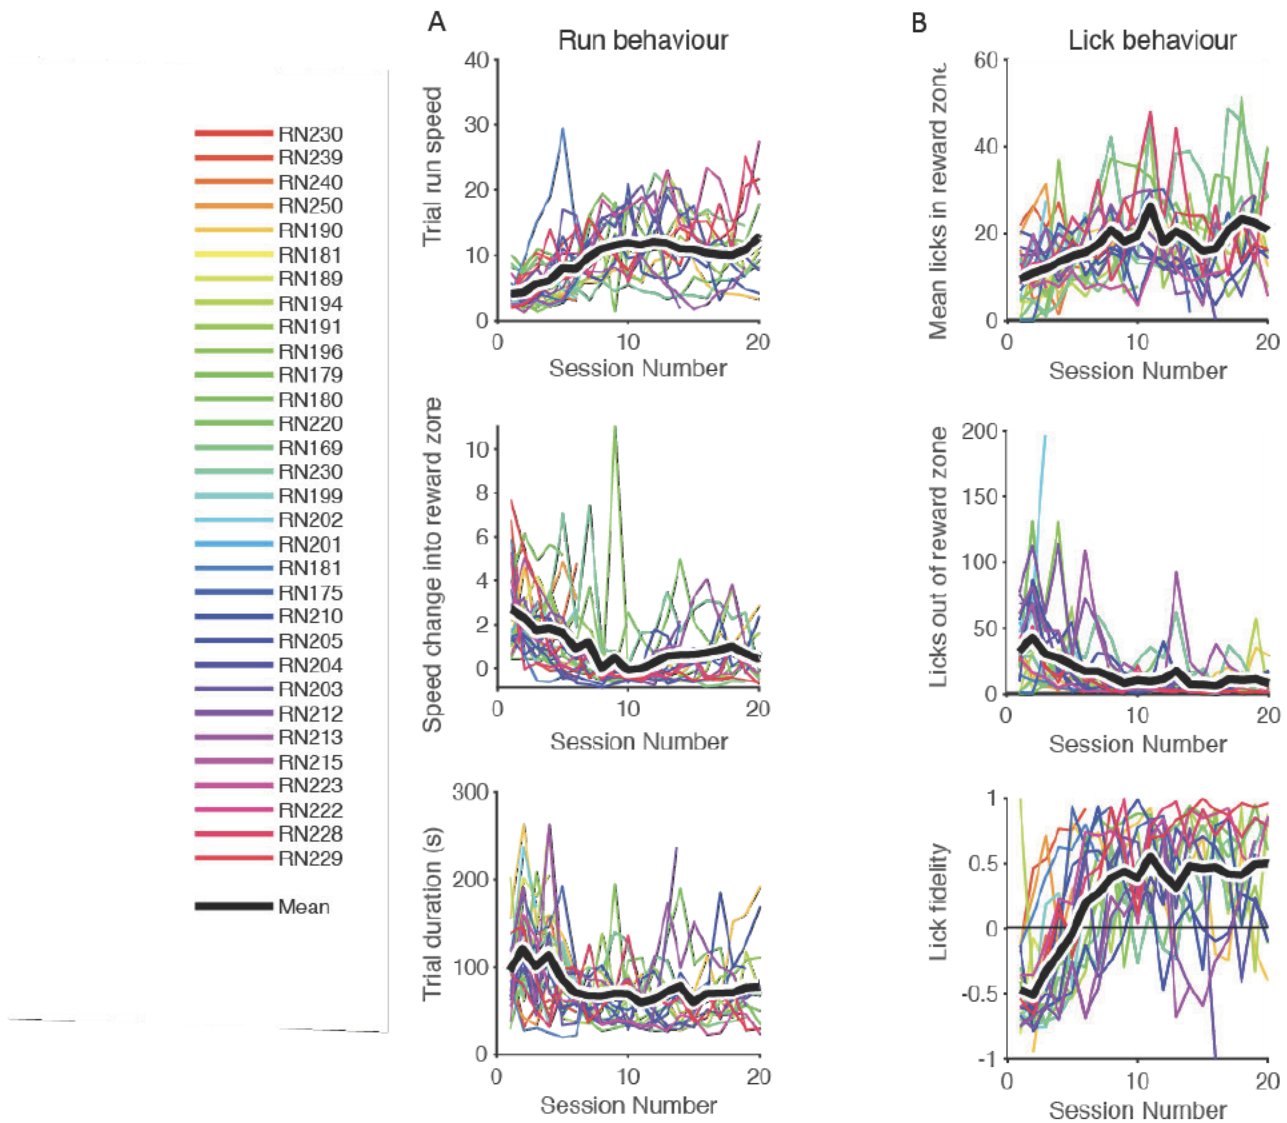
Figure 4-3 Learning curves for run and lick behaviours across animals trained in the mini VR.

A) Run behaviour: Top: Running speed (cm/s) and along the virtual track (before the reward zone) for all mice (coloured lines) over all the sessions completed, up to session 20. Middle) Speed change into the reward zone, defined as the difference between the running speed in the middle of the track vs. running speed in the reward zone. Bottom: Average trial duration (seconds) over sessions.

B) Lick behaviour: Top: Licks in the reward zone before the reward delivery across animals and over sessions. Middle: Licks outside of the reward zone and along the virtual track for all animals over sessions. Bottom) Lick fidelity, defined as the differences in licks inside the reward zone before the reward delivery and outside the reward zone as a proportion for total licks. Each coloured represents a different mouse, black line is the mean. Not all animals competed 25 sessions hence the individual correspond to the number of sessions completed by that animal. N = 31.

A

B er reward zone changes

Larger reward zone changes

Lick
Correct trial
Run trajectory
Reward zone location

Larger reward zone changes

C Smaller reward zone changes

D Larger reward zone changes

Mean
RN191
RN239
RN248
RN240
RN300
RN301
RN299
RN322
RN296
RN324
RN309
RN321
RN310
RN308

Mean
RN240
RN300
RN301
RN322
RN296
RN321
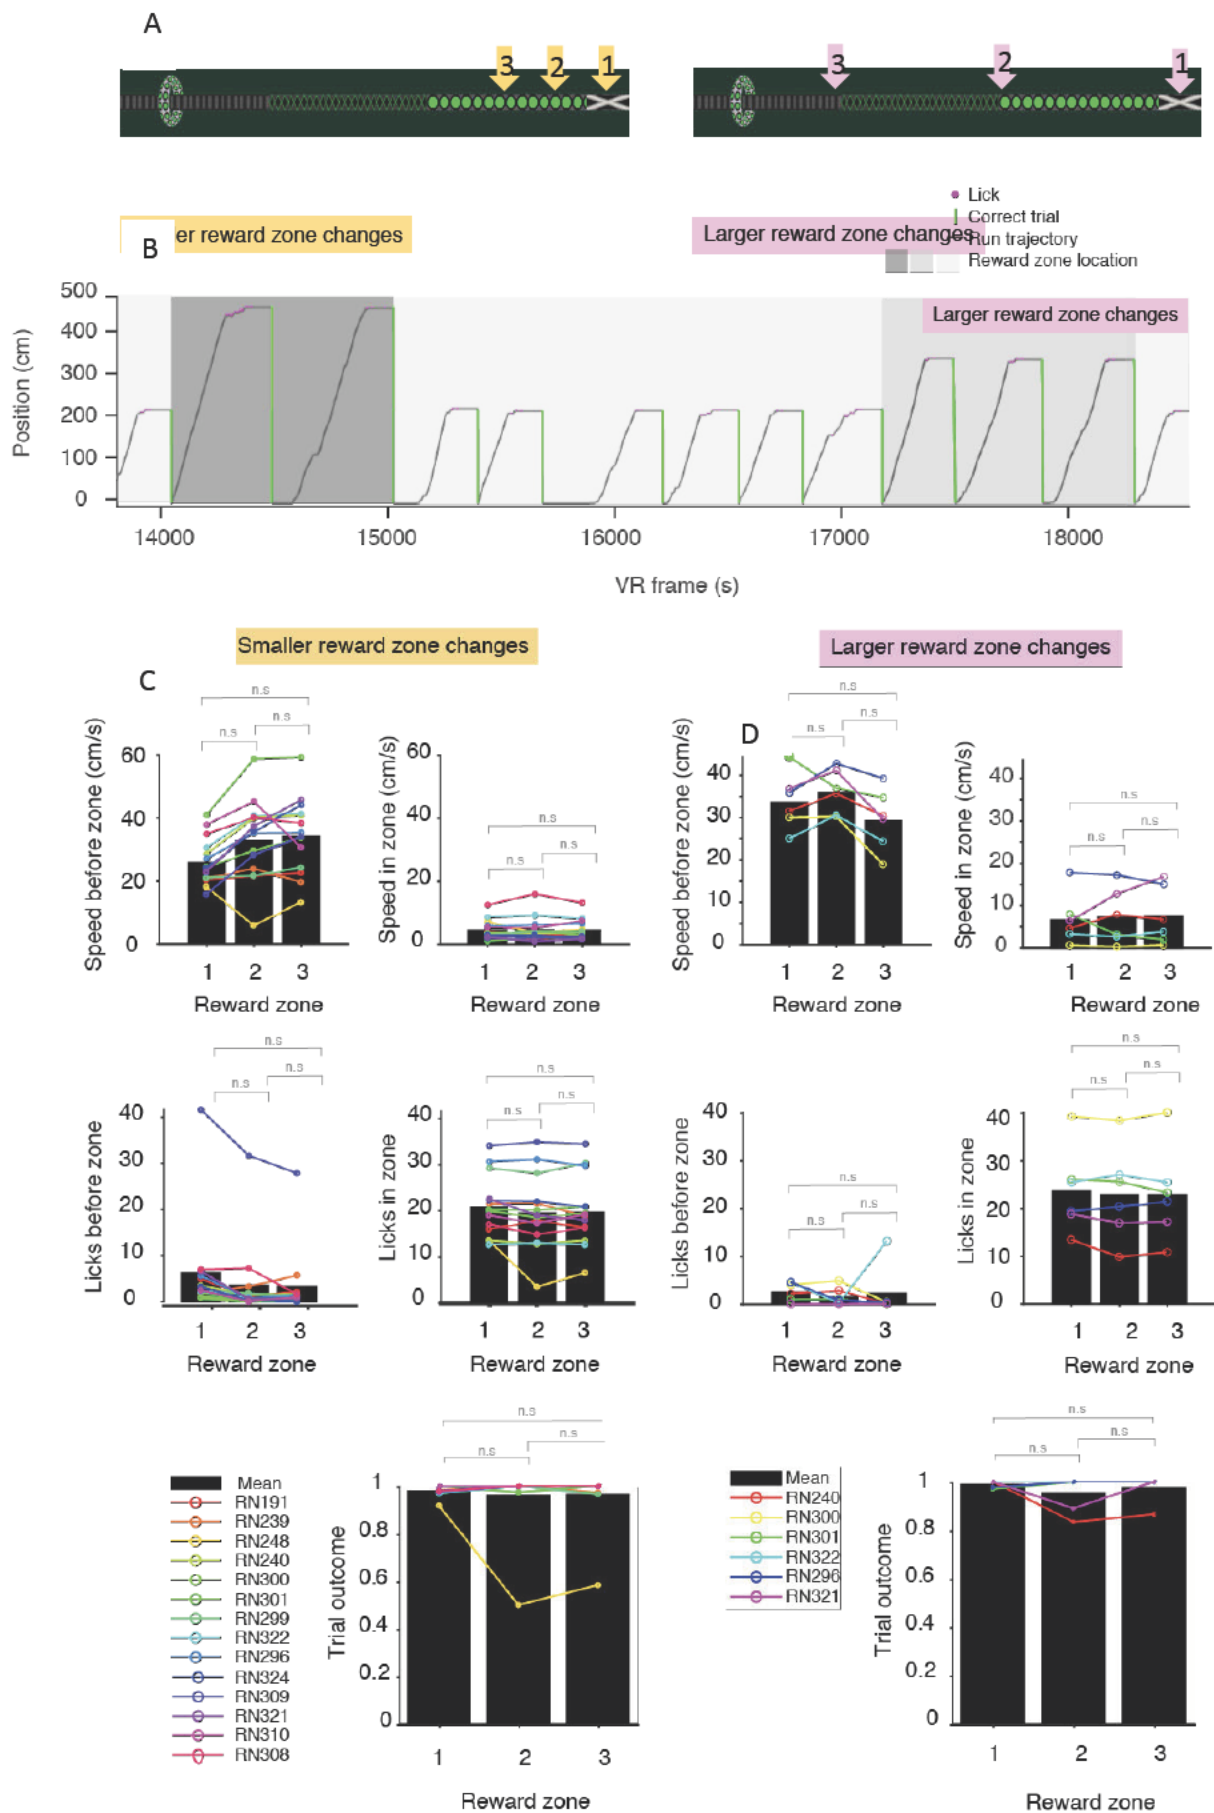
Figure 4-4. Mice are actively engaged during virtual navigation and use visual cues to complete the task.

A) Track schematics with labelled rewarded locations (coloured arrows) for each version of the task. Numbers identify the reward zone.

B) Example run (grey trajectory) and lick (pink dots) behaviours for 13 trials in one session. Green lines indicate correct trials and grey shaded areas represent the rewarded. Left) Run trajectory (grey line) through virtual space as a function of time (seconds). Green lines indicate correct trials and teleportation to the beginning of the track. Grey shaded areas indicate the reward zone position (1, 2, or 3).

C top) Mouse running speed (cm/s) before the reward zone (left) and in the reward zone (right). C middle) Mouse licking behaviour (lick count), before the reward zone (left) and in the reward zone (right). C bottom) Trial outcome, defined as proportion of trials correct for each of the reward zone locations. N = 14.

D) Same as in C, except the rewarded locations were in different location, spread along the track (see A). N = 6. Black bars are mean values and coloured lines represent data from individual mice. All Wilcoxon rank sum tests with Bonferroni correction were not significant (n.s.; $p > 0.05/3$).

4.2.8 Place cells

For our mini VR system to be useful experimentally, we need animals to not only be able to navigate in experimenter controlled virtual worlds, but we also need to be able to detect location specific activity in the hippocampus. Hence, we used 2P imaging to record from > 21 animals in CA1 and CA3 (Chapter 4). As a proof of principle, we were able to detect place cells as defined by criteria similar to those used by other authors (Dombeck et al., 2010; Bartos & Hainmueller , 2018). We identified 500 ROIs and found 69.4% (n = 347/500) met our place cell criteria (Figure 4.5). We define place cell criteria as those cells with place field width ≥ 20cm, place fields on ≥ 50% trials, and a mean in field response that is 2 x greater than the mean out of place field response (Dombeck et al., 2010). Cells that did not pass all 3 criteria were not defined as place cells. We found place cells mapped our virtual track and were more numerous at salient locations i.e. the end of the track (Figure 4.5 B). For a more detailed quantification of optically identified place cells see Chapter 5.

A

B

C

ROI #290

ROI #386
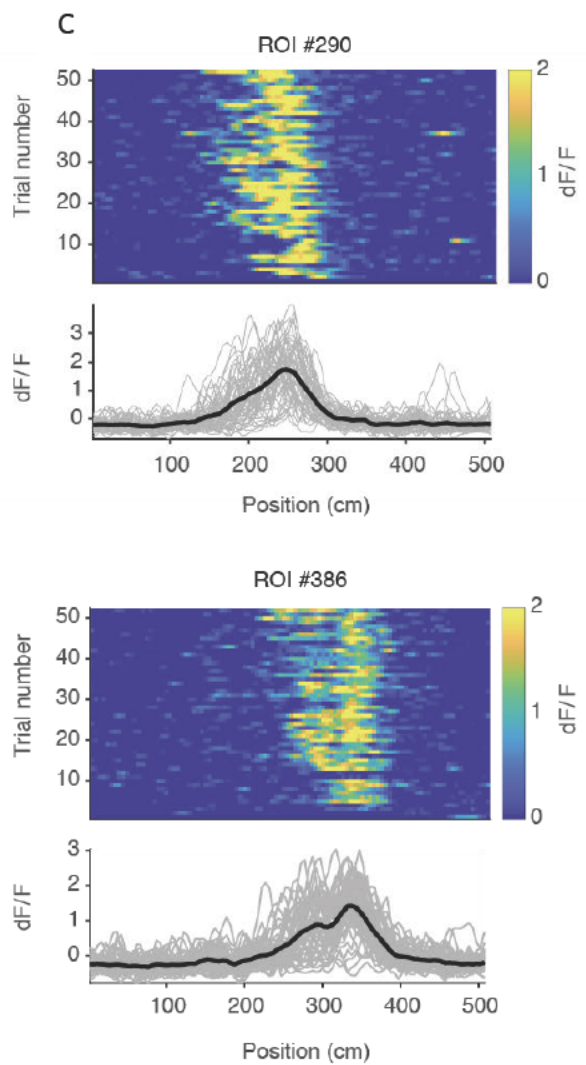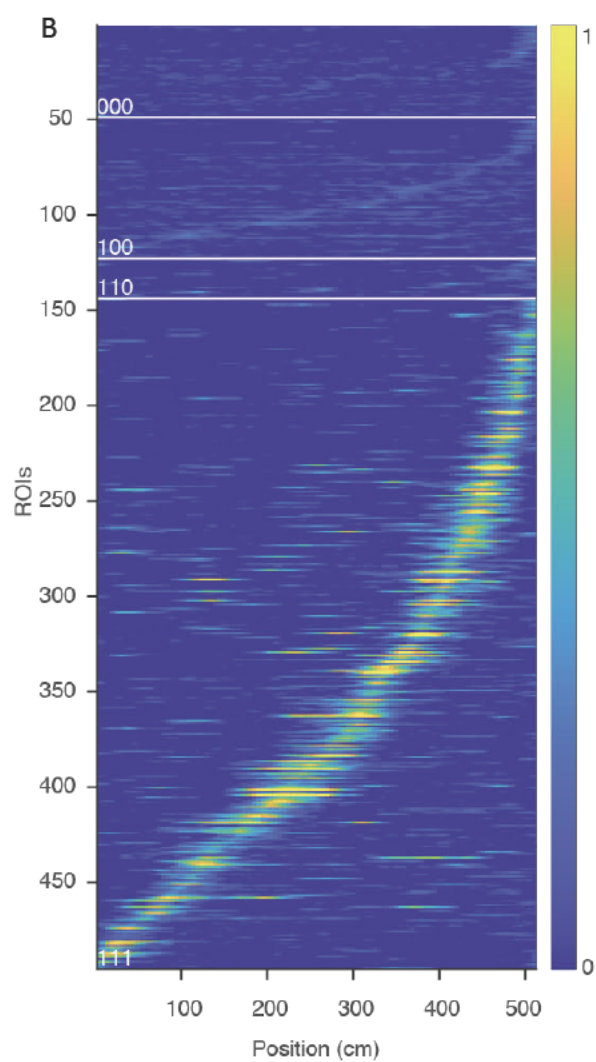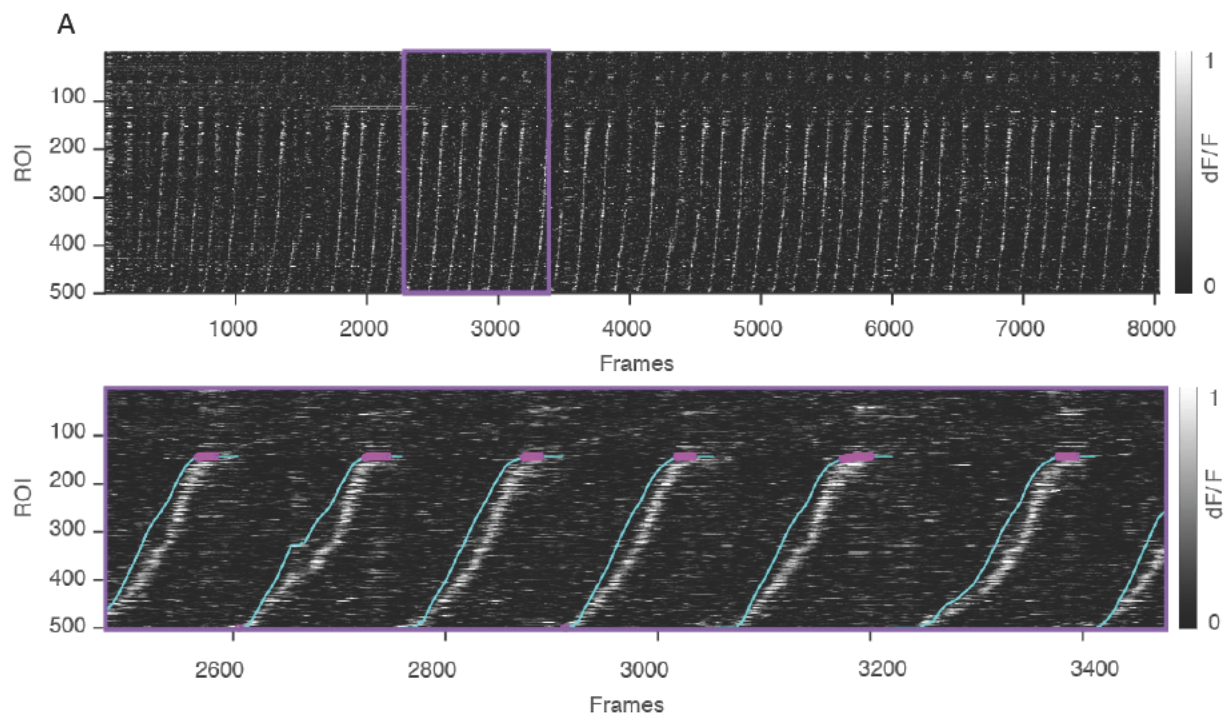
Figure 4-5 Optically identified place cells from navigation in a mini VR.

A)   Top: Heat map of ΔF/F traces for all detected ROIs (n = 500) across an hour recording. Lighter colour represent increased fluorescence and traces are sorted according to their peak preferred location. Purple square outlines the zoom in area below. Below) same as above but with running trajectory (blue) and lick behaviour (pink rings) overlaid.

B)   Heat map of ΔF/F traces for all detected ROIs (N = 500) sorted by place field preference along the 500cm virtual track. Warmer colours represent increased fluorescence. ROIs are also sorted according to the place cell criteria they meet (see white dividing lines), from those ROIs meeting 0 criterion (top; 000) to those meeting all 3 criteria (bottom; 111). Field size, response ratio (in field vs our field reliability, (% trials with response; Dombeck et al., 2010).

C)   Randomly selected place cells ROI 290, ROI 368. Top: Heatmap of normalized ΔF/F traces from that ROI across trials. Warmer colours represent increased fluorescence. Bottom) Single trial normalized ΔF/F traces (grey) with averaged traces (black) as a function of virtual position (cm).

4.3 Discussion

We developed a miniaturized VR system, capable of flexibly fitting on our microscope configuration (see photo in Figure 3.1 A) with the hardware (Figure 3.1 A insert). We found high variability in our 31 trained animals in both their run and lick behaviour, but on average an animal's performance plateaus after 10 sessions. This session threshold was similar to that previously reported (Harvey et al., 2009) and enabled us to create a training time scale, (10 sessions). We were also able to define expert behaviour along two metrics; as lick fidelity and run speed change into the reward zone. We found that our task was learnable for all animals, with most reaching experimental utility.

We were also able to optically identify place cells, in animals concurrently imaged during virtual navigation. These place cells were defined according to existing criteria and are similar to other studies (Dombeck et al., 2010; Hainmueller & Bartos, 2018; Robinson et al., 2020). Identification of place cells combined with task engaged behaviour indicate our animals are actively 'navigating' in our virtual world. Therefore, head-restrained mice can perform visually guided spatial behaviours in a mini VR environment. Our miniature system offers an affordable, space efficient alternative to typical more immersive VR configurations. It is modular and useful for a number of experiments beyond those described in this thesis. Full code is available opensource, for others in the science community to use and amend to their experimental needs.

Code repository: https://github.com/BNutbrown/mouseVR

5

CA1 & CA3 response dynamics to parametrically morphed worlds

5.1 Introduction

The heterogenous hippocampal subregions differ in their anatomical inputs, in addition to their internal circuitry. Together these differences in architecture support context-dependent learning and episodic memory. Although these differences, notably the recurrent connections between CA3 pyramidal neurons, have inspired multiple theoretical models of differential processing capacities of these regions, differences in the spatial properties of place cells in CA1 and CA3 in behaving animals are less well described. Electrophysiological recordings of CA3 population activity responded more coherently than activity in CA1 when real environment manipulations resulted in local-global reference frame conflicts (Lee et al., 2004). Simultaneous single unit recordings in CA3 and DG during a similar behavioural protocol found the CA3 representation of the altered environment was more similar to the familiar environment, compared to its DG inputs (Neunebel & Knierim, 2014). These results suggest that CA3 produced an output pattern closer to the originally stored representation than the degraded input patterns from the DG in the cue conflict environment, providing evidence for pattern completion of degraded inputs in the DG-CA3 circuit (Neunebel & Knierim, 2014).

5.1.1 Population properties in CA1 and CA3 networks

Pattern completion is generally considered to be a property of an attractor network (Rolls et al., 2013). Briefly, it refers to the network's ability to reinstate a previously stored pattern of activity from activation when cued with partial or degraded input. Coherent population dynamics, indicative of pattern completion have also been observed in CA1 during real world navigation in environments morphed between two end points (Wills et al., 2005). However, Leutgeb et al. found conflicting results suggesting gradual and abrupt changes to neural representations coexist in simultaneously recorded CA3 populations, such that progressive representation transformations were more likely when environments were manipulated incrementally. These studies highlight the impact of different training protocols, specifically whether morphed worlds

113

were presented incrementally or in a random order. Hence, the specific features of contexts driving remapping are debated. New work suggests that a more general principle for remapping of neural representations.

Other factors contributing to this apparent contradiction in CA3 attractor dynamics may stem from intra CA3 heterogeneity. Leutgeb et al. recorded from a small number of CA3 neurons (n = 38), from multiple CA3 intra-subregions, with most data likely from CA3c given the advancement of electrodes from the dorsal surface (Leutgeb et al., 2005). The opposing proximo-distal gradients of decreasing mossy fibre excitation (Claiborne, Amaral, & Cowan, 1986) and increasing excitation from cortical inputs (Ishizuka et al., 1995), combined with more abundant recurrent synaptic connections in distal CA3 (CA3a and CA3b; (Ishizuka et al., 1990), are consistent with the notion that direct cortical drive, recurrent excitation, and weak mossy fibre excitation, enables distal CA3 to participate in pattern completion (Lee & Kesner, 2004; Sun et al., 2017) Furthermore, models suggests weakly recurrent networks are less subject to attractor collapse, as population activity tends to track ongoing inputs, resulting in linear transitions along a morphing sequence (Papp, Witter, & Treves, 2007). Hence, it remains uncertain whether hippocampal assemblies can exist along continua, with stable or intermediate states emerging when incoming input is incongruent with stored representations. CA3 circuits, at least in CA3a and b have the anatomical attributes to serve encoding and retrieval of memory (Rolls, 2013).

To experimentally assess the interpretation that CA3 functions as an associative attractor network, we must probe recall properties, through titrated input to the network and assess the local as well as downstream CA1 responses. To this end, we will record from a large population of genetically defined CA3 neurons and CA1 neurons in animals trained to navigate in a virtual world. We will assess CA1 and CA3 network responses (simultaneously in some animals) during navigation in a novel world, in addition to intermediate environments precisely morphed between these end point (familiar and novel) worlds, to probe different hippocampal subregion spatial representations to different degrees of overlap in input patterns.

5.2 Results

5.2.1 Defining hippocampal subregions

To assess the possibility of non-linear response dynamics in different hippocampal subregions, we used 2P volume imaging to record hippocampal neurons transgenically expressing GCaMP6s in subregion CA1 (n = 5, CA1 mice) and CA3 (n = 7 CA3 mice) in vivo during virtual navigation. We also imaged CA1 and CA3 simultaneously in a subset of animals (n = 5 dual area mice; see Figure 4.1 B & C for example dual FOV). To define our 'dual area' FOV we used the morphological characteristics of CA3 neurons (characterised in Chapter 2) to identify this subregion and positioned our FOV such that extending 400μm axially over 4 planes, it included that CA3 FOV and subregion CA1 on the more superficial planes (Figure 4.1 A). We manually curated our FOVs into separate the two subregions (Figure 4.1 C − white tack lines). We excluded a 150μm boundary region and found that the average somatic CA3 ROI size in 'dual' FOV animals was similar, but slightly smaller to that of genetically defined CA3 neurons (mean ± STD; dual: 17.71μm ± 1.11 vs. CA3 transgenic; 22.06μm ± 2.67, P = 0.005, Wilcoxon rank sum test, Figure 4.1 D). For dual area mice, we included these curated CA1 and CA3 ROIs to their respective areas for subsequent group analysis. The somatic diameters of CA3 ROIs were significantly larger than those in CA1 (CA1: 14.98μm ± 0.68 vs. CA3: 20.25μm ± 3.06; Figure 4.1 E). We recorded from more CA1 than CA3 neurons (mean ROIs per animal in CA1 mice 2188.1 ROIs ± 617.47, CA3 mice, 315.5 ROIs ± 169.23; Figure 4.1 F), given the angled nature of CA3 anatomy relative to the imaging plane (Chapter 2.2.2; Bartos & Hainmueller, 2018; Schoenfeld et al., 2021).
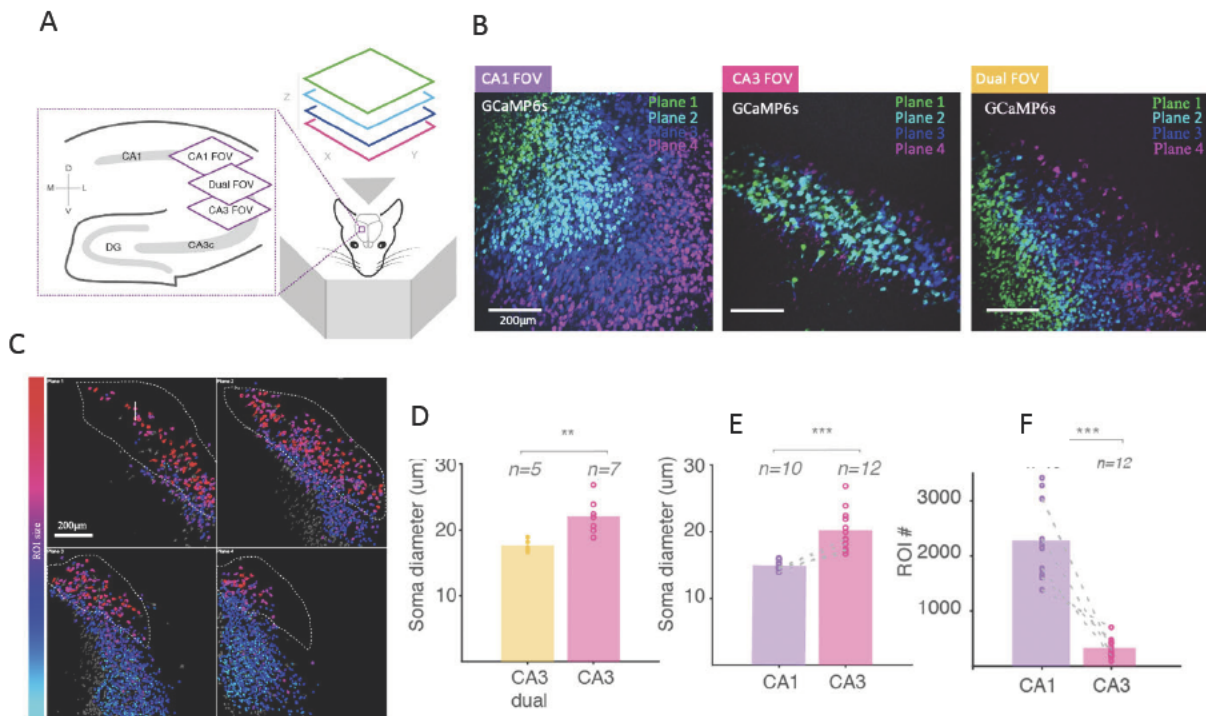
Figure 5-1. Two-photon volume imaging in hippocampal subregions CA1 and CA3 in vivo.

A) Schematic showing 4 plane 2P volume imaging in the hippocampus during virtual navigation, left inset (purple track line zoom in), schematic coronal section with FOV locations marked. Not to scale.

B) Example FOVs from CA1 (left- purple), CA3 (middle- pink) and dual area (right- yellow) mice. Expression image: colour = plane imaged (dorsal to ventral: plane 1-green, 2-light blue, 3- dark blue, 4- pink).

C) Example animal with a 'dual' area FOV. ROIs are coloured by size, where warmer colours indicate larger ROIs and overlaid white track lines ring the area manually curated as subregion CA3.

D) Mean soma diameter (µm) of CA3 ROIs for ROIs defined in CA3 transgenic animals (pink rings) and CA3 ROIs defined from the CA3 subregion in 'dual area' animals (yellow rings). Bars = mean soma diameters.

E) Mean ROI soma diameter size in CA1 (purple rings, n = 10) and CA3 (pink rings, n = 12). Bars represent mean soma diameter size for each subregion (CA1 and CA3, including ROIs from dual FOV animals).

F) Mean ROIs for CA1 (purple rings) and CA3 (pink rings) mice. Bars represent mean ROIs per subregion.

Grey track lines labels mean ROI numbers in dual mice. All scale bars are 200µm. All Wilcoxon rank sum tests were significant (**, $P > 0.01$, ***, $P > 0.001$).

116

To complete the task, animals navigated down a virtual linear track (Figure 4.2 B) for a water reward, initially in a familiar green world (world 1; W1) and then a novel blue world (world 5; W5), followed by 3 worlds derived from parametric interpolation of W1 and W5. We termed these 'intermediate worlds' (worlds 2-4; W2-4; Figure 4.2 A & B), and all worlds were presented in a randomised manner. For the full training procedure see Methods & Materials (Chapter 5.6).
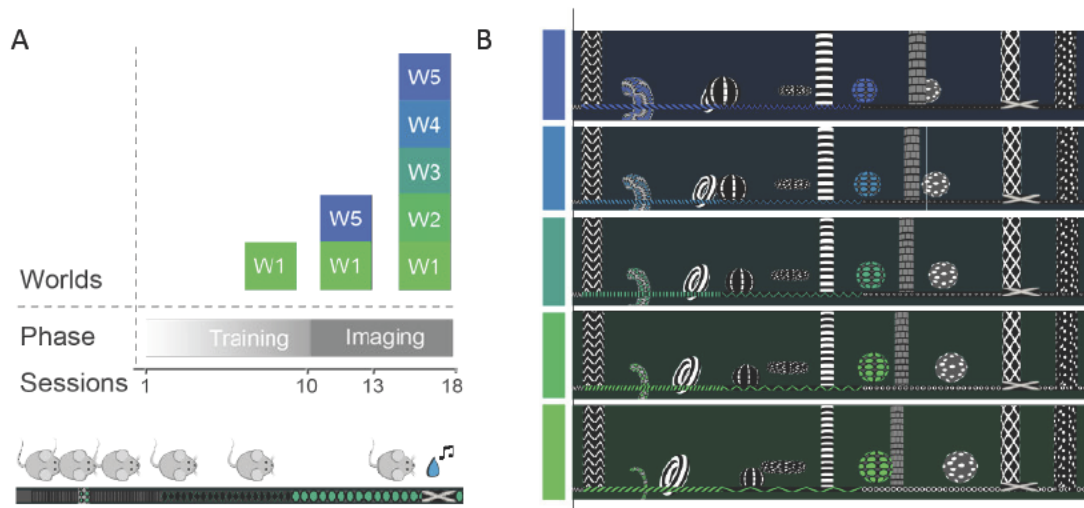
Figure 5-2. VR morphing worlds behaviour protocol.

A) Schematic of the training and experiment protocol (phase). Briefly, after 10 training sessions, mice entered the experimental phase, where concurrent 2P imaging was performed during navigation for a water reward in worlds 1-5 as shown.

B) Side profile in experimental worlds 5-1 shown showing morphed colour, object position, wall textures. Reward Zone is marked by a cross and animals were required to stop and lick in this zone to initiate a reward. Morphing floor textures are not shown in this view (see Chapter 2 for more details on world design).

5.2.2 Experiment 1: Single World.

5.2.2.1 Comparing the spatial properties of neurons in CA1 and CA3 within and between animals

There were no marked differences in the level of performance between the three groups of mice (Figure 4.4 A) in terms of trial durations (CA1 mice = 50.88s ± 13.48 vs. CA3 48.98s ± 14.23, vs dual mice = 64.78s ± 38.83, P = 0.89, Wilcoxon rank sum test, Figure 4.4 C). The animals also exhibited no statistical differences in behavioural performance, a correct trial was defined as the mouse slowing down and licking in the reward zone to initiate the reward (CA1 mice: 94% trials correct ± 1.2 vs. CA3 93 ± 0.06, vs Dual 0.91 ± 0.05, P = 0.89, Wilcoxon rank sum test; Figure 4.4 D). Interestingly, we found mean CA1 place cell population activity was positively correlated with running speed, whereas CA3 place cell population activity was negatively correlated (CA1, r = 0.06 ± 0.03, CA3, r = -0.04 ± 0.03, P < 0.001, Wilcoxon rank sum test; Figure 4.3), suggesting that at higher running speeds, the CA3 population is less active and vice versa in CA1.

To classify place cells - neurons with clearly defined spatial modulation - we created a null distribution of spatial information scores for each cell from their circularly permuted time series (we shifted the position relative to the activity trace within each trial). We shuffled 1000 times per ROI and ROIs exceeding 99% of permutations were considered to be significant 'place cells' (see Giocomo & Plitt, 2019 for details). Additionally, those 'place cells' were required to have at least one clearly defined place field (see Methods & Materials) in order for them to be included in our subsequent analyses (see Figure 4.4 for example place maps [E] and place cells[F]).
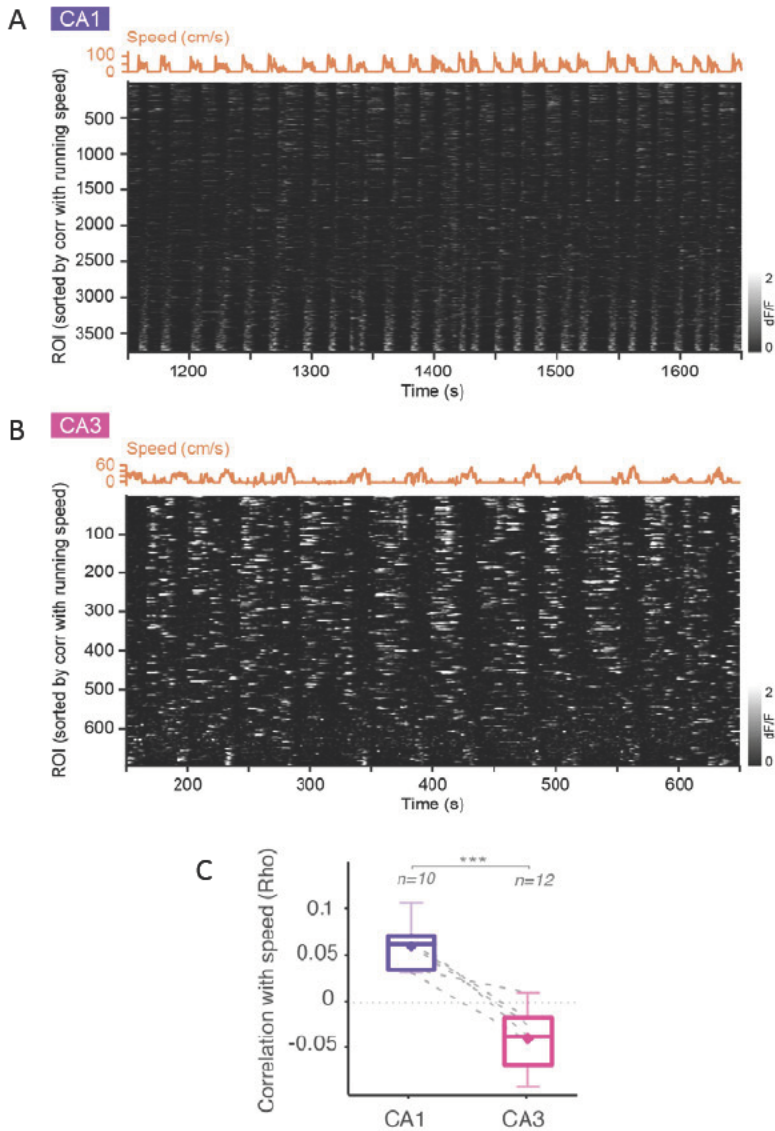
Figure 5-3 CA1 and CA3 neurons show an opposite correlation with animal running speed.

A) Top: Running speed trace (cm/s; orange) over time in example CA1 mouse. Below: Heatmap of neural activity per ROI over time (lighter colour = higher activity; normalised mean ΔF/F), sorted by increasing correlation with animal running speed (R).

B) Same as in A) but for example CA3 activity.

C) Neural activity correlation with running speed presented as box-and-whisker plots displaying median (line), mean (dot) interquartile (box) and 90% ranges (whiskers) for each subregion.

Grey track lines label dual mice (n = 5). All *s indicate Wilcoxon rank sum tests (***, *P* < 0.001).
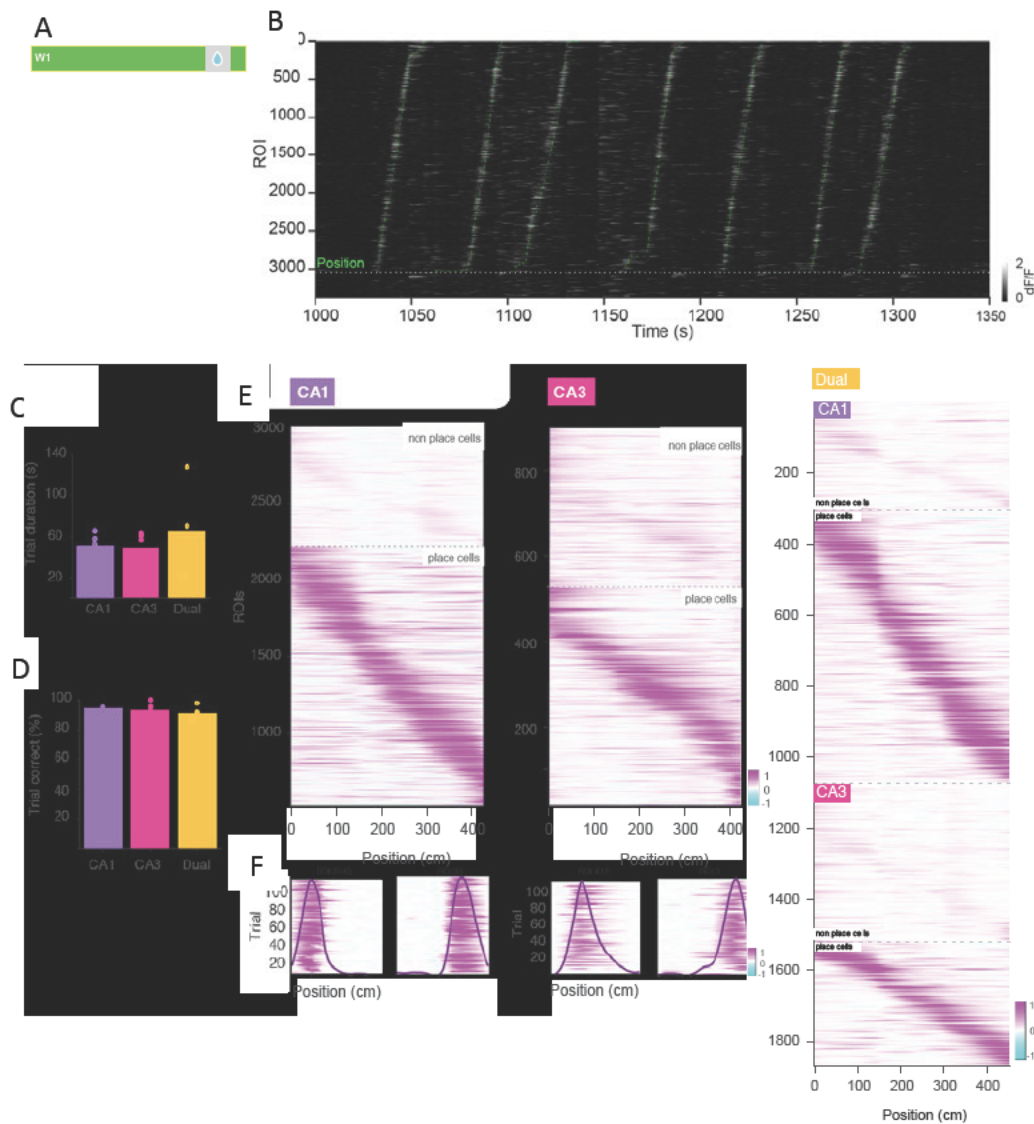
Figure 5-4. Animals' spatial behaviour and neural activity in subregions CA1 and CA3 during virtual navigation.

A) Task schematic: animals navigated down a virtual track for a water reward (W1 = green).

B) Example heatmap of all ROI activity (normalised ΔF/F) over time (s) during virtual navigation in W1 where lighter colour = higher activity. Animal position trajectories over time are overlaid (green track lines).

C) Average trial duration for each mouse category (CA1 [purple], n= 5, CA3= [pink] n = 7, dual [yellow] n = 5), defined as the time taken (in seconds) to run one lap of the virtual track for a water reward. All Wilcoxon rank sum tests were non-significant (n.s.; *P* > 0.05).

D) Mean trial outcome (% correct) for each mouse group (CA1 [purple], n= 5, CA3= [pink] n = 7, dual [yellow] n = 5). All Wilcoxon rank sum tests were non-significant (n.s.; *P* > 0.05).

E) Example place maps for each mouse category (CA1 left, CA3 middle, dual right). Plots show all ROIs detected in an example animal. Grey dotted line dividing (above) 'non place cells' and (below) 'place cells'. Each row in the heatmap indicates mean ΔF/F of that cell as a function of position in virtual space, traces are sorted by place cell peak preferred location on the track.

F) Example mean ΔF/F trace (right: CA3, left CA1) as a function of position (dark purple line) overlaid on a normalised ΔF/F heatmap per trial for that ROI (warmer colours indicate higher ΔF/F).

Using our place cell detection method, we identified more place cells as a proportion of total recorded ROIs in CA1 compared to CA3 (0.58 ± 0.14 vs. 0.47 ± 0.16, *P* < 0.05, Wilcoxon rank sum test; Figure 4.5 A). Place cells in CA1 and CA3 mapped the track, with different place fields preferring different locations along the virtual track (Figure 4.5 B). Although the ends of the track were overrepresented by CA3 place cells (Figure 4.5 B), i.e., fewer CA3 place cells had place fields in the middle of the track than at the beginning or end of the track. To quantify this we took the mean proportion of cells with fields at the ends (the first 5 cm spatial bin from each end) and compared that to (by subtracting from) the proportion of cells with fields in the middle of the track (middle 15 cm; CA1, 0.01 difference in proportion of end-vs-middle place cells ± 0.03, CA3, 0.12 ± 0.13, *P* < 0.005, Wilcoxon rank sum test; Figure 4.5B, insert plot). CA1 and CA3 place cells were also similar in that they mostly had one place field (CA1: 97.4% ± 1.56, CA3: 98.1% ± 1.46 place cells have 1 place field, *P* = 0.07, Wilcoxon rank sum test; Figure 4.5 C) with only a few having multiple fields (CA1, 2.58% ± 1.54 CA3, 1.30% ± 1.46 place cells have 2 place fields, *P* = 0.07, Wilcoxon rank sum test; CA1, 1.37 ± 2.91, CA3, 0, place cells with 3 place fields, *P* = 0.07, Wilcoxon rank sum test).

When averaging within each subregion (CA1 and CA3) across animals, we found no statistical differences in the average place field width, reliability or magnitude between subregions (Figure 4.5 D-F). However, when pooling all cells across animals we observed that CA3 place cells had slightly smaller place field widths, compared to CA1 place cells (CA3, 21.46cm ± 7.09, CA1, 22.50cm ± 6.83, P < 0.001, Wilcoxon rank sum test; Figure 4.5 G) but larger field amplitudes (CA3, 2 .27 ΔF/F ± 1.56, vs. CA1, 1.68 ΔF/F ± 0.92, P < 0.001, Wilcoxon rank sum test; Figure 4.5 H). We

121

found CA3 place cells to be less reliable (less spatially stable) than CA1 place cells, defined as the correlation between the place cell preferred position in the first half of the trial vs. the second half of that trial. CA3 place cells were present on a lower proportion of trials (CA3, 0.34 r ± 0.2, CA1, 0.38 r ± 0.23, P < 0.001, Wilcoxon rank sum test; Figure 4.5 I). Although spatial information scores (see Methods & Materials 5.17 for calculation details) were not statistically different between CA1 and CA3 place cells (CA1, 0.08 bits per event ± 0.05, CA3, 0.09 bits per event ± 0.05, P = 1, Wilcoxon rank sum test; Figure 4.5 J).

Finally, to investigate the spatial clustering of functionally similar place cells in order to assess the potential for topographic place maps, we looked at pairwise correlations of place cell activity. We found CA3 place cells to be more correlated with neighbouring cells, whereby the mean pairwise correlation between the closest cells is highest and declines with distance (CA1, r = 0.05 ± 0.01 vs. CA3, r = 0.15 ± 0.09, *P* < 0.001, Wilcoxon rank sum test; Figure 4.5 K insert plot) although, the fall off is slower in CA3. To quantify this, we fitted an exponential curve to the correlation vs distance profiles and compared the full width at half maximum (FWHM) of these lines (CA1, 1.43µm ± 0.05, [n = 10], CA3. 1.54µm ± 0.12, *P* < 0.01, Wilcoxon rank sum test; Figure 4.5 K inset plots).
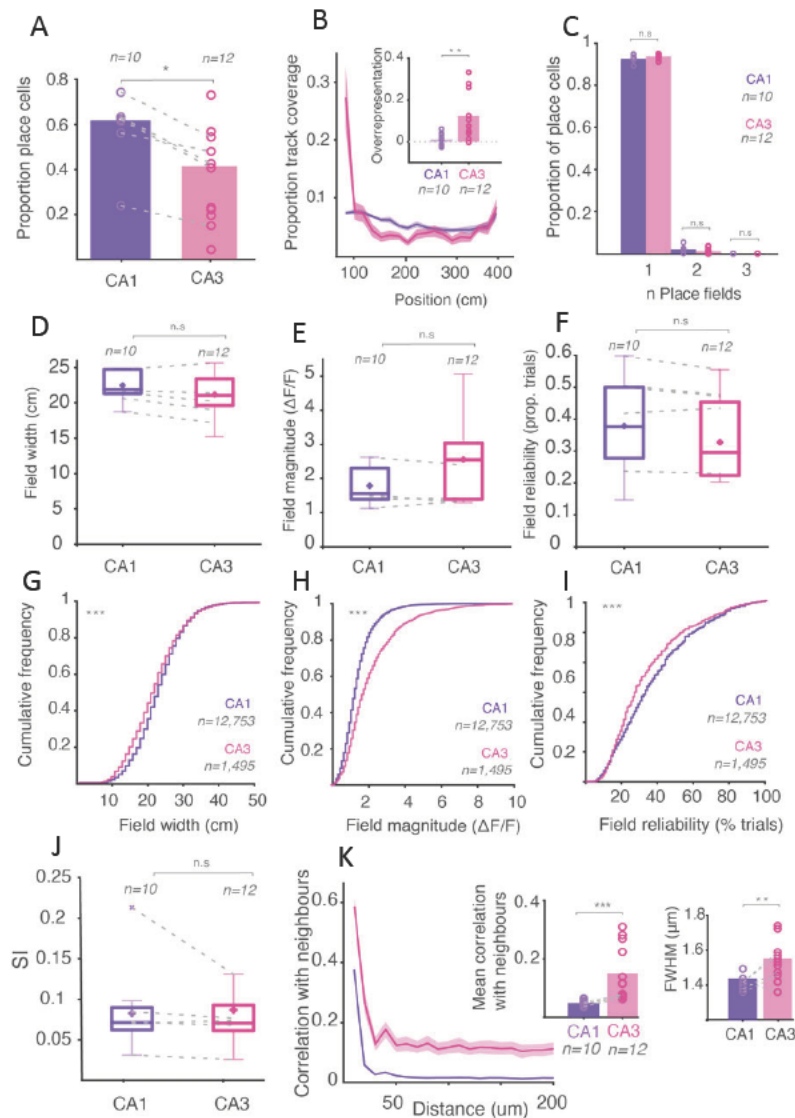
Figure 5-5 Comparing the spatial properties of neurons in subregions CA1 and CA3 within and between animals.

A) Proportion of ROIs defined as place cells for CA1 (purple rings) and CA3 (pink rings), bars indicate means.

B) Proportion of place cells with preferred peak locations at a given animal position along the track for CA1 (purple) and CA3 (pink) place cells, shaded areas = S.E.M. Inset) Mean end of track overrepresentation in CA1 and CA3.

C) Proportion of place cells as a function of number of place fields for CA1 (purple rings) and CA3 (pink rings) subregions, bars indicate mean proportion of place cells with that field number.

D) Place cell field widths presented as box-and-whisker plots displaying median (line), mean (dot) interquartile (box) and 90% ranges (whiskers) for each subregion

E) Same as in D) but for place field response magnitude.

F) Same as in D) but for place field reliability over trials.

G) Cumulative frequency distribution as a function of ROI field width for CA1 (purple) and CA3 (pink).

H) Cumulative frequency distribution as a function of place field ROI response magnitude for CA1 (purple) and CA3 (pink).

I) Cumulative frequency distribution as a function of place field ROI reliability across trials for CA1 (purple) and CA3 (pink).

J) Spatial information (SI) scores presented as box-and-whisker plots displaying median (line), mean (dot) interquartile (box) and 90% ranges (whiskers) for each subregion.

K) Mean pairwise correlation of place maps as a function of distance from neighbouring ROIs, inset 1) Mean pairwise correlation averaged across all distance bins for each subregion, inset 2) full width at half maximum (FWHM) per subregion.

Grey track lines label dual mice in each subregion category (n = 5). All scale bars are 200μm. All *s indicate Wilcoxon rank sum tests (n.s. non-significant, *$P < 0.05$ **, $P < 0.01$, ***, $P < 0.001$).

5.2.3 Experiment 1: Two Worlds.

5.2.3.1   Comparing the spatial properties of neurons in CA1 and CA3 in the familiar and a novel environment.

When introduced to a novel distinct blue world (W5), that was different from W1 in colour, texture and object locations, changing the visual properties of the environment but not the size or reward location (Figure 4.6 A), animals' behavioural performance was not affected. Animals had no statistical differences in trial durations (CA1 W1: 51.41s ± 16.03 vs CA1 W5: 46.76s ± 20.78, $P = 0.43$, CA3 W1: 53.65s ± 16.50, vs CA3 W5: 50.38s ± 23.52, P = 0.44, CA1 W1 vs CA3 W1, $P = 0.60$, CA1 W5 vs CA3 W5 $P = 0.80$, Wilcoxon rank sum test, Figure 4.6 B) and trial outcomes (% trials correct; CA1 W1: 90.60% ± 6.04, CA1 W5: 86.06% ± 10.88, $P = 0.45$, CA3 W1: 91.17% ± 5.66, CA3 W5: 85.20% ± 10.01, $P = 0.15$, CA1 W1 vs CA3 W1, $P = 0.87$, CA1 W5 vs CA3 W5 $P = 0.95$ Wilcoxon rank sum test; Figure 4.6 C).
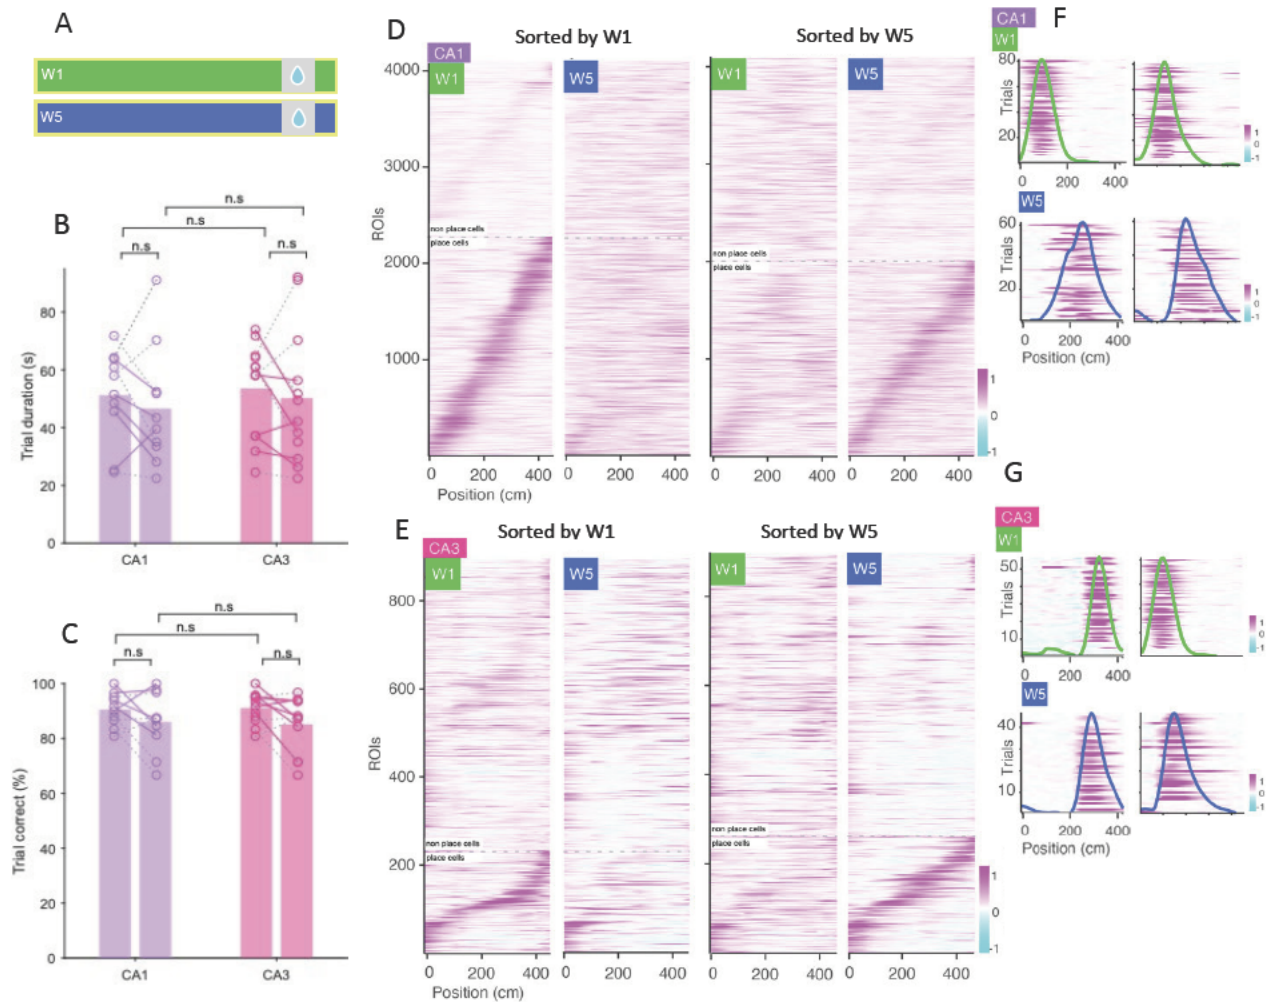
Figure 5-6 Spatial behaviour and neural activity in CA1 and CA3 in novel and familiar virtual worlds.

A) VR track schematic to represent W1 and W5 experimental phase.

B) Average trial duration for each mouse category (CA1 [purple], n= 10, CA3= [pink] n = 12), defined as time taken (s) to run 1 lap for a water reward. All Wilcoxon rank_sum tests were non-significant (n.s.; *P* > 0.05).

C) Mean trial outcome (% correct) for each mouse group (CA1 [purple], n= 10, CA3= [pink] n = 12). All Wilcoxon rank_sum tests were non-significant (n.s.; *P* > 0.05).

D) Example ROI activity maps (normalised mean ΔF/F) in W1 (left) and W5, sorted by peak fluorescence (ΔF/F) as a function of VR position in W1 (left two plots) and sorted by peak activity as a function of VR position in W5 (right two plots). Grey dotted line divides (above) 'non place cells' and (below) 'place cells'.

E) Same as in A, but for CA3 ROI activity.

F) Example trial-wised average (ΔF/F) trace as a function of position in world 1 (green line) and world 5 (blue line) for the same CA1 ROI across worlds.

G) Same as in C but for CA3 ROIs.

We recorded neural activity in CA1 and CA3 while animals where navigating through the two worlds (W1 and W5) in a random interleaved trial structure. We constructed average place activity maps for the population activity in each of the two worlds, and when we sorted one world's map by the peak activity locations in the other world. We observed the activity in W5 appeared qualitatively different to in W1 for both CA1 and CA3 populations (see Figure 4.6 D & E). To explore this, we found a higher proportion of ROIs detected to be place cells in CA1 than CA3 in both W1 and W5 (CA1 W1: 55% ± 9, CA1 W5: 39% ± 12, CA3 W1: 29% ± 13, CA3 W5: 24% ± 14. 2-way ANOVA main effect of area $F_{(1,1)}$ = 29.9, $P < 0.001$. Figure 4.7A) and both subregions had a higher proportion of place cells in W1 than in W5 (2-way ANOVA main effect of world $F_{(1,1)}$ = 7.7, $P = 0.009$; Figure 4.7 A). However, the cumulative proportion of place cells detected increased in W5 for both CA1 and CA3 animals (CA1 W1: 55% ± 9, CA1 W5 69% ± 9, CA3 W1: 29% ± 13, CA3 W5: 43% ± 16. 2-way ANOVA main effect of world $F_{(1,1)}$ = 50.4, $P < 0.001$; Figure 4.7 B) indicating that W5 is represented by a different, or at least additional, population of place cells (see Figure 4.8 for further analysis).

We found CA3 place fields overrepresented the beginning of the virtual track, whereas CA1 place fields uniformly tiled the whole track (Proportion of place cells preferring locations <= 50 cm; CA1 W1: 6% ± 1, CA1 W5: 7% ± 2, CA3 W1: 11% ± 7, CA3 W5: 12% ± 6; Figure 4.7 C). These place field coverage distributions were consistent between novel (W5) and familiar (W1) worlds (Figure 4.7 C). The average S.I. scores across all neurons in each subregion were not statistically different between worlds and region (CA1 W1: 0.11 bits per event ± 0.05, CA1 W5: 0.11 ± 0.06, CA3 W1: 0.11 ± 0.05, CA3 W5: 0.10 ± 0.04; Figure 4.7 D). Taken together, these findings suggest that the novel world (W5) is potentially represented with a similar fidelity.

However, at the level of single cells, we found CA3 place fields to be less reliable than CA1 place fields, in addition to lower reliability scores in both subregions in the novel world (CA1 W1: 38%

± 22, CA1 W5: 33% ± 18, CA3 W1: 41% ± 24, CA3 W5: 34% ± 20; Figure 5.7 E). CA3 place cells had slightly narrower place fields, compared to CA1 place cells, with both subregions having narrower place fields in the novel world (CA1 W1: 23.64cm ± 7.07, CA1 W5: 24.74cm ± 7.06, CA3 W1: 22.75cm ± 7.20, CA3 W5: 24.12cm ± 7.96; Figure 5.7 G). We found CA3 had higher in field activity magnitudes than CA1 place cells and both subregions have smaller field magnitudes in the novel world (CA1 W1: 1.56 ΔF/F ± 0.84, CA1 W5: 1.84 ΔF/F ± 0.98, CA3 W1: 1.96 ΔF/F ± 1.41, CA3 W5: 2.33 ΔF/F ± 1.55; Figure 5.7 F).
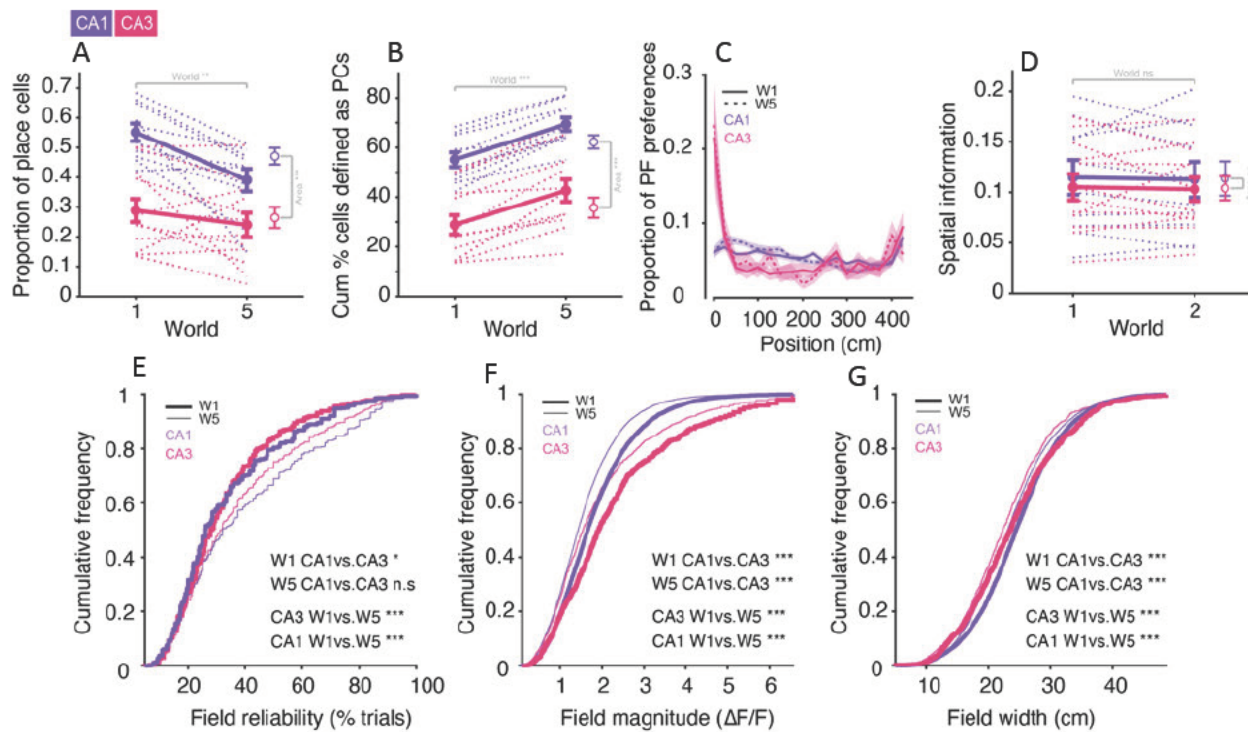
Figure 5-7. Comparing the spatial properties of neurons in subregions CA1 and CA3 in novel and familiar worlds.

A) Proportion of ROIs defined as place cells for CA1 and CA3 animals.

B) Cumulative proportion of place cells in W1 in CA1 and CA3 and in W1 and W5 (W5) for CA1 and CA3.

C) Proportion of place cells with preferred peak locations at a given animal position along the track for CA1 and CA3 across W1 (solid lines) and W5 (track lines).

D) Mean SI scores for CA1 and CA3 across W1 and W5 (thick line).

A – D CA1 (purple), n= 10, CA3 (pink) n = 12.

E) Cumulative frequency distribution as a function of place field ROI reliability across trials for CA1 (purple) and CA3 (pink). CA1 W1 n = 12,222, CA1 W5 n = 9,066, CA3 W1 n = 1,130, CA3 W5 n = 786 neurons.

F) Cumulative frequency distribution as a function of place field ROI response magnitude for CA1 (purple) and CA3 (pink). CA1 W1 n = 12,222, CA1 W5 n = 9,066, CA3 W1 n = 1,130, CA3 W5 n = 786 neurons.

G) Cumulative frequency distribution as a function of ROI field width for CA1 (purple) and CA3 (pink). CA1 W1 n = 12,222, CA1 W5 n = 9,066, CA3 W1 n = 1,130, CA3 W5 n = 786 neurons.

Shaded areas and error bars = S.E.M. Track lines = individual animals. Thicker solid lines in top row = means. A-D all *s indicate 2-way ANOVA main effects. E-G all *s indicate Wilcoxon rank sum tests (n.s. = non-significant, *$P < 0.05$ **, $P < 0.01$, ***, $P < 0.001$).

To further investigate the extent to which CA1 and CA3 populations represented the novel and familiar worlds distinctly, we quantified the proportion of cells defined exclusively in each world; cells that were only active in W1, only active in W5 or active in both. As before CA1 had more place cells than CA3 and we found there were more W1 (familiar) preferring cells than W5 (novel) preferring cells (Figure 4.8 A & B; 2-way ANOVA main effect of world: $F_{(1,1)} = 9.8$ $P < 0.001$, main effect of area: $F_{(1,1)} = 20.6$ $P < 0.001$). We next asked if cells that were active in both worlds were active in the same location (Figure 4.8 C for an example). We found only a weak correspondence between the preferred firing locations in W1 and in W5 for both CA1 and CA3 (correlation of CA1 W1 field locations with W5 field locations: 0.20 r ± 0.04, CA3 W1 with W5: 0.26 r ± 0.12; Figure 4.8 D & E).

To look at this in another way, we identified the preferred firing locations in W1 and looked at each cells' activity in that same location in W5 (Figure 4.8 F for examples). We found the response magnitude of CA1 place cells averaged across all laps (trials) was larger than that for CA3 place cells (response in W1 fields in W1 CA1: 0.31 ΔF/F ± 0.13, CA3: 0.19 ± 0.10; 2-way ANOVA W1 field locations main effect of area $F_{(1,1)} = 7.9$, $P = 0.008$; Figure 4.8 G & H; response in W5 fields in W5 CA1: 0.27 ΔF/F ± 0.12, CA3: 0.16 ΔF/F ± 0.12; 2-way ANOVA for responses in W5 field locations main effect of area $F_{(1,1)} = 6.12$, $P = 0.018$; Figure 4.8 G & H), which is explained by the greater reliability of CA1 place cells (CA1 W1: 0.38 ±0.22, W5: 0.33 ±0.18 vs CA3 W1: 0.41 ±0.24, CA3 W5: 0.34 ±0.20, proportion of trials with significant response) despite the larger in-field responses of CA3 cells (CA1 W1: 1.56 ±0.84, CA1 W5: 1.84 ±0.98 vs. CA3 W1: 1.96 ±1.41 CA3 W5: 2.33 ±1.55,

128

ΔF/F on significant response trials; see Figure 4.7 E). Both CA1 and CA3 place cells were not active in the same location across worlds (response in W1 fields in W5 in CA1: -0.01 ΔF/F ± 0.04, CA3: -0.04 ΔF/F ± 0.05. Response in W5 fields in W1 in CA1: -0.00 ΔF/F ± 0.04, CA3 W1: -0.02 ΔF/F ± 0.04. 2-way ANOVA for response in W1 field locations main effect of world $F_{(1,1)} = 110.5$, $P < 0.001$; Figure 4.8 G & H).

To quantify the degree of similarity of the world population representations, we correlated all single trial place maps with the trial-averaged map from W1 and W5 (Figure 4.8 I). Again, we found that both CA1 and CA3 formed distinct place map representations of the two worlds whereby the W5 trials were distinct from the W1 map and vice versa (average correlation for W5 trial maps with W1 average map, CA1: 0.07 r ± 0.06, CA3: 0.09 r ± 0.03. Average correlation of W1 trial maps with W5 average map, CA1: 0.08 r ± 0.06, CA3: 0.09 r ± 0.03. 2-way ANOVA for correlations with W1 average map $F_{(1,1)} = 110.4$, $P < 0.001$). We found no difference in spatial representation similarity between CA1 or CA3 (2-way ANOVA for correlations with W1 average map, no main effect of area $F_{(1,1)} = 2.1$ $P = 0.15$). We also developed a metric to assess the probability a given trial is more similar to the W1 representation than W5; on each trial we computed the correlation of that trials place map to the average W1 map, and from that subtracted that trials correlation to the average W5 map. From this difference, trials with values greater than 0 are more similar to the W1 representation, and scored as 1 (i.e. they are using the W1 map), and values less than zero are more similar to the W5 representation, and scored as 0 (i.e. using the W5 map). We then average these binarized trial scores to get an overall probability that trials in a given world have a neural representation more similar to the W1 world or the W5 world. We term this metric the 'binarized end-point similarity' and in our implementation it can be thought of as akin to a world 'decoder'. We found that areas respond similarly (2-way ANOVA no main effect of area $F_{(1,1)} = 0.2$, $P = 0.77$; Figure 4.8 L), in that as worlds become more dissimilar from W1, the probability that the neural representation matches W1 decreases (2-way ANOVA main effect of world $F_{(1,1)} = 10027.4$, $P < 0.001$; Figure 4.8 L).
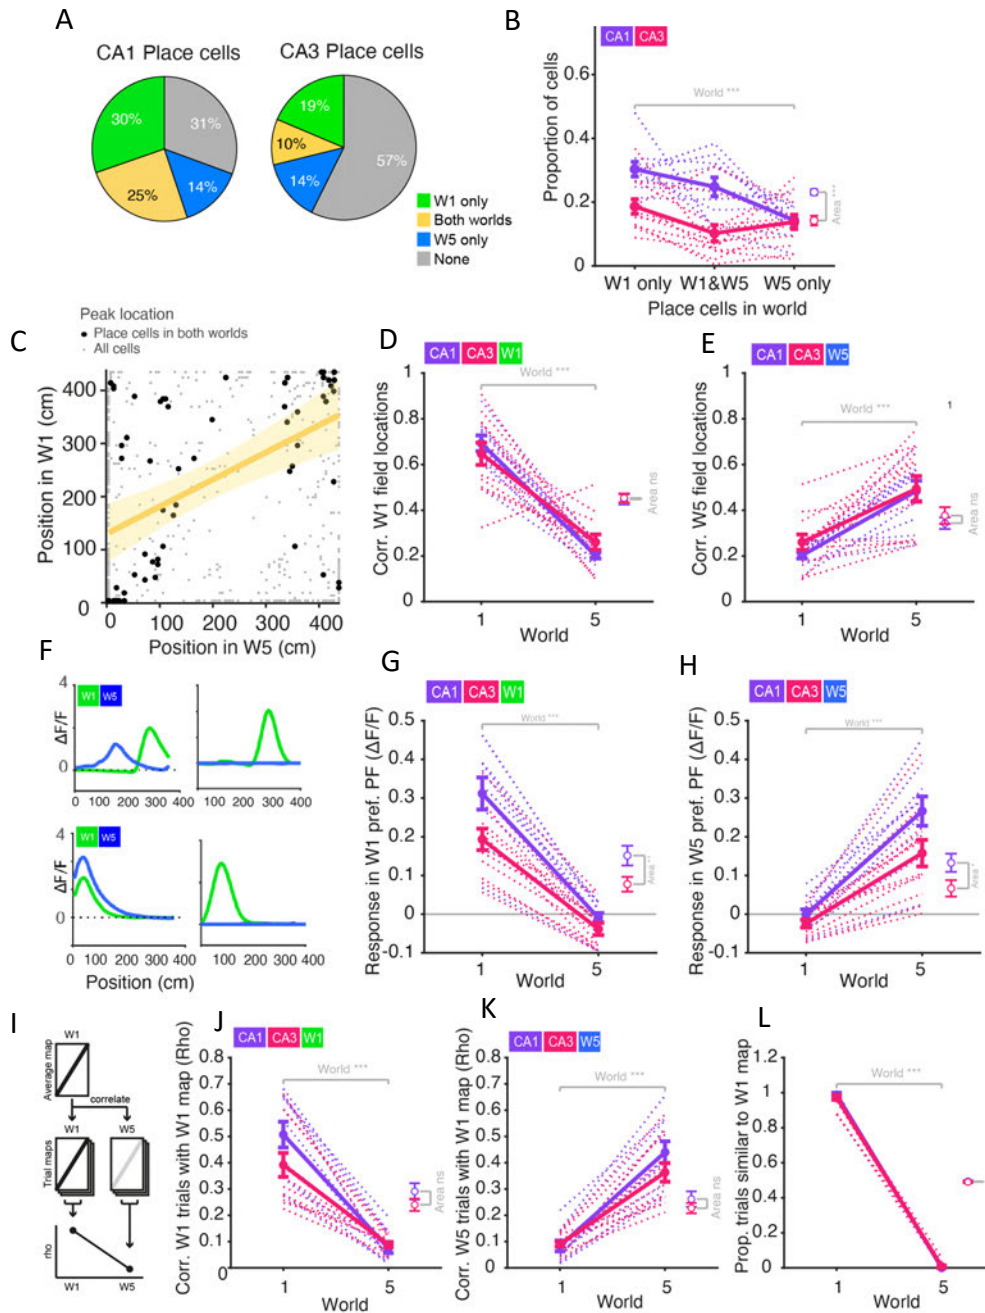
Figure 5-8. Both CA1 and CA3 hippocampal networks form distinct neural representations of different virtual worlds.

A) Pie chart to visualise averaged overlapping place cell populations in CA3 (right) and CA1 (left) animals. Grey represents ROI detected that do not meet place cell criteria. Colour represents ROIs that meet place criteria in W1 only (green), W5 only (blue) or are present in both W1 and W5 (yellow).

B) Mean proportion of cells determined to be place cells in world 1 only, both W1 and W5 and W5 only in CA1 (purple) and CA3 (pink).

C) Scatter plot of ROI peak location in W1 as a function of the same ROI's peak location in W5, for all ROIs (grey dot) and place cells (black dot). Linear fit (yellow line), shaded is the 95% confidence interval.

D) Cross validated correlation of W1 place field locations with place field location in W1 and W5 for CA1 (purple), and CA3 (pink).

E) Same as in D) but cross validated correlations of W5 place fields.

F) Example trial averaged mean ΔF/F trace as a function of position for the given ROI in world 1 (green line) and world 5 (blue line).

G) Magnitude of response (ΔF/F) of W1 field locations in W1 and response in the same location in W5

H) Same as in G) but for magnitude of response in W5 place field locations.

I) Analysis schematic to demonstrate how we quantify the similarity of single trials in a given world to the average W1 (or W5) place map.

J) Correlation of trial-averaged W1 population place map with all single trial population place maps in W1 and W5.

K) Same as in J) but with correlation of trial-averaged W5 population place map.

L) Proportion of trials using W1 map (defined as trials with a larger correlation to the W1 average map than to the W5 average map) in W1 and W5.

Shaded areas and error bars = S.E.M (except C). Track lines = individual animals. Thicker solid lines = means. CA1 (purple), n= 10, CA3 (pink) n = 12. All *s indicate 2-way ANOVA main effects. (n.s. non-significant, *$P < 0.05$ **, $P < 0.01$, ***, $P < 0.001$).

5.2.4 Experiment 1: Five Morphing Worlds.

5.2.4.1   Comparing CA1 and CA3 responses to parametrically morphed worlds

After experiencing the novel world for 2 days (2 sessions), animals became familiar with both W1 and W5. With 'end point' worlds established, we introduced 3 intermediate worlds (W2-4; Figure 4.9 A) that were proportional linear morphs of the end point worlds. Thus, this allowed us to investigate how each subregion represents gradual changes to environmental inputs. We found no significant difference in animal behaviour across W1-5 in terms of average trial duration and trial outcome (2-way ANOVA no main effect of world, $F_{(4,1)} = 0.5$, P = 0.755; Figure 4.9 B & C). Although, CA3 animals had a shorter trial duration on average than CA1 animals (2-way ANOVA main effect of area, $F_{(4,1)} = 5.8$, P = 0.018; Figure 4.9 B).

131
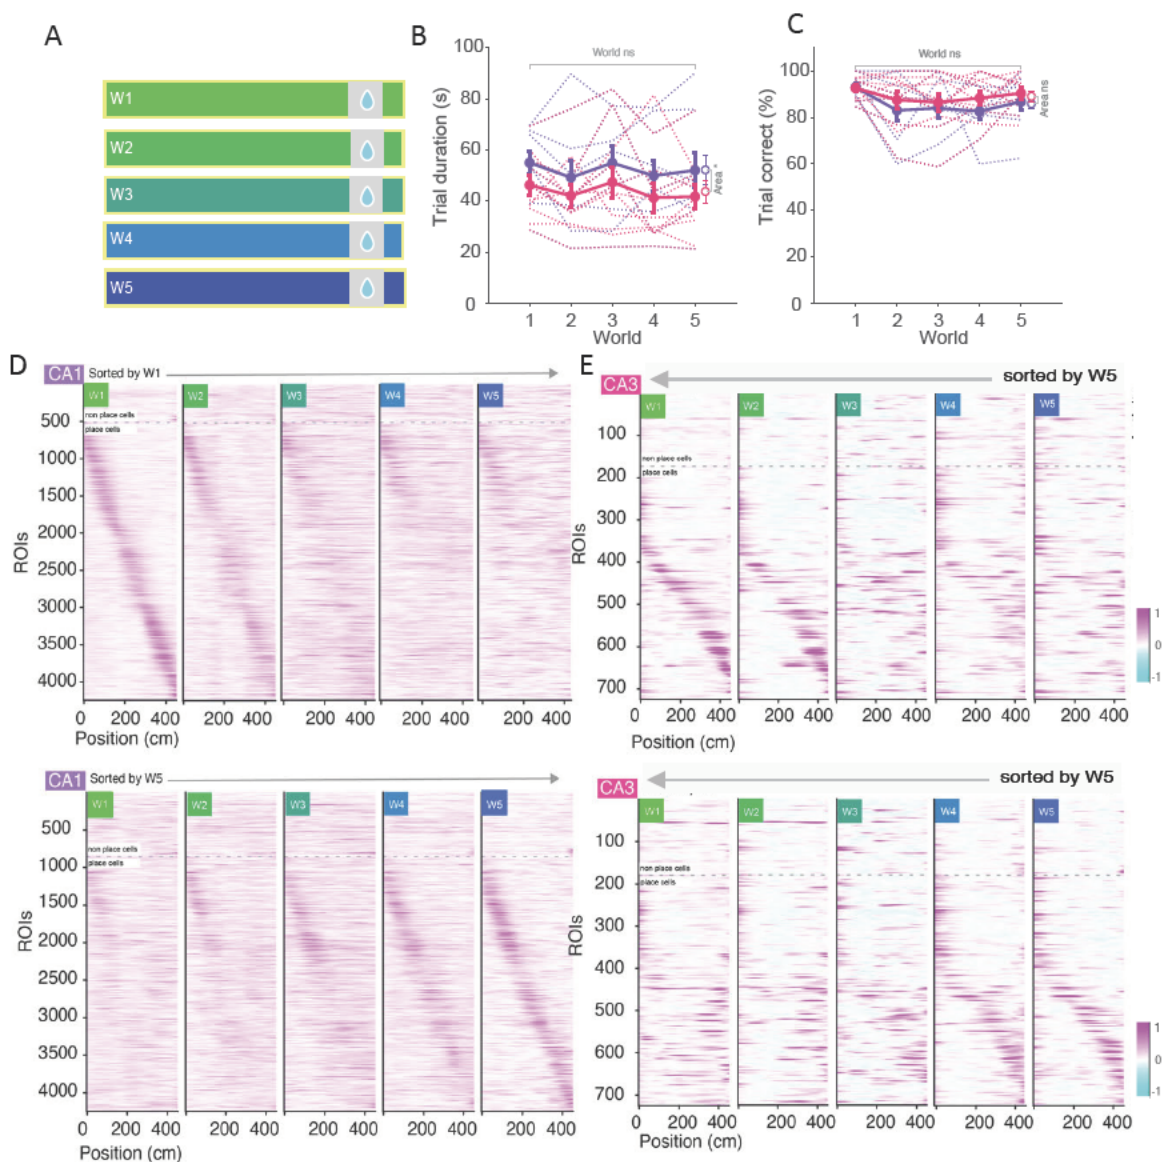
Figure 5-9. Spatial behaviour and neural activity in CA1 and CA3 across 5 morphed VR environments.

A) Schematic to represent morphing 5 worlds experimental phase.

B) Average trial duration for each mouse category (CA1 [purple], n= 10, CA3 [pink] n = 12), defined as the time taken (in seconds) to run one lap of the virtual track for a water reward. Dual area mice are included in both CA1 and CA3 categories (CA1 [purple], n= 10, CA3 [pink] n = 12).

C) Same as in B) but for trial outcome, defined as % trials correct.

D) Example CA1 ROI activity maps (normalised mean ΔF/F) in W1 (left) to W5, sorted by peak fluorescence (ΔF/F) as a function of VR position in W1 (top plot) and then sorted by peak activity as a function of VR position in W5 (bottom plot). Grey dotted line divides (above) 'non place cells' and (below) 'place cells'.

E)   Same as D) but for CA3 activity traces.

Error bars = S.E.M. Track lines = individual animals. Thicker solid lines = means. All *s indicate 2-way ANOVA main effects. tests (n.s. non-significant).

Interestingly, in both subregions we found a higher proportion of ROIs were classified as place cells in W1 and subregion CA1 to have an overall higher average proportion of place cells across worlds than CA3 (2-way ANOVA main effect of world, $F_{(4,1)} = 5.7$, $P < 0.001$. Main effect of area $F_{(4,1)} = 49.5$, $P < 0.001$; Figure 4.10 A). The cumulative proportion of place cells in both subregions increase as we considered more worlds, suggesting new place cells are recruited to represent each world (2-way ANOVA main effect of world, $F_{(4,1)} = 7.4$, $P < 0.001$; Figure 4.10 B). We found no significant differences in these place cell populations across worlds and subregions in terms of their average spatial information scores (2-way ANOVA, no main effect of world or area; Figure 4.10 C).

To begin to assess graded place cell response changes in different worlds, we first looked at the average response in a given location across worlds (Figure 4.10 D & E). To do this we identified each cell's preferred firing location in W1 and measured the response in that location across all worlds. We saw, as worlds become more dissimilar, activity in the W1 field location was reduced in both CA1 and CA3 (2-way ANOVA of responses in W1 fields, main effect of world $F_{(4,1)} = 39.1$, $P < 0.001$; (Figure 4.10 F & G). Consistent with the notion that different place populations represent different worlds, we also found a decrease in the correlation between place cell peak firing locations in W1 with their peak firing locations in the other worlds as those worlds became less similar to W1, for both CA1 and CA3. We found the same relationship when comparing the peak firing locations to their preferred location in W5 (2-way ANOVA of correlation with W1 field locations, main effect of world $F_{(4,1)} = 81.3$, $P < 0.001$; Figure 4.10 H & I).
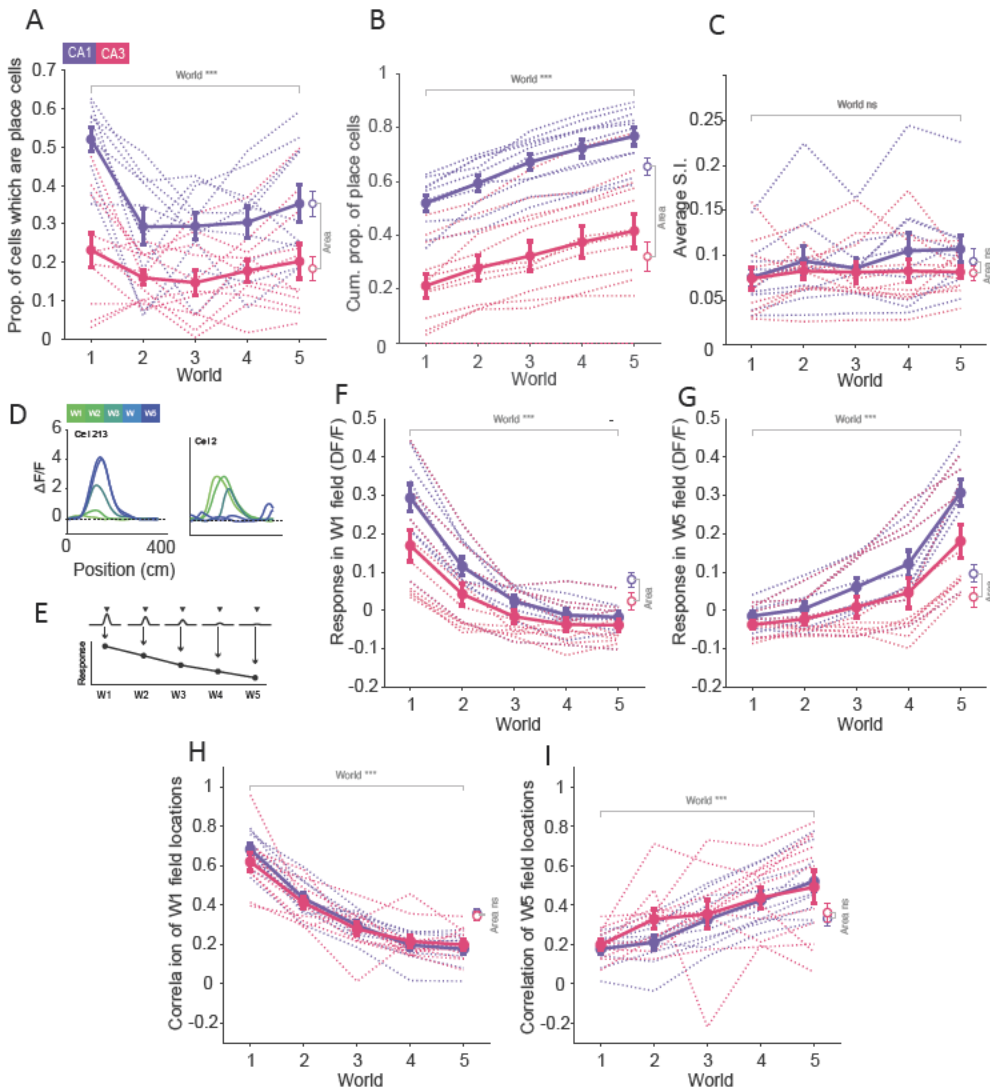
Figure 5-10. CA1 and CA3 place cell properties across of parametrically morphed VR worlds.

A) Proportion of ROIs defined as place cells across W1-5 for subregion CA1 and CA3.

B) Cumulative proportion of ROIs defined as place cells across worlds 1-5.

C) Average SI scores per animal across the 5 worlds.

D) Example trial averaged mean ΔF/F trace as a function of position for the same given ROI in world 1-5.

E) Schematic to show the method used to calculate F) & G). Briefly, ROIs preferred field is identified in W1 and the response in that location is measured across worlds (W1-5).

F) Response magnitude (ΔF/F) of W1 field locations in W1 and responses in the same location across W2-5.

G) Same as in F). but for magnitude of response in W5 place field locations across W1-4.

H) Cross validated correlation of W1 place field locations with place field location in W1 to W5 for CA1 (purple), and CA3 (pink) animals.

I) Same as in H) but for cross validated correlations for W5 place field locations.

As before, we sought to assess the degree of similarity of each world's representation relative to a reference of the average W1 or W5 map. Here, for each animal we correlated the activity of the entire neural population in a given world, position bin by position bin to build to a profile of similarity to the end points tracks, across distance down each track (Figure 4.11 A-D). We found that spatial representations were most similar at the start of the track but diverged as animals advanced along it (Figure 4.11 A-D). For this reason, we excluded the beginning of the track (0 to 40 cm) from subsequent analyses in order to focus on the maximally different representations of the tracks.

To quantify the overall similarity between world representations at a population level, we correlated all single trial place maps across worlds with the average map in W1 or W5 (Figure 4.11 E). We found that single trial average place map correlations to the end point maps (W1 and W5) decreases across worlds, suggesting decreasing similarity in neural representations as world similarity decreases for both subregions (2-way ANOVA for correlation to W1 map, main effect of world $F(4,1) = 71.4$, $P < 0.001$. 2-way ANOVA for correlation to W5 map, main effect of world $F(4,1) = 50.1$, $P < 0.001$; Figure 4.11 F & G). We then used our metric (discussed earlier) to assess the probability a given trial is more similar to the W1 representation than W5. We found that CA1 and CA3 respond similarly (2-way ANOVA main effect of area $F(4,1) = 0$, $P = 0.97$; Figure 4.11 H), in that as worlds become more dissimilar from W1, the probability that the neural representation matches W1 decreases (2-way ANOVA main effect of world $F= 157.7$, $P < 0.001$). Notably, while this decoding approach revealed a relatively linear response in CA3, CA1 responses appeared to be non-linear.

135
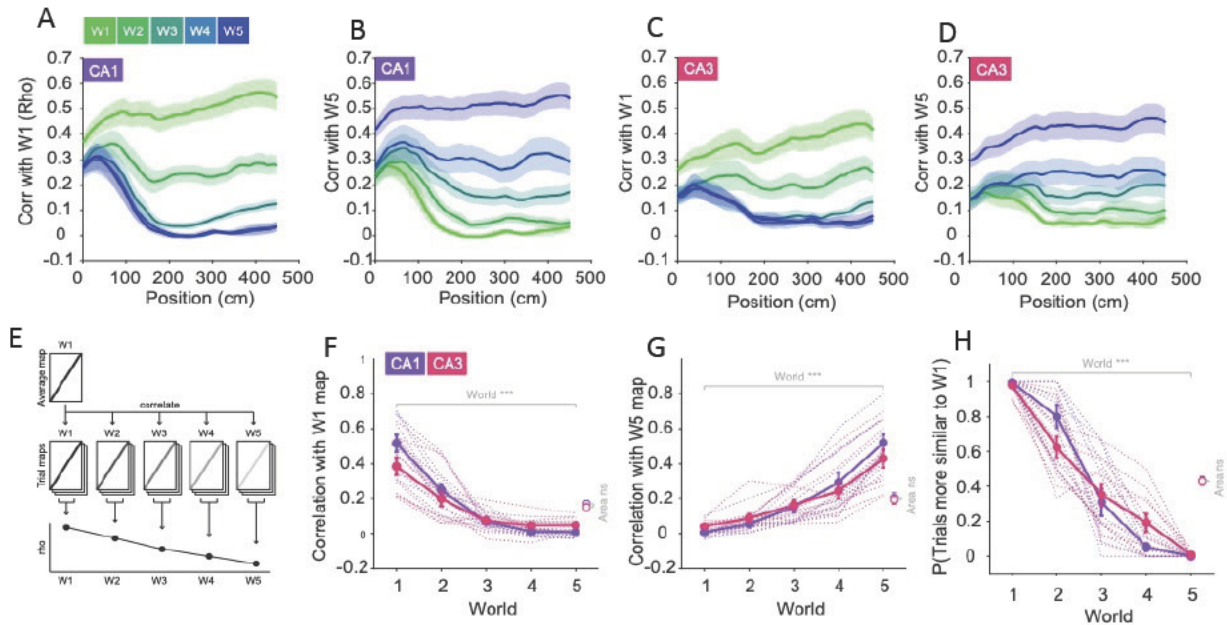
Figure 5-11. Populations of neurons in CA1 and CA3 similarly map parametrically morphed virtual worlds.

A) Correlation of trial-averaged W1 population place map with activity (ΔF/F) along at different positions along the track in W1 (green) to W5 (blue), for CA1 (purple) and CA3 (pink) mice.

B) Same as in A) but correlating to a trial-averaged W5 population place map.

C) Same as in A) but for CA3 neural activity.

D) Same as in B) but for CA3 neural activity.

E) Schematic to demonstrate the analysis method used to assess neural representation similarity to the end point world maps. Briefly, we correlated all single trial place maps across worlds with the average map in W1 or W5

F) Correlation of trial-averaged W1 population place map with all single trial population place maps in W1-5

G) Same as in F) but with correlation of trial-averaged W5 population place map.

H) Proportion of trials using W1 map (defined as trials with a larger correlation to the W1 average map than to W2-4 average maps) in in W1-5.

Shaded areas and error bars = S.E.M. Track lines = individual animals. Thicker solid lines = means. F-H: All *s indicate 2-way ANOVA main effects (n.s. non-significant, $*P < 0.05$ **, $P < 0.01$, ***, $P < 0.001$). CA1 (purple), n= 10, CA3 (pink) n = 12.

136

5.2.5 Experiment 2: Two Worlds

5.2.5.1 Context discrimination task

We reasoned that having animals actively discriminate their environment may strengthen the differences between neural representation when it is behaviourally important to treat the environments differently. In order to link neural activity to behaviour, we developed a task to provide a behavioural read out from which we could infer which world the animal perceived itself to be present in (Figure 4.12 A & B). In each of the training worlds (worlds representing the two extremes of the morph continuum) there were two reward zones (Figure 4.12 C). Animals were taught to lick (for a reward) in only 1 of the 2 reward zones for a particular world contingency (see Methods & Materials 5.7). We found no statistical difference between behavioural performance of CA1 and CA3 animals in terms of trial duration and trial outcomes across worlds (2-way ANOVA no main effect of world, $F_{(1,1)} = 1.5$, $P = 0.23$ or area $F_{(1,1)} = 0.2$, $P = 0.68$; Figure 4.12 D & E). We found that animals were able to discriminate between worlds in order to select the appropriate reward zone (rather than licking in both), as we found significantly more mean licks in the world appropriate reward zones for both CA1 and CA3 mice (W1 appropriate zone: 2-way ANOVA main effect of world, $F_{(1,1)} = 9.9$, $P = 0.004$; Figure 4.12 G and W5 appropriate zone: $F_{(1,1)} = 11.6$, $P = 0.002$; Figure 4.12 F).

137
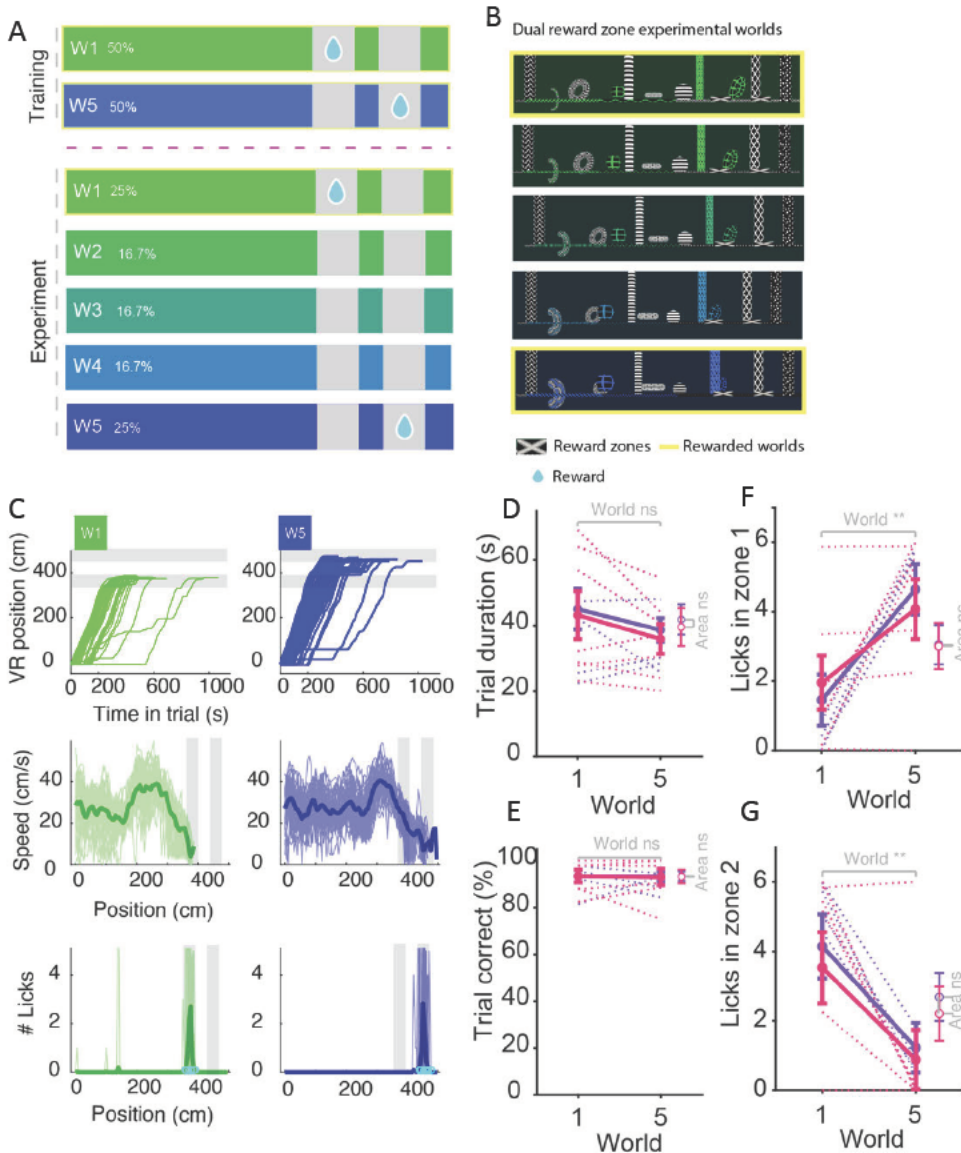
Figure 5-12. Animals can learn different lick contingencies when navigating in distinct virtual worlds.

A) Task schematic to demonstrate the two training worlds prior to 5 experimental worlds. W1 and W5 were rewarded in opposite reward zones (grey squares), on probe (unrewarded trials)

B) Side view of VR worlds complete in the experiment, rewarded worlds are outlined in yellow, dual reward zones are marked with crosses and object locations, texture design frequency and colour are shown in a parametric manner.

C) Example behavioural performance across W1 (green) and W5 (blue) in the training phase. Top row: VR position over time in trial (s; W1- green, W5- blue). Middle row: Animal speed along the virtual track per trial, thicker line represents mean speed over the track. Bottom row: number of licks across the track.

Thicker line represents the mean lick count as a function of position across trials. Blue rings indicate rewards received. Grey bars indicate reward zones.

D) Average trial duration for each mouse category (CA1 [purple] n = 8, CA3 [pink] n = 7), defined as the time taken (in seconds) to run one lap of the virtual track for a water reward. Dual area mice are included in both CA1 and CA3 categories.

E) Same as in E) but for trial outcome, defined as % trials correct.

F) Mean licks in appropriate zone for W5 (Licks after reward delivery were discarded).

G) Same as F) but for appropriate licks in W1

Reward zones marked in grey. Error bars = S.E.M. Pale lines and track lines = individual animals. Thicker solid lines = means. Blue rings indicate water rewards received. All *s indicates 2-way ANOVA main effects (n.s. non-significant, *$P < 0.05$ **, $P < 0.01$, ***, $P < 0.001$).

We recorded neural activity from CA1 and CA3 populations whilst animals performed this discrimination task and saw place modulated activity (Figure 4.13 B). We found some place cells were exclusively in W1 (CA1, w1: 0.19 ±0.10 vs. CA3, w1: 0.14 ±0.08), or W5 (CA1, w5: 0.20 ±0.10 vs. CA3, w5: 0.16 ±0.10) and some cells that had place fields in both worlds (CA1, w1+5: 0.18 ±0.13 vs. CA3, w1+5: 0.09 ±0.07; Figure 4.13 A). We quantified the proportion of place cells present in W1 and W5 and found no differences between worlds (2-way ANOVA main effect of world, $F_{(1,1)} = 0.1$, $P = 0.79$; Figure 4.14 C) but CA3 animals had significantly lower proportions of place cells overall in both worlds (2-way ANOVA main effect of world, $F_{(1,1)} = 11.6$, $P = 0.002$; Figure 4.13 C). We found no statistical differences in CA1 and CA3 place cells spatial information scores across subregions and worlds in terms (2-way ANOVA main effect of world, $F_{(1,1)} = 0.1$, $P = 0.8$; main effect of area, $F_{(1,1)} = 0.1$, $P = 0.81$; Figure 4.13 D). Importantly, as before in Experiment 1, the presence of world exclusive place cells and the low similarity of W5 trial maps to the W1 average place maps suggests the two worlds have distinct neural representations in both CA1 and CA3 (2-way ANOVA main effect of world, $F_{(1,1)} = 0.1$, $P = 0.8$; main effect of area, $F_{(1,1)} = 0.1$, $P = 0.81$; Figure 4.13 E & F).

139
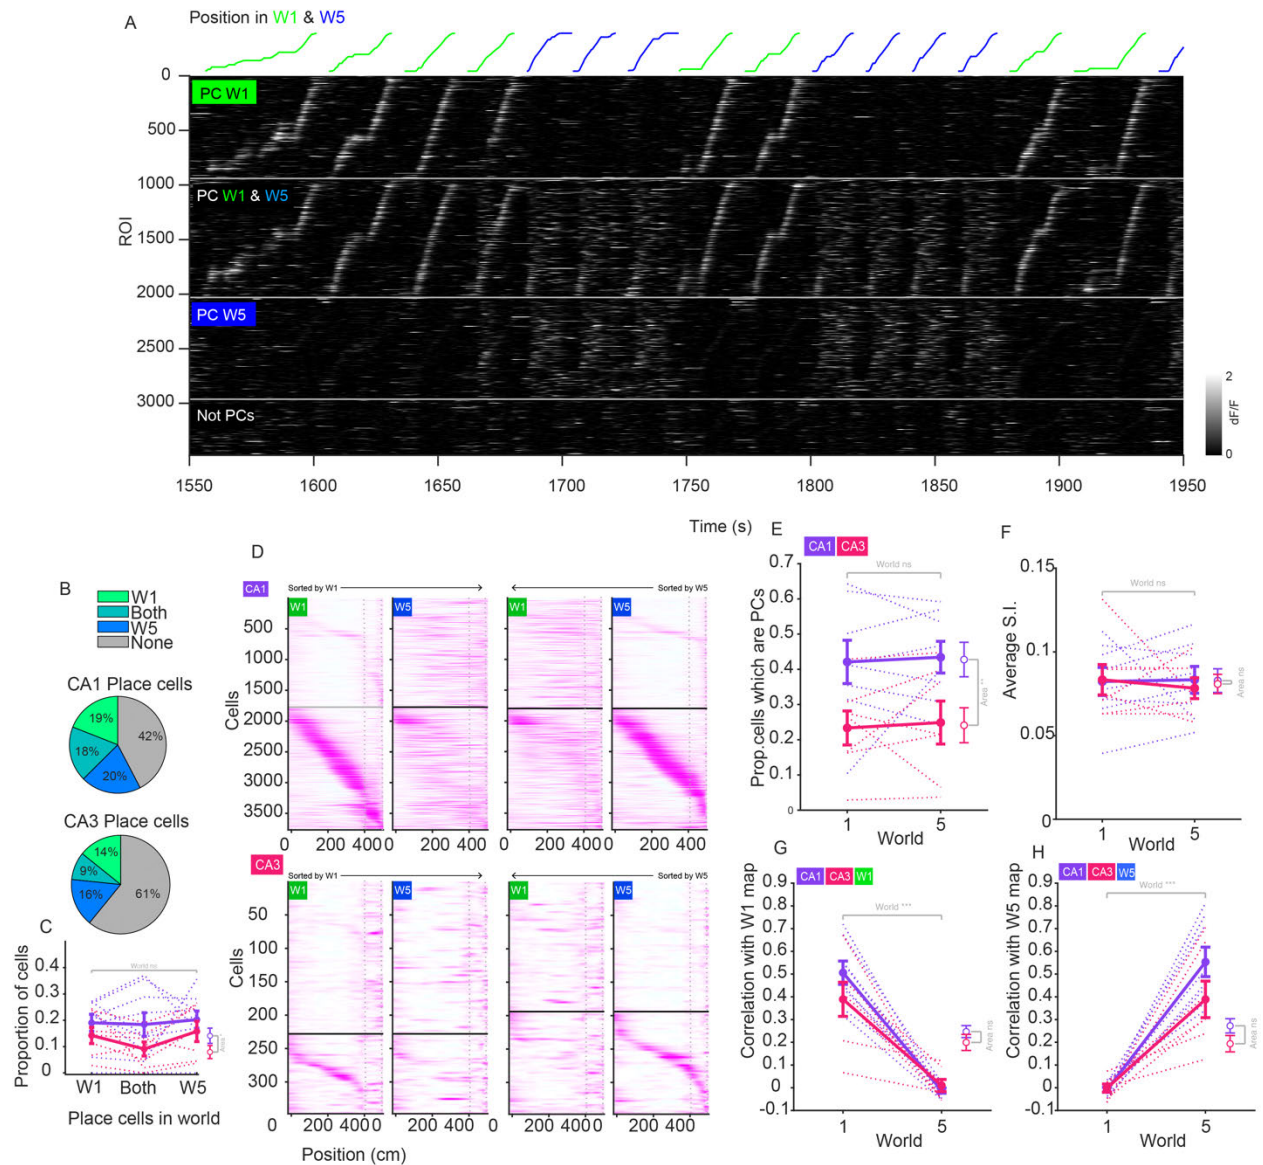
Figure 5-13. CA3 and CA1 hippocampal neurons form distinct neural representations of different virtual worlds with different behavioural contingencies.

A) Top: Example trial trajectories per animal over time, coloured by world presented (W1 = green, W5 = blue). Below: Heatmap of neural activity per ROI over time (lighter colour = higher activity; normalised mean ΔF/F), grouped by world preference and sorted by preferred place field in W1

B) Pie chart to visualise averaged overlapping place cell populations in CA3 (bottom) and CA1 (top) animals. Grey represents ROI detected that do not meet place cell criteria. Colour represents ROIs that meet place criteria in W1 only (green), W5 only (blue) or are present in both W1 and W5 (teal).

C) Mean proportion of cells determined to be place cells in W1 only, both worlds and W5 only in CA1 (purple) and CA3 (pink).

D) Example place maps for each mouse category (CA1 top, CA3 bottom). Plots show all ROIs detected in an example animal. Black line divides (above) 'non place cells' and (below) 'place cells'. Each row in the heatmap indicates mean ΔF/F of that cell as a function of position in virtual space, traces are sorted by place cell peak preferred location on the track in W1 (left two plots) and W5 (right two plots).

E) Proportion of ROIs defined as place cells across W1-5 for subregion CA1 and CA3.

F) Same as in C) but for average SI scores for each subregion In W1 and W5.

G) Correlation of trial-averaged W1 population place map with all single trial population maps in W1 and W5.

H) Same as in E) but for correlation with average W5 place map.

Reward zones marked by grey track lines. Error bars = S.E.M. Pale lines and track coloured lines = individual animals. Thicker solid lines = means. All *s indicate 2-way ANOVA main effects (n.s. non significant, $*P < 0.05$ **, $P < 0.01$, ***, $P < 0.001$), CA1 [purple] n = 8, CA3 [pink] n = 7.

5.2.6 Experiment 2: Five Morph Worlds

As before, we introduced parametrically morphed worlds (W2-4) as probe trials where the two reward zones were present but not rewarded (Figure 4.14 A). W2 was designed to be most similar to W1, W3 was designed to be as ambiguous as possible between W1 and W5, while W4 was designed to be most similar to W5. Animals' licking was recorded but not rewarded in the intermediate environments - yielding a behavioural readout of their beliefs contingent on prior training (W1 and W5; Figure 4.14 A). We recorded behaviour and neural activity when animals were introduced to these intermediate worlds on a subset of trials (Figure 4.2; 17.5 % intermediate trial types). We found no significant difference between animals behavioural performance in terms of trial duration (2-way ANOVA main effect of world, $F(4,1) = 1.4$, P = 0.27; no main effect of area, $F(4,1) = 0.6$, P = 0.45; Figure 4.14 B) and trial outcome (restricted to only rewarded worlds; 2-way ANOVA no main effect of area, $F(4,1) = 1.3$, P = 0.26; Figure 4.14 E). Both CA1 and CA3 mice showed no statistical differences in licking behaviour, as intermediate worlds became more different from W1 the animals licked less in the W1 associated reward zone (2-way ANOVA main effect of world, $F(4,1) = 49.2$, P = 0.000; no main effect of area, $F(4,1) = 0.0$, P = 0.98; Figure 4.14 D), and more in W5 reward zone (2-way ANOVA main effect of world, $F(4,1) = 60.2$, P = 0.000; no main effect of area, $F(4,1) = 0.6$, P = 0.46; Figure 4.14 C).
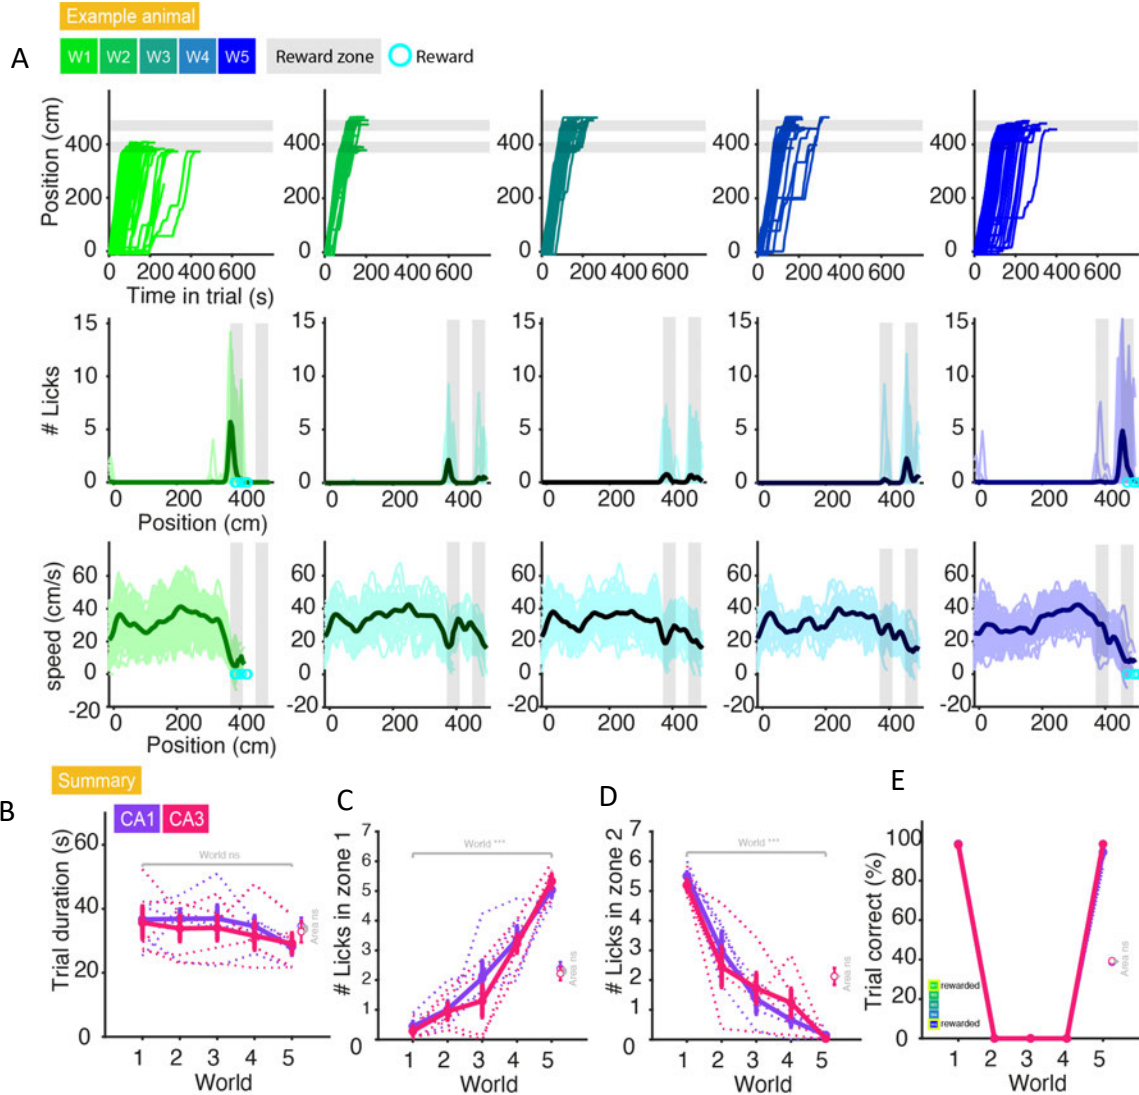
Figure 5-14. Animals show graded behavioural responses to intermediate virtual worlds.

A) Example behavioural performance across W1 (green) through to W5 (blue). Top row: VR position over time in trial (s; W1-W5 [green to blue]). Middle row: number of licks across the track. W1 and W5 were rewarded contingent on licking in the correct location for that world and W2-4 were probe trials, no reward was given (see Chapter 4 for details). Thicker line represents the mean lick count as a function of position across trials Bottom row: Animal speed along the virtual track per trial, thicker line represents mean speed over the track. Grey shaded bars = rewarded locations. Blue rings indicate rewards received (only in W1 and W5).

B) Average trial duration for each mouse category across W1-5, defined as the time taken (in seconds) to run one lap of the virtual track for a water reward. Dual area mice are included in both CA1 and CA3 categories.

C) Mean licks in appropriate zone for W5 (zone 1 here. Licks after reward were discarded) for CA1 and CA3.

D) Same as E) but in appropriate reward zone in W1

E) Mean trial outcome defined as % correct trials. W1 and W5 were the only rewarded trials i.e. W2-4 were probe trials and not rewarded.

Reward zones marked by grey bars. Error bars = S.E.M. Pale lines and track coloured lines = individual animals. CA1 [purple], n= 6, CA3= [pink] n = 5. Thicker solid lines = means. All *s indicate 2-way ANOVA main effects tests (n.s. non-significant, *P < 0.05 **, P < 0.01, ***, P < 0.001).

We found the greatest proportion of place cells in W1 for both animal groups, and as before, the proportion of ROIs determined to be place cells was greater in CA1 across all worlds (2-way ANOVA main effect of area, $F_{(4,1)}$ = 12.8, P = 0.001; Figure 4.15 A). Similarly, the cumulative proportion of place cells was greater in CA1 than in CA3 and increased for both areas over worlds suggesting distinct populations of place cells might represent each world in each area (2-way ANOVA main effect of world, $F_{(4,1)}$ = 7.1, P = 0.000 and main effect of area $F_{(4,1)}$ = 33.6, P = 0.000; Figure 4.15 B). We also found the mean magnitude of responses in W1 place field locations decreased over worlds for both CA1 and CA3 regions (W1 field preferences: 2-way ANOVA main effect of world, $F_{(4,1)}$ = 26.5, P = 0.000 and no main effect of area $F_{(4,1)}$ = 0.0, P = 0.93; W5 field preferences: $F_{(4,1)}$ = 46.6, P = 0.000 and no main effect of area $F_{(4,1)}$ = 0.0, P = 0.89; Figure 4.15 E & F) hinting at a cellular-level remapping of place field firing.
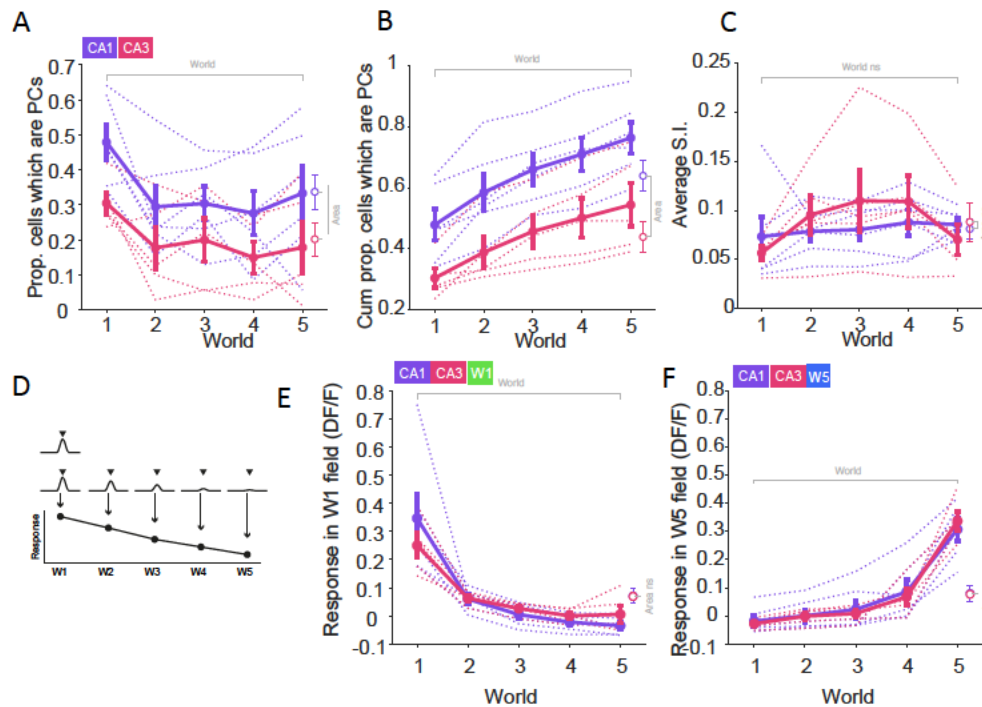
Figure 5-15. CA1 and CA3 place cells show similar properties and modulation over 5 morphed virtual worlds.

A) Proportion of ROIs defined as place cells across W1-5 for subregion CA1 and CA3.

B) Same as in A) but for the cumulative proportion of place cells

C) Same as in A) but for average S.I. scores.

D) Schematic to show the method used to calculate E) & F). Briefly, ROIs preferred field is identified in W1 and the response in that location is measured across worlds (W1-5).

E) Magnitude of response (ΔF/F) of W1 field locations in W1 and in the same location across W2-5.

F) Same as in E) but for W5 field locations.

Error bars = S.E.M. Pale track-coloured lines = individual animals. Thicker solid lines = means. CA1 [purple], n= 6, CA3= [pink] n = 5. All *s indicate Wilcoxon rank_sum tests (n.s. non-significant, *$P < 0.05$ **, $P < 0.01$, ***, $P < 0.001$).

To compare the population place code across worlds we computed the similarity of all trials in a given world relative to the average map in one of the end point worlds (W1 or W5; Figure 4.16 A,B). We found that the neural representation of W2 (adjacent to W1) was most similar to the W1 representation and the representation of W4 (adjacent to W5) was most similar to that of W5, suggesting a graded switch of neural representations as the world is morphed between two familiar end points (similarity to W1: main effect of world $F_{(4,1)} = 34.3$, $P = 0.00$; similarity to W5: main effect of world, $F_{(4,1)} = 22.9$, $P = 0.00$; Figure 4.16 D & E). At this level of population analysis

we did not observe a difference in terms of how CA1 or CA3 represented the morphed intermediate worlds (similarity to W1: main effect of area $F_{(4,1)} = 0.4$, $P = 0.52$; similarity to W5: main effect of area, $F_{(4,1)} = 0.7$, $P = 0.40$; Figure 4.16 D & E).
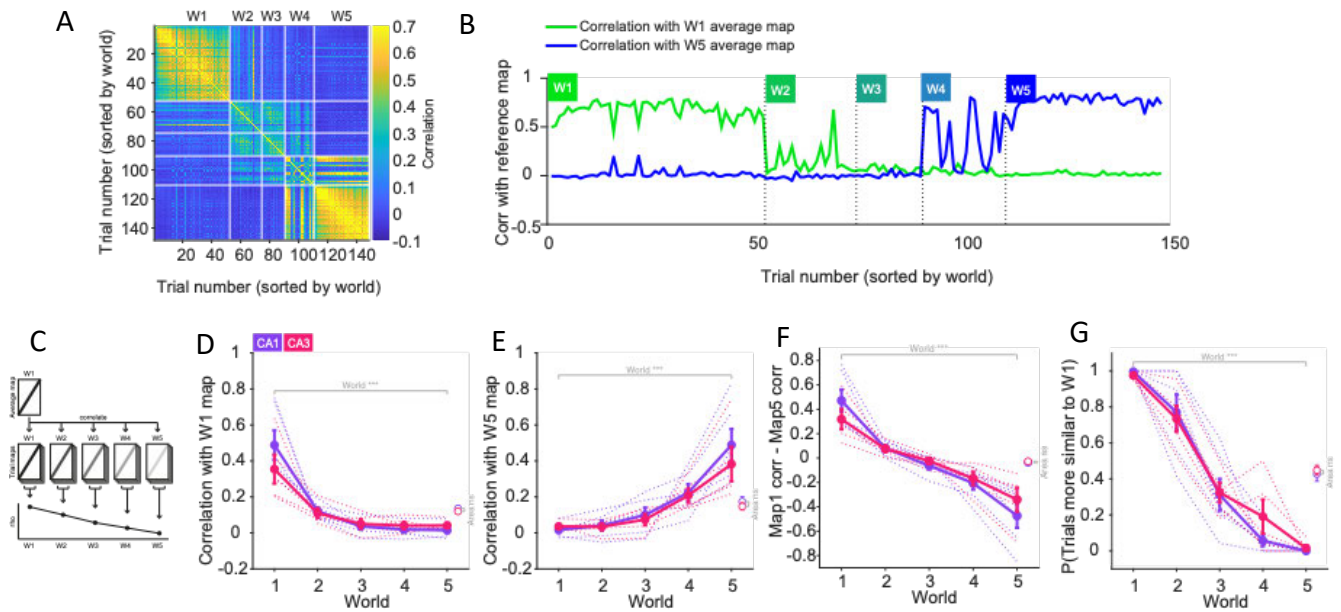
Figure 5-16. CA1 and CA3 place cells show similar modulation of spatial representations over 5 morphed virtual worlds.

A) Example trial by trial correlation matrix of population place cell activity. Warmer colours = higher correlations (R).

B) Example neural activity correlation (R) with average place map in W1 (green) and W5 (blue) across trials. Trials are sorted by world number and not in the order presented during the experiment.

C) Schematic to demonstrate the analysis method of correlating single trial maps to the average map in W1

D) Correlation of trial-averaged W1 population place map with all single trial population place maps in W1-5

E) Same as in D) but with correlation of trial-averaged W5 population place map.

F) The difference between D) and E) i.e. the difference in correlation to the W1 and W5 average place map across worlds 1-5

G) Proportion of trials using W1 map (defined as trials with a larger correlation to the W1 average map than to the W5 average map) across W1-5.

Error bars = S.E.M. Dotted lines = individual animals. Thicker solid lines = means. CA1 [purple], n= 6, CA3= [pink] n = 5. All *s indicate 2-way ANOVA tests (n.s. non-significant, *$P < 0.05$ **, $P < 0.01$, ***, $P < 0.001$).

146

5.2.7 Experiment 2: 5 Morph Worlds

5.2.7.1 Linking neural activity to behaviour

Next, since we had simultaneous access to the animals' behaviour and neural activity, we examined how these covaried in the morph environments. First, comparison of licking zone preference with the relative neural similarity of intermediate environments to endpoint worlds revealed a clear relationship. Behaviour at reward zones reflected the activity patterns seen as the animal was traversing the track prior to that point (Figure 4.17 A & B). We found when the neural representation most closely resembled one of the end point worlds, the animal preferentially licked in the corresponding (contingency appropriate) reward zone. Conversely in intermediate worlds, with ambiguous neural representations (less similar to either of the end points), the behaviour too was ambiguous (Figure 4.17 A & B). We found this activity-behaviour relationship took the form of a steep sigmoid function, in that small increases or decreases in the similarity of the neural representations, relative to the learned endpoints, caused large increases or decreases in the animal's behaviour, in the same direction as the learned associations in those end point worlds. This steep relationship persists even when we exclude the end point worlds from analysis and focus only on intermediate worlds (Figure 4.17 C). While we did not observe a difference between CA1 or CA3, we did find a direct link between activity in the hippocampal circuit and the ultimate reward-seeking navigation behaviour of the animal.
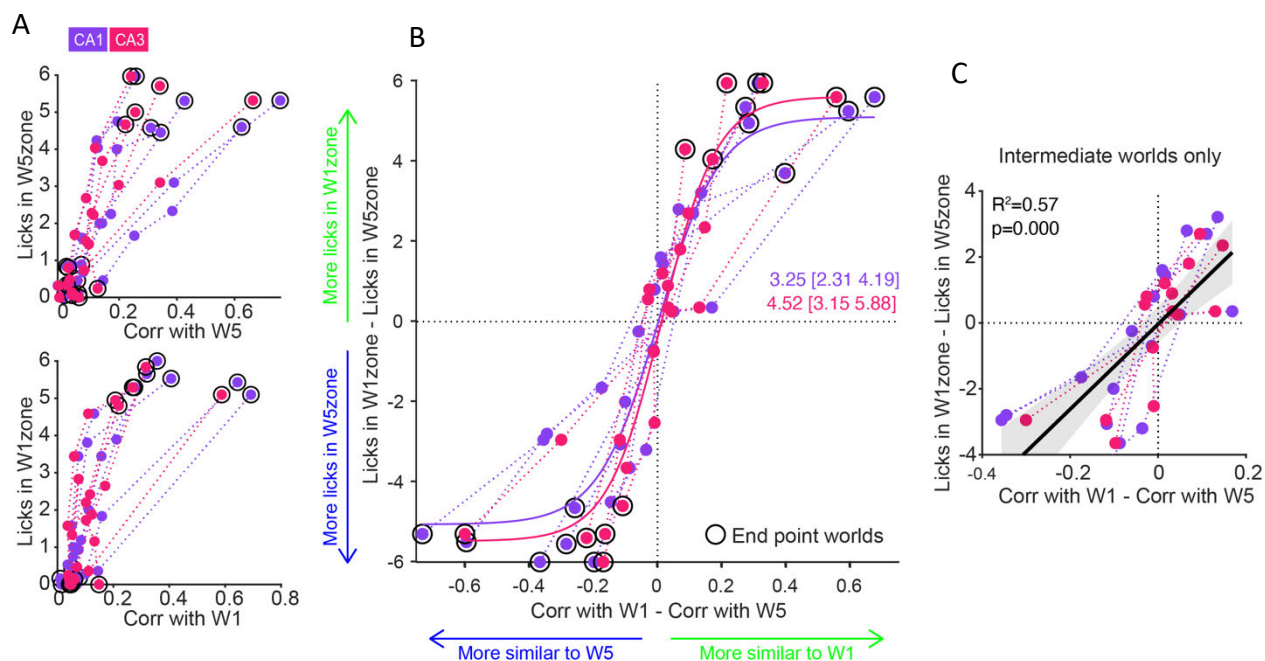
Figure 5-17. Animals' spatial behaviour in parametrically morphed virtual worlds drives the neural representation of those worlds in CA1 and CA3.

A) Top: Relationship between licks in reward zone 1 and the correlation of the average W5 map. Bottom) Same as in A top but for licks in reward zone 2 and plotted as a function of correlation with the average place map in W1.

B) Relationship between the lick difference between reward zone 1 and 2 and the difference in place map correlation to the W1 and W5 average place maps. Solid line represents the sigmoid fit for CA1 (purple) and CA3 (pink). The steepness parameter estimate = CA1: 3.25 (CI: 2.31 4.19), CA3: 4.25 (CI: 3.15 5.88).Each point represents 1 of the 5 worlds per animal (animals are connected by the track lines) and end point worlds (W1, W5) are indicated by the open black circles. Purple = CA1. Pink = CA3.

C) Same as in B) but for intermediate worlds only (end point worlds excluded). Black line = linear fit, grey shading = 95% confidence intervals.

CA1 [purple], n= 6, CA3= [pink] n = 5.

The average activity maps we use to compare all single trials maps to in the previous analysis may not be the true representations animals use, for example if population drift occurs or the representations are unstable. To combat this, we employed a new approach utilising K-means clustering to extract two maximally different average place maps. These maps by definition are the two most distinct representation seen throughout the recordings and we then used them as reference maps for assessment of the single trial maps. By projecting single trial activity into the

space of the k-means clustered, maximally distinct, maps we were able to ask on a given trial which map is the current activity most similar to (Figure 4.18 A). We then averaged the proportion of trials using a given map in a given world and aligned the maps relative to the endpoint world most frequently using them in order to enable comparison across animals.

When looking at the proportion of trials across each world that used the W1 dominant activity map, we noted a nonlinear function on average, whereby the frequency map 1 is used decreases across the worlds. CA3 animals tend to use the world 1 map more frequently than CA1 animals (Fig 4.18 B; 2-way ANOVA main effect of area, $F_{(4,1)} = 5.8$, $P = 0.02$, main effect of world, $F_{(4,1)} = 7.8$, $P < 0.001$). We noted that some animals show very steep relationships while for others it was shallower (track lines in Figure 4.18 B). This suggests heterogeneity in terms of how individual animals are representing the intermediate worlds relative to the learnt familiar end points.

To quantify the level of attractor dynamics, we assessed how well a sigmoid function explains the neural representation across worlds. We devised a method to quantify the goodness of fit of a sigmoid, relative to a simple linear fit which we termed the sigmoid score (Figure 4.18 C). Surprisingly, we found that on average CA3 is no different from CA1 in terms of how it binds similar worlds into similar representations (Figure 4.18 D; CA1 population sigmoid score (AIC sigmoid fit – AIC linear fit): -98.36 ±110.45, n=6. CA3 population sigmoid score (AIC sigmoid fit – AIC linear fit): -28.42 ±57.48, n=5; P = 0.177 Wilcoxon rank sum test). However, we did observe a range of sigmoid scores, highlighting that some individual animals display pronounced attractor dynamics, while others do not. To explore the factors governing this, we related the average sigmoid score to various variables of interest. We looked animals' behaviour including how long the animals were trained for and how fast the animals were traversing the tracks but found no relationship (Figure 4.18 E-H). However, when we looked at the potential fidelity of how the population was representing the environments, we discovered a relation to the strength of the attractor states. As the populations were made up of more place cells (and to a lesser extent, better encoding place cells) and so as the encoding quality increased so too did the degree of binding to the familiar end points ($R^2 = 0.66$, P= 0.002, n = 11 animals). Although this relationship

149

appears to be driven by two data points, we found a similar result when this method was applied to data from Experiment 1 (N = 17 animals and r values Appendix, Chapter 7)
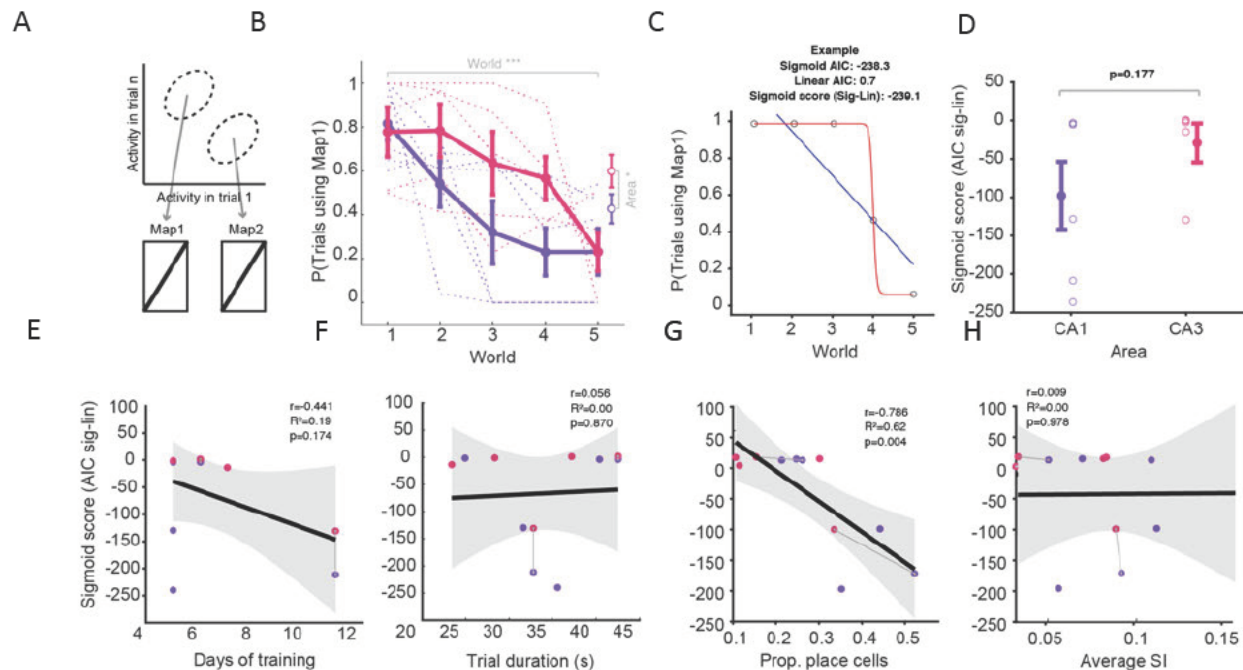
Figure 5-18. Using K-means clustering to identify the 'optimal' neural representation, reveals nonlinear mapping of morphed worlds.

A) Schematic to demonstrate the k-means method of unsupervised identification of maximally different representations (clusters).

B) The proportion of trials using map (cluster) 1 across worlds and by subregion

C) Example to demonstrate the simple linear (blue) and sigmoid (red) fitting procedure for data points (proportion of trials using k-means map 1) across worlds per for one animal.

D) Sigmoid score for each subregion, defined as AIC of the non-linear minus the AIC of the linear fit. Open circles = animals, closed circles with S.E.M error bars = average across animals

E) Relationship between the average sigmoid score and the number of training days per animal. Grey connecting lines = dual mice, i.e area CA1 and CA3 are from the same animal.

F) Same as in E) but for trial duration

G) Same as in E) but for proportion of place cells

H) Same as in E) but for average S.I.

Grey shaded area & Error bars = S.E.M. Pale dotted coloured lines = individual animals. Thicker solid lines = means. CA1 [purple] n= 6, CA3 [pink] n = 5.

5.2.7.2 Heterogenous single cell world representations

After finding varying degrees of non-linearity of world representation at the neural population level, and concluding that the proportion on place cells in the network determined the degree of population sigmoid strength, we next explored the composition of those populations. We investigated how single cells behaved in terms of how similar their representation across worlds was relative to their optimal maps for the end point worlds (Figure 4.19A). We found a heterogenous population (Figure 4.19 B) in that some place cells are best described by a sigmoid fit (Figure 4.17C; 24.1 ± 7.3 % of place cells classed as sigmoidal in CA1 (n=6), 20.9 ± 7.4 % in CA3 (n=5). P = 0.329 Wilcoxon rank sum test), indicating an abrupt switch whereby some worlds were most similar to W1, and others most similar to W5. Other place cells were best fit by a linear fit (Figure 4.19 D; 17.7 ± 8.1 % of place cells classed as linear in CA1 (n=6), 11.1 ± 6.5 % in CA3 (n=5). P = 0.637 Wilcoxon rank sum test), indicating a gradual change in representation across worlds. We did not observe any significant difference in single cell composition of population between subregions. We noted that with increasing numbers of place cells in the population, more of those place cells were sigmoidal in their world representation behaviour (Figure 4.19 E). Finally, to relate the sigmoidal character of single cells back to the population average nonlinearity, we investigated the relationship between the two. We observed that with increasing numbers of sigmoidal single cells, the population nonlinearity increases (Figure 4.19 F).
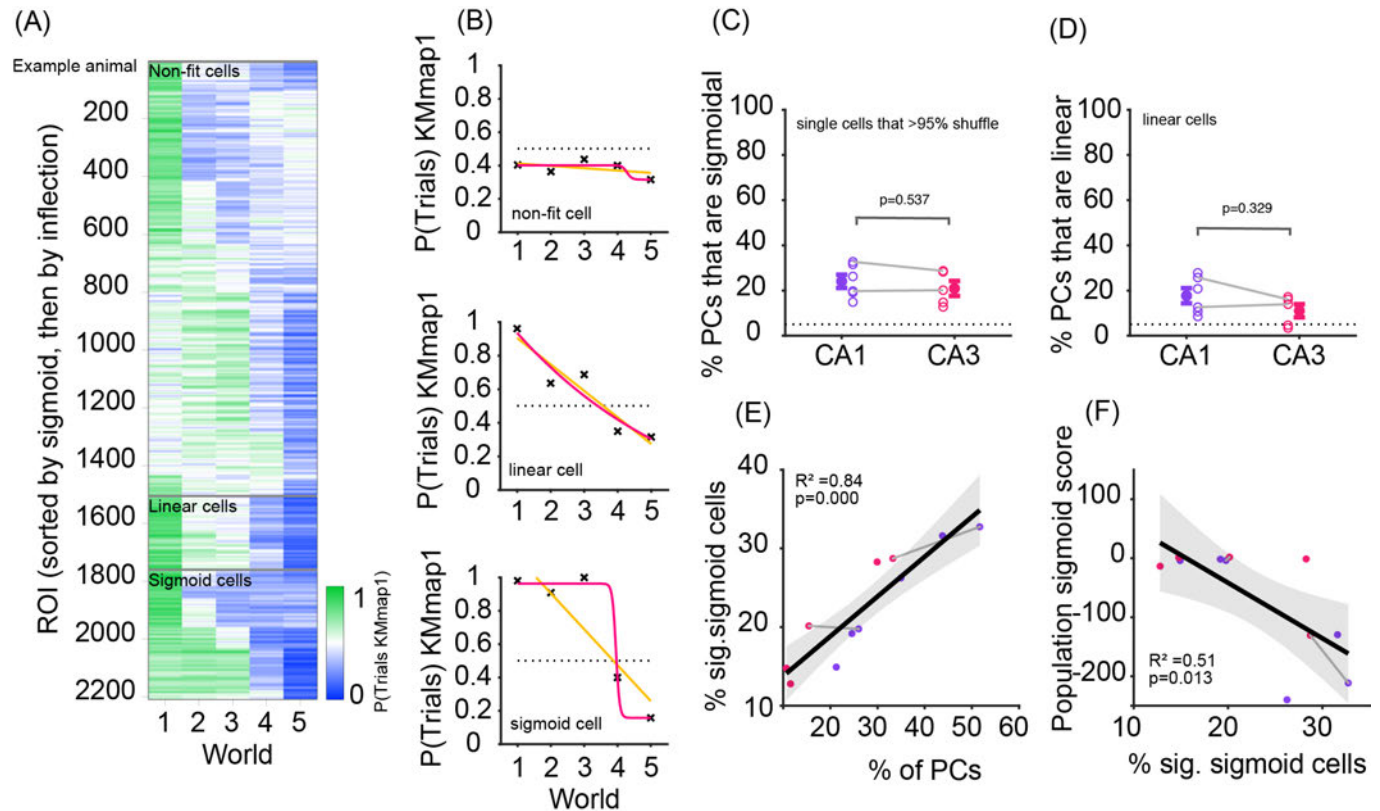
Figure 5-19. The proportion of single cells with nonlinear response profiles determines the population level nonlinearity.

A) Example heat map of all place cells in one session, where colour indicates which map (W1 or W5) each cell is using across worlds. Cells are sorted by those that are best fit by a sigmoid, or a linear fit or not fit by either.

B) Examples of the fitting procedure, where the proportion of trials a single cell is using the optimal map 1 across worlds is fit by a straight line and a sigmoid. yellow line = linear fit, pink line = sigmoid fit. Top= a cell that is not fit by either, middle= cell that is fit by a linear fit bottom = cell that is fit by a sigmoidal fit.

C) Percentage of place cells defined as sigmoidal within each animal for each subregion. Open circles = animals, closed circles with S.E.M error bars = average across animals.

D) Same as in C) but for cells defined as linear.

E) Relationship between the percentage of significant sigmoidal cells in an animal and percentage of place cells in that animal.

F) Relationship between the population sigmoid score (Figure 4.18 D) and the percentage of significantly sigmoidal cells in that population

Grey shaded area & Error bars = S.E.M. grey lines in C & D = dual animals. CA1 [purple], n= 6, CA3= [pink] n = 5.

5.3 Discussion

How are distinct memories formed and used for behaviour? To probe this question, we titrated the external environmental inputs by parametrically morphing virtual worlds and recording from large populations of CA1 and CA3 neurons. We related neuronal and behavioural discrimination during memory formation in virtual worlds. Our data describe similarities and differences in CA1 and CA3 place cell properties within the same animals for the first time. We show that both CA1 and CA3 neural population representations of the external world can be characterised by a non-linear relationship to gradually changing environmental inputs at both a neural population and single cell level. We provide evidence of a direct link between activity in hippocampal CA1 and CA3 circuits and reward seeking navigation behaviour, suggesting non-linear hippocampal place codes drive spatial behavioural discrimination.

5.3.1 Comparing CA1 and CA3 response properties in W1

5.3.1.1   Relationship between running speed and neural population activity in CA1 and CA3

While striking, our finding that animal running speed is positively correlated with network activity in CA1 and negatively correlated with CA3 network activity might be expected given evidence for an opposite relationship between population oscillations in CA1 and CA3 and running speeds (Czurko et al., 1999). Hippocampal network states include rhythms at theta (5–10Hz) and gamma (40–100Hz) frequencies. Standard models of hippocampal function posit that these two states are associated with distinct memory functions: encoding and consolidation, which occur during periods of either activity (encoding, Chapter 1.3.2) or quiescence (consolidation; Chapter 1.3.3). Running speed modulates the frequency of hippocampal gamma oscillations, notably slow (~20–50 Hz) and fast (~50–100 Hz) gamma frequency ranges in CA1 correspond respectively to CA3 or ERC coherence with CA1 (Ahmed & Mehta, 2012; Colgin et al., 2009; Kemere et al., 2013). Specifically, as running speeds increases, CA3 firing rates reduce, and ERC firing rates increase, suggesting that CA1 transitions from CA3 inputs to ERC inputs as running speed increases (Colgin et al., 2009). While we do not have the temporal resolution to resolve the frequency of such

events using 2P imaging, we were able to show changes in population activity in thousands of cells that would be consistent with supporting this interpretation.

Furthermore, this interpretation is in line with slow and fast gamma rhythms reflecting functionally distinct states in the hippocampal network. During slow gamma, hippocampal place cells tend to represent spatial locations separate from the current location (Zheng et al., 2016). This can occur as part of the sharp wave-triggered reactivation process (Carr, Karlsson, & Frank, 2012) or as maps of upcoming locations are activated during exploration (Bieri, Bobbitt, & Colgin, 2014). As running speed increases, animals likely rely more on self-motion cues, conveyed by MEC (Jacob, Poucet, Liberge, Save, & Sargolini, 2014; McNaughton et al., 2006; Van Cauter et al., 2013), to accurately keep track of where they are. Under these conditions, MEC spiking increases, and MEC-driven fast gamma becomes the dominant gamma activity in CA1. Hence, fast gamma frequency increases with running speed, allowing place cells to accurately keep track of rapid transitions across spatial locations (Jacob, et al., 2014; McNaughton et al., 2006; Van Cauter et al., 2013).

5.3.1.2 CA1 & CA3 assemblies

The manner in which environments are mapped by neural circuit is influenced by local variations in network structure, given connectivity determines the spatial arrangement of postsynaptic activation. Hence, pairwise activity correlations between neurons captures much of the statistical properties of a single neuron and provide a measure for studying the properties of population activity (Dettner, et al., 2016; Helias, Tetzlaff, & Diesmann, 2014). Weak pairwise correlations can give rise to emergent spatiotemporal structures in population activity (Halliday, 2000; Renart et al., 2010; Schneidman et al., 2006; Kriener et al., 2009; Yu, Hendrickson, Song, & Berger, 2018) and therefore we characterised the topographic organisation activity correlations in CA1 and CA3. Our finding of a similar spatial structure of pairwise correlations in CA1 and CA3 but higher neighbour correlations in CA3 fits with the notion that CA3 pyramidal neurons are more bursty than pyramidal neurons in CA1 (Hunt et al., 2016; Rauš Balind et al., 2020). While both CA1 and CA3 pyramidal neurons exhibit burst discharges for example, associated with the theta rhythm

and SWRs (Buzsáki, 1989; Csicsvari et al., 2000), it is known that CA3 neurons exhibit more spikes per burst, greater spike frequency, and larger amplitude spikes than CA1 neurons (Dong & Sheffield, 2021; Sneider et al., 2006). In line with this, we also found CA3 place cells had larger in field response amplitudes. These results support recent work that showed large calcium events are associated with complex spike bursting (Schoenfeld, Carta, Rupprecht, Ayaz, & Helmchen, 2021), and might lead to plasticity in the recurrently connected CA3 network to support the formation of functional engrams (Rauš Balind et al., 2019; Schoenfeld et al., 2021). The emergence of co-active CA3 ensembles and their relevance for hippocampus-dependent behaviours warrant further investigations using longitudinal calcium imaging (Schoenfeld et al., 2021) or functional manipulation (Carillo-Reid et al., 2018; Packer et al., 2015; Russell et al., 2019; Robinson et al., 2020).

5.3.1.3 Place coding in CA1 and CA3

Place cells have been found in all hippocampal subregions and regional specific differences in their coding properties have already been documented (Oliva et al., 2016; Jung & McNaughton, 1993, Leutgeb et al., 2007, Kjelstrup et al., 2008; Skaggs et al., 1996). Most notably Oliva et al, (2016) performed large-scale silicon probe recordings simultaneously across all layers in CA1, CA2, and CA3 in rats during real world spatial navigation (Oliva et al., 2016). While most of our findings agree, in that we both found fewer principal cells displayed space-dependent responses in CA3 than in CA1, and bigger place field sizes in CA3, we found opposing results on other spatial properties (Oliva et al., 2016). Notably, we found similar spatial information scores between CA1 and CA3 place cells, whereas Oliva et al. found CA3 place cells code more spatial information (Oliva et al., 2016). This might be explained by differences in our recording sites. On closer inspection, Oliva et al. found CA3a and b and proximal CA1 (CA1 closest to CA2) to have similar spatial information scores and it is likely that these are the parts of the subregion we had optical access too rather than distal CA1 and CA3c (Figure 1 A & K Oliva et al., 2016).

Moreover, we used different methodologies (electro- vs opto-physiology), and recently other optically identified place cells in CA1 and CA3 have been shown to be similar in their spatial coding

(Hainmueller & Bartos, 2018). It is possible that more heterogeneity in spatial tuning of place cells within subregions exists and is not detected using optophysiological methods as we currently do not consider hippocampal layers differences (Cherubini & Miles, 2015) and ignore the known proximo-distal gradients (Chapter 1.2.4; Lu et al., 2015). Hence, while similar to the other optically identified place cells, our between subregion comparisons were also similar given the existing proximodistal hippocampal gradients (Lu et al., 2015). Given the high spatial resolution provided by 3D 2P imaging, our future work could include mapping cells spatial position in the imaging volume to reconstruct the layers in the subregion and investigate subregion place coding at a finer spatial resolution. Ultimately, whether different behavioural correlates of these place cell properties exist across subregions was a point of interest for our later comparisons.

5.3.2 CA1 and CA3 response properties in a novel environment

5.3.2.1   Place cell properties in novel environments

The reorganization in the place field locations results in unique population-wide representations for different environmental contexts (Colgin, Moser & Moser, 2008). However, how different hippocampal subregions remap has not been fully explored. In line with previous work, we found a higher number of place cells on the familiar track than on the novel track (Hainmueller & Bartos, 2018). As many neurons were active in both contexts, we investigated whether active neurons were likely to have a place field in both contexts and found place cells appeared to form independent subgroups within the active cell population. This was similar to previous reports of place cells in CA1 and CA3 (Hainmueller & Bartos, 2018; Figure 1f) indicating that a separate place-code or representation exists for each context. Further comparison of spatial coding properties revealed both subregions similarly coded spatial information in both novel and familiar worlds. This differs from recent work, where an increase in spatial coding was found in novel worlds for both subregions, (Hainmueller & Bartos, 2018).

Notably in our experiments, CA3 place cells were less reliable (spatially stable) than CA1 place cells in novel contexts, which is surprising given some models of CA1 place field formation require

spatially tuned presynaptic inputs from CA3 (Bittner et al., 2017; Magee & Grienberger, 2020). Spatially tuned CA3 inputs coincident with ERC inputs generate dendritic plateau potentials in CA1 cells to drive synaptic potentiation (Bittner, et al., 2015). However, recently Dong and Sheffield (2021) found in a novel context, CA1 instantly forms a well-organized map with 30% place cells remapping on the first track lap, whereas CA3 forms a new population representation gradually with experience, with only 9% CA3 place cells remapping instantly (Dong & Sheffield, 2021; Figure 1h, i; Extended Data Fig. 2c, d). Hence, delayed CA3 place fields might explain our reduced trial-wise CA3 place cell reliability measures. Following this, instant CA1 place fields might not be formed through the mechanism previously described, but instead might form through other direct ERC inputs or inputs from Nucleus Reunions (Dolleman-Van Der Weel et al., 2019). Dong and Sheffield (2021) suggest, delayed CA3 place fields might form through a process of local dendritic spikes that occur on the initial laps in the absence of somatic firing (Sheffield & Dombeck, 2019; Sheffield et al., 2017; Sheffield & Dombeck, 2015; Bittner et al., 2017), potentiating a local clusters of synapses that become strong enough to drive firing, and possibly plateau potentials after a delay to ultimately form novel place fields (Sheffield & Dombeck, 2019; Sheffield et al., 2017; Sheffield & Dombeck, 2015; Bittner et al., 2017). These interpretations suggest CA1 does not simply inherit spatial information from CA3 during this form of spatial learning and suggests that CA3 place fields are less stable than CA1 place cells, contrary to the traditional assumption.

5.3.3 Morphing worlds

5.3.3.1 Gradual and abrupt transitions at different points along the morph sequence

Pattern completion and pattern separation are two algorithms thought to be implemented in the hippocampus (Chapter 1.5.2). How information undergoes such transformations between downstream (CA3) and upstream (CA1) areas is little understood beyond theoretical predictions (Rolls et al., 1998; Rolls 2007). According to current attractor network theory (Hopfield, 1984; Marr, 1971; Rolls et al., 1998; Chapter 1.5.2) discrete attractor dynamics for contextual memory are expected to manifest as an abrupt shift in the neural representation due to pattern

completion properties, as one context is morphed into the next, as is the case for global remapping (Colgin et al., 2008). We found non-linear response dynamics, indicative of attractor behaviour in both CA1 and CA3 networks. Hence, in familiar and novel environments that are sufficiently different, both CA1 and CA3 population activity shows separate neural representations of these worlds. However, in similar worlds (intermediate worlds), CA3 and CA1 representations were more coherent with the end point world they most closely resemble. Although we noted heterogeneity in the remapping response non-linearity between different animals.

The heterogeneity in neural representation dynamics we reported reflects the debate (mentioned in Chapter 1.5.3) between conflicting evidence suggesting average population activity transitions gradually (Leutgeb et al., 2007) or abruptly (Wills et al., 2005) from one familiar context to the next similar context. These contradictory results confuse the interpretation of network properties associated with the stored representations (Leutgeb et al., 2007). Interestingly, at the level of individual CA3 cells Leutgeb et al. showed a mix of gradual and abrupt representation transitions at different points along the morph sequence (Leutgeb et al., 2007). They also found some cells displayed hysteresis, a signature of attractor dynamics. While with calcium imaging we do not have the temporal resolution to resolve such dynamics, we were also able to show heterogeneity in representation transitions at the level of single cells, suggesting sub populations of cells likely drive non-linear population events.

Moreover, theoretical work has developed neural network models of CA3 with attractors for both position (continuous) and discrete contexts, to account for the possibility that CA3 memories might consist of local attractor states embedded in a continuous spatial map (Leutgeb et al. 2017; Solstad, Yousif, & Sejnowski, 2014). Solstad et al. found, position-dependent context attractors made transitions at different points along the morph sequence (Solstad et al., 2014). Subsequently, smooth transition curves arose from averaging across the population, while heterogeneous responses were observed on the single cell (Solstad et al., 2014). In contrast, orthogonal memories led to abrupt and coherent transitions on both population and single unit

levels as experimentally observed when remapping between two independent spatial maps (Solstad et al., 2014).

5.3.3.2   Similarity in neural population behaviour in subregions CA1 and CA3

Surprisingly, we found population activity in both subregions to be similar in that both CA1 and CA3 neural representations showed non-linear response dynamics to morphing VR worlds. We found this neural behaviour to more apparent in CA1 populations, contrary to our hypothesis that CA3 neurons would show a more abrupt transition in population activity given the differences in connectivity between subregions, in particular the abundant recurrent collaterals in CA3 and the functional properties such an architecture endows (Chapter 1.5.2; Rolls, 2007; Rolls et al., 2013; Treves & Rolls, 1992; Tsodyks, 1999). Our results are inconsistent with previous electrophysiology studies that showed CA3 population activity was less sensitive than CA1 activity to subtle environmental cue alterations, evidenced by a more coherent population representations between environments in CA3 than CA1 (Lee, Yoganarasimha, Rao & Knierim, 2004). Hence, Lee et al. concluded that CA3 displayed more homogeneity in the ensemble response to a subtly different novel environment than CA1 (Lee et al., 2004). It is possible, given their recording sites extending from the dorsal surface through CA1, tha Lee and colleagues target CA3c, whereas our work records from CA3a and b. Anatomical differences along the CA3 transverse axis may serve different computational roles (Chapter 1.5.3; Lu et al., 2005; Sun et al., 2017). Specifically, a functional gradient exists making CA3c maximally different from CA1 (Lu et al., 2005).

The functional gradient within CA3 is important, as more recent work has shown using a local-global cue mismatch (double-rotation) experiment, recordings from CA3c have revealed these place cells remapped (become silent or gain a new place field), whereas CA3a and b cells rotate their place fields to follow the local cues (Neunuebel & Knierim, 2014; Knierim & Neunuebel, 2016). Hence CA3c, where recurrent collaterals are weakest, showed behaviour reflective of the pattern separation property of the DG (Allegra et al., 2020; Neunuebel & Knierim, 2014; Knierim & Neunuebel, 2016). Whereas CA3a and b, where recurrent collaterals are strongest, maintained

coherent representations resembling the classic attractor network system (Neunuebel & Knierim, 2014; Knierim & Neunuebel, 2016). Hence, when considering the intra subregion recording sites, and the global environmental changes, it is less surprising Lee et al., found that 'CA3' remaps. These results highlight the importance of considering behavioural variables and functional heterogeneity within subregion when interpreting findings from these brain areas.

Furthermore, recent optophysiology studies have also found evidence of attractor behaviour in CA1 networks. Allegra, et al. (2019), used 2P imaging during virtual navigation of familiar, familiar' (similar) and novel environments. They found place cell correlations between similar environments were higher in CA1 than in the DG, suggesting that CA1 minimised and the DG amplified small context differences, providing functional evidence of pattern completion in CA1 and pattern separation in DG in vivo (Allegra et al., 2020). Adding to this, Plitt and Giocomo (2021) found that activity in CA1 is strongly driven by the animal's prior beliefs about the frequency of certain context. They recorded CA1 activity during a similar morphing experiment to ours and found the more familiar animals were with the intermediate morph worlds, the more gradual the neural representation transition (Plitt & Giocomo, 2021). Whereas the less familiar animals were with the intermediate morph worlds the more their neural representations showed non-linear dynamics (Plitt & Giocomo, 2021). These findings help contextualise and explain our own results. Our experimental procedures (Figures; 4.2, and 4.12), particularly in experiment 2, were most similar to Plitt & Giocomo's rare morph condition in which 50% of trials were in the end point (maximally different) worlds and the other 50% split between the intermediate worlds. It follows that we would be more likely to see non-linear neural response dynamics under these conditions and that such network effect are also expected in CA1 populations (Plitt & Giocomo, 2021).

5.3.4 Neural discrimination drives behavioural discrimination during virtual spatial navigation

5.3.4.1 Linking hippocampal activity to behaviour

While place cells have long been assumed to form the neural basis of a cognitive map of space, the direct influence of place cell activity on spatial navigation behaviour has only recently been demonstrated by our lab (Robinson et al., 2020). Using an 'all-optical' approach (Chapter 1.6.4), we identified and selectively activated place cells that encode behaviourally relevant locations in VR. Remarkably, targeted optogenetic stimulation of a small number of place cells was sufficient to bias the behaviour of animals toward that associated with the location of their place fields during a spatial memory task (Robinson et al., 2020). These findings demonstrate a causal role for place cell activity in guiding spatial navigation (Robinson et al., 2020), thereby providing direct support for long-standing theories about the behavioural function of the hippocampal cognitive map (O'Keefe and Nadel, 1978; Eichenbaum et al., 1999; McNaughton et al., 2006).

Following on from Robinson et al.'s (2020) demonstration of place cells driving behaviour, we still know relatively little about how other subregions encode distinct memories of similar objects and events (Danielson et al., 2016) and more specifically how these distinct representations are used for behavioural decisions (Leal, & Yassa, 2018). How information undergoes transformations between downstream (CA3) and upstream (CA1) areas is little understood. Supporting Allgera et al., (2020) we found a close relationship between behavioural and neural discrimination in CA1. We also found a similar relationship between CA3 neural discrimination and behaviour. In light of recent work, these data are compatible with the notion that behavioural relevance and saliency (given the rich nature of our VR environment) might trigger formation of new representations in CA3/1, perhaps driven by decorrelated information from the DG (Allgera et al; 2020). It follows that attractor dynamics in CA3 transition from pattern separation to completion (reflecting heterogeneity found within subregion), possibly as the differences in subregion inputs exceed a threshold (Plitt & Giocomo, 2021). CA1 then integrates CA3 inputs as well as other inputs from ERC and surrounding areas and the final representation in CA1 is then used by downstream neocortical circuits to guide learning and drive behavioural decisions (Frankland & Bontempi, 2005; Guzman et al., 2016; Hasselmo et al., 1995; Marr, 1971; McNaughton and Morris, 1987).

Collectively, our findings show single cells display non-linear neural dynamics in response to parametrically morphed virtual environments, and this activity likely underpins the non-linear neural population dynamics and also explains the variability in linearity. Furthermore, the similarity in response dynamics between CA1 and CA3 networks suggests the hierarchy of information flow in the trisynaptic loop model of the hippocampal may not be as pronounced as previously assumed. We show that activity in CA1 and CA3 reflects discrimination behaviour and hence, this decision-making information must already be present in CA3, again supporting the notion that rather than a trisynaptic chain, the hippocampus may be organised as a set of parallel loops. Hainmueller and Bartos (2020) suggest each loop could be formed by a hippocampal subfield receiving direct PP synapses and inputs from upstream hippocampal areas (Hainmueller & Bartos, 2018; Hainmueller & Bartos, 2020 Fig. 1c; McNaughton et al., 1989; Kitamura, et al., 2014), as there are multiple instances where correct memory storage and retrieval is achieved based on the PP–CA1 or PP–CA3–CA1 loop alone (Hainmueller & Bartos, 2020).

6

General Discussion

The neural firing of hippocampal place cells often strongly correlates with an animal's current spatial location in an environment, providing a neural basis for the brain's code for space O'Keefe & Dostrovsky, 1971; O'Keefe & Nadel, 1979). Between environmental contexts, place cell firing fields can reorganise in phenomena collectively referred to as remapping. The extent to which different hippocampal subregions remap and under which environmental contexts has been an area of intense debate (Chapter 1.4.3; Kentros et al., 1998; Cacucci, et al., 2007; Mankin et al., 2012). Different hippocampal subregion connectivity suggests specialised subregion computations (Rolls, 2007). However, methodological difficulties in accessing large neural populations across subregions and challenges in precisely controlling environmental inputs has limited experiments assessing this question. In this thesis we devise novel methods and experiments to probe questions arising from this body of work.

6.1.1 Summary of thesis

This thesis describes new methods that:

1) Contribute to founding the emerging field of in vivo 2P calcium imaging of CA3 pyramidal neurons, by establishing a method for imaging genetically defined CA3 pyramidal neurons in vivo in a 3D volume of hippocampus, optically accessing the largest population of CA3 cells reported to date.

2) Demonstrate the first protocol for dual hippocampal subregion imaging based on unique markers for CA3 pyramidal cells and differences in soma size, opening up cross subregion imaging experiments within the same animals.

3) Provide the first demonstration of an all optical approach involving simultaneous 2P imaging and 2P photoactivation of neurons in subregion CA3, demonstrated in (Russell et al., in prep).

4) Provide open source material to quickly build an affordable and space efficient mini modular VR systems, simplifying complex experiments.

This thesis also describes experiments using these methods that suggest the following scientific findings:

164

1) CA1 and CA3 network properties are differently modulated by brain state. Increased running speed increased CA1 network activity and vice versa for the CA3 network.

2) Fewer principal cells displayed space-dependent responses in CA3 than in CA1, but place cells in CA3 had bigger place field sizes, despite similar spatial information scores.

3) CA1 and CA3 place cell networks behave similarly demonstrating non-linear neural response dynamics during virtual navigation through parametrically morphed VR environments.

4) Non-linear neural population behaviour in CA1 and in CA3 drives behavioural discrimination during spatial navigation.

6.1.2 Future perspectives

6.1.2.1 CA3 and dual subregion 2P imaging

At the time of project conception (2016) there were no reported 2P studies of subregion CA3, a testament to the difficulty in accessing this region with an imaging objective. However, recently a few imaging studies have been conducted in CA3 and use a similar method to ours (Bartos & Hainmueller, 2019; Schoenfeld et al., 2020). The similarity in our FOVs, and findings reflect the robustness of the surgical method developed. One limitation of the current published CA3 imaging studies is that they access a relatively small population of neurons given the angled nature CA3 anatomy. We circumvented this limitation by using volume imaging to access 3 additional imaging planes and thus access a much larger subsample of CA3 neurons. One future direction might include identifying layers in CA3 (Amaral et al., 2007). For example, by imaging at subcellular resolution (see Figure 2.3), the recently identified athorny CA3 cells, known to exist deep layer SP in CA3a and b in vitro (Hunt et al., 2016) could be confirmed in vivo. Athorny CA3 cells are bursty (Hunt et al., 2016) and thought to play a key role in triggering slow waves, hence establishing this function in vivo is paramount.

The new dual hippocampal subregion imaging protocol demonstrated in this thesis opens up opportunities to study the propagation of information between downstream (CA1) and upstream

(CA3) regions. While much is known regarding the information flow around the trisynaptic loop, it is only recently we have been able to study at the level of large neural populations the different subregions individually (Dombeck et al., 2010; Bartos & Hainmueller, 2018) and our method represents the first time two subregions were recorded simultaneously in vivo. Although exciting, our approach does have caveats as it ignores the smallest subregion, CA2 (Lorente De Nó, 1934; Dudek, Alexander, & Farris, 2016). Although, there is debate over whether CA2 is truly a distinct region from CA3 (Gaarskjaer, 1986). CA2 has been implicated as fundamental to social memory processes and olfactory cues, given CA2's non-spatial information inputs from ERC (Chevaleyre, & Siegelbaum 2010; McHugh et al., 2007; Mankin et al., 2015; Dudek et al., 2016; Smith et al., 2016). It has also been suggested that CA3 neurons closest to CA2 have firing properties more similar to those in area CA2 than to the rest of the CA3 (Caruana, Alexander, & Dudek, 2012; Dudek et al., 2016). Hence, to exclude CA2 neurons in our dual subregion preparations, we use an exclusion boundary between subregions, but we cannot be absolutely certain we are not imaging some CA2 cells in our dual region experiments. CA2 cells would be difficult to identify using our current method, given CA2 pyramidal cell bodies are larger than those in CA1, and similar in size to CA3 pyramidal cell bodies (Caruana et al., 2012). One solution would be to post hoc identify CA2 pyramidal neurons using a Purkinje cell protein 4 (PCP-4) stain (San Antonio et al., 2014). However, this too could be difficult given the challenge in identifying reconstructed imaging FOVs. Hence, an exclusion zone between boundaries, even with its caveats should suffice and similar exclusion methods have been used to sperate other cross region imaging cortical preparations.

6.1.2.2   All optical activation of CA3 neurons in vivo

In this thesis we demonstrate the first all optical activation of CA3 pyramidal neurons. While the reliability of our expression strategy remains an area of continued work, our efforts contribute to the development of this strategy from initially photoactivating neurons in superficial cortex (Packer et al., 2015), to deep cortex, (Russell et al., 2019) to subcortical areas (CA1), and now to deeper subcortical areas (CA3; Russell et al., 2021 in prep). This all-optical method (see Russell et al., 2021) offers the ability to record from thousands of neurons in multiple brain areas and probe

166

network function at single cell resolution, combined with behavioural read outs. The extension of this methodology to deep subcortical regions described in this thesis, takes us closer to being in a position to crack the hippocampal neural code, linking different subregion circuit function directly to behaviour.

The notion that CA3 networks endowed with recurrent connectivity and function as an attractor network was an idea put forward decades ago (Bhatia, Moza, & Bhalla, 2019; Marr, 1971; Rolls, 2007). Work in this thesis utilised our novel all optical approach to assess CA3 circuit function with single cell resolution and provides the first demonstration of potential attractor dynamics in CA3 in vivo (Chapter 2.3.3). While neural circuits are the substrate for information processing, their dynamics are governed by a complex interplay between input activity and the responses of excitatory and inhibitory neurons (Bhatia et al., 2019). Thus, the organization of patterns of synaptic connectivity is an important determinant of the dynamics, and therefore one of the main caveats of our experiment is that we were not able to resolve connectivity motifs in vivo. We could only infer that our boosted network response is the result of CA3-CA3 recurrent connections.

Furthermore, we have similar issues clarifying our interpretation of suppressed CA1 responses to CA3 stimulation as indicative of feedforward inhibition reflecting the canonical EPSP-IPSP sequence (Alger & Nicoll, 1982; Lovett-Barron et al., 2012; Miles, 1990; Pouille & Scanziani, 2001; Schwartzkroin, 1975). To confirm this, we need to resolve underlying connectivity motifs to reveal monosynaptic excitation and disynaptic inhibition (Guzman et al., 2016). This has not yet been demonstrated in vivo, but recent work developing in vivo voltage imaging (Adam et al., 2019; Knöpfel & Song, 2019) and improvements in somatically restricted opsin expression (Marshel et al., 2019) could potentially make connectivity mapping at this resolution a real possibility in the near future. Nevertheless, experimentally linking neural microcircuit function to emergent properties of the brain requires fine-scale manipulation and measurement of neural activity during behaviour, where each neuron's coding and dynamics can be characterized (Carrillo-Reid et al., 2020). Our proof-of-principle experiments takes the first steps towards this

goal by demonstrating all optical manipulation of subcortical hippocampal circuits in subregion CA3 (Russell et al., 2021 in prep).

6.1.2.3 Population activity and spatial coding differences between CA1 and CA3 networks

One stark contrast between CA1 and CA3 circuits is their population activity correlation with behavioural state, specifically running speed. Our findings reflect the notion that there is continuous modulation of information processing in the hippocampal circuit as a function of behavioural state, transitioning from strong CA3 drive of CA1 activity at low running speeds to ERC drive of CA1 activity at higher speeds. It would be interesting to probe the impact of environment novelty on this correlation (Kemere et al., 2013). Typically, as the environment become more familiar to the animals, they run faster, hence speed modulation should be stronger in novel environments. Such a finding would strengthen our interpretation that CA3 inputs predominate at slow gamma where place cells tend to represent spatial locations separate from the current location. In contrast, at fast running speeds animals likely need to rely on self-motion cues, conveyed by MEC (McNaughton et al., 2006; Van Cauter et al., 2013; Jacob et al., 2014), to accurately track position. Under these conditions, MEC spiking increases, and MEC-driven fast gamma becomes the dominant gamma activity in CA1 (Zheng et al., 2015). Hence, fast gamma frequency increases as running speed increases, enabling place cells to accurately keep track of rapid transitions across spatial locations (Zheng et al., 2015).

The present pattern of results regarding spatial properties of place cells in CA1 and CA3 highlight the importance of considering the subregions beyond homogenous entities. As mentioned throughout this thesis (Chapter, 1.2.2, Chapter 6.1.2.1), molecular, (Vogel et al., 2020) connectivity (Qiongling et al., 2020) and functional gradients in place cell properties across subregions exist (Colgin et al., 2010; Lee et al., 2015; Lu et al., 2015). Our work presented in this thesis neglects these subtleties. However, future work could use alternative methods for example, a GRIN lens (Moretti et al., 2016) or excavating CA1 to image deeper CA3 areas (CA3c and DG). These differences in methodological approaches might be useful in revealing differences between CA1 and CA3(c) that are more comparable with traditional electrophysiological studies

that compare and contrast these specific subregions. While such methods would enable a comparison of CA3c to CA1, they are not without their own caveats, most notably increased damage to the hippocampus than the methods we used. A better understanding of the specific part of the subregion imaged would enable a deeper understanding of the conditions in which subnetworks, specifically the CA3c – distal CA1 'non-spatial' network and the proximal CA1 - CA3a/b 'spatial' networks are recruited to comprehend how information is processed in the hippocampal network (Beer et al., 2018). At present, performing recording simultaneously across subregions still remains a major challenge for both electrophysiological and optophysiological studies since the coordinates of the proximal and distal parts of CA1 and CA3 vary greatly along the transverse axis of the hippocampus due to its folding (Beer et al., 2018). Our work takes the first steps to opening up cross subregion imaging and combined with the alternative methods discussed, will begin to unpick circuit functions at a level not previously possible.

6.1.2.4 Manipulating environments to probe hippocampal subregion dynamics

To clearly represent similar objects and events that have different behavioural significance, the hippocampus must transform similar inputs into well-separated neuronal memory representations (Allgera et al., 2020). The input region of the hippocampus, the DG, is thought to implement this by orthogonalizing cortical inputs through sparse firing activity patterns (Gilbert et al., 2001; GoodSmith et al., 2017; Leutgeb et al., 2007; McNaughton & Nadel, 1990; Neunuebel & Knierim, 2014; O'Reilly & McClelland, 1994; Rolls & Treves, 1998; Treves & Rolls, 1992). In contrast, the downstream circuits CA3 and CA1 are thought to retrieve memorized patterns from incomplete or degraded input via attractor dynamics and to transfer these memory representations to the neocortex, where they are processed to drive behaviour (Frankland & Bontempi, 2005; Guzman et al., 2016; Hasselmo et al., 1995; Marr, 1971; McNaughton & Morris, 1987). Our findings support these interpretations. While average population activity in CA3 was reported to transition nonlinearly from one familiar context to the next, suggesting some attractive forces associated with the two stored representations (Leutgeb et al., 2005), although this finding was less abrupt than predicted based on attractor network theory (Chapter 1.5; Rolls,

2007). It is possible that a mix of gradual and abrupt network transitions occur at different points along the morph sequence and might underpin this result (Solstad et al., 2014).

Alternatively, it is possible that our intermediate VR worlds were not as similar as we assumed them to be to the end point worlds. Our parametric morphing procedure involved arbitrarily choosing maximally different features (e.g. blue vs. green) to morph; as a result the intermediate worlds created might be more different from the end point worlds than optimal. To probe how subregions truly discriminate very similar inputs and account for this caveat, one could make an additional 5 morph worlds using the 3rd intermediate world as one end point world, thus creating more subtly morphed worlds, that would ultimately be more similar. Finally, it is also thought that CA1 representations might diverge over time, Lever et al. (2002) found CA1 place cell representations, after repeated expose to two subtly differently shaped environments, gradually and incrementally diverged (Lever, Burgess, Cacucci, Hartley, & Keefe, 2002). In this thesis, we have only considered representations at one point in time and thus it is possible that CA1 spatial representations diverge over repeated exposure to the intermediate worlds, such that each world is represented differently. Perhaps this would result in a more seemingly gradual transition in CA1 spatial representations between the end point worlds, fitting the initial predictions from attractor theory (see Chapter 1.5). However, much work is needed to dissect the role of specific hippocampal subregion place cells after place field emergence, and how the place field properties (position, width, shape) evolve with familiarization. Studying these dynamics is a proxy for studying the synaptic plasticity mechanisms supporting experience-dependent activity patterns in the hippocampus at a system level

6.1.2.5 Neural discrimination behaviour underlies behavioural discrimination behaviour for both subregions

By correlating the behavioural readout of which world the animal thought itself to be in with the neural representation, we were able to investigate the extent to which behaviour reflects neural activity patterns seen as the animal was traversing through the track prior to the reward point.

While not the first example of place activity in CA1 driving behaviour, given we showed that targeted stimulation of specific place cells was sufficient to bias behaviour towards that associated with the location of their place fields (Robinson et al., 2020). Our current findings are the first example of CA3 neural representations reflecting behavioural discrimination and take the first steps towards suggesting a causal role for CA3 place cell activity in guiding spatial navigation and decision making. Further work is needed to confirm this statement, ideally through direct manipulation of CA3 place cell activity (see below 6.1.3).

Alternatively, an interesting future avenue would be to assess the role of neuromodulators such as acetylcholine that can reconfigure neural networks (Hasselmo, 2006; Prince, Tsaneva-Atanasova, Clopath, & Mellor, 2017). Acetylcholine has been proposed as a filter or gate mechanism to select which experiences should be encoded (Hasselmo et al., 1995). Acetylcholine is thought to promote encoding of novel information by facilitating NMDA receptor function and induction of synaptic plasticity (Buchanan, Petrovic, Chamberlain, Marrion, & Mellor, 2010; Dennis et al., 2016; Marino, Rouse, Levey, Potter, & Conn, 1998; Markram & Segal, 1992) and also selectively suppressing recurrent activity representing stored information, in favour of feed-forward activity representing novel information (Hasselmo, 1995). Hence, salience is indicated by increasing the release of acetylcholine in the hippocampus (Dayan & Yu, 2006). In light of our experiments, optogenetic acetylcholine release could be used to increase the saliency of the end point worlds, potentially improving their encoding as two distinct worlds. It follows then, that we might find a more abrupt shift in the neural representations such that intermediate worlds are represented as more similar to one or the other end point worlds. This is particularly true if acetylcholine release is also blocked during presentation of intermediate worlds to reduce the suppression of recurrent activity potentially containing the stored end point world representations. Manipulating this neurotransmitter would enable us to probe its function and it provides testable predictions for population dynamics under uncertain conditions (subtle environment changes). This would enable us to strengthen our interpretation of neural population activity driving behavioural decision making.

6.1.3 A dream experiment

The hippocampus shows three main classes of rhythms: theta (∼4–12 Hz), sharp wave–ripples (∼150–200 Hz ripples superimposed on ∼0.01–3 Hz sharp waves; SWs) and gamma (slow and fast∼25–100 Hz; Colgin et al., 2016). Sharp wave/ripple (SWR) events in the hippocampus are thought to coordinate the reactivation of stored memory traces and direct their reinstatement in cortical circuits (O'Niell et al., 2010; van de Ven et al., 2016) and the planning of future behaviours (Foster & Wilson 2006; Ólafsdóttir et al., 2018; Pfeiffer, 2020). SWs originate from the CA3a and CA2 regions and are observed in a diverse range of mammalian species from primates to rodents (Csicsvari et al., 2000). These phylogenetically conserved rhythms that occur during periods of quiescence, and slow-wave sleep (Colgin et al., 2016). SW attractor states reflect pattern completion computations vital to memory function (Hunt et al., 2018; Rolls 2007). Evidence in support of this relationship comes from the observation that cellular activity patterns occurring during recent experiences are reactivated in temporally compressed synchronous epochs (SW replay; Ólafsdóttir et al., 2018). In the context of our findings, information about cell assembly patterns stored at recurrent synapses can be reactivated during SWs when driven by internally generated or extrinsic partial patterns, (readout of these partial patterns might be performed by athorny cells; Hunt et al., 2018) promoting pattern completion that manifests as a SW attractor. Hence one could imagine an experiment, similar to those described in Chapter 4, but combined with an all-optical approach to enable user controlled optogenetic activation of part or the complete neural representation of one of the end point worlds when the animal is in the intermediate world. This approach would enable us to potentially bias behaviour and demonstrate the functional relevance of neuronal ensembles containing spatial information underpinning behavioural decisions. Ultimately, an experiment like this would allow us to understand and probe the neural code for spatial memory.

This dream experiment would also enable us to assess the functional relevance of attractor behaviour by activating part of known neural representations and read out the network response while simultaneously recording conditional spatial behaviour. While exciting, these types of

experiments would be technically difficult on multiple levels. Firstly, not all neurons in the ensemble are 'pattern completing' neurons (Carrilo-Reid et al., 2019; see below), and therefore we would need to select neurons best positioned to elicit reactivation of their functionally connected neighbours, particularly if we hope to drive behaviour with these inputs. Secondly, these experiments would also enable us to probe the precise number of neurons in a functional unit required to drive behaviour. While proponents of emergent network theory (Chapter 1.5.1) believe in the distributed yet collective activity underpinning cognitive functions, it would be experimentally challenging to capture such distributed populations or titrate the number of neurons needed to drive behaviour, particularly when connectivity motifs are not known. While challenging, these experiments are not out of reach, as recently similar studies to those described above have been conducted in visual cortex (Carrilo-Reid et al., 2019; Carrilo-Reid et al., 2020; Marshel et al., 2019; Russell et al., 2019). With our imaging protocol giving us access to CA3 neural populations, we are in a prime position to extend these experiments to the hippocampus. A particularly interesting step given the wealth of computational literature suggesting the CA3 circuit functions as an auto associative network (Chapter 1.5.2), that has yet to be directly experimentally probed with sufficient resolution (Knierim & Zhang, 2012; McNaughton et al., 1996; Rolls, 2007).

One factor outstanding in these dream experiments is provision for the temporal resolution of neural activity, both at a macro level (population oscillations) and a micro level (firing rates). While 2P imaging has slow kinematics, other recent examples of artificially activating neurons with 2P optogenetics has been shown to drive behaviour (Carrilo-Reid et al., 2019; Carrilo-Reid et al., 2020; Marshel et al., 2019; Russell et al., 2019). In light of population dynamics, if network wide oscillations could be detected, it is possible behaviour change might be more likely during these events (Colgin et al., 2016; Csicsvari et al., 2000). Hence if we could detect SW using the LFP, then activate some cells from a stored memory representation, we might be more likely to see pattern completion than if we were to activate the same cells at a different point in time. While the functional relevance of network wide oscillations has been shown particularly during replay, in light of the functional consequences of disturbed replay, it has not been experimentally

linked to network dynamics like pattern completion beyond theory or demonstrated to directly to link to behaviour. Adding this layer of complexity to the experiment would be technically difficult given the combination of electrophysiological and optophysiological methodologies and interfering electrical artefacts. However, the combination of LFP detection, to select potentially more plastic periods of concerted network activity during large-scale imaging and to then use these times to inform artificial stimulation in vivo, would be key in confirming the role of neuronal ensembles or groups of coactive neurons, as the functional building blocks of cortical circuits and units of the neural code (Carrilo-Reid et al., 2020).

6.1.3.1 Future experiments summary

The notion that recurrently connected ensembles of neurons work as functional units was developed by Hebb, who argued that these "cell assemblies", arise from synaptic learning rules where pre and postsynaptic cells would be coactive, and could be used to implement associative memories (Hebb, 1949; Chapter 1.3). Hopfield subsequently provided a mathematical foundation for these ideas, arguing that recurrently connected neural circuits settle on dynamical "attractors", i.e., stable states of network activity, (Chapter 1.5; Hopfield, 1982). The settling dynamics endow these attractors with pattern completion properties, i.e., activation of a few of the neurons could trigger the entire pattern. These stable states implement memory representations. Collectively these pioneers predicted essentially the same idea: that the neural code is written with multicellular units, built by the coordinated activity of groups of neurons (Carrilo-Reid et al., 2020).

Combining our advancements in imaging with the all-optical approach provides the first opportunity to experimentally assess Hebbian learning in the hippocampus and specifically in subregions CA3 where attractor behaviour has been extensively theorised. Recent results using holographic optogenetics have demonstrated that activation of ensembles, by activating individual pattern completion neurons, can control visually evoked behaviour (Chapter 1.6.3; Carrillo-Reid et al., 2019; Marshel et al., 2019). Stimulation of as few as two pattern completing neurons was sufficient to change the behaviour (Carrillo-Reid et al., 2019; Marshel et al., 2019),

providing a causal link between neuronal ensemble activity in V1 and behaviour (Carrillo-Reid et al., 2019). These results have also demonstrated that targeting neurons with pattern completion capabilities can be an effective strategy to control behaviour through indirect recruitment of functionally connected neurons (Carrillo-Reid et al., 2016; 2019). Although the idea of recalling a learned behaviour by the stimulation of single neurons is not new (Brecht et al., 2004; Houweling & Brecht, 2007; Romo et al., 1998) the precise targeting of neurons paves the way to understanding circuit mechanisms underlying different behaviours with single cell precision. In light of these recent experiments and work demonstrated in this thesis, understanding the functional unit of cognitive behaviours is within reach. We have the experimental tools, and can build on preceding work to finally probe Tolman's proposition that the hippocampus functions as a cognitive map.

7

Appendix
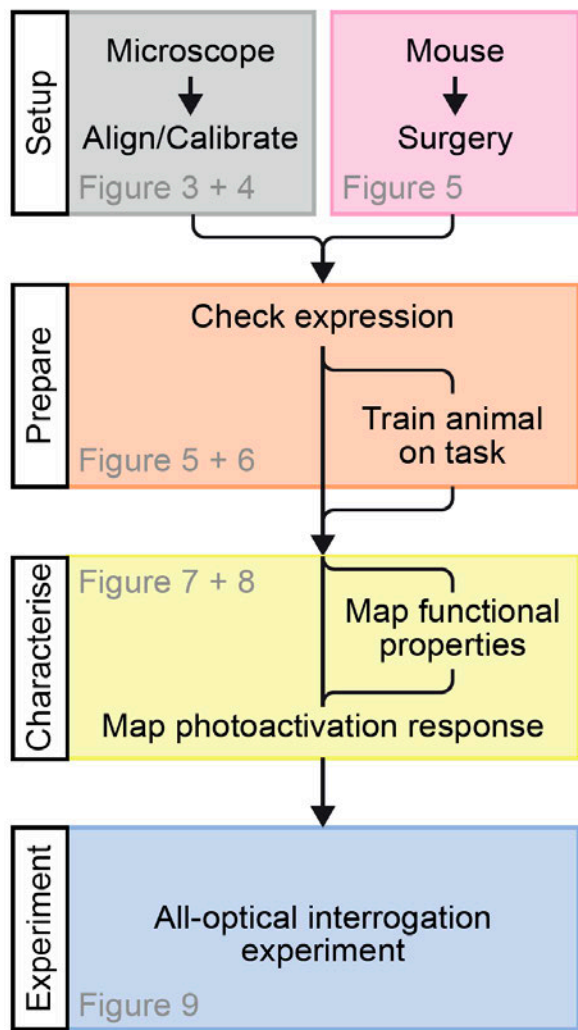
Figure 7-1. Experimental workflow for all optical experiments.

Stages involved in setting up and running an all-optical experiment. For full details see Russell et al., 2021. Taken from Russell et al., 2021 in prep).
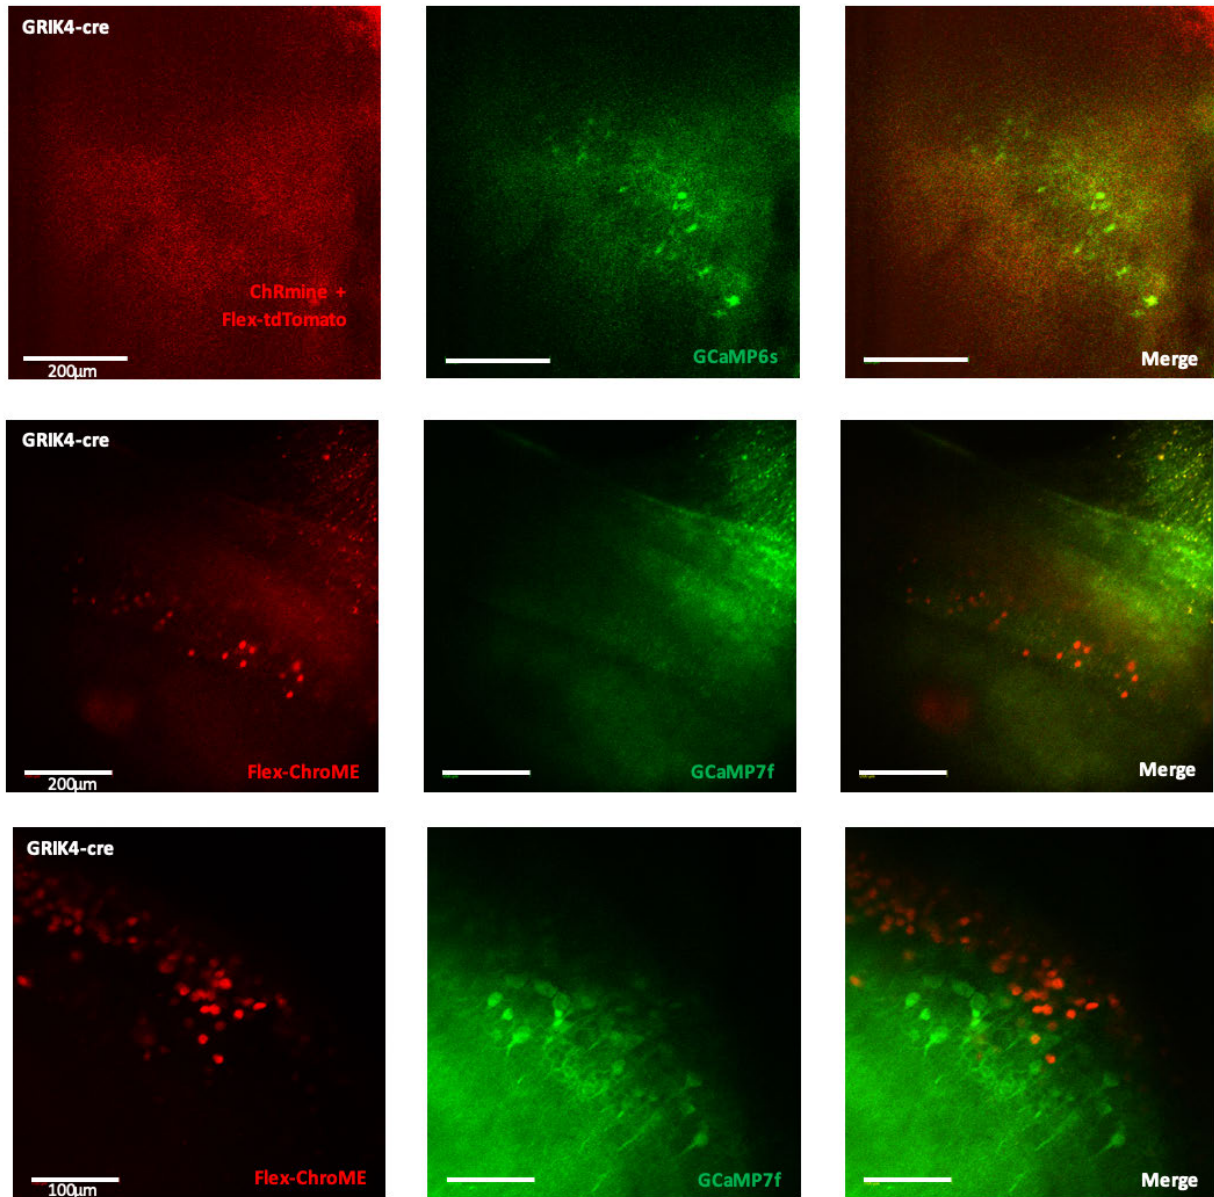
Figure 7-2. Co-expression problems for activity indicators and optogenetic actuators in CA3 in vivo.

Two photon FOVs showing expression of opsin (red), GCaMP6s (green) and co-expression (yellow; merge) in CA3 of multiple mice (genotype listed in top left corner, specific opsin/GECI listed bottom right). Top row) GCaMP expression with failed opsin expression, Middle row) opsin expression with failed GCaMP expression. Bottom row) epxression of both GCaMP and opsin but mostly not co expressed in the same cell. Scale bars 200 or 100 µm as listed.

GRIK4-cre

GCaMP6s +
ChRmine + Flex-

200μ

Stim Groups
1
20

GRIK4-cre

ChroME
GCaMP6s

200μ

Stim Groups
1
20

GRIK4-cre

Flex-tdTomato
GCaMP6m-ChRmine

200μ

Stim Groups
1
14

EMX-cre

ChRmine
GCaMP6s

200μ

Stim Groups
1
16
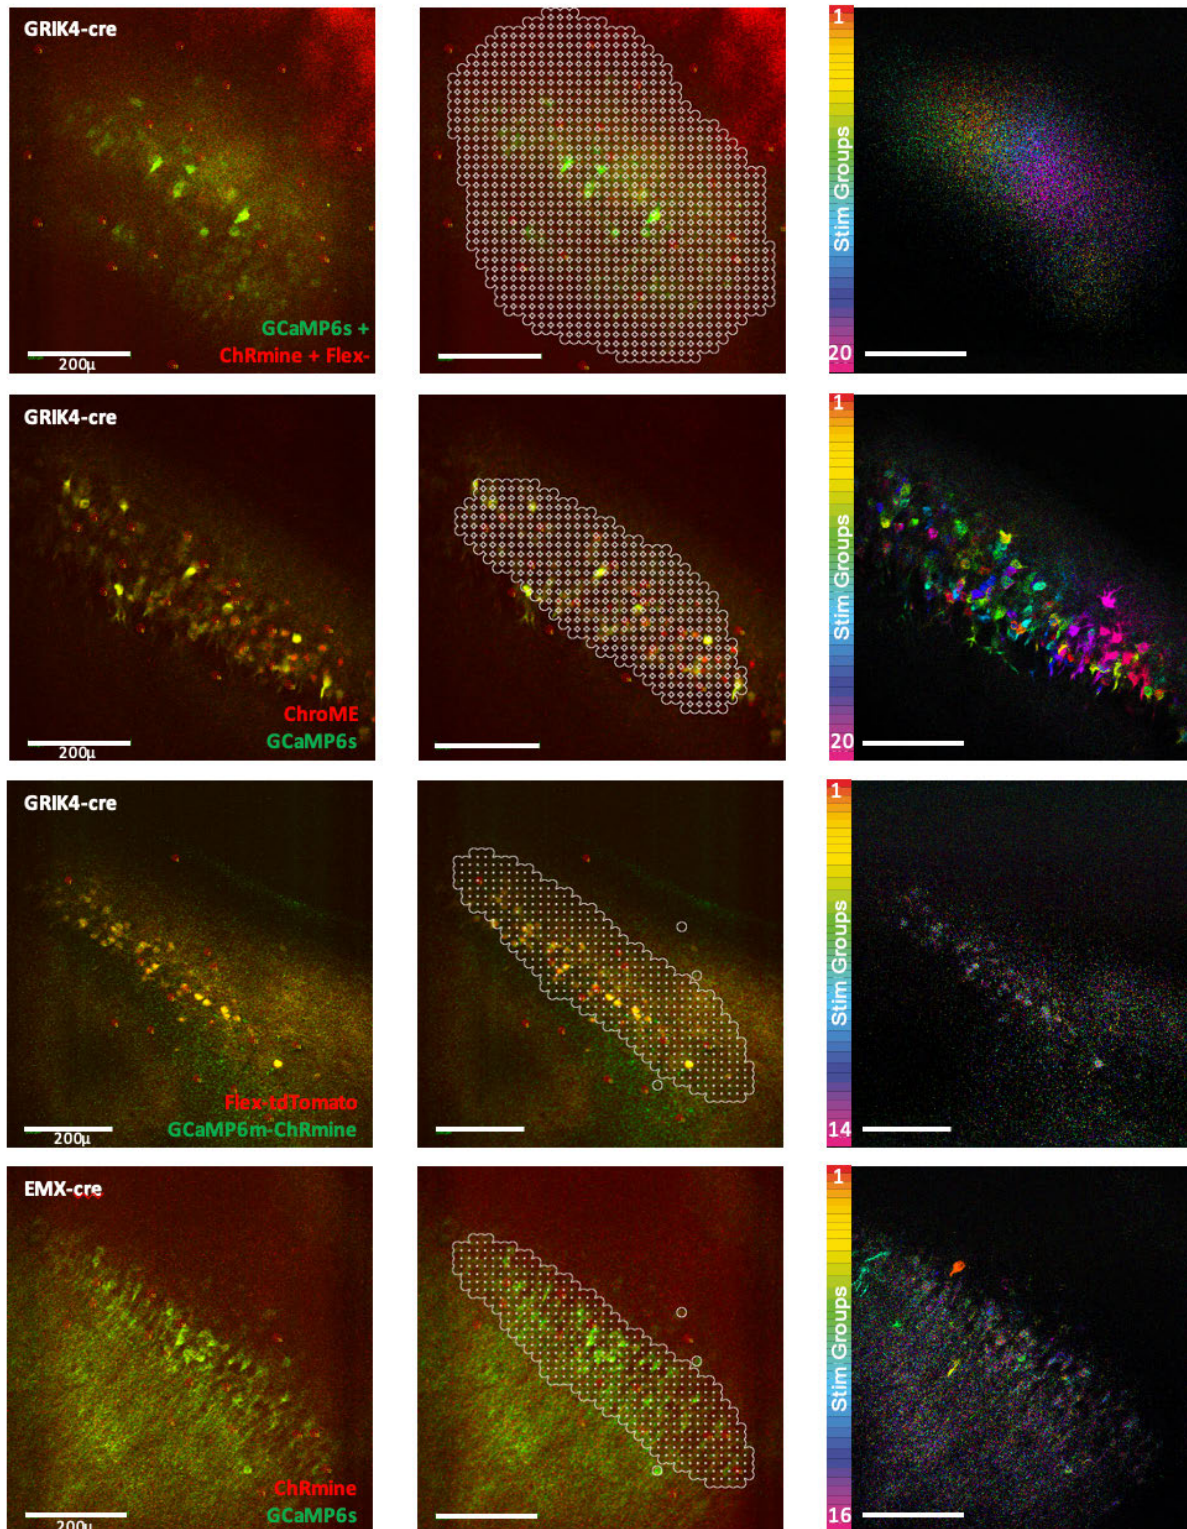
Figure 7-3. Non photoresponsive FOVs in CA3.

CA3 FOV in various mice (genotype listed top left corner), expressing GCaMP6 (green) and opsin (red; specific opsin/GECI combination listed bottom right corner and showing failed photostimulation responses. Next, ROIs selected for photostimulation (white rings). Next, target sites were grouped into different patterns of 20 cells each and targeted for sequential photostimulation to confirm responsivity. Pixel intensity indicates the change in florescence caused by photostimulation. Colour corresponds to the photostimulation cluster (numbered – left) which caused the largest change in activity (all scale bars, 200µm).
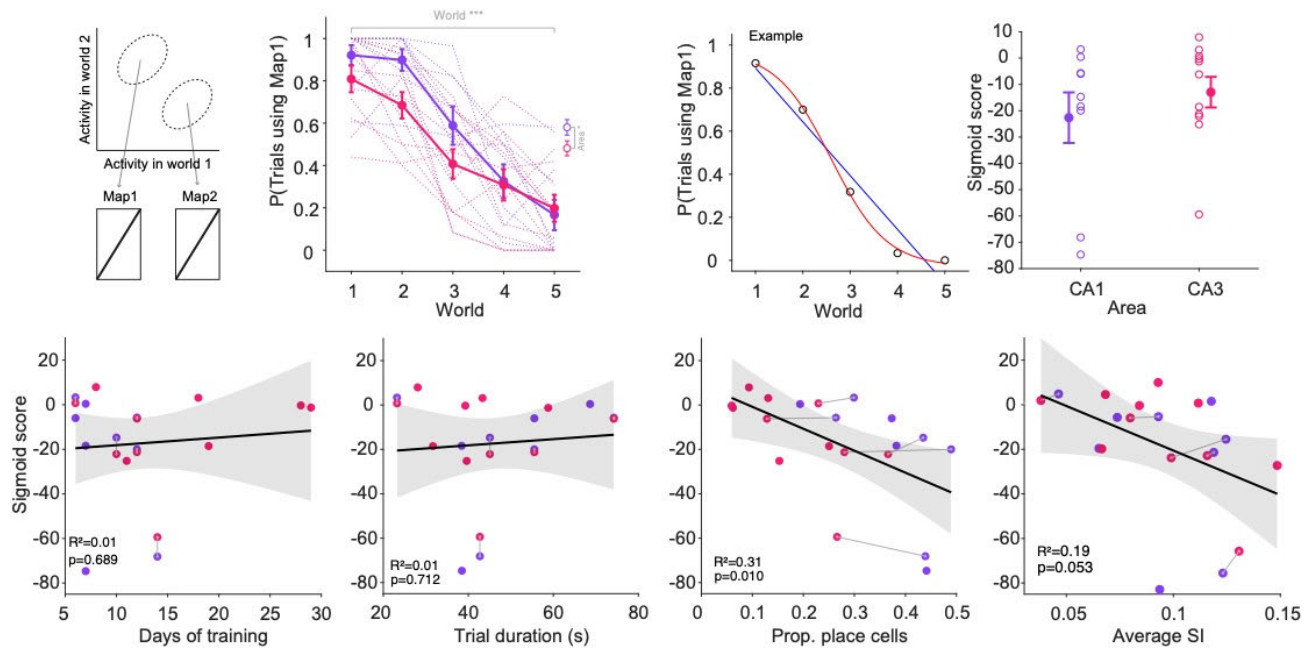
Figure 7-4. K-means based identification of optimal world maps reveals nonlinear mapping of world representations.

A) Schematic to demonstrate the k-means method of unsupervised identification of maximally different representations (clusters).

B) The proportion of trials using map (cluster) 1 across worlds and by subregion

C) Example to demonstrate the simple linear (blue) and sigmoid (red) fitting procedure for data points (proportion of trials using k-means map 1) across worlds per for one animal.

D) Sigmoid score for each subregion, defined as AIC of the non-linear minus the AIC of the linear fit. Open circles = animals, closed circles with S.E.M error bars = average across animals

E) Relationship between the average sigmoid score and the number of training days per animal. Grey connecting lines = dual mice, i.e. area CA1 and CA3 are from the same animal.

F) Same as in E) but for trial duration

G) Same as in E) but for proportion of place cells

H) Same as in E) but for average S.I.

Grey shaded area & Error bars = S.E.M. Pale track coloured lines = individual animals. Thicker solid lines = means. CA1 [purple] n= 6, CA3 [pink] n = 5.

181

8 References

Abbott, L., (2000). Synaptic plasticity: taming the beast. *Nature.Com*.

Ahmed, O. J., & Mehta, M. R. (2012). Running speed alters the frequency of hippocampal gamma oscillations. *Journal of Neuroscience*, *32*(21), 7373–7383. https://doi.org/10.1523/JNEUROSCI.5110-11.2012

Alger, B. E., & Nicoll, R. A. (1982). Pharmacological evidence for two kinds of GABA receptors on rat hippocampal pyramidal cells studied in vitro. *The Journal of Physiology*, *328*(1), 125–141. https://doi.org/10.1113/jphysiol.1982.sp014256

Aliñanes, G. R. L. G., Archand, P. A. U. L. J. M., Aphaël, R., Urcotte, T., Ellat, S. E. P., Uber, D. A. H., … Servet, R. M. (2018). *Optical alignment device for two-photon microscopy*. *9*(8), 15741–15746. https://doi.org/10.1364/BOE.9.003624

Allegra, M., Posani, L., & Schmidt-hieber, C. (2019). *The hippocampus as a perceptual map : neuronal and behavioral discrimination during memory encoding*.

Allen, K., Nick, J., Rawlins, P., Bannerman, D. M., & Csicsvari, J. (2012). Behavioral/Systems/Cognitive Hippocampal Place Cells Can Encode Multiple Trial-Dependent Features through Rate Remapping. *Soc Neuroscience*. https://doi.org/10.1523/JNEUROSCI.6175-11.2012

Alme, C. B., Miao, C., Jezek, K., Treves, A., Moser, E. I., & Moser, M. B. (2014). Place cells in the hippocampus: Eleven maps for eleven rooms. *Proceedings of the National Academy of Sciences of the United States of America*, *111*(52), 18428–18435. https://doi.org/10.1073/pnas.1421056111

Amaral, D. G. (1993). Emerging principles of intrinsic hippocampal organization. *Current Opinion in Neurobiology*, *3*(2), 225–229. https://doi.org/10.1016/0959-4388(93)90214-J

Amaral, D. G., Ishizuka, N., & Claiborne, B. (1990). Chapter Neurons, numbers and the hippocampal network. *Progress in Brain Research*, *83*(C), 1–11. https://doi.org/10.1016/S0079-6123(08)61237-6

Amit, D. J. (1989). *Modeling brain function: The world of attractor neural networks*. Cambridge

University Press.

Andersen, P., Bliss, T. V. P., & Skrede, K. K. (1971). Lamellar organization of hippocampal excitatory pathways. *Experimental Brain Research*, *13*(2), 222–238. https://doi.org/10.1007/BF00234087

Anderson, M. I., & Jeffery, K. J. (2003). Behavioral/Systems/Cognitive Heterogeneous Modulation of Place Cell Firing by Changes in Context. In *Soc Neuroscience*. Retrieved from https://www.jneurosci.org/content/23/26/8827.short

Aronov, D., Nevers, R., & Tank, D. W. (2017). Mapping of a non-spatial dimension by the hippocampal-entorhinal circuit. *Nature*, *543*(7647), 719–722. https://doi.org/10.1038/nature21692

Aronov, D., & Tank, D. W. (2014). Engagement of Neural Circuits Underlying 2D Spatial Navigation in a Rodent Virtual Reality System. *Neuron*. https://doi.org/10.1016/j.neuron.2014.08.042

Bains, J., Longacher, J. (1999). Reciprocal interactions between CA3 network activity and strength of recurrent collateral synapses. *Nature.Com*.

Baker, P. F., Hodgkin, A. L., & Ridgway, E. B. (1971). Depolarization and calcium entry in squid giant axons. *The Journal of Physiology*. https://doi.org/10.1113/jphysiol.1971.sp009641

Barrionuevo, G., & Brown, T. H. (1983). Associative long-term potentiation in hippocampal slices. *Proceedings of the National Academy of Sciences of the United States of America*, *80*(23 I), 7347–7351. https://doi.org/10.1073/pnas.80.23.7347

Barry, C., Hayman, R., Burgess, N., & Jeffery, K. J. (2007). Experience-dependent rescaling of entorhinal grids. *Nature neuroscience*, *10*(6), 682-684.

Barry, C., Lever, C., Hayman, R. (2006.). The boundary vector cell model of place cell firing and spatial memory. *Ncbi.Nlm.Nih.Gov*. Retrieved from https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2677716/

Bear, M. F., & Abraham, W. C. (1996).Long term depression in the hippocampus. In *Annual Reviews Inc. All rights reserved* (Vol. 19). Retrieved from www.annualreviews.org

Behrens, T. E. J., Muller, T. H., Whittington, J. C. R., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018, October 24). What Is a Cognitive Map? Organizing Knowledge for Flexible Behavior. *Neuron*, Vol. 100, pp. 490–509.

https://doi.org/10.1016/j.neuron.2018.10.002

Bhatia, A., Moza, S., & Bhalla, U. S. (2019). Precise excitation-inhibition balance controls gain and timing in the hippocampus. *ELife*, *8*. https://doi.org/10.7554/eLife.43415

Bi, G.-Q., & Poo, M.-M. (1998). Synaptic Modifications in Cultured Hippocampal Neurons: Dependence on Spike Timing, Synaptic Strength, and Postsynaptic Cell Type. In *Soc Neuroscience*. Retrieved from https://www.jneurosci.org/content/18/24/10464.short

Bieri, K. W., Bobbitt, K. N., & Colgin, L. L. (2014). Slow and Fast Gamma Rhythms Coordinate Different Spatial Coding Modes in Hippocampal Place Cells. *Neuron*, *82*(3), 670–681. https://doi.org/10.1016/j.neuron.2014.03.013

Bingman, V. P., Gagliardo, A., & Loal~, P. (1996). Hippocampal participation in the sun compass orientation of phase-shifted homing pigeons. In *J Comp Physiol A* (Vol. 179). Springer-Verlag.

Bittner, Katie, Health, A., & Grienberger, C. Magee, J. (2018). Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science.Sciencemag.Org*. https://doi.org/10.1126/science.aan3846

Bittner, KC, Grienberger, C., Vaidya, S., Magee., J. (2015). Conjunctive input processing drives feature selectivity in hippocampal CA1 neurons. *Nature.Com*. R

Bliss, T. V.P., Collingridge, G. L., & Morris, R. G. M. (2014). Synaptic plasticity in health and disease: Introduction and overview. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *369*(1633). https://doi.org/10.1098/rstb.2013.0129

Bliss, T V P, & Collingridge, G. L. (1993). A synaptic model of memory: long-term potentiation in the hippocampus. In *nature.com*. Retrieved from https://idp.nature.com/authorize/casa?redirect_uri=https://www.nature.com/articles/36 1031a0&casa_token=WgHrjC4b9DEAAAAA:S1YWyCaDWtWEpIkGwGJNXjaSBXydn-BN0l8475ihJvhRGJATs8X1uJKjcJdWveBsrGA7DVetMNupHHf2_w

Bliss, Tim V.P., & Collingridge, G. L. (2013). Expression of NMDA receptor-dependent LTP in the hippocampus: bridging the divide. *Molecular Brain*, Vol. 6. https://doi.org/10.1186/1756-6606-6-5

Boccara, C. N., Sargolini, F., Hult Thoresen, V., Solstad, T., Witter, M. P., Moser, E. I., & Moser, M.-B. (2010). Grid cells in pre- and parasubiculum. *Nature Publishing Group*.

https://doi.org/10.1038/nn.2602

Boettiger, C., (2001). Developmentally restricted synaptic plasticity in a songbird nucleus required for song learning. *Elsevier*. Retrieved from https://www.sciencedirect.com/science/article/pii/S0896627301004032

Bostock, E., Muller, R. U., & Kubie, J. L. (1991). Experience-dependent modifications of hippocampal place cell firing. *Hippocampus*, *1*(2), 193–205. https://doi.org/10.1002/hipo.450010207

Brun, V. H., Solstad, T., Kjelstrup, K. B., Fyhn, M., Witter, M. P., Moser, E. I., & Moser, M. B. (2008). Progressive increase in grid scale from dorsal to ventral medial entorhinal cortex. *Hippocampus*, *18*(12), 1200–1212. https://doi.org/10.1002/hipo.20504

Buchanan, K. A., Petrovic, M. M., Chamberlain, S. E. L., Marrion, N. V., & Mellor, J. R. (2010). Facilitation of Long-Term Potentiation by Muscarinic M1 Receptors Is Mediated by Inhibition of SK Channels. *Neuron*, *68*(5), 948–963. https://doi.org/10.1016/j.neuron.2010.11.018

Bush, D., Barry, C., & Burgess, N. (2014, March 1). What do grid cells contribute to place cell firing? *Trends in Neurosciences*, Vol. 37, pp. 136–145. https://doi.org/10.1016/j.tins.2013.12.003

Buzsáki, G. (1989). Two-stage model of memory trace formation: a role for "noisy" brain states. *Neuroscience*, *31*(3), 551–570. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/2687720

Buzsaki, G., Horvath, Z., & Urioste, R. (1992). *High-Frequency Network Oscillation in the Hippocampus Author (s): György Buzsáki , Zsolt Horváth , Ronald Urioste , Jamille Hetke and Kensall Wise Published by : American Association for the Advancement of Science Stable URL : https://www.jstor.org/stab*. *256*(5059), 1025–1027.

Cacucci, F., Wills, T. J., Lever, C., Giese, K. P., & O'keefe, J. (2007). Brief Communications Experience-Dependent Increase in CA1 Place Cell Spatial Information, But Not Spatial Reproducibility, Is Dependent on the Autophosphorylation of the-Isoform of the Calcium/Calmodulin-Dependent Protein Kinase II. *Soc Neuroscience*. https://doi.org/10.1523/JNEUROSCI.1704-07.2007

Carpenter, F., & Barry, C. (2016, March 1). Distorted Grids as a Spatial Label and Metric. *Trends*

in Cognitive Sciences, Vol. 20, pp. 164–167. https://doi.org/10.1016/j.tics.2015.12.004

Carr, M. F., Karlsson, M. P., & Frank, L. M. (2012). Transient Slow Gamma Synchrony Underlies Hippocampal Memory Replay. *Neuron*, *75*(4), 700–713. https://doi.org/10.1016/j.neuron.2012.06.014

Carrillo-Reid, L., Yang, W., Bando, Y., Peterka, D. S., & Yuste, R. (2016a). Imprinting and recalling cortical ensembles. *Science*. https://doi.org/10.1126/science.aaf7560

Carrillo-Reid, L., Yang, W., Bando, Y., Peterka, D. S., & Yuste, R. (2016b). Imprinting and recalling cortical ensembles. *Science*, *353*(6300), 691–694. https://doi.org/10.1126/science.aaf7560

Carrillo-Reid, Luis, Han, S., Yang, W., Akrouh, A., & Yuste, R. (2019). Controlling Visually Guided Behavior by Holographic Recalling of Cortical Ensembles. *Cell*, *178*(2), 447-457.e5. https://doi.org/10.1016/j.cell.2019.05.045

Carrillo-Reid, Luis, & Yuste, R. (2020, October 1). Playing the piano with the cortex: role of neuronal ensembles and pattern completion in perception and behavior. *Current Opinion in Neurobiology*, Vol. 64, pp. 89–95. https://doi.org/10.1016/j.conb.2020.03.014

Caruana, D. A., Alexander, G. M., & Dudek, S. M. (2016). New insights into the regulation of synaptic plasticity from an unexpected place: Hippocampal area CA2. *Learnmem.Cshlp.Org*. https://doi.org/10.1101/lm.025304.111

Chen, T. W., Wardill, T. J., Sun, Y., Pulver, S. R., Renninger, S. L., Baohan, A., … Kim, D. S. (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature*. https://doi.org/10.1038/nature12354

Cherubini, E., & Miles, R. (2015, February 5). The CA3 region of the hippocampus: How is it? What is it for? How does it do it? *Frontiers in Cellular Neuroscience*, Vol. 9. https://doi.org/10.3389/fncel.2015.00019

Chevaleyre, V.,(2010). Strong CA2 pyramidal neuron synapses define a powerful disynaptic cortico-hippocampal loop. *Elsevier*. Retrieved from https://www.sciencedirect.com/science/article/pii/S0896627310002849

Claiborne, B. J., Amaral, D. G., & Cowan, W. M. (1986). A light and electron microscopic analysis of the mossy fibers of the rat dentate gyrus. *Journal of Comparative Neurology*, *246*(4), 435–458. https://doi.org/10.1002/cne.902460403

Colgin, L. L., Denninger, T., Fyhn, M., Hafting, T., Bonnevie, T., Jensen, O., … Moser, E. I. (2009). Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature*, *462*(7271), 353–357. https://doi.org/10.1038/nature08573

Collingridge, G. L., Kehl, S. J., & McLennan, H. (1983). Excitatory amino acids in synaptic transmission in the Schaffer collateral-commissural pathway of the rat hippocampus. *The Journal of Physiology*, *334*(1), 33–46. https://doi.org/10.1113/jphysiol.1983.sp014478

Constantinescu, A. O., O'Reilly, J. X., & Behrens, T. E. J. (2016). Organizing conceptual knowledge in humans with a gridlike code. *Science*, *352*(6292), 1464–1468. https://doi.org/10.1126/science.aaf0941

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

Csicsvari, J., Hirase, H., Czurkó, A., Mamiya, A., & Buzsáki, G. (1999). Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving rat. *Journal of Neuroscience*, *19*(1), 274–287. https://doi.org/10.1523/jneurosci.19-01-00274.1999

Czurkó, A., Hirase, H., Csicsvari, J., & Buzsáki, G. (1999). Sustained activation of hippocampal pyramidal cells by 'space clamping'in a running wheel. *European Journal of Neuroscience*, *11*(1), 344-352.

Danielson, N. B. B., Kaifosh, P., Zaremba, J. D. D., Lovett-Barron, M., Tsai, J., Denny, C. A. A., … Kheirbek, M. A. A. (2016). Distinct Contribution of Adult-Born Hippocampal Granule Cells to Context Encoding. *Neuron*, *90*(1), 101–112. https://doi.org/10.1016/j.neuron.2016.02.019

Danjo, T., Toyoizumi, T., & Fujisawa, S. (2018). Spatial representations of self and other in the hippocampus. *Science*, *359*(6372), 213–218. https://doi.org/10.1126/science.aao3898

Dayan, P., & Yu, A. J. (2006). Phasic norepinephrine: A neural interrupt signal for unexpected events. *Network: Computation in Neural Systems*, *17*(4), 335–350. https://doi.org/10.1080/09548980601004024

De Almeida, L., Idiart, M., & Lisman, J. E. (2007). Memory retrieval time and memory capacity of the CA3 network: Role of gamma frequency oscillations. *Learning and Memory*, *14*(11), 795–806. https://doi.org/10.1101/lm.730207

Debanne, D., Gähwiler, B. H., & Thompson, S. M. (1998). Long-term synaptic plasticity between

pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. *Journal of Physiology*, *507*(1), 237–247. https://doi.org/10.1111/j.1469-7793.1998.237bu.x

Deco, G., & Rolls, E. T. (2003). Attention and working memory: a dynamical model of neuronal activity in the prefrontal cortex. *European Journal of Neuroscience*, *18*(8), 2374–2390. https://doi.org/10.1046/j.1460-9568.2003.02956.x

Delatour, B., & Witter, M. P. (2002). Projections from the parahippocampal region to the prefrontal cortex in the rat: Evidence of multiple pathways. *European Journal of Neuroscience*, *15*(8), 1400–1407. https://doi.org/10.1046/j.1460-9568.2002.01973.x

Dennis, S. H., Pasqui, F., Colvin, E. M., Sanger, H., Mogg, A. J., Felder, C. C., … Mellor, J. R. (2016). Activation of Muscarinic M1 Acetylcholine Receptors Induces Long-Term Potentiation in the Hippocampus. *Cerebral Cortex*, *26*(1), 414–426. https://doi.org/10.1093/cercor/bhv227

Denny, C., Kheirbek, M., Alba, E., (2014). Hippocampal memory traces are differentially modulated by experience, time, and adult neurogenesis. *Elsevier*.

Dettner, A., Münzberg, S. (2016). Temporal pairwise spike correlations fully capture single-neuron information. *Nature.Com*.

Diamantaki, M., Coletta, S., Nasr, K., Zeraati, R., Carandini., M. (2018). Manipulating hippocampal place cell activity by single-cell stimulation in freely moving mice. *Elsevier*.

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

Do, V. H., Martinez, C. O., Martinez, J. L., & Derrick, B. E. (2002). Long-term potentiation in direct perforant path projections to the hippocampal CA3 region in vivo. *Journal of Neurophysiology*, *87*(2), 669–678. https://doi.org/10.1152/jn.00938.2000

Dolleman-Van Der Weel, M. J., Griffin, A. L., Ito, H. T., Shapiro, M. L., Witter, M. P., Vertes, R. P., & Allen, T. A. (2010). The nucleus reuniens of the thalamus sits at the nexus of a hippocampus and medial prefrontal cortex circuit enabling memory and behavior. *Learnmem.Cshlp.Org*. https://doi.org/10.1101/lm.048389.118

Dombeck, D. A., Harvey, C. D., Tian, L., Looger, L. L., & Tank, D. W. (2010). Functional imaging of hippocampal place cells at cellular resolution during virtual navigation. *Nature Neuroscience*, *13*(11), 1433–1440. https://doi.org/10.1038/nn.2648

Donal Hebb. (1949). *The organization of behavior: A neuropsychological approach*. John Wiley & Sons.

Dong, C., & Sheffield, M. E. J. (2021). Title: Distinct place cell dynamics in CA1 and CA3 encode experience in new contexts 1. *Biorxiv.Org*. https://doi.org/10.1101/2020.09.10.292177

Dudek, S. M., & Bear, M. F. (1993). Bidirectional Long-Term Modification of Synaptic Effectiveness in the Adult and Immature Hippocampus. In *The Journal of Neuroscience* (Vol. 13). Retrieved from https://www.jneurosci.org/content/13/7/2910.short

Ebbinghaus, H. (1885). *Memory. A Contribution to Experimental Psychology*. New York: Teachers College, Columbia University.

Eichenbaum, H. (2013). What H.M. taught us. *Journal of Cognitive Neuroscience*, *25*, 14–21.

Eichenbaum, Howard. (2017, April 7). The role of the hippocampus in navigation is memory. *Journal of Neurophysiology*, Vol. 117, pp. 1785–1796. https://doi.org/10.1152/jn.00005.2017

Ekstrom, A. D., Kahana, M. J., Caplan, J. B., Fields, T. A., Isham, E. A., Newman, E. L., & Fried, I. (2003). Cellular networks underlying human spatial navigation. *Nature*, *425*(6954), 184–187. https://doi.org/10.1038/nature01964

Emiliani, E., Cohen, A., Deisseroth, D., & Hausser, M. (2012). Optical interrogation of neural circuits. *Neuromethods*, *67*(41), 3–20. https://doi.org/10.1007/7657_2011_17

Foster, D. J., & Wilson, M. A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, *440*(7084), 680–683. https://doi.org/10.1038/nature04587

Fyhn, M., Hafting, T., Treves, A., Moser, M.-B., & Moser, E. I. (2007a). Hippocampal remapping and grid realignment in entorhinal cortex. *Nature*, *446*(7132), 190–194. https://doi.org/10.1038/nature05601

Fyhn, M., Hafting, T., Treves, A., Moser, M. B., & Moser, E. I. (2007b). Hippocampal remapping and grid realignment in entorhinal cortex. *Nature*, *446*(7132), 190–194. https://doi.org/10.1038/nature05601

Gauthier, J., Neuron, D. T.-, & 2018, undefined. (2018). A dedicated population for reward coding in the hippocampus. *Elsevier*. Retrieved from

https://www.sciencedirect.com/science/article/pii/S0896627318304768

Girardeau, G., Benchenane, K., Wiener, S. I., Buzsáki, G., & Zugaro, M. B. (2009). Selective suppression of hippocampal ripples impairs spatial memory. *Nature Neuroscience*, *12*(10), 1222–1223. https://doi.org/10.1038/nn.2384

Gonzales, R. B., Galvan, C. J. D. E. L., Rangel, Y. M., & Claiborne, B. J. (2001). *Distribution of Thorny Excrescences on CA3 Pyramidal Neurons in the Rat Hippocampus*. *368*(October 2000), 357–368.

Goymer, P. (2007). Colour vision for mice. *Nature Reviews Genetics*, *8*(5), 324–325. https://doi.org/10.1038/nrg2106

Guzman, S. J., Schlogl, A., Frotscher, M., & Jonas, P. (2016). Synaptic mechanisms of pattern completion in the hippocampal CA3 network. *Science*, *353*(6304), 1117–1123. https://doi.org/10.1126/science.aaf1836

Hafting, T., Fyhn, M., Molden, S., Moser, M. B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, *436*(7052), 801–806. https://doi.org/10.1038/nature03721

Hainmueller, T., & Bartos, M. (2018). Parallel emergence of stable and dynamic memory engrams in the hippocampus. *Nature*, *558*(7709), 292–296. https://doi.org/10.1038/s41586-018-0191-2

Hainmueller, T., & Bartos, M. (2020, March 1). Dentate gyrus circuits for encoding, retrieval and discrimination of episodic memories. *Nature Reviews Neuroscience*, Vol. 21, pp. 153–168. https://doi.org/10.1038/s41583-019-0260-z

Halliday, D. M. (2000). Weak, stochastic temporal correlation of large-scale synaptic input is a major determinant of neuronal bandwidth. *Neural Computation*, *12*(3), 693–707. https://doi.org/10.1162/089976600300015754

Hardcastle, K., Maheswaranathan, N., Neuron, S. G.-,. (2017). A multiplexed, heterogeneous, and adaptive code for navigation in medial entorhinal cortex. *Elsevier*.

Hasselmo, M. E. (1995). *BEHAVIOURAL Review article Neuromodulation and cortical function : modeling the physiological basis of behavior*. *67*.

Hasselmo, M. E. (2006, December). The role of acetylcholine in learning and memory. *Current*

Opinion in Neurobiology, Vol. 16, pp. 710–715. https://doi.org/10.1016/j.conb.2006.09.002

Helias, M., Tetzlaff, T., & Diesmann, M. (2014). The Correlation Structure of Local Neuronal Networks Intrinsically Results from Recurrent Dynamics. *PLoS Computational Biology*, *10*(1). https://doi.org/10.1371/journal.pcbi.1003428

Henze, D. A., Borhegyi, Z., Csicsvari, J., Mamiya, A., Harris, K. D., & Buzsaki, G. (2000). Intracellular features predicted by extracellular recordings in the hippocampus in vivo. *Journal of neurophysiology*, *84*(1), 390-400.

Henze, D. A., Wittner, L., & Buzsáki, G. (2002). Single granule cells reliably discharge targets in the hippocampal CA3 network in vivo. *Nature Neuroscience*, *5*(8), 790–795. https://doi.org/10.1038/nn887

Hölscher, C., Schnee, A., …(2005). Rats are able to navigate in virtual environments. *Jeb.Biologists.Org*. https://doi.org/10.1242/jeb.01371

Hopfield, J. J. (1982a). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the United States of America*, *79*(8), 2554–2558. https://doi.org/10.1073/pnas.79.8.2554

Hopfield, J. J. (1982b). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, *79*(8), 2554–2558. https://doi.org/10.1073/PNAS.79.8.2554

Hopfield, J. J. (1984). Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences of the United States of America*, *81*(10 I), 3088–3092. https://doi.org/10.1073/pnas.81.10.3088

Høydal, Ø., Skytøen, E., Andersson, S (2019). Object-vector coding in the medial entorhinal cortex. *Nature.Com*.

Hunt, D. L., Linaro, D., Si, B., Romani, S., & Spruston, N. (2018). A novel pyramidal cell type promotes sharp-wave synchronization in the hippocampus. *Nature Neuroscience*, *21*(July). https://doi.org/10.1038/s41593-018-0172-7

Hunt, D. L., Puente, N., Grandes, P., & Castillo, P. E. (2013). Bidirectional NMDA receptor plasticity controls CA3 output and heterosynaptic metaplasticity. *Nature Neuroscience*, *16*(8), 1049–1059. https://doi.org/10.1038/nn.3461

Ishizuka, N., Cowan, W. M., & Amaral, D. G. (1995). A quantitative analysis of the dendritic organization of pyramidal cells in the rat hippocampus. *Journal of Comparative Neurology*, *362*(1), 17–45. https://doi.org/10.1002/cne.903620103

Ishizuka, N., Weber, J., & Amaral, D. G. (1990). Organization of intrahippocampal projections originating from CA3 pyramidal cells in the rat. *Journal of Comparative Neurology*, *295*(4), 580–623. https://doi.org/10.1002/cne.902950407

Jacob, P. Y., Poucet, B., Liberge, M., Save, E., & Sargolini, F. (2014). Vestibular control of entorhinal cortex activity in spatial navigation. *Frontiers in Integrative Neuroscience*, *8*(JUNE), 1–9. https://doi.org/10.3389/fnint.2014.00038

Jeffery, K. J. (2018, February 1). The Hippocampus: From Memory, to Map, to Memory Map. *Trends in Neurosciences*, Vol. 41, pp. 64–66. https://doi.org/10.1016/j.tins.2017.12.004

Jung, M. W., & McNaughton, B. L. (1993). Spatial selectivity of unit activity in the hippocampal granular layer. *Hippocampus*, *3*(2), 165–182. https://doi.org/10.1002/hipo.450030209

Karlsson, M., neuroscience, L. F.-N., & 2009, undefined. (2009). Awake replay of remote experiences in the hippocampus. *Nature.Com*. L0S4FtPECAQDN3ekg-9q_ZUyf4O-YDDI7A-bVFrq_jlOj4_0bSB82nb98bvxfL6UA_IffQ

Kemere, C., Carr, M. F., Karlsson, M. P., & Frank, L. M. (2013). Rapid and continuous modulation of hippocampal network state during exploration of new places. *PloS One*, *8*(9). https://doi.org/10.1371/journal.pone.0073114

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

Kentros, C., Hargreaves, E., Hawkins, R. D., Kandel, E. R., Shapiro, M., & Muller, R. V. (1998). Abolition of long-term stability of new hippocampal place cell maps by NMDA receptor blockade. *Science*, *280*(5372), 2121-2126.

Kesner, R. P. (2013). A process analysis of the CA3 subregion of the hippocampus. *Frontiers in Cellular Neuroscience*, *7*(MAY). https://doi.org/10.3389/fncel.2013.00078

Kesner, R. P., & Rolls, E. T. (2015, January 1). A computational theory of hippocampal function, and tests of the theory: New developments. *Neuroscience and Biobehavioral Reviews*, Vol.

48, pp. 92–147. https://doi.org/10.1016/j.neubiorev.2014.11.009

Kheirbek, M., Drew, L., Burghardt, N., Neuron, (2013). Differential control of learning and anxiety along the dorsoventral axis of the dentate gyrus. *Elsevier*. Retrieved from https://www.sciencedirect.com/science/article/pii/S0896627313000469

Kirkwood, A., & Bear, M. F. (1994). Hebbian Synapses in Visual Cortex. In *The Journal of Neuroscience* (Vol. 14). Retrieved from https://www.jneurosci.org/content/14/3/1634.short

Kjelstrup, K. B., Solstad, T., Brun, V. H., Hafting, T., Leutgeb, S., Witter, M. P., … Moser, M. B. (2008). Finite scale of spatial representation in the hippocampus. *Science*, *321*(5885), 140–143. https://doi.org/10.1126/science.1157086

Knierim, J. J., & Neunuebel, J. P. (2016). Tracking the flow of hippocampal computation: Pattern separation, pattern completion, and attractor dynamics. *Neurobiology of Learning and Memory*, *129*, 38–49. https://doi.org/10.1016/j.nlm.2015.10.008

Knierim, J. J., & Zhang, K. (2012). Attractor Dynamics of Spatially Correlated Neural Activity in the Limbic System. *Annual Review of Neuroscience*, *35*(1), 267–285. https://doi.org/10.1146/annurev-neuro-062111-150351

Knight, H. M., Walker, R., James, R., Porteous, D. J., Muir, W. J., Blackwood, D. H. R., & Pickard, B. S. (2012). GRIK4/KA1 protein expression in human brain and correlation with bipolar disorder risk variant status. *American Journal of Medical Genetics Part B: Neuropsychiatric Genetics*, *159B*(1), 21–29. https://doi.org/10.1002/ajmg.b.31248

Knöpfel, T., & Song, C. (2019, December 1). Optical voltage imaging in neurons: moving from technology development to practical tool. *Nature Reviews Neuroscience*, Vol. 20, pp. 719–727. https://doi.org/10.1038/s41583-019-0231-4

Kondo, H., Lavenex, P., & Amaral, D. G. (2009). Intrinsic connections of the macaque monkey hippocampal formation: II. CA3 connections. *The Journal of Comparative Neurology*, *515*(3), NA-NA. https://doi.org/10.1002/cne.22056

Krupic, J., Bauza, M., Burton, S., & O'Keefe, J. (2016). Framing the grid: effect of boundaries on grid cells and navigation. *Journal of Physiology*, *594*(22), 6489–6499. https://doi.org/10.1113/JP270607

Krupic, J., Burgess, N., & O'Keefe, J. (2012). Neural representations of location composed of spatially periodic bands. *Science*, *337*(6096), 853–857. https://doi.org/10.1126/science.1222403

Lashley, K. S. (1950). In search of the engram.itle. *Symposia of the Society for Experimental Biology*.

Lashley, K. S. . (1929). *Brain mechanisms and intelligence* (U. of C. Press, Ed.).

Lassalle, J.-M., Bataille, T., & Halley, H. (2000). Reversible Inactivation of the Hippocampal Mossy Fiber Synapses in Mice Impairs Spatial Learning, but neither Consolidation nor Memory Retrieval, in the Morris Navigation Task. *Neurobiology of Learning and Memory*, *73*, 243–257. https://doi.org/10.1006/nlme.1999.3931

Latuske, P., Kornienko, O., Kohler, L., & Allen, K. (2018). Hippocampal Remapping and Its Entorhinal Origin. *Frontiers in Behavioral Neuroscience*, *11*, 253. https://doi.org/10.3389/fnbeh.2017.00253

Lauer, M., & Senitz, D. (2006). *Dendritic excrescences seem to characterize hippocampal CA3 pyramidal neurons in humans*. 1469–1475. https://doi.org/10.1007/s00702-005-0428-8

Leal, S., neuroscience, M. Y.-N., & 2018, undefined. (2018). Integrating new findings and examining clinical applications of pattern separation. *Nature.Com*.

Lee, I., & Kesner, R. P. (2004). Encoding versus retrieval of spatial memory: Double dissociation between the dentate gyrus and the perforant path inputs into CA3 in the dorsal hippocampus. *Hippocampus*, *14*(1), 66–76. https://doi.org/10.1002/hipo.10167

Leutgeb, J. K., Leutgeb, S., Treves, A., Meyer, R., Barnes, C. A., McNaughton, B. L., … Moser, E. I. (2005). Progressive transformation of hippocampal neuronal representations in "morphed" environments. *Neuron*, *48*(2), 345–358. https://doi.org/10.1016/j.neuron.2005.09.007

Leutgeb, S. (2004). Distinct Ensemble Codes in Hippocampal Areas CA3 and CA1. *Science*, *305*(5688), 1295–1298. https://doi.org/10.1126/science.1100265

Lever, C., Burgess, N., Cacucci, F., Hartley, T., & Keefe, J. O. (2002). *What can the hippocampal representation of environmental geometry tell us about Hebbian learning ? 372*, 356–372. https://doi.org/10.1007/s00422-002-0360-z

Lever, C., Burton, S., Jeewajee, A., O'Keefe, J., & Burgess, N. (2009). Boundary vector cells in the

subiculum of the hippocampal formation. *Journal of Neuroscience*, *29*(31), 9771–9777. https://doi.org/10.1523/JNEUROSCI.1319-09.2009

Levy, W., Neuroscience, O. S.-, & 1983, undefined. (1983). Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Elsevier*. Retrieved from https://www.sciencedirect.com/science/article/pii/0306452283900106

Li, X.-G., Somogyi, P., Ylinen, A., & Buzsáki, G. (1994). The hippocampal CA3 network: An in vivo intracellular labeling study. *Journal of Comparative Neurology*, *339*(2), 181–208. https://doi.org/10.1002/cne.903390204

Linden, D. J., & Connor, J. A. (1995). Long-Term Synaptic Depression. *Annual Review of Neuroscience*, *18*(1), 319–357. https://doi.org/10.1146/annurev.ne.18.030195.001535

Lorente De Nó, R. (1934). Studies on the structure of the cerebral cortex. II. Continuation of the study of the ammonic system. *Journal Für Psychologie Und Neurologie*, *46*, 113–177.

Lovett-Barron, M., Turi, G. F., Kaifosh, P., Lee, P. H., Bolze, F., Sun, X. H., … Losonczy, A. (2012). Regulation of neuronal input transformations by tunable dendritic inhibition. *Nature Neuroscience*, *15*(3), 423–430. https://doi.org/10.1038/nn.3024

Madroñal, N., … J. D.-G.-N., (2016). Rapid erasure of hippocampal memory following inhibition of dentate gyrus granule cells. *Nature.Com*. Retrieved from https://www.nature.com/articles/ncomms10923

Magee, J. C., & Grienberger, C. (2020). Synaptic Plasticity Forms and Functions. *Annual Review of Neuroscience*, Vol. 43, pp. 95–117. https://doi.org/10.1146/annurev-neuro-090919-022842

Magee, J., Science, D. J.-, & 1997, undefined. (1997). A synaptically controlled, associative signal for Hebbian plasticity in hippocampal neurons. In *science.sciencemag.org*. Retrieved from https://science.sciencemag.org/content/275/5297/209.short

Maguire, E., Gadian, D., Johnsrude, I., Good, C., Ashburner, J., Frackowiak, R., & Frith, C. (2000). Navigation-related structural change in the hippocampi of taxi drivers. *Proceedings of the National Academy of Sciences*, *97*, 4398–4403.

Makay, D. (n.d.). Information theory, inference, and learning algorithms : MacKay, David J. C : Free Download, Borrow, and Streaming : Internet Archive.

Mankin, E. A., Diehl, G. W., Sparks, F. T., Leutgeb, S., & Leutgeb, J. K. (2015). Hippocampal CA2

Activity Patterns Change over Time to a Larger Extent than between Spatial Contexts. *Neuron*, *85*(1), 190–201. https://doi.org/10.1016/j.neuron.2014.12.001

Mankin, E. A., Sparks, F. T., Slayyeh, B., Sutherland, R. J., Leutgeb, S., & Leutgeb, J. K. (2012). *Neuronal code for extended time in the hippocampus*. https://doi.org/10.1073/pnas.1214107109

Marino, M. J., Rouse, S. T., Levey, A. I., Potter, L. T., & Conn, P. J. (1998). Activation of the genetically defined m1 muscarinic receptor potentiates N-methyl-D-aspartate (NMDA) receptor currents in hippocampal pyramidal cells. *Proceedings of the National Academy of Sciences of the United States of America*, *95*(19), 11465–11470.

Markram, H., Lübke, J., Frotscher, M. (1997). Henry Markram and EPSPs Regulation of Synaptic Efficacy by Coincidence of Postsynaptic APs. *Science.Sciencemag.Org*. https://doi.org/10.1126/science.275.5297.213

Markram, H., & Segal, M. (1992). Activation of protein kinase C suppresses responses to NMDA in rat CA1 hippocampal neurones. *The Journal of Physiology*, *457*(1), 491–501. https://doi.org/10.1113/jphysiol.1992.sp019389

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

Marshel, J. H., Kim, Y. S., Machado, T. A., Quirin, S., Benson, B., Kadmon, J., … Deisseroth, K. (2019). Cortical layer-specific critical dynamics triggering perception. *Science*, *365*(6453). https://doi.org/10.1126/science.aaw5202

Marshel, J. H., Kim, Y. S., Machado, T. A., Quirin, S., Kadmon, J., Raja, C., … Deisseroth, K. (2019). *HHMI Author Manuscript HHMI Author Manuscript HHMI Author Manuscript Cortical layer – specific critical dynamics triggering perception*. *365*(6453), 1–30.

McCormack, S. G., Stornetta, R. L., & Zhu, J. J. (2006). Synaptic AMPA Receptor Exchange Maintains Bidirectional Plasticity. *Neuron*, *50*(1), 75–88. https://doi.org/10.1016/j.neuron.2006.02.027

McHugh, T. J., Jones, M. W., Quinn, J. J., Balthasar, N., Coppari, R., Elmquist, J. K., … Tonegawa, S. (2007). Dentate gyrus NMDA receptors mediate rapid pattern separation in the hippocampal network. *Science*, *317*(5834), 94–99.

https://doi.org/10.1126/science.1140263

McNaughton, B. L., Barnes, C. A., Gerrard, J. L., Gothard, K., Jung, M. W., Knierim, J. J., … Weaver, K. L. (1996). Deciphering the hippocampal polyglot: the hippocampus as a path integration system. *Journal of Experimental Biology*, *199*(1).

McNaughton, B. L., & Morris, R. G. M. (1987). Hippocampal synaptic enhancement and information storage within a distribution memory system. *Trends in Neurociences*, *10*(10), 408–415.

Mercer, A., Trigg, H. L., & Thomson, A. M. (2007). *Characterization of Neurons in the CA2 Subfield of the Adult Rat Hippocampus*. https://doi.org/10.1523/JNEUROSCI.1829-07.2007

Miles, R. (1990). Synaptic excitation of inhibitory cells by single CA3 hippocampal pyramidal cells of the guinea-pig in vitro. *The Journal of Physiology*, *428*(1), 61–77. https://doi.org/10.1113/jphysiol.1990.sp018200

Miller, T. D., Chong, T. T. J., Davies, A. M. A., Johnson, M. R., Irani, S. R., Husain, M., … Rosenthal, C. R. (2020). Human hippocampal CA3 damage disrupts both recent and remote episodic memories. *ELife*, *9*. https://doi.org/10.7554/eLife.41836

Moser, M. B., Rowland, D. C., & Moser, E. I. (2015). Place cells, grid cells, and memory. *Cold Spring Harbor Perspectives in Biology*, *7*(2), a021808. https://doi.org/10.1101/cshperspect.a021808

Mulkey, R., Neuron, R. M.-, & 1992, undefined. (1992). Mechanisms underlying induction of homosynaptic long-term depression in area CA1 of the hippocampus. *Cell.Com*.

Muller, R. U., & Kubie, J. L. (1987). The Effects of Changes in the Environment on the Spatial Firing of Hippocampal Complex-Spike Cells. In *The Journal of Neuroscience* (Vol. 7).

Nakazawa, K., Quirk, M. C., Chitwood, R. A. R. A., Watanabe, M., Yeckel, M. F. M. F., Sun, L. D. L. D., … Tonegawa, S. L. B. science. 297. 2002. 21. (2002). Requirement for Hippocampal CA3 NMDA Receptors in Associative Memory Recall. *Science*, *297*(5579), 211–218. https://doi.org/10.1126/science.1071795.Requirement

Neunuebel, J. P., & Knierim, J. J. (2014). CA3 retrieves coherent representations from degraded input: Direct evidence for CA3 pattern completion and dentate gyrus pattern separation. *Neuron*, *81*(2), 416–427. https://doi.org/10.1016/j.neuron.2013.11.017

O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, *34*(1), 171–175. https://doi.org/10.1016/0006-8993(71)90358-1

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

Ólafsdóttir, H. F., Carpenter, F., & Barry, C. (2017). Task Demands Predict a Dynamic Switch in the Content of Awake Hippocampal Replay. *Neuron*, *96*(4), 925-935.e6. https://doi.org/10.1016/j.neuron.2017.09.035

Omer, D. B., Maimon, S. R., Las, L., & Ulanovsky, N. (2018). Social place-cells in the bat hippocampus. *Science*, *359*(6372), 218–224. https://doi.org/10.1126/science.aao3474

Packer, A. M., Russell, L. E., Dalgleish, H. W. P., & Häusser, M. (2014). Simultaneous all-optical manipulation and recording of neural circuit activity with cellular resolution in vivo. *Nature Methods*, *12*(2), 140–146. https://doi.org/10.1038/nmeth.3217

Papp, G., Witter, M. P., & Treves, A. (2007). *The CA3 network as a memory store for spatial representations The CA3 network as a memory store for spatial representations Introduction: A comparative perspective*. 732–744. https://doi.org/10.1101/lm.687407

Park, S., Kramer, E., Mercaldo, V., … A. R.. (2016). Neuronal allocation to a hippocampal engram. *Nature.Com*. Retrieved from https://www.nature.com/articles/npp201673

Pfeiffer, B. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature.Com*.

Plitt, M. H., & Giocomo, L. M. (2019). Experience dependent contextual codes in the hippocampus. *BioRxiv*, 864090. https://doi.org/10.1101/864090

Plitt, M. H., & Giocomo, L. M. (2021). Experience-dependent contextual codes in the hippocampus. *Nature Neuroscience*, 1–10. https://doi.org/10.1038/s41593-021-00816-6

Pouille, F., & Scanziani, M. (2001). Enforcement of temporal fidelity in pyramidal cells by somatic feed-forward inhibition. *Science*, *293*(5532), 1159–1163. https://doi.org/10.1126/science.1060342

Prince, L. Y., Tsaneva-Atanasova, K., Clopath, C., & Mellor, J. R. (2017, October 11). Acetylcholine disinhibits hippocampal circuits to enable rapid formation of overlapping memory ensembles. *BioRxiv*, p. 201699. https://doi.org/10.1101/201699

Rauš Balind, S., Manojlović-Stojanoski, M., Šošić-Jurjević, B., Selaković, V., Milošević, V., & Petković, B. (2020). An Extremely Low Frequency Magnetic Field and Global Cerebral Ischemia A¡ect Pituitary ACTH and TSH Cells in Gerbils. *Bioelectromagnetics*, *41*(2), 91–103. https://doi.org/10.1002/bem.22237

Rebola, N., Carta, M., Lanore, F., Blanchet, C., & Mulle, C. (2011). NMDA receptor-dependent metaplasticity at hippocampal mossy fiber synapses. *Nature Neuroscience*, *14*(6), 691–693. https://doi.org/10.1038/nn.2809

Rebola, N., Carta, M., & Mulle, C. (2017). Operation and plasticity of hippocampal CA3 circuits: Implications for memory encoding. *Nature Reviews Neuroscience*, *18*(4), 209–221. https://doi.org/10.1038/nrn.2017.10

Redish, A. D., Battaglia, F. P., Chawla, M. K., Ekstrom, A. D., Gerrard, J. L., Lipa, P., … Barnes, C. A. (2001). *Independence of Firing Correlates of Anatomically Proximate Hippocampal Pyramidal Cells*. Retrieved from http://www.jneurosci.org/cgi/content/full/5093

Reilly, R., Storage, J. M.-, Trade-off, . (1994). *Hippocampal Conjunctive Encoding*.

Renart, A., Rocha, J. D. La, Bartho, P. (2010). The asynchronous state in cortical circuits. *Science.Sciencemag.Org*. Retrieved from https://science.sciencemag.org/content/327/5965/587.abstract

Rickgauer, J. P., Deisseroth, K., & Tank, D. W. (2014). Simultaneous cellular-resolution optical perturbation and imaging of place cell firing fields. *Nature Neuroscience*, *17*(12), 1816–1824. https://doi.org/10.1038/nn.3866

Rolls, E. T. (2007). An attractor network in the hippocampus: Theory and neurophysiology. *Learning & Memory*, *14*(11), 714–731. https://doi.org/10.1101/lm.631207

Rolls, Edmund T. (2013a). A quantitative theory of the functions of the hippocampal CA3 network in memory. *Frontiers in Cellular Neuroscience*, *7*(June 2013). https://doi.org/10.3389/fncel.2013.00098

Rolls, Edmund T. (2013b). A quantitative theory of the functions of the hippocampal CA3 network in memory. *Frontiers in Cellular Neuroscience*, *7*(June), 1–24. https://doi.org/10.3389/fncel.2013.00098

Rolls, Edmund T. (2018, September 1). The storage and recall of memories in the hippocampo-

cortical system. *Cell and Tissue Research*, Vol. 373, pp. 577–604. https://doi.org/10.1007/s00441-017-2744-3

Rolls, Edmund T., & Kesner, R. P. (2016). Pattern separation and pattern completion in the hippocampal system. Introduction to the Special Issue. *Neurobiology of Learning and Memory*, *129*, 1–3. https://doi.org/10.1016/j.nlm.2016.02.001

Rolls, Edmund T., & Treves, A. (1990). The relative advantages of sparse versus distributed encoding for associative neuronal networks in the brain. *Network: Computation in Neural Systems*, *1*(4), 407–421. https://doi.org/10.1088/0954-898X_1_4_002

Rolls, Edmund T., Treves, A., Robertson, R. G., Georges-François, P., & Panzeri, S. (1998). Information About Spatial View in an Ensemble of Primate Hippocampal Cells. *Journal of Neurophysiology*, *79*(4), 1797–1813. https://doi.org/10.1152/jn.1998.79.4.1797

Rosenbaum, R. S., Priselac, S., Köhler, S., Black, S. E., Gao, F., Nadel, L., & Moscovitch, M. (2000). *Remote spatial memory in an amnesic person with extensive bilateral hippocampal lesions*. *3*(10).

Russell, L. E., Yang, Z., Tan, L. P., Fişek, M., Packer, A. M., Dalgleish, H. W. P., … Häusser, M. (2019). *The influence of visual cortex on perception is modulated by behavioural state*. 1–30.

Sarel, A., Finkelstein, A., Las, (2007). Vectorial representation of spatial goals in the hippocampus of bats. *Science.Sciencemag.Org*. Retrieved from https://science.sciencemag.org/content/355/6321/176.abstract

Sargolini, F., Fyhn, M., Hafting, T., … B. Moser, E. (2006). Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science.Sciencemag.Org*. Retrieved from https://science.sciencemag.org/content/312/5774/758.short

Sasaki, T., Piatti, V., Hwaun, E., … S. A.-N., & 2018, undefined. (2018). Dentate network activity is necessary for spatial working memory by supporting CA3 sharp-wave ripple generation and prospective firing of CA3 neurons. *Nature.Com*.

Schneidman, E., Berry, M., Segev, R., Nature, W. B.-, & 2006, undefined. (2006). Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature.Com*.

Schoenfeld, G., Carta, S., Rupprecht, P., Ayaz, A., & Helmchen, F. (2021). In vivo calcium imaging of CA3 pyramidal neuron populations in adult mouse hippocampus. *BioRxiv*,

2021.01.21.427642. https://doi.org/10.1101/2021.01.21.427642

Schomburg, E. W., Anastassiou, C. A., Buzsáki, G., & Koch, C. (2012). The Spiking Component of Oscillatory Extracellular Potentials in the Rat Hippocampus. *Soc Neuroscience*. https://doi.org/10.1523/JNEUROSCI.0656-12.2012

Schwartzkroin, P. A. (1975). Characteristics of CA1 neurons recorded intracellularly in the hippocampal in vitro slice preparation. *Brain Research*, *85*(3), 423–436. https://doi.org/10.1016/0006-8993(75)90817-3

SCOVILLE, W. B., & MILNER, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of Neurology, Neurosurgery, and Psychiatry*, *20*(1), 11–21. https://doi.org/10.1136/jnnp.20.1.11

Sheffield, M. E. J., Adoff, M. D., & Dombeck, D. A. (2017). Increased Prevalence of Calcium Transients across the Dendritic Arbor during Place Field Formation. *Neuron*, *96*(2), 490-504.e5. https://doi.org/10.1016/j.neuron.2017.09.029

Sheffield, M. E. J., & Dombeck, D. A. (2015). Calcium transient prevalence across the dendritic arbour predicts place field properties. *Nature*, *517*(7533), 200–204. https://doi.org/10.1038/nature13871

Shi, S. H., Hayashi, Y., Petralia, R. S., Zaman, S. H., Wenthold, R. J., Svoboda, K., & Malinow, R. (1999). Rapid spine delivery and redistribution of AMPA receptors after synaptic NMDA receptor activation. *Science*, *284*(5421), 1811–1816. https://doi.org/10.1126/science.284.5421.1811

Siegelbaum, S., neurobiology, E. K.-C. (1991). Learning-related synaptic plasticity: LTP and LTD. *Elsevier*.

Sik, Atilla, Ylinen, Aarne, Penttonen, M, Buzsáki, B. (1994). *Inhibitory CA1-CA3-Hilar Region Feedback in the Hippocampus Author (s): Attila Sik , Aarne Ylinen , Markku Penttonen and György Buzsáki Published by : American Association for the Advancement of Science Stable URL : https://www.jstor.org/stable/2884619*. *265*(5179), 1722–1724.

Sjöström, P., Turrigiano, G., (2001). Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Elsevier*. Retrieved from https://www.sciencedirect.com/science/article/pii/S0896627301005426

Skaggs, W. E., McNaughton, B. L., M., G. K., & Markus J., E. (1993). An Information-Theoretic Approach to Deciphering the Hippocampal Code. *Proceedings of the IEEE*. https://doi.org/10.1109/PROC.1977.10559

Smetters, D., Majewska, A., & Yuste, R. (1999). Detecting action potentials in neuronal populations with calcium imaging. *Methods: A Companion to Methods in Enzymology*, *18*(2), 215–221. https://doi.org/10.1006/meth.1999.0774

Smith, A. S., Williams Avram, S. K., Cymerblit-Sabba, A., Song, J., & Young, W. S. (2016). Targeted activation of the hippocampal CA2 area strongly enhances social memory. *Molecular Psychiatry*, *21*(8), 1137–1144. https://doi.org/10.1038/mp.2015.189

Sneider, J., Chrobak, J., Quirk, M., Oler, J., (2006). Differential behavioral state-dependence in the burst properties of CA3 and CA1 neurons. *Elsevier*.

Solstad, T, Boccara, C., Kropff, E. (2008). Representation of geometric borders in the entorhinal cortex. *Science.Sciencemag.Org*. Retrieved from https://science.sciencemag.org/content/322/5909/1865.short

Solstad, Trygve, Yousif, H. N., & Sejnowski, T. J. (2014). Place Cell Rate Remapping by CA3 Recurrent Collaterals. *PLoS Computational Biology*, *10*(6), e1003648. https://doi.org/10.1371/journal.pcbi.1003648

Stackman, R. W., Clark, A. S., & Taube, J. S. (2002). Hippocampal spatial representations require vestibular input. *Hippocampus*. https://doi.org/10.1002/hipo.1112

Steemers, B., Vicente-Grabovetsky, A., Barry, C., Smulders, P., Schröder, T. N., Burgess, N., & Doeller, C. F. (2016). Hippocampal Attractor Dynamics Predict Memory-Based Decision Making. *Current Biology*, *26*(13), 1750–1757. https://doi.org/10.1016/j.cub.2016.04.063

Sun, Q., Sotayo, A., Cazzulino, A. S., Snyder, A. M., Denny, C. A., & Siegelbaum, S. A. (2017). Proximodistal Heterogeneity of Hippocampal CA3 Pyramidal Neuron Intrinsic Properties, Connectivity, and Reactivation during Memory Recall. *Neuron*, *95*(3), 656-672.e3. https://doi.org/10.1016/j.neuron.2017.07.012

Tamamaki, N., & Nojyo, Y. (1995). Preservation of topography in the connections between the subiculum, field CA1, and the entorhinal cortex in rats. *The Journal of Comparative Neurology*, *353*(3), 379–390.

Taube, J. S., & Burton, H. L. (1995). Head direction cell activity monitored in a novel environment and during a cue conflict situation. *Journal of Neurophysiology*, *74*(5), 1953–1971. https://doi.org/10.1152/jn.1995.74.5.1953

Taube, Jeffrey S, Muller, R. U., & Ranck, J. B. (1990). Head-Direction Cells Recorded from the Postsubiculum in Freely Moving Rats. I. Description and Quantitative Analysis. In *The Journal of Neuroscience*.

Tavares, R., Mendelsohn, A., Grossman, Y., Neuron, C. W.-, & 2015, undefined. (n.d.). A map for social navigation in the human brain. *Elsevier*. Retrieved from https://www.sciencedirect.com/science/article/pii/S0896627315005243

Tayler, K., Tanaka, K., Reijmers, L., Biology, B. W.-C., & 2013, undefined. (n.d.). Reactivation of neural ensembles during the retrieval of recent and remote memory. *Elsevier*. Retrieved from https://www.sciencedirect.com/science/article/pii/S0960982212013280

Thompson, L. T., & Best, P. J. (1990). Long-term stability of the place-field activity of single units recorded from the dorsal hippocampus of freely behaving rats. In *Brain Research* (Vol. 509).

Toni, N., & Schinder, A. F. (2016, January 1). Maturation and functional integration of new granule cells into the adult hippocampus. *Cold Spring Harbor Perspectives in Biology*, Vol. 8. https://doi.org/10.1101/cshperspect.a018903

Treves, A., & Rolls, E. T. (1992a). Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. *Hippocampus*, *2*(2), 189–199. https://doi.org/10.1002/hipo.450020209

Treves, A., & Rolls, E. T. (1992b). *Treves et al. 1992.pdf*. *2*(2).

Treves, A., & Rolls, E. T. (1994). Computational analysis of the role of the hippocampus in memory. *Hippocampus*, *4*(3), 374–391. https://doi.org/10.1002/hipo.450040319

Tsodyks, M. (1999). Attractor neural network models of spatial maps in hippocampus. *Hippocampus*, *9*(4), 481–489. https://doi.org/10.1002/(SICI)1098-1063(1999)9:4<481::AID-HIPO14>3.0.CO;2-S

Tsukamoto, M., Yasui, T., Yamada, M. K., Nishiyama, N., Matsuki, N., & Ikegaya, Y. (2003, February 1). Mossy fibre synaptic NMDA receptors trigger non-Hebbian long-term potentiation at entorhino-CA3 synapses in the rat. *Journal of Physiology*, Vol. 546, pp. 665–675.

https://doi.org/10.1113/jphysiol.2002.033803

Ulanovsky, N., neuroscience, C. M.-N., & 2007, undefined. (n.d.). Hippocampal cellular and network activity in freely moving echolocating bats. *Nature.Com*.

Van Cauter, T., Camon, J., Alvernhe, A., Elduayen, C., Sargolini, F., & Save, E. (2013). Distinct roles of medial and lateral entorhinal cortex in spatial cognition. *Cerebral Cortex*, *23*(2), 451–459. https://doi.org/10.1093/cercor/bhs033

van de Ven, G. M., Trouche, S., McNamara, C. G., Allen, K., & Dupret, D. (2016). Hippocampal Offline Reactivation Consolidates Recently Formed Cell Assembly Patterns during Sharp Wave-Ripples. *Neuron*, *92*(5), 968–974. https://doi.org/10.1016/j.neuron.2016.10.020

Van Strien, N. M., Cappaert, N. L. M., & Witter, M. P. (2009, April). The anatomy of memory: An interactive overview of the parahippocampal- hippocampal network. *Nature Reviews Neuroscience*, Vol. 10, pp. 272–282. https://doi.org/10.1038/nrn2614

Verwer, R. W. H., Meijer, R. J., Van Uum, H. F. M., & Witter, M. P. (1997). Collateral projections from the rat hippocampal formation to the lateral and medial prefrontal cortex. *Hippocampus*, *7*(4), 397–402. https://doi.org/10.1002/(SICI)1098-1063(1997)7:4<397::AID-HIPO5>3.0.CO;2-G

Vyleta, N. P., Borges-Merjane, C., & Jonas, P. (2016). Plasticity-dependent, full detonation at hippocampal mossy fiber-CA3 pyramidal neuron synapses. *ELife*, *5*(OCTOBER2016). https://doi.org/10.7554/eLife.17977

Wagatsuma, A., Okuyama, T., Sun, C., Smith, L. M., Abe, K., & Tonegawa, S. (2017). Locus coeruleus input to hippocampal CA3 drives single-trial learning of a novel context. *Proceedings of the National Academy of Sciences of the United States of America*, *115*(2), E310–E316. https://doi.org/10.1073/pnas.1714082115

Wills, T. J. (2005). Attractor Dynamics in the Hippocampal Representation of the Local Environment. *Science*, *308*(5723), 873–876. https://doi.org/10.1126/science.1108905

Wilson, M., Science, B. M.-, & 1993, undefined. (1993). Dynamics of the hippocampal ensemble code for space. *Science.Sciencemag.Org*. Retrieved from https://science.sciencemag.org/content/261/5124/1055.short

Yartsev, M. M., & Ulanovsky, N. (2013). Representation of three-dimensional space in the

hippocampus of flying bats. *Science*, *340*(6130), 367–372. https://doi.org/10.1126/science.1235338

Yeckel, M. F., & Berger, T. W. (1990). Feedforward excitation of the hippocampus by afferents from the entorhinal cortex: Redefinition of the role of the trisynaptic pathway. In *Proc. Natl. Acad. Sci. USA* (Vol. 87). Retrieved from https://www.pnas.org/content/87/15/5832.short

Ylinen, A., Bragin, A., Nbdasdy, Z., Jand6, G., Szab6, L., Bc, ', … Buzshkil, G. (1995). Sharp Wave-Associated High-Frequency Oscillation (200 Hz) in the Intact Hippocampus: Network and Intracellular Mechanisms. In *The Journal of Neuroscience* (Vol. 15). Retrieved from https://www.jneurosci.org/content/15/1/30.short

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

Yu, G. J., Hendrickson, P. J., Song, D., & Berger, T. W. (2018). *Spatiotemporal Patterns of Granule Cell Activity Revealed by a Large-Scale, Biologically Realistic Model of the Hippocampal Dentate Gyrus*.

Yuste, R. (2015). from the neural doctrine to neural networks. *Nature Reviews Neuroscience*.

Yuste, R., & Katz, L. C. (1991). Control of postsynaptic Ca2+ influx in developing neocortex by excitatory and inhibitory neurotransmitters. *Neuron*, *6*(3), 333–344. https://doi.org/10.1016/0896-6273(91)90243-S

Zhang, F., Vierock, J., Yizhar, O., Fenno, L. E., Tsunoda, S., Kianianmomeni, A., … Deisseroth, K. (2011). The microbial opsin family of optogenetic tools. *Cell*, *147*(7), 1446–1457.

Zhang, Z., Russell, L. E., Packer, A. M., Gauld, O. M., & Häusser, M. (2018). Closed-loop all-optical interrogation of neural circuits in vivo. *Nature Methods*, *15*(12), 1037–1040.

Zhu, J. J., Qin, Y., Zhao, M., Van Aelst, L., & Malinow, R. (2002). Ras and Rap control AMPA receptor trafficking during synaptic plasticity. *Cell*, *110*(4), 443–455.