UNIVERSITY OF CALIFORNIA,
IRVINE

Neocortical Dynamics with and without a Hippocampus

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Biological Sciences

by

Scott Kilianski

Dissertation Committee:
Distinguished Professor Bruce McNaughton, Ph.D.
Professor Craig Stark, Ph.D.
Professor Michael Yassa, Ph.D.

2021

DEDICATION

To all those who are kind and curious,
let your inquisitive minds be your flashlights and open hearts be your lightbulbs.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

VR – virtual reality

RSC – retrosplenial cortex

M2 – secondary motor cortex

HPC – hippocampus

CA1 – Cornu Ammonis 1

CA3 – Cornu Ammonis 3

NMDA – n-methyl-d-aspartate

mPFC – medial prefrontal cortex

NC – neocortex

LTP - Long-term potentiation

DG – dentate gryus

EC – entorhinal cortex

SWS – slow wave sleep

SO – slow oscillation

REM – rapid eye movement

SWR – sharp wave-ripple

EEG - electroencephalogram

LFP – local field potential

ACKNOWLEDGEMENTS

I would like to thank all of my colleagues in the McNaughton lab, first and foremost, Justin Shobe and Rajat Saxena. Without their mentorship, guidance, advising, training, and assistance, I would not have completed these final experiments and presented this work to the world. They were absolutely essential to this project. More importantly than that, however, through our numerous hours-long discussions and arguments, often extending beyond our individual capacities, I gained the theoretical and conceptual experience that I felt I needed to complete a PhD. Without these debates, and without the dedication that each one of these individuals has to uncovering the honest truth in its purest form, I would not have the expertise nor the confidence that I now have. For that, and for the personal emotional support that both of them provided, I cannot thank them enough.

I also thank all other members of the McNaughton lab that were involved in my growth as a scientist, regardless of whether they worked on the projects in my dissertation or not. Rachael Hokenson provided much needed technical support, intellectual insight, and organization to much of my early work in the McNaughton lab. I wish them the greatest of successes in their scientific career (because they are so greatly deserved) and in life in genera (because they are also greatly deserved). I thank Zaneta Navratilova, Dhruba Banerjee, Aida Andujo, Elham Ghanbarian, Ivan Skelin, Aaron Wilber, Soyun Kim, and Geoffrey Shafer for all their assistance, support, and advice over the years.

I thank my dissertation committee, Drs. Bruce McNaughton, Mike Yassa, and Craig Stark. Together, the three of you shepherded me through the final stages of lengthy graduate studies. Without their support, I could not have completed this dissertation. Dr. McNaughton can synthesize such vast research fields and distill complicated subjects into a digestible form. I have learned so much from him: not just about brain physiology, but, more importantly, how to hypothesize, experiment, and teach. Dr. Stark always provides excellent feedback about experimental design and data analysis and interpretation. Those contributions were much needed. Dr. Yassa provided me with extraordinary personal guidance and support through the final stages of my dissertation. My deepest personal thanks go out to him for his help, and I have the greatest respect of him, not just as a scientist, but as a human being.

I also thank my advancement committee members, Drs. Norbert Fortin and Jeffrey Krichmar, for valuable feedback on earlier versions of my dissertation project. Other UCI professors, including Drs. Marcelo Wood and Steve Mahler, provided me with essential advising on how to complete my dissertation and effectively communicate it. The Neurobiology and Behavior staff, specifically Naima Louridi and Sally Dabiri, are true heroes and allies to all graduate students in our department. I extend my thanks to them, too.

I thank my family, Mom, Dad, and Mark, for all the love and support and confidence they provided during some of the most difficult times of my life in graduate school. They have shaped me into the person that I am today, and I am proud and honored to have them as my family forever.

My greatest thanks go to my wonderful partner, Morgan Coburn. She has been a source of inspiration and hope for me since the moment I met her. Her experiences academically and personally have changed the way I think about the world, and for that, I am profoundly grateful. She supported me in so many ways: listening to me, encouraging me, taking care of our cats, cleaning the house, feeding me. She even gave me technical pointers on lab procedures that directly affected this dissertation project. She taught me about the importance of microglia in the nervous system, which, before knowing her, I virtually ignored. She is my greatest love and my greatest inspiration. Without her, none of this would be possible.

CURRICULUM VITAE

Scott Kilianski

Education

2014 – 21021	**University of California, Irvine** Ph.D. in Neurobiology and Behavior GPA: 3.96	**Irvine, CA**
2010 – 2014	**Rutgers, The State University of New Jersey** Bachelors of Arts, School of Arts and Sciences Honors Program Double Major in Psychology (Honors) and English **GPA:** 3.96	**New Brunswick, NJ**

Research Experience

2014 – 2021	Graduate Student Researcher, Neurobiology and Behavior, UC Irvine Advisor: Dr. Bruce McNaughton, Ph.D., Distinguished Professor (brucemcn@uci.edu)
2011-2014	Undergraduate Researcher, Behavioral and Systems Neuroscience, Rutgers University Advisor: Dr. Tim Otto, Ph.D., Principle Investigator (totto@rci.rutgers.edu)

Teaching Experience

2016	Teaching Assistant for Introductory Neurobiology, University of California Irvine Advisor: Dr. Michael Leon, Ph.D., Professor (mleon@uci.edu)
2017- 2018	Course Coordinator – Neurobiology Lab , UC Irvine Advisors: Dr. Andrea Nicholas, Ph.D., Lecturer (acnichol@uci.edu) Dr. Audrey Chen Lew, Ph.D., Lecturer (lewac@uci.edu)
2016	Teaching Assistant for Neurobiology Lab, UC Irvine Advisor: Dr. Andrea Nicholas, Ph.D., Lecturer (acnichol@uci.edu)
2015 & 2019	Teaching Assistant for Introductory Biology, UC Irvine Advisor: Dr. Marcelo Wood, Ph.D., Professor (mwood@uci.edu)
Fall 2012	First-year Interest Group Seminar (FIGS) Peer Instructor, Rutgers University Advisor: Lyn R. Krueger, Director of New Student Programs (lyn.krueger@rutgers.edu)

Publication

2018 Skelin, I., **Kilianski, S.**, & McNaughton, B. L. (2018). Hippocampal coupling with cortical and subcortical structures in the context of memory consolidation. *Neurobiology of learning and memory*.

Presentations

2018 **Kilianski, S.**, Hokenson, R., Sahagian, A. & McNaughton, B.L. (2018) Testing the hippocampus' role in a rapidly-acquired, novel spatial sequence task in rats. Neuroscience 2018, San Diego, CA

2016 Skelin, I., Wilber, A.A., Montz, C.J., **Kilianski, S.**, & McNaughton, B.L. (2016). "Sequential multineuronal reactivation in rat parietal cortex". Neuroscience 2016, San Diego, CA

2014 **Kilianski, S.** & Otto, T. "Hippocampal Arc Protein in the Contextual Modulation of Fear Extinction". Poster presentation at the Rutgers Department of Psychology Honors Research Symposium: Rutgers University, Piscataway, NJ. April, 2014.

2013 Hudgins, C.*, Ben-Harush, T., **Kilianski, S.**, Otto, T. "Learning induced protein expression within subfields of dorsal and ventral hippocampus". Poster presentation at Society for Neuroscience: New Orleans, LA.

Other Experience and Professional Memberships

2019 - 2021	Graduate Student Representative for Neurobiology and Behavior (UC Irvine)
2019	Mentor, Brain Explorer Academy (UC Irvine)
2018	Activate to Captivate Presentation Workshop (UC Irvine)
2018	Improvisation for Teaching Workshop (UC Irvine)
2017	Peer Mentor, Competitive Edge Research Program (UC Irvine)
2016	Trainee, Becoming an Effective Mentor Workshop (UC Irvine)
2015	Trainee, Introductory Workshop on Computational Methods in Neuroscience (ULeth)
2015 - 2021	Graduate Member, GPS-BIOMED (funded by NIH-BEST award)
2015 - 2018	Graduate Member, Data Science Initiative (UC Irvine)
2014 -	Graduate Student Member, Society for Neuroscience

Honors

2021	Division of Teaching Excellence and Innovation Summer Fellowship (UC Irvine)
2020	John W. Haycock Award from Center for Neurobiology of Learn and Mem (UC Irvine)
2017	Edward Steinhaus Teaching Award (UC Irvine)
2017	Ralph W. Gerard Prize in the History of Neuroscience, runner-up (UC Irvine)
2014	Francisco J. Ayala Graduate Fellowship (UC Irvine)
2014	Graduated *Summa Cum Laude*, Honors in Psychology, SAS Honors Program
2014	Henry Rutgers Award for Outstanding Undergraduate Thesis (Rutgers University)
2014	Alice M. and Walter F. Philips Award for Outstanding Honors Thesis: "Hippocampal Arc Protein in the Contextual Modulation of Fear Extinction" (Rutgers University)
2012 & 2013	Cooper Summer Research Fellowship (Rutgers University)
2013	Phi Beta Kappa Academic Honors Society, Rutgers Chapter (Rutgers University)
2012 - 2014	Jerome and Lorraine Aresty Undergraduate Research Grants (Rutgers University)
2011 - 2014	Edward J. Bloustein Distinguished Scholar (New Jersey)
2011 - 2014	Rutgers School of Arts and Sciences Honors Program (Rutgers University)
2011-2014	Scarlet Award/Scholarship (Rutgers University)

ABSTRACT OF THE DISSERTATION

Neocortical Dynamics with and without a Hippocampus

by

Scott Kilianski

Doctor of Philosophy in Biological Sciences

University of California, Irvine, 2021

Distinguished Professor Bruce McNaughton, Chair

Sparse, spatially selective activity in neuronal populations is believed to be reflective of an “indexing” system that stores the patterns of activity corresponding to memories. Such activity has been observed in neurons in the hippocampus (HPC) and across many regions of the neocortex, particularly in the superficial layers. Generally, this kind of place-cell like activity in the neocortex (NC) seems to be dependent on the hippocampus, as hippocampal lesions greatly reduce the spatial selectivity of superficial neocortical neurons. Questions about these spatially selective neocortical neurons remain outstanding though, including the extent to their distributions across layers and how their activity is modulated by stimuli in different sensory domains. To better understand what shapes spatial selectivity in neocortical neurons, neuronal ensembles in the secondary motor (M2) and retrosplenial (RSC) cortices of head-fixed mice with hippocampal or sham lesions were recorded using linear electrode arrays. Mice ran through a visual virtual reality environment for reward during the recording. Behavioral results showed that sham controls slowed down before upcoming rewards, indicating they likely remembered the reward locations. Hippocampus lesioned mice did not. This

is consistent with many previous accounts of deficits in spatial memory following hippocampal lesion. Physiological results showed that in M2 neurons, in both sham and lesioned mice, there was a significant increase in activity around reward areas. In sham mice, there was greater M2 activity ramping up before the reward, whereas in lesioned mice, the increase in M2 activity occurred after reward administration. While many neurons responded around the reward sites in both a new and old VR environment, their firing rates at those rewards often changed. Firing rate differences between VR environments were observed, which may be analogous to the 'rate remapping' seen in the hippocampus when a rodent's spatial environment changes without a corresponding change in its path integration system. Neurons in lesioned mice showed greater firing rate differences, indicating that this phenomenon may be driven more by changing visual inputs when the hippocampus is lesioned. This study is the first to examine neocortical responses to different contexts (real or virtual) in hippocampal lesioned animals.

PREFACE

The objective of this dissertation was to test predictions made by one of the most widely cited theories of how the brain achieves memory: the hippocampal memory indexing theory (Teyler & DiScenna, 1986). In short, this theory asserts that the hippocampus (HPC), generates an “index code” that stores the identity of neocortical areas activated by experiential events. According to this theory, information about memory attributes is not stored in HPC; only the index to the proper neocortical areas, which actually contain the experiential information about specific memory attributes, is stored in HPC. This theory of memory has existed for decades, and although it has gone through minor updates (Teyler & Rudy, 2007; McNaughton, 2010), it remains widely influential. Despite its broad influence, some of its predictions are still untested. One of the most fundamental is the prediction that HPC somehow orchestrates reactivation of stored memory attributes in areas that are broadly distributed across the neocortex (NC) to successfully retrieve a whole memory. There is a fair amount of correlational evidence supporting this, but no causal relationship between HPC and reactivation of memories across NC has been demonstrated.

Another derivative of the theory is that, in order to activate NC areas that are so widely distributed, anatomically speaking, HPC must work through a compressed transmission scheme that necessitates decompression, and thus a “second index”, in NC (McNaughton, 2010). There is accumulating evidence that an NC “index” does indeed exist. Superficial neocortical cells seem to behave like hippocampal place cells (Mao et al., 2017; Esteves et al., 2021.), but the extent of similarity between the HPC index and the NC index is untested.

This dissertation will inform the reader in more detail about the hippocampal memory indexing theory and its biological plausibility. It will introduce the reader to brain structures and processes that perform the functions described in the theory. It will present an experiment designed to test the aforementioned predictions of the theory and outstanding corollaries. It will present several findings, including that the NC seems capable of a ‘rate remapping’ function that may be driven primarily by visual stimuli in this experimental protocol. These findings will be interpreted in the context of hippocampal memory indexing theory, and by its conclusion, this dissertation should convince the reader that the theory still stands on solid evidential ground. It will also highlight one outstanding, untested prediction of the theory that the presented experiment may be able to answer with further analysis: that HPC coordinates the reactivation of precise, broadly distributed spatiotemporal NC sequences, which corresponds to memory retrieval.

CHAPTER 1:
Background and Significance

1.1 Hippocampal Memory Indexing Theory

Ever since Brenda Milner and William Scoville first documented that human patients with bilateral medial temporal-lobe resection had severe memory loss, HPC has been the concentration of much research in the neurobiology of learning and memory field (Scoville, 1954; Scoville & Milner, 1957). After several decades of further research about the function of HPC, the mechanisms by which it could store and retrieve memories became clearer. Specifically, a strengthening of synaptic input from perforant path fibers to the dentate gyrus (DG) in HPC was discovered (Bliss & Lomo, 1973). This modification of synaptic strength was a concrete biological phenomenon that could conceivably correspond to the brain modification or 'engram' that was theorized as necessary for memory earlier in the 20th century (Semon, 1921; Konorski, 1948; Hebb, 1949; Schacter et al., 1978). The input-specificity, cooperativity, and associativity of LTP were all physiological features that made it appealing as a biological mechanism for storing memories (Andersen et al., 1977; Lynch et al., 1977; McNaughton et al., 1978;

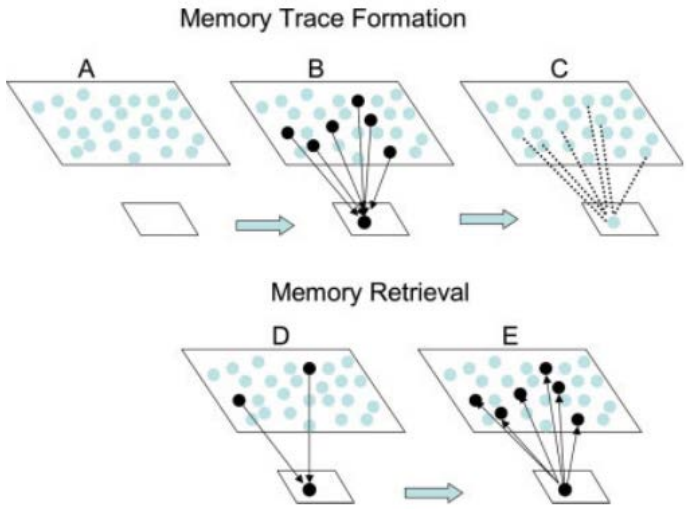

Fig 1.1. Hippocampal Memory Indexing Theory. Memory Trace Formation: A) Circles in the large top rectangle represent NC neurons. The bottom rectangle represents HPC. B) Activated NC neurons (dark) project to a unique set of HPC neurons, the "index". C) Dotted lines indicate synaptic strengthening between NC and HPC. Memory Retrieval: D) A subset of the initial NC pattern is activated. E) Because of earlier strengthening, the NC input is sufficient to activate the HPC index, which in turn activates the remaining NC neurons in the initial pattern. Adapted from Teyler & Rudy, 2007

Levy & Steward, 1979). Not only was HPC the first brain region where this was discovered, it also had the strongest and longest lasting LTP relative to other limbic forebrain pathways (Racine et al., 1983). Most evidence, behavioral and physiological, was pointing towards HPC as a structure critically involved in memory.

As insights into the function of HPC emerged, theories about exactly *how* HPC was involved in memory started to develop. The main theory with which this dissertation is concerned is the so-called “hippocampal memory indexing theory” (Teyler & DiScenna, 1986). In short, the theory states that HPC acts as an indexing system that stores the spatiotemporal patterns of activity across NC, which correspond to original sensory experiences themselves. In other words, activity across disparate areas in NC produces a sensory experience; HPC simply stores the spatiotemporal pattern of that activity and can reactivate it later. That subsequent reactivation manifests cognitively as retrieval of a memory. Teyler and DiScenna asserted that because HPC has bidirectional connectivity with associational areas of NC, it is well positioned to integrate sensory information. Furthermore, because LTP is so readily inducible in HPC, they argued, it would be an ideal system for rapidly storing new experiences. Finally, at the behavioral level, HPC lesions in humans and nonhuman animals lead to memory profound deficits in memory retrieval, leading them to conclude that HPC is somehow critical for normal mnemonic function. Numerous experiments have confirmed predictions derived from it, upholding it as a leading theory of HPC function to the present day (Teyler and Rudy, 2007).

1.2 What Makes a Good Memory Indexing System?

While Teyler and DiScenna focused on the advantageous reciprocal connectivity between HPC and NC and the striking plasticity of HPC synapses in the form of LTP, there are several other features of HPC that are ideal for performing a memory indexing function. These desirable features will be described in the following section from a theoretical and conceptual perspective. The proposed anatomical and/or physiological aspects of HPC that correspond to these features will also be explained.

Pattern Completion and Separation

Although they don't explicitly use the phrase "pattern completion", Teyler and DiScenna (1986) recognized its necessity in the indexing theory. They write, "If the hippocampus exceeds some threshold level, the remainder of the hippocampal index will be activated, and will, in turn, reactivate the entire neocortical array". Such a pattern completion function likely works in HPC through the recurrent collateral system within CA3 (Marr, 1971). The recurrent connections from pyramidal cells in CA3 back onto CA3 make up the majority input, at least in terms of number of fibers, to CA3 (Fig 1.2; Amaral, 1990; Treves & Rolls, 1992).

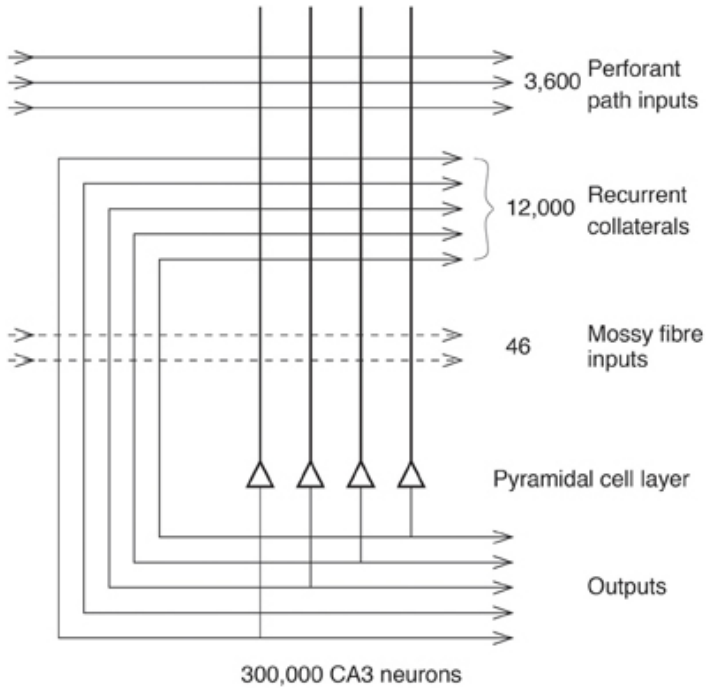

Fig 1.2. Connectivity of CA3 Pyramidal Cells. The largest number of inputs to CA3 come from CA3 itself ("Recurrent collaterals"). The perforant path and the exceptionally strong mossy fibers of the DG also input to CA3 pyramidal cells. Adapted from Treves & Rolls, 1992; Rolls, 2013.

This proposed system works by having inhibitory interneurons set the level of overall activity in CA3 and pyramidal cells having modifiable recurrent connections with one another. A simple, conceptual version of the proposed system works as follows: when an incompletely matching input arrives at CA3, the tonic inhibition effectively increases the threshold for activation, such that each CA3 pyramidal cell needs a significant amount of excitatory synaptic input to overcome the inhibition and activate. The recurrent connectivity allows the cells that are active to excite connected cells and recruit more into the active subset until a sufficient match with a previous pattern is made. CA3 outputs to CA1, subiculum, and ultimately, the neocortex where the actual information about a memory is stored (Marr, 1971; McNaughton & Morris, 1987; McNaughton & Nadel, 1990; Willshaw & Buckingham, 1990).

The corollary to pattern completion is pattern separation, the process by which two sets of overlapping input are made more distinct. It has been proposed that the perforant path, projecting from entorhinal cortex (EC) to DG, is a prime candidate for performing pattern completion (McNaughton & Nadel, 1990). The idea is that by projecting activity, or the subset of cells active (what Marr called a “codon”; Marr, 1969), divergently from one level onto a level with higher dimensionality, the distinctiveness or “orthogonality” of the pattern can be increased. In the case of HPC, the higher dimensionality comes from the fact that there is a nearly tenfold increase in the number of pyramidal cells from EC to DG. This is one feature of the EC to DG perforant path projection that makes it a candidate anatomical unit for pattern separation. In theoretical terms, an ideal indexing system must be capable of both pattern completion and separation so that it can retrieve items in a content-addressable fashion and can store

partially overlapping items sufficient uniquely to avoid interference. In anatomical reality, the recurrent connectivity of CA3 and dimensionality expansion from EC to DG are the paths through which these respective functions can be performed.

The Advantage of Sparse, Distributed Coding

Projecting patterns from a lower into a higher dimensional space, for the purposes of pattern separation, implies that, at the high-dimensional level, the system will use a relatively sparse coding scheme. This is the whole purpose of moving into the higher-dimensional space: patterns can become more orthogonal (i.e., minimally correlated) and, therefore less susceptible to interference with one another, by moving to a sparser coding scheme. In more concrete biological terms, the information being stored will be represented with a more unique set of cells at the higher dimensional level. In practice, this means that cells will have lower overall firing rates and more restricted receptive fields because they are firing to fewer sets of inputs. DG granule cells show exactly these features: they exhibit very sparse activity as a population (Chawla et al., 2005), and individual cells have smaller place fields than other HPC subregions (Jung & McNaughton, 1993; Park et al., 2011). DG granule cells also show lower firing rates during spatial navigation tasks than CA1, subiculum, and EC (Barnes et al., 2000).

Another factor affecting the orthogonality of representations in HPC is the divergence of its afferents. Although, the perforant path accounts for much of the input to DG, it is quite divergent in itself: it is estimated that each granule cell in DG receives about 6,000 inputs from EC and most are from unique EC cells (McNaughton & Nadel,

1990). As described above, orthogonality is achieved in part by moving into higher-dimensional space, but also by this divergence wherein the EC inputs needed to activate DG granule cells get “scrambled” such that initially random combinations of EC activity trigger unique, orthogonalized patterns in DG. There is further divergence in very strong, but relatively few, mossy fibers from DG to CA3. These so-called “detonator” synapses have an overwhelming influence on CA3 pyramidal cell firing (Urban et al., 2001) and can thereby further randomize the activity in CA3 and preserve the sparsity present in DG (Rolls, 1989b). Since sparse representations are more likely to be decorrelated with one another (Rolls, 2008), this EC>DG>CA3 series of projections produces enhances storage capacity while keeping interference between stored patterns minimal. There is physiological evidence for this random assignment principle, too. One compelling recent example, Rich et al. (2014), showed that the process by place fields are formed for a CA1 pyramidal cell (i.e., have a place field at a particular spot on a track/maze) is best modeled as a Poisson process. This means that the population activity in CA1 at any given location is indeed random. This is key to reducing interference between patterns and thus maximizing the capacity of the index.

Transmitting Back to the Neocortex

A necessary quality of any indexing system is to have access to the content addressed by the index. In the case of the hippocampal memory indexing theory, the content storage site is NC. The relay from an index in HPC to content in NC presents a problem though: the sparse activity across a large neuronal population in HPC, needs access to all these widely distributed NC areas. An all-to-all connectivity between the

HPC and NC would be prohibitively expensive, energetically speaking. It would require too many neurons and too many synapses between them at too long a distance. As Teyler and DiScenna presciently noted, HPC does have access to the entire NC, but it's not simply all-to-all. In some cases, this access is present directly via outputs from CA1 and subiculum to NC areas like retrosplenial (RSC; Wyss & van Groen, 1992) and medial prefrontal (mPFC; Swanson, 1981) cortices. In others, the HPC to NC exists indirectly through associational areas like RSC, mPFC, and EC, which then project to other, more primary cortical areas lower in the cortical hierarchy. Transitive connectivity between HPC and NC, by way of the hierarchical organization of the cortex, provides a partial solution to the index-content connection problem.

An additional element of the solution comes in the form of non-sparse coding in the subiculum and EC. This effectively compresses the index that was created and orthogonalized upstream in HPC and efficiently transmits it back out to NC. Then in NC, the corresponding patterns can be reactivated and memories retrieved. This is part of how the brain overcomes the biological constraint of an energetically expensive, all-to-all connectivity problem. There is physiologically support for this compression concept: subiculum and EC, the primary output structures of HPC, use denser coding schemes than DG, CA1, and CA3 as inferred by higher firing rates and lower spatial information per spike (Barnes et al., 2000). In further support of the idea that subiculum and EC function as transmitters of information rather than storage sites are their relatively low concentrations of NMDA receptors that mediate LTP (Monaghan & Cotman, 1986). This suggests they are more likely static pathways used to transmit information rather than modifiable ones used to store it.

HPC has the features of an ideal indexing system: 1) pattern completion and separation capabilities 2) sparse, orthogonal codes for partially related experiences 3) access to the locations of the stored information in NC. Several aspects of hippocampal memory indexing theory have not been answered in this introductory section. Namely: if/how is the compressed HPC index decompressed in NC? To what extent is an NC index similar or different from the one in stored in HPC? Most fundamental to the theory, is neocortical reactivation actually orchestrated by HPC during memory retrieval/reactivation? The additional introductory sections below will partly address the former two questions. The experiments described in the later chapters of this dissertation will address all three.

1.3 An Index Code Outside the Hippocampus

The fact that the non-sparse code in the subiculum and EC is used to transmit information back out to NC means that it must be decompressed somewhere in NC. Evidence from hippocampectomized patients and animals also suggests that there is an NC index, too. In these groups of people and animals, semantic memory is largely intact. One can easily explain this pattern of findings by assuming that there is an autoassociative system outside of HPC. Since this autoassociative function of pattern completion is an essential component of an indexing system, it is reasonable to think that an HPC-independent index must exist. Combined with the fact that this dense code from subiculum and EC must be “decompressed” to maintain the orthogonality of the index, it makes even more sense that there’s an NC index.

The NC, specifically the superficial layers (2/3), seem to exhibit properties that conform to the idea that it functions as an indexing system like HPC (McNaughton,

2010). Specifically, it has been shown that superficial cortical neurons across NC show more context-dependency than their deep layer counterparts. In other words, superficial layers are more sensitive to multimodal (conjunctive information) and deep layers to less complex stimuli. An exemplary case of this differential responding to individual features and conjunctions across cortical depths comes from Burke et al. (2005). In this study, the authors trained rats to turn counterclockwise or clockwise in either the same or different contexts, all while receiving a food reward. They then quantified the distribution of Arc mRNA expression (as a proxy for activity in the different conditions) in the deep and superficial layers of several cortical areas. In gustatory cortex, they found that deep layer neurons didn't differentiate between rooms or turns and were highly active, likely responding to the same food reward across all conditions. In the superficial gustatory cortex however, neurons did differentiate between both rooms and turn directions. Posterior parietal cortex findings were similar in that the superficial layer cells distinguished between the turns and rooms while the deep layer cells only responded uniquely to turn direction. This simply experimental design and corresponding set of findings supply compelling evidence that superficial layers are more responsive to conjunctive inputs, or "context ", while deep layers neurons encode less complex inputs, or "content".

Similar findings have been made by observing the *in vivo* physiology of cells across NC layers. Mao et al. (2017) recorded neural activity using 2-photon imaging of a Ca^{2+} indicator, GCaMP6m, in different parts of RSC while mice ran head-fixed in a treadmill apparatus. By imaging at different cortical depths, they found that while about 15% of cells in the superficial layers of granular or agranular RSC passed place cell

detection criteria, only 6% of cells in the deep layers did. The superficial layers in other NC areas seem to show spatial selectivity in the same behavioral protocol, too, but comparisons to deep layer were not made (Esteves et al., 2021). Interestingly, there are several earlier reports that contradict these more recent findings of place-cell-like activity in NC, especially in prefrontal cortex areas (Poucet, 1997; Jung et al., 1998; Gemmell et al., 2002; Euston & McNaughton, 2005). It is critical to emphasize, however, that these recordings primarily, if not exclusively, recorded activity from deep layers of the prefrontal cortex. This may explain why some report a lack of complex selectivity in NC that is characteristic of superficial layers and HPC and which is ideal for an index code.

There are other reports that in the prefrontal cortex, some neurons, specifically those in layer 2/3, discriminate between behavioral conditions. For example, Fujisawa et al. (2008) show that on an odor-cued match-to-sample task, layer 2/3 neurons, more than layer 5 neurons, showed differential firing depending on the upcoming choice. So, while these can't necessarily be considered "place cells" per se, they are still capable of performing an indexing function because they have orthogonal responses across behavioral conditions. Another group showed in an auditory-motor mapping task, neuronal ensembles in layer 2/3 differentiate between sensorimotor combinations, not simply stimulus, motor response, or trial outcome (Siniscalchi et al., 2016). There also exists a disparity in simple firing rate metrics between superficial and deep NC layers, with the deep layers showing significantly higher firing rates (De Kock et al., 2007; Wallace and Palmer, 2008; Sakata and Harris, 2009; Petersen and Crochet, 2013). Having sparser, highly conjunctive, coding that is orthogonal even between very similar

behavioral conditions is very similar to what is seen in HPC, and of course, advantageous for an indexing system. This makes very seductive the hypothesis that layer 2/3 acts as an “index” in NC.

1.4 Index-like Activity in Head-fixed and Nonspatial Behaviors

One important consideration for this dissertation is the degree to which sparse, orthogonal activity (e.g., sequences of HPC place cells or NC “index cells” with unique fields) can be generated in head-fixed behavioral protocols like the one used here. Such protocols are useful for recordings from large ensembles of neurons because they allow for recording arrays and microscopes that would otherwise be impossible to mount to the heads of freely moving rodents. These large recordings then allow us to answer new questions about the areas from which we are recording that would otherwise be unanswerable. For example, this dissertation is primarily concerned with looking at sparse, orthogonal codes outside of HPC. To assess the sparsity of a given NC region, it is valuable, and to a certain extent necessary, to record from large ensembles of cell so one can get an idea of how much of the population is active at any one time. Again, if one has questions about similar patterns are across different behaviors, environments, brain states, etc., it is valuable to know what large proportions of the population are doing.

Because HPC is implicated in spatial information processing and 2-dimensional spatial tuning is impaired in VR (Aghajan et al., 2015) and reliable spatial tuning in HPC relies on input from the vestibular system (Stackman et al., 2002), it is reasonable to question whether the sparse, orthogonal codes characteristic of HPC place cells can

still develop in head-fixed VR behavior protocols. Fortunately, there is ample evidence of such activity in many brain regions, including HPC and neocortex, in head-fixed behaviors. This evidence will be reviewed below.

'Index Cells'' in HPC in Head-fixed Tasks

The earliest demonstration of place cell-like activity in a head-fixed behavior is from Harvey et al. (2009). Using whole-cell patch clamping and extracellular single unit recording, they found that hippocampal neurons fired reliably at discrete locations along the linear track, often preferentially in one direction over the other, like the way in which they respond in a freely moving linear track environment (Nakazawa et al., 2003; Navratilova et al., 2012) or down individual arms of a radial maze (McNaughton & Barnes, 1983). Along with demonstrating that hippocampal neurons have circumscribed fields in a virtual environment, they also showed that these cells precess relative to the local theta rhythm like place cells in freely moving behaviors (O'Keefe & Reece, 1993). This group replicated the finding in the same kind of head-fixed virtual navigation paradigm with 2-photon imaging of a fluorescent Ca2+ indicator in the hippocampus (Dombeck et al., 2010; Rickgauer et al., 2014). This second study also found that the 2-photon imaging in the same behavioral paradigm, led to a broadening of place fields by about ~20%, which is expected because of the slower temporal dynamics of intracellular Ca2+ relative to those of membrane voltage.

Geiller et al. (2017) found 'context-modulated' and 'landmark vector' neurons preferentially in the superficial and deep CA1 cell layer, respectively, in a primarily tactile 1D navigation task with head-fixed mice. The 'context-modulated' neurons are

those that have a single field and are sensitive to the reordering of cues on the treadmill belt. These types of neurons have high spatial information because of their sparse activity patterns; they are like archetypal HPC ‘place cells’. The ‘landmark vector’ cells (originally reported in Deshmukh & Knierim, 2013) have fields at a particular distance relative to specific landmarks (in this case, textured patches attached to the belt) and fire reliably at that distance even if the greater context around those landmarks has changed. These can still have high spatial information scores because they respond only where the landmark is present. So, for instance, if only 1 landmark is present, they will only respond in one place. If two or more are present, they will fire at each of those locations, and spatial information will be progressively reduced with an increase in landmark occurrences. 2-photon Ca^{2+}-imaging in HPC and several dorsal NC regions also shows ‘place cell’-like activity in a head-fixed tactile belt behavioral protocol (Mao et al., 2017; Esteves et al., 2021). Considered altogether, these results clearly demonstrate that even in a head-fixed behavioral task with no explicit mnemonic demands, HPC neurons, especially those in deep CA1 (Geiller et al., 2017), develop discrete responsive fields reminiscent of canonical ‘place cells’ in classical 2D free foraging experiments.

Additionally, several studies have shown that well-defined place fields develop in 2D visuospatial virtual environments, too. These environments are created by either a body-restricting harness that allows for vestibular changes by turning around the azimuth (Aronov & Tank, 2014) or a pressure-sensitive head-fixation apparatus that allows for horizontal head movement (Chen et al., 2018, 2019). Spatial tuning in these VR environments is sharpened by the animal’s ability to make horizontal movements

(Aghajan et al., 2015) adding to earlier evidence that vestibular input greatly influences spatial selectivity of HPC neurons.

Index Cells in HPC in Nonspatial Tasks

As described above and elsewhere throughout this dissertation, sparse, orthogonal activity patterns of 'place cells' are active while an animal navigates through an environment, whether it be 1D, 2D, head-fixed, or open to unrestricted movement. However, all experiences important to an organism won't necessarily unfold in the spatial dimension. For example, as a student listening to a lecture in a classroom or a friend having a long, intimate discussion over coffee, it is possible to remember the content and sequence of the experience even though one was virtually immobile throughout it. If the index code was purely spatial in nature, the brain would not be able to preserve the temporal structure of extended experiences that occurred in only one location. On the contrary, there is evidence that the brain, and specifically the hippocampus, can still index parts of an experience that involves nonspatial sequences. Though there are more, I will discuss two striking examples of this below.

Sequences of HPC activity have been observed during behavior for decades. A clear example of such sequences can be observed in behaving rats when they are required to time an action or wait for an event to occur. In all these early examples of HPC sequences during waiting or delay periods of a behavioral task, the rats from which activity was recorded were moving continuously either freely in space or in-place on a treadmill or running wheel (Pastalkova et al., 2008; Gill et al., 2011; MacDonald et al., 2011; Kraus et al., 2013). To decouple the delay phase of the behavior and the

animals' movement, MacDonald et al. (2013) used a head-fixed delayed matching to sample task while recording activity in HPC. In short, they observed that HPC neurons were reliably active at brief portions of the delay phase, together generating a reliable sequence that spanned the entire delay phase. This is very similar to the way in which \ 'place cells' are sequential active all along a linear track that a rat is running. These neuronal sequences could even be the basis for the temporal organization of episodic memory. MacDonald et al. also observed that the sequences corresponding to different odors were unique, which could allow the brain to distinguish memories for experiences that happen in the same place. This bears striking resemblance to remapping in HPC that occurs when an animal enters a different spatial environment (Leutgeb et al., 2004).

Another attempt to resolve the question of whether HPC indices could be mapped onto nonspatial dimensions used a behavioral task in which rats had to hold and release a lever that increased the frequency of an auditory pure tone until it reached a target frequency (Aronov et al., 2017). More than 30% of CA1 neurons had discrete bouts of activity at specific sound frequencies even when the speed at which the sound passed through the frequency dimension was slowed down or sped up varied. Only 1.7% of all CA1 cells responded during a passive version of the frequency sweeps that didn't require any action from the rats. This is akin to the reduction of spatially selective firing when a rat is restrained and passively moved through a previously experienced environment (Foster et al., 1989). Aronov et al. also found there was no correlation between cells that had place fields in a free foraging task and

‘frequency fields’ in the auditory task indicating orthogonality between the tasks, which would be expected of a memory indexing system.

These two examples of nonspatial index-like responses in HPC demonstrate that HPC activity can be mapped onto sequences that are temporal or auditory in nature, not just spatial. In fact, this mapping of neuronal sequences or “indexes” has been advanced as a general-purpose function of HPC (Buzsaki and Tingley, 2018). This is, at the very least, an important feature of HPC function because it can potentially explain how HPC, and the brain at large, can temporally arrange memories for experiences during which the organism’s location was constant.

1.5 Review of RSC Structure and Function

Anatomy

In the rat, the retrosplenial cortex can be divided into at least four regions: the dysgranular, agranular, granular a, and granular b regions. van Groen & Wyss did many of the early studies mapping the connectivity of RSC, exploring both afferent and efferent connectivity of subdivisions of the RSC (van Groen & Wyss, 1990; van Groen & Wyss, 1992; van Groen & Wyss, 2003). Broadly speaking, the RSC receives input from the following structures: anterior cingulate cortex (ACC), subiculum, postsubiculum, visual areas 17 and 18b, dorsal thalamic regions, claustrum. Subcortical efferents from the RSC go to the caudate and anterior and lateral thalamus. Cortical efferents go to postsubiculum, parietal, visual, and anterior medial cortices. van Groen & Wyss repeatedly assert that the RSC’s position anatomically between the hippocampus and neocortex makes it a prime candidate to mediate information flow to and from those two

brain structures. A more detailed, but still brief, summary of the RSC subdivisions can be found in the following 2 paragraphs.

The granular regions are known as such because they contain a layer containing granule cells, while the dysgranular and agranular areas have relatively few and no granule cells, respectively. The granular a region is more ventral and lateral to the granular b region. Although the two do share inputs from several brain structures, they are primarily distinguished based on unique connectivity: granular b RSC is targeted by the anteroventral thalamic nucleus, dorsal subiculum, and visual cortical areas 18a and 18b (van Groen & Wyss, 1992). Granular a RSC is targeted by the laterodorsal thalamus, ventral subiculum, and presubiculum (van Groen & Wyss, 1992).

The boundary between the dysgranular and agranular regions of RSC is less clear. There is a gradual reduction of the granule cell layer as the RSC moves caudally and laterally. These regions, considered together, receive more input, relative to a and b subdivisions, from the claustrum, orbitofrontal cortex, and ventromedial thalamus (van Groen & Wyss, 2003). Importantly, the dysgranular region has dense reciprocal connectivity with the granular a region, which itself has dense projections from the subiculum, an output structure of the HPC (van Groen & Wyss, 1990). In addition, there is a long-range GABAergic projection to the granular RSC from HPC (Jinno et al., 2007; Miyashita & Rockland, 2007). The experimental work in this dissertation is concentrated on the dysgranular subdivision of the RSC as it is the most dorsal part of the RSC, and its laminae lie normal to the probes' channel orientation, maximizing our ability to record from multiple cortical layers.

A circuit not mentioned by van Groen & Wyss, but potentially very important for this dissertation, is the connection between RSC and M2 (Yamawaki et al., 2016). In this paper, Yamawaki et al. show that there exists a monosynaptic, reciprocal connection between RSC and M2. Using optogenetics, single-cell electrophysiology, and fluorescent tracers, they show an RSC → M2 projection that excites M2 pyramidal cells across most layers that themselves have a diverse range of downstream targets including RSC, thalamus, pons, and superior colliculus. They also found that M2 → RSC projections primarily targeted the dysgranular RSC and that these M2 → RSC neurons also received input from RSC themselves, thus creating a reciprocal loop. They suggest that this circuitry allows for systems involved in spatial memory (e.g., dorsal hippocampus, RSC) to directly influence diverse brain processes such as motor control, and decision making. This of particular interest to this dissertation because I am recording single unit activity in both areas and potentially capturing long-range, monosynaptically connected pairs of neurons across regions.

Function

Several studies have shown that RSC lesions impair allocentric spatial memory while preserving egocentric spatial memory. For example, Cain et al. (2006) showed that even RSC-lesioned rats that had been pre-trained to swim directly to a refuge platform in the water maze task were impaired relative to sham-lesioned rats when tested on a hidden-platform, random start location, and therefore allocentric memory-taxing, version of the task. In a different study, rats with extensive retrosplenial cortex lesions were impaired in radial arm maze, water maze, and object-in-place tasks, all of

which rely on allocentric spatial memory (Vann and Aggleton, 2002). Importantly, these same rats were not impaired on object recognition and an egocentric spatial discrimination task in which they were reinforced to always choose the left or right turn at a T-intersection. Other studies using various lesion approaches report spatial memory deficits as well (Keene and Bucci, 2008), but there are a few reports of intact spatial memory following RSC lesions. In general, it seems that most discrepancies about the necessity, or lack thereof, of RSC in spatial memory can be explained by lesion size (Vann and Aggleton, 2002). Complete lesions of the RSC unequivocally impair allocentric spatial memory while preserving other kinds of memory like egocentric discrimination and object recognition.

The physiology of the RSC is multifaceted, however, single unit recordings and 2-photon Ca^{2+}-imaging show a clear correspondence between various kinds of spatial information processing and unit activity. For example, 10% of RSC cells are 'head-direction' cells and found throughout the region, granular and dysgranular regions included. Some of these RSC 'head-direction' cells are also sensitive to non-angular movement, too (Chen et al., 1994). Similarly, another study reported that nearly 20% of cells recorded in RSC showed a response to particular combinations of location and direction (Cho & Sharp, 2001). These various directional, positional, and movement correlates of RSC activity strongly suggest a role in spatial processing which accords with the lesion and inactivation findings described in the foregoing paragraph.

Additional studies have shown spatial information correlates in the RSC. Among those listed above, there are also conjunctions of internal and external spaces (Alexander & Nitz, 2015), different levels of sub-spaces and distances traveled

(Alexander & Nitz, 2017), and unique and reliably activated ensemble of neurons across similar trajectories (Czajkowski et al., 2014). Most related to the present dissertation is the finding that in a head-fixed setup in which the mice have to walk on a treadmill belt with tactile cues attached, both RSC and HPC show similarly sparse, orthogonal representations at different positions along the belt (Mao et al., 2017). In this study, RSC 'place cells' have nearly identical distributions across the belt, field widths, and number of fields as CA1 'place cells'. The RSC population vector correlation, a measure of ensemble similarity, also decreases as a function of distance at the same rate as that of the CA1. Mao et al. (2017) had two additional findings that are highly relevant to this dissertation: 1) there was a much greater percentage of 'place cells' in the superficial RSC (15% for both granular and agranular regions) relative to the deep layer (6%) 2) a noticeable number of RSC neurons preferred firing location changed in accordance with a change in the reward position. In other words, there was an apparently global increase in RSC activity around the reward site, such that when the site changed, the increased activity went with it. Both of these findings will be discussed later in relation to the results of the experiments in this dissertation.

Other studies report reward-related activity in the RSC, too. Tabuchi et al. (2005) trained head-fixed rats to discriminate between different conditional stimuli alone and in combination while recording single units from RSC. 58% of neurons recorded responded to at least one of the rewards, either sucrose solution or intracranial self-stimulation, far more than responded to during CS presentation, CS-US delay, or non-rewarded trials combined. Furthermore, the activity of some RSC neurons was correlated with the actual lick action as well as reward type. Other studies have also

reported reward-selective neurons along with the standard spatial conjunctive responses seen in other studies (Smith et al., 2012; Vedder et al., 2017). In one case, RSC neurons also gradually shaped their responses to the light-on CS cue over several days of training (Vedder et al., 2016). The results of these studies, taken together with the others summarized above, highlight the flexibility of RSC in general. The highly flexible and conjunctive representations in RSC made it an ideal candidate to store a neocortical 'index code'.

As is described above, the connectivity of RSC and HPC suggests a bidirectional communication of some sort between the two structures. It appears that inactivation or lesion of one structure does indeed have consequences for activity in the other. For example, Cooper & Mizumori (2001) found, among other things, that 1) RSC inactivation impaired learning of a spatially-taxing radial maze task specifically with the lights on and 2) HPC place cells change field locations when RSC is temporarily inactivated. The spatial memory deficit result contributes to the abundance of evidence that RSC is involved in integrating spatial and visual information. The changing of fields of HPC place cells finding indicates that RSC informs HPC place cells about locations in which to fire.

In the other direction, hippocampal lesions greatly affect the spatial coding properties of RSC. Mao et al. (2018) showed that bilateral, but not unilateral, excitotoxic NMDA lesions of HPC dramatically reduce the spatial selectivity of superficial RSC neurons. The fraction of neurons meeting the 'place cell' criteria is dramatically reduced, and position reconstruction based on neural activity is greatly impaired. And while measures of spatial selectivity and stability of place fields increased over multiple days

of training in the control mice, they did not in HPC-lesioned mice. This effect appears to be more global, not just limited to RSC; several other neocortical regions across the dorsal neocortex including PPC, M2, M1, and SS1 show a reduction in spatial information in mice with bilateral HPC lesion (Esteves et al., 2021). This finding that HPC lesions reduce spatial information in the neocortex is clear and consistent, but it has only been shown in this particular tactile-cued treadmill apparatus with 2-photon Ca2+ imaging, and ideally it should be replicated in a variety of behavioral paradigms and with multiple recording methods.

1.6 Review of M2 Structure and Function

Anatomy

M2 (sometimes called MOs) afferents come from a variety of subcortical and cortical regions. M2 receives extensive input from many thalamic nuclei (Reep & Corwin, 1999) and a diverse array of cortical regions including visual, somatosensory, auditory, parietal, retrosplenial, and fronto-orbital areas (Reep et al., 1990; Hoover & Vertes, 2007, Yamawaki et al., 2016). M2 is often considered part of the prefrontal cortex specifically because it, along other PFC regions, receives projections from the mediodorsal thalamus and therefore is part of the so-called ‘MD-projection” network (Leonard, 1969; Oh et al., 2014). In general, rostral M2 receives more input from motor cortex areas while caudal M2 receives more long-range input from sensory cortices like V2, RSC, and auditory cortex (van Eden et al., 1992; Zhang et al., 2016).

The efferents of M2 go to areas associated with direct motor control, sensation, and higher-order associational areas. The M2 → thalamus projections include midline nuclei such as the reticular, anteromedial, and anteroventral nuclei. Furthermore, these

primarily target the rostral end of these nuclei (Sesack et al., 1989). Other subcortical efferents from M2 include the superior colliculus, pons, dorsal striatum, and through the corticospinal tract to the spinal cord (Donoghue & Wise, 1982; Stuesse & Newman, 1990; Berendse et al., 1992; Gabbott et al., 2005). The cortical targets of M2 efferents are quite various and are much denser than from those from primary motor cortex (M1). They include fronto-orbital, insular parietal, and retrosplenial areas (Reep et al., 1987; Jeong et al., 2016).

The hodology of M2, with its convergence of multimodal sensory and motor information and its direct output to motor areas, has led to the assertion that it is highly involved in at least the following processes: 1) direct modulation of motor output 2) the sensorimotor integration necessary for decision making and adaptive behavioral control and 3) ‘corollary discharge’ or ‘motor efference copy’ to reafferent sensory associational areas with information about ongoing action and decisions (Barthas & Kwan, 2017).

Function

As would be expected from the prodigious connectivity of M2, it has been implicated in many different cognitive processes. Below, I will briefly review findings from microstimulation, lesion/inactivation, and physiological recording experiments that generally point to M2 being involved in sensorimotor integration and decision making.

Although M2 borders the primary motor cortex (M1), microstimulation mapping supports the idea that M2 is a nonprimary motor area since only large currents can evoke a motor response in that area whereas in primary motor areas, relatively smaller currents can generate a movement (Donoghue & Wise, 1982). Electrical stimulation of

M2 results in the movement of a group of adjacent whiskers in unison (Neafsey et al., 1986). Otherwise, microstimulation of M2 causes combinations of whisker, eye, and head/neck movements altogether which is characteristic of an orienting response (Hall and Lindholm, 1974; Donoghue & Wise, 1982; Tennant et al., 2011).

Rats with M2 lesions or inactivations show a bias for turning to the ipsilateral side (Cowey & Bozek, 1974; Erlich et al., 2015) which suggests either an inability to turn in the contralateral direction, a neglect of that direction, or both. With additional training however, rats can overcome this neglect and turn in the contralateral direction for reward (Cowey & Bozek, 1974). Unilateral M2 lesions also lead to impaired orientation to cues originating in the contralateral sensory field, and this is true for tactile, visual, and auditory cues (Crowne & Pathria, 1982). It has also been shown that temporary unilateral inactivation of M2 leads to a strong ipsilateral bias in a free-choice task (Erlich et al., 2015). Interestingly, lesion studies also show that M2 is also necessary for 'arbitrary' sensory-motor associations, for which HPC is also necessary (Wise and Murray, 2000). Taken altogether, evidence suggests that M2 is involved in orienting responses, selection of a motor output, and creating arbitrary sensory-motor mapping rather than actual action execution or spatial attention.

Many *in vivo* electrophysiological recordings in M2 have been carried out while rodents are in a forced-choice behavioral paradigm. M2 neuronal activity in these tasks is predictive of upcoming decisions, more predictive than other frontal and striatal areas (Sul et al., 2010, 2011), suggesting that it may be involved in planning or selecting upcoming actions. Though some neurons prefer chosen contralateral actions while others prefer ipsilateral ones, the direction of upcoming action can be predicted by M2

activity up to 500ms in advance for some M2 neurons (Erlich et al., 2011; Siniscalchi et al., 2016). When examined in even closer detail, it seems that two populations of M2 neurons can be distinguished based on their activity leading up to action selection. One of these populations shows a ramping up of activity, and the other shows brief bouts of activity, circumscribed in time, during a waiting task (Murakami et al., 2014).

Finally, M2 neurons can be selective for specific actions and conditions. Neurons respond differentially to a nose-poke and a lever-press even though the reward contingencies are similar (i.e., wait the same interval to either nose-poke or lever-press for the same reward; Murakami et al., 2014). The receipt of reward alone also seems to affect neuronal activity in M2 in some cases (Kargo et al., 2007; Sul et al., 2011), although it is difficult from these studies to decouple the reward from the consummatory behavior. In conclusion, M2 is clearly active before and during decisions, but there is a heterogeneity of single unit responses within the population during these time periods.

1.7 Historical Perspective on Sleep's Effect on Memory

An important component to the foregoing theory of memory is that it is the repeated reactivation of memory traces that strengthens them, or, in other words, makes them resist to forgetting. Of course, every time a memory is actively retrieved there is an opportunity for additional strengthening. Memory reactivation, and thereby strengthening, has also been theorized to occur during sleep. Marr (1971) was the first to hypothesis that HPC could store up a day's experiences and replay back to the rest of the brain during sleep when there is no interference from external stimuli. Even before Marr, however, a strong, beneficial role for sleep in memory has been

documented (see Rasch & Born, 2013 for exhaustive review). The following two sections discuss this relationship: the first will provide semi-chronological recounting of foundational experiments uncovering the relationship between sleep and memory. The second will introduce specific sleep physiology phenomena that are uniquely related to memory.

Researchers have proposed a relationship between sleep and memory since the early 20th century (Jenkins and Dallenbach, 1924). Although there was evidence already available to suggest a sleep-memory relationship in the late 19th century, not much consideration was given to it (Ebbinghaus, 1913). Furthermore, then contemporary conceptions of memory processes, such as the "perseveration-consolidation" (Müller and Pilzecker, 1900), lead many researchers to conclude that sleep acted as a time window free from so-called "retroactive interference" that disrupted memory consolidation during waking hours (Heine, 1914; Jenkins and Dallenbach, 1924). In other words, memory benefited from sleep because the brain was in a passive state, during which no new information could interrupt the perseverative neural activity that strengthened memory *post hoc*. Eventually, the advent and widespread adoption of EEG in the early to mid-20th century, which allowed for the discovery and characterization of several different unique sleep stages (Hartman, 1967), triggered a shift in theorizing about the relationship between sleep and memory: the theory that sleep was actually an active brain state, facilitating memory reprocessing emerged and supplanted older conceptions that sleep is a purely passive state. Ultimately, this new idea of sleep as an active state inspired further investigation of the neurobiological phenomena occurring during sleep that might contribute to memory

consolidation. Those phenomena will be discussed in subsequent sections; the present section is centrally concerned with the history of research on the sleep-memory relationship. This is not an exhaustive review of the literature demonstrating a beneficial effect of sleep on memory. Rather, it is intended to serve as a historical perspective on sleep and memory research, with an emphasis on how theories of active memory processing during sleep evolved. The purpose is to contextualize and justify HPC-orchestrated replay during sleep as an important mechanism for memory consolidation.

The history of research on sleep and memory is rich and old. Many current reviews (Walker and Stickgold, 2006; Diekelmann and Born, 2010 for two examples) of the subject attribute the first study of sleep's effect on memory to John Jenkins and Karl Dallenbach (1924). However, there were several studies preceding Jenkins and Dallenbach that deserve mention even though they were not explicitly designed with the intention of investigating a relationship between sleep and memory. Rather, while studying memory retention over time, this handful of studies in the nascent field of memory research incidentally used training-testing retention intervals during which experimental subjects slept.

The first of these memory experiments comes from the pioneering memory researcher, Hermann Ebbinghaus, and is considered the first true experimental study of memory as opposed to earlier medical reports and case studies. In his monograph, Ebbinghaus very briefly considers, but eventually dismisses, the possibility that sleep may have had an effect on his memory (Ebbinghaus, 1913). While interpreting his finding that the rate of forgetting seems to decrease in the 9-24 hour training-testing interval but then increase in subsequent intervals, he writes:

> *"Such a condition is not credible, since in the case of all the other numbers the decrease in the after-effect is greatly retarded by an increase in time. It does not become credible even under the plausible assumption that night and sleep, which form a greater part of the 15 hours but a smaller part of the 24, retard considerably the decrease in the after- effect. Therefore, it must be assumed that one of these three values* [memory at 9, 24, and 48-hour delays] *is greatly affected by accidental influences"*

This is likely the first quantitative evidence suggesting a possible relationship between sleep and memory, but Ebbinghaus concluded that the data must be "greatly affected by accidental influences". In subsequent replications of Ebbinghaus's experiments, other independent researchers consistently found a similar forgetting curve including that reduced rate of forgetting between 9 and 24 hours post-learning that Ebbinghaus originally considered "not credible" (Bean, 1912; Finkenbinder, 1913; Foucault, 1913). One study even showed an enhancement of memory between an 8-hour train-test interval, which did not contain any sleep, and a 24-hour train-test interval, which did contain sleep (Radossawljewitsch, 1907). Although no deliberate effort was made to investigate sleep in these historic experiments, Ebbinghaus's original finding and those generated in replication studies are essential to the history of sleep and memory research. They serve as the inspiration for future investigations of the topic.

Another important idea influencing memory research was the perseveration-consolidation" hypothesis, which stated that neural processes must "perseverate" uninterrupted for some time after initial learning to eventually stabilize and stay in long-term memory. The hypothesis implied that memory exists in a temporarily labile state

until it is stored more permanently or “consolidated” and no longer vulnerable to interruption (Müller and Pilzecker, 1900). Indeed, their experiments revealed that memory for lists of nonsense syllables was susceptible to “retroactive interference”, a phenomenon in which additional learning of new material after initial learning impairs subsequent recall of the material. One experiment specifically intended to vary the level of retroactive interference by learning immediately before sleep or upon waking showed that sleep reduced the rate of forgetting (Heine, 1914). It was hypothesized that because sleep is a time during which the mind is free from any external input, no retroactive interference can disrupt the ongoing perseveration process whereas during waking hours, new learning can disrupt the temporarily labile memory. The results of Heine (1914) supported this hypothesis.

It was not until 1924 that this hypothesis would be tested more rigorously (Jenkins & Dallenbach, 1924). In this study, subjects were required to learn lists of nonsense syllables immediately upon waking or at night before sleep. However, while Heine used only 8-hour retention intervals, Jenkins and Dallenbach used multiple intervals: 1, 2, 4 and 8 hours after learning in both sleeping and waking conditions. Memory retention at every interval was superior in the sleep condition relative to the waking condition. While retention became progressively lower over the course of 8 hours in the waking condition, it did not change over time after the first 2 hours of sleep (i.e., retention decreased from baseline in the first 2 hours of sleep but stayed constant thereafter). The results from this foundational experiment indicate that sleep, especially the first 2 hours of sleep, have a beneficial effect on memory. They, like the researchers mentioned above, also concluded that this seemingly protective effect of sleep on

memory is due to the brain existing in an offline, passive state during sleep, in which no new input can interfere with the learning that occurred just prior to sleep.

The body of research on sleep's beneficial effect on memory continued to grow after Jenkins and Dallenbach (1924). Although results varied slightly from study to study, the general trend of a beneficial sleep effect on memory was consistent. For example, in a replication of Jenkins and Dallenbach's study, results again indicated that sleep protected against forgetting, but the difference on retention scores between the sleeping and waking conditions only became clear after 4 hours, rather than the 2 originally reported by Jenkins and Dallenbach (Van Ormer, 1932). Efforts to ensure that subjects were sufficiently awake at the various time intervals were also made by comparing initial learning at the time of retention testing and non-mnemonic cognitive abilities at the time of testing to control for any effects of a lingering "hypnagogic" state upon waking (Van Ormer, 1933). Memory retention was also tested at longer intervals (i.e., over multiple days) that included multiple sleep epochs. Results of these experiments similarly indicated that initial learning just before sleep benefitted memory relative to learning upon waking, although the effect only became clear after 3 days (Graves, 1936). Again, in all of these studies, the benefit of sleep on memory was assumed to be caused solely by the absence of retroactive interference during the unconscious state of sleep.

A groundbreaking technical development that would vastly change the study of sleep and memory was gaining momentum while all this research was going on during the late 18th and early 19th centuries: electroencephalography (EEG). Originally applied in non-human mammals by Richard Caton (Caton, 1875) and eventually used on

humans (Berger, 1929), the EEG was instrumental in categorizing sleep into multiple unique stages (Loomis et al., 1935a; 1935b). In addition to the pioneering work of Alfred Loomis's lab, the observation of a recurring sleep state during which changes in EEG activity, ocular motility, and overt bodily activity, resulted in the definition of yet another sleep stage, rapid eye movement (REM) or "paradoxical" sleep (Aserinsky and Kleitman, 1953).

Over a decade later, using experiments in which subjects had to learn two separate lists of associations to experimentally induce interference, it was demonstrated that sleep didn't simply facilitate memory, it benefited memory for the 1st more than that for the 2nd list (Ekstrand, 1967). That pattern of results is the exact opposite of what a traditional "perseveration-consolidation" hypothesis would predict: only retention of the 1st list would be impaired by retroactive interference, and the 2nd list would greatly benefit from the interference-free nature of sleep following learning. This was the first in a series of studies that would provide compelling evidence against a passive role of sleep with respect to memory.

To further test the perseveration-consolidation hypothesis and its implications about sleep benefiting memory, a new wave of studies all relied on the observation that there is an uneven distribution of REM and slow-wave sleep (SWS) across an entire night's sleep. Sleep stages are distributed such that there is much more SWS sleep in the first half of sleep and not much REM. In the second half of sleep, there is little SWS and a relatively larger quantity of REM (Hartman, 1967). This is important because if sleep protected against forgetting solely because no retroactive interference could occur, then the relative time spent in certain stages of sleep (i.e., REM vs SWS) should

not affect memory retention. Only the presence or absence of sleep and total time spent sleeping should affect memory. However, results indicated that the SWS-dominated first half of sleep was beneficial for memory while the relatively REM-heavy second half of sleep had no effect on memory relative to controls that were awake between learning and testing (Yaroush et al., 1971). A follow up study controlled for time-of-day effects (i.e., learning after at 2 am, after 3.5-4 hours of sleep as opposed to learning at 10 pm, just before sleep) and produced similar results, while showing time-of-day had no effect on initial learning or memory retention (Barrett and Ekstrand, 1972). Another study in this initial wave demonstrated the specific effect of the first half of sleep was consistent for memory across modalities: verbal paired-associates and visual forms (Fowler et al., 1973). Furthermore, it also confirmed the original observation that there is significantly more SWS in the first half of sleep and more REM in the second half, reinforcing the conclusion that the benefit of sleep on memory is not due solely to a lack of interference. This series of findings opened the door for theories about sleep as a time for active memory processing to emerge.

Currently, many theories about memory functions of SWS and REM exist (see Rasch & Born, 2013; Ackermann & Rasch, 2014 for review). Some highly speculative theories suggest that REM sleep is essential for “reverse learning,” to avoid recalling inappropriate or unlikely associations (Crick & Mitchison, 1983; Crick & Mitchison, 1986). Others posit that REM is primarily involved in procedural memory (Diekelmann & Born, 2010). Most current thinking indicates an active role of SWS in the consolidation of declarative (Gais & Born, 2004) and episodic memory (Walker, 2009), but this is certainly not without opposition (Siegel, 2001; Vertes, 2004; Vertes & Siegel, 2005).

Now, many researchers think of sleep, and non-REM sleep in particular, as a time during which memory traces (patterns of neural activity correspond to a memory) are replayed and strengthened (Fig 1.3). As mentioned earlier, this was first conceived by Marr (1971) and has now been formalized by others memory trace reactivation theory (Sutherland & McNaughton, 2000; McNaughton, 2010). These modern ideas about memory replay during sleep are based on neurophysiological observations and interventions during sleep that seem to be related with memory consolidation. They will be discussed in the section below.

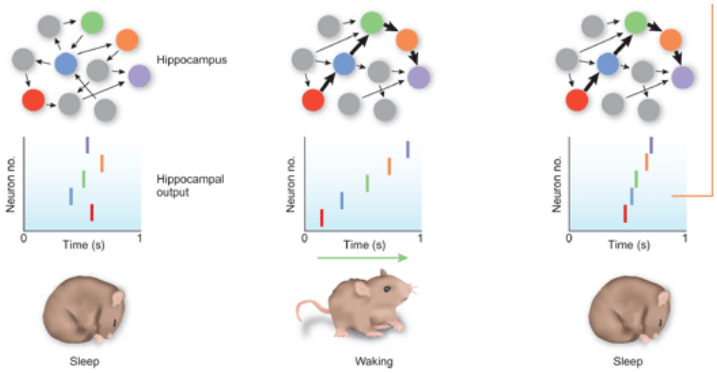

Figure 1.3. Hippocampal activity during waking and sleep. A) During sleep before behavior there is random activity during SWRs in HPC. B) During behavior there are repetitive, sequenced inputs while mice run through the neurons' "place fields". This sequential activity asymmetrically strengths connections between HPC neurons. C) During post-task sleep, ensembles are again activated, but this time, with the sequence from the task preserved. Adapted from Mehta (2007).

1.8 The Role of Neural Oscillations During Sleep on Memory

Sleep is most often defined as a temporary condition during which an organism is unconscious, relatively motionless, and has a clearly diminished ability to respond to external stimuli (Cirelli and Tononi, 2008). Along with these three general criteria, there are certain patterns of field potential oscillatory activity in the brain that consistently appear during sleep which have also been implicated in memory processing, at least in mammals: the slow oscillation of SWS (SO, <1Hz), cortical sleep spindles (7-20Hz), and hippocampal sharp wave-ripples (SWRs, 100-300Hz). There are two additional field

potential patterns, namely the theta oscillation and ponto-geniculo-occipital waves, which are present during REM sleep and will not be discussed in any great detail here. These oscillations have all been linked to memory independently. Researchers have reported that they all either increase in frequency, density, and/or amplitude following learning events, and in some cases, have even demonstrated a causal role for them in consolidating memory. Importantly, the spatiotemporal relationship between these three field potential oscillations is likely at least as important for memory consolidation as any of them considered in isolation.

That sleep is comprised of several unique stages with signature electrical events and oscillations has been known for decades (Dement and Keitmann, 1957). The neural processes generating, maintaining, and eventually disrupting sleep to awaken an animal are complex but relatively well characterized. The functions of generating and maintaining the characteristic oscillations in brain EEG associated with different sleep stages are generally ascribed to reciprocal interactions between the thalamus and cortex (McCormick and Bal, 1997). Awakening and maintaining a conscious, aroused state is produced by the apparent disruption of such highly synchronous oscillatory activity, and is controlled by a very localized brainstem region, originally called the reticular activating formation (Moruzzi and Magoun, 1949), and connected areas collectively known now as the ascending reticular activating system (Paus, 2000). Mechanisms regulating sleep and arousal states are highly conserved across mammalian species (Nicolau et al., 2000).

As described above, there are many unique field potential patterns occurring during sleep that have been implicated in memory processing. Before describing the

SO, it is important to mention that there are two related field potential events: K-complexes and down-to-up state transitions. These three events all represent similar patterns of cortical activity but occupy different sleep stages and may have different generative mechanisms (Amzica and Steriade, 1997). In this dissertation, SO is to be understood as all down-to-up state transitions in the cortex during sleep. SO dominates the cortical field potential during deep non-REM sleep and reflects recurring transitions from down- to up-states in the cortical field potential (Steriade et al., 1993a; Isomura et al., 2006). Down-states are periods of relative silence of spiking activity in the cortical neuronal population, and, on the level of single cells, a characteristic membrane hyperpolarization is also present (Steriade et al., 1993b). Up-states are essentially the opposite: the first 100ms of an up-state reflect synchronous activation of cortical neuronal populations, often in a very stereotyped, temporally structured manner (Luczak et al., 2007). Extended periods of hyperpolarization in cortical neuronal populations, followed by brief periods of temporary disinhibition, have been proposed as a mechanism causing synchrony of the SO over long cortical distances (Amzica and Steriade, 1995). Together, this results in the characteristically large deflections in local field potential that are capable of synchronous oscillation across long distances in the cortex, which may be important in information transfer and synaptic modification as will be discussed below.

Some evidence suggests a relationship between SO, recent learning experiences, and memory processing. Using an implicit learning task, Huber et al. (2004) showed that learning, compared to a similar kinematic task that involved no learning, produced an increase in SO during non-REM sleep immediately following the

task. Moreover, the results of that study also show a correlation between increases in SO amplitude and improved performance on the task, indicating enhanced learning and suggesting a positive relationship between SO and learning, although causality cannot be inferred. Moreover, this increased SO amplitude and the correlation with learning were specific to brain regions activated by the task within the parietal cortex. Furthermore, many studies demonstrate a beneficial effect of the 1st half of sleep, which is rich in SO, on declarative memory while the 2nd half of sleep, which has comparatively less SO, provides no benefit to declarative memory (Yaroush et al., 1971; Fowler et al., 1973; Plihal and Born, 1997; Gais et al., 2000). Perhaps more compelling than those earlier studies is a relatively recent report demonstrating that transcranial application of current within the SO frequency range (0.75Hz) enhanced memory retention (Marshall et al., 2006). Stimulating at a higher frequency (5Hz) had no effect. Such findings all support the hypothesis that the SO is important for memory processing during sleep, although these findings are primarily correlative (Hubert al., 2004) or come from experiments using preparations that do not specifically target the SO, but temporal halves of a night's sleep instead (Yaroush et al., 1971, Fowler et al., 1973; Gais et al., 2000).

Sleep spindles are yet another class of field potential oscillations present during non-REM sleep that have been implicated in memory processing. Spindles are 7-20Hz field potential oscillations occurring in both the thalamus and cortex (Loomis et al., 1935). Unlike the SO, the thalamic reticular nucleus is essential in generating spindle oscillations in the cortex (Steriade et al., 1985); isolated cortical slices do not exhibit spontaneous or evoked spindling, but, similar to SO, intracortical connectivity is

essential for synchronization of spindles across anatomically distinct regions of the cortex (Contreras et al., 1997; Kandel and Buzsaki, 1997). Spindles have been characterized as local oscillations (Andrillon et al., 2011) that are capable of traveling and synchronizing across long distances within the cortex (de Souza et al., 2016). The firing of many cortical neurons highly active during wakefulness seems to be phase-locked to spindle troughs, which may reflect the role of spindles in sustaining that enhanced level of activity during sleep (Steriade, 1999; Steriade et al., 2003; Peyrache et al., 2009). In fact, at least one computational model of different sleep stages asserts that spindles can induce Ca2+ entry into dendrites of cortical neurons which can, in turn, induce a cascade of intracellular biochemical events that culminate in synaptic plasticity (Sejnowski, 1998; Sejnowski and Destexhe, 2000). It appears that spindles, because of their presence in the cortex during sleep, their ubiquity and synchrony across disparate cortical regions, and their ability to induce favorable conditions for plasticity, are an important oscillation in modifying cortical networks, a necessary operation in memory consolidation.

Complementing the physiological evidence that sleep spindles are involved in modifying connectivity in the cortex are several studies indicating a relationship between spindles, learning, and memory. For example, following a motor memory task, specific deprivation of stage 2 sleep, in which spindling is densest (Carskadon and Dement, 2005), produces subsequent impairments relative to deprivation of other sleep stages (Smith and Macneill, 1994). Similar impairments following stage 2 sleep deprivation have been shown using a simple tracing task (Aubrey et al., 1999). Cortical EEG showed more spindling in subsequent sleep when human subjects were required to

learn word-pair associations as opposed to during sleep after a nonlearning task. In this same report (Aubrey et al., 1999), increases in spindling were correlated with better memory recall. Furthermore, only those subjects that showed significantly greater spindle activity after the association task showed improved recall the following day (Schabus et al., 2004). A similar increase in spindle density has also been reported following a procedural learning task (Fogel and Smith, 2006). With respect to memory in other modalities, sleep spindle density following learning of face-name associations and a Rey-Osterrieth Complex Figure Test (visuospatial) can be used as a predictor of subsequent memory performance on those tasks (Clemens et al., 2005; Clemens et al., 2006). Additionally, a study administering zolpidem (a GABAa agonist sleep aid) or placebo after a verbal learning task and before sleeping reported that sleep spindle density increased after zolpidem, and verbal memory was improved. Importantly, spindle amplitude and frequency were unaffected, suggesting that the spindle density (number of occurrences during sleep) is a critical factor influencing subsequent memory recall (Mednick et al., 2013). Although, much of this work has been performed in humans, some of these phenomena have been replicated in rodents. Specifically, it has been shown that spindle density increases after learning and retrieval in rats (Eschenko et al., 2006). Johnson et al. (2010) also show that memory trace reactivation, as measured by explained variance and template matching, is enhanced during spindles and down-to-up state transitions. Clearly, there is a substantial body of evidence, across experimental protocols and organisms, linking sleep spindle activity to memory.

SWRs will be the final category of field potential oscillations discussed here. SWRs are unique from SO and spindles because they have a much higher frequency

(100-300Hz) and do occur during wakefulness, although they occur with much greater frequency during quiescence and with the greatest frequency during SWS (Buzsaki, 1986; Buzsaki, 2015). The term “ripples” is first used in O’Keefe and Nadel (1978) with respect to brief, fast hippocampal oscillations during periods of rest and immobility in rats. Ripples reflect a period of tremendous population synchrony in the brain; there is an approximate 6-fold activity gain in hippocampal unit activity and a high degree of synchrony in the structure, although some SWRs are localized to dorsal, ventral, or intermediate HPC while others appear more global (Csicsvari et al., 1999; Patel et al., 2013). Similar to cortical up-states mentioned earlier, there is a highly synchronous activation of entire ensembles of hippocampal units during SWRs, and during these activations, cell pairs with correlated spike trains tend to be coactive (Wilson and McNaughton, 1994; Kudrimoti et al., 1999; O’Neill et al., 2010). Like the SO, and sleep spindles mentioned earlier, the high synchronicity of activity among units during the SWR make it a candidate mechanism for inducing synaptic plasticity between those units, which may ultimately underlie learning and memory (Sadowski et al., 2016).

In addition to the physiological characteristics that make SWRs candidate temporal windows for plasticity, and thus, memory processing, there is substantial behavioral evidence that cannot go unnoted. Like the other field potential oscillations discussed, it has been reported that SWRs increase after learning in rats (Eschenko et al., 2008; Ramadan et al., 2009) and humans (Axmacher et al., 2008). SWRs are very discrete events and can be detected quite reliably. As such, methods of online detection and closed-loop stimulation have been developed so that SWRs can be disrupted as they develop in real time. Using this technique, two labs, working independently have

shown that SWRs play a role in memory consolidation during sleep (Girardeau et al., 2009; Ego-Stengel and Wilson, 2010) and presumably memory recall when awake (Jadhav et al., 2012). Along with the correlative evidence showing increased SWR density following learning, these loss-of-function type experiments support the theory that SWRs are normally involved in memory processing.

These three field potential oscillations, SO, sleep spindles, and SWRs, have garnered much attention individually, but they also have relationships with one another that are likely relevant to memory processing during sleep (Staresina, 2015). More specifically, coherence between cortical regions has been hypothesized to support plasticity between them (Steriade et al., 1999). Similarly, interactions between hippocampal SWRs and cortical oscillations is thought to support information transfer and plasticity between those participating structures (Buzsaki, 1986; Buzsaki, 1996; Maingret et al., 2016). SWRs are coupled to SOs: SWRs tend to occur at the onset or offsets of cortical up-states and rarely occur during cortical down-states (Battaglia et al., 2004; Peyrache et al., 2009; Wilber et al., 2017; Karimi Abadchi et al., 2020). It has even been proposed that cortical population bursts associated with sleep spindles and SOs are responsible for triggering SWR-related discharges, although this may not always be the case (Sirota et al., 2003). As that proposal would predict, cortical spindles and hippocampal SWRs are indeed temporally correlated (Siapas and Wilson, 1998).

The temporal relationships of these oscillations across the cortex and HPC are not just epiphenomena of active systems during sleep; they appear to have functional significance with respect to memory. Using a word-pair association task, Mölle et al. (2004) reported an increase in EEG coherence across multiple frequency bands during

learning. Specifically, during the depolarized phase of the SO in sleep, EEG coherence in several frequency bands (<1Hz, 1-4Hz, 8-10Hz, and 25-40Hz) was increased following the learning. Importantly, such SO-locked coherence was not seen after a nonlearning control task. Their findings suggest that coincident field potentials in spatially distinct cortical regions may support memory processing during sleep. A recent study used a closed-loop stimulation protocol to detect the onset of hippocampal SWRs and stimulate prefrontal cortex at the SO frequency (Maingret et al., 2016). This closed-loop stimulation enhanced coupling between SWRs and cortical SO and spindles generated by the exogenous stimulation. This artificially induced SWR-SO-spindle coupling had the effect of enhancing memory on an object-place association task, whereas delayed stimulation that did not produce coupling had no effect on memory performance. Along with SWRs in the HPC, there seem to be bursts of transient high-frequency activity (100-150Hz) in association cortices, but not primary cortical regions that appear very similar to hippocampal SWRs. These cortical “ripples”, in addition to SOs and spindles, are reported to be more tightly coupled to SWRs following learning of a cheeseboard maze task than following a non-mnemonic maze exploration period (Khodagholy et al., 2017). These cortical ripple-like events are yet another example of interacting field potential oscillations that may facilitate communication between memory structures during sleep.

SWS appears to play an important role in processing of recently acquired memories. SOs are correlated with the subsequent memory performance, and artificially induced enhancements of the SO produced better memory retention (Huber et al., 2004; Marshall et al., 2006). Similarly, sleep spindle density is influenced by prior

learning before sleep, and deprivation of sleep stages densest in spindles impairs memory retention for a variety of tasks (Schabus et al., 2004; Clemens et al., 2006; Mednick et al., 2013). SWRs are widely recognized as important field potential oscillations for learning and memory. They, like the SOs and sleep spindles, have been reported to increase during sleep after learning (Axmacher et al., 2008; Eschenko et al., 2008). More compellingly, the closed-loop online detection and disruption of SWRs during sleep significantly impaired memory consolidation (Girardeau et al., 2009; Ego-Stengel and Wilson, 2010). While these field potential oscillations can and have been considered in isolation, it may be more appropriate to investigate their temporal coupling, or lack thereof, with respect to memory (Sirota et al., 2003).

1.9 Combining NC recording and HPC Lesion

The novelty of the experimental approach described herein should be noted. Although HPC lesions have been a widely used experimental procedure for decades, and single-unit electrophysiology has been used for nearly as long, the two are rarely used in conjunction. In fact, the only studies combining these two experimental techniques use neonatal ventral hippocampal lesion as a model for schizophrenia and record single neurons with intracellular or patch clamp electrophysiology (Tseng et al., 2009). These studies often concentrate on changes in GABAergic signaling, often in the frontal cortical areas, and results are interpreted in the context of schizophrenia, but they do not examine changes at the level of ensembles (or populations) of neurons and the HPC lesions are targeted only to the ventral HPC.

Two other studies, also from the McNaughton lab, recorded cortical activity in HPC lesioned mice (Mao et al., 2018; Esteves et al., 2021). These two studies found reductions in the spatial selectivity of NC neurons in HPC lesioned mice while they ran on a treadmill with tactile cues, but they did not address changes in other behavioral states like sleep, rest, or anesthesia. Furthermore, because those studies used 2-photon imaging of fluorescent Ca^{2+} indicators to observe neuronal activity, they could not assess changes in LFP. There may be alterations of frequency, power, or coherence in certain LFP frequency bands within or between NC regions. Because of the relatively long decay time constants of Ca^{2+} indicators (Wei et al., 2020) and lower sampling rates, these earlier studies had limited ability to capture temporal relationships between pairs or ensembles of neurons. The high-density extracellular electrophysiological recordings used in these experiments can answer a variety of questions about whether the general physiology and dynamics of NC change when HPC is lesioned. In that respect, this study is the first of its kind.

1.10 Background Summary and Conclusion

The hippocampal memory indexing theory has proven to be a durable explanation for how the brain could possibly achieve its impressive memory functions. In short, the theory asserts that HPC can serve as a store not for the experiential information itself but for indices to the appropriate NC patterns storing the actual information. HPC anatomy, hodology, and propensity for plasticity support this theory. Experimental results generally accord with predictions of the theory.

Predictions of the theory remain untested. This dissertation is primarily concerned with a prediction derived from the theory: there should be an “index”. HPC

lesions should affect the response properties of this “index”. It is now clear that there does appear to be an NC “index” which may be localized to the superficial cortex. Only one experiment thus far has systematically questioned whether there are differences in spatial coding properties between the superficial and deep NC (Mao et al., 2017). This study was limited only to the RSC. This dissertation furthers that analysis by looking at spatial coding properties across cortical layers in RSC and M2 using a similar head-fixed, 1D navigation behavioral paradigm.

CHAPTER 2:

Behavioral Effects of Hippocampal Lesion on Virtual Reality Task

Mice were trained to run through a visual virtual reality task while head-fixed. The task required mice to simply run on a styrofoam ball to update their position in a 1D, annular-track VR environment and receive food rewards at fixed locations on the track. Because mice needed at least 2-3 weeks of pre-training on this task to run consistently, and then re-training following the HPC Lesion or Sham surgery, we gathered data from mice learning in two VR environments: one before the HPC Lesion or Sham surgery, VR1 (Old) and another following the surgery, VR2 (Interim). Since HPC lesions have been widely documented to cause spatial memory impairments in freely moving animals, we decided to check whether we saw similar memory impairments in these head-fixed mice running in 1D annular VR environments.

Methods

Experimental Subjects

Experimental protocols were approved by the Institutional Animal Care and Use Committee of the University of California, Irvine (Protocol No. AUP-19-181). 10–12-week-old C57BL/6J mice were procured from Jackson Labs (000664). All mice were housed individually following the initial headbar surgery and unless stated otherwise, they were on a 12hr on/off light cycle and received daily scheduled and limited food to maintain their body weight between 75%-80% of their *ad libitum* weight as measured prior to headbar surgery.

Headbar Implant Surgery

After at least 1 week after the mice were delivered, stainless steel headbars were implanted on the skull. To begin, after a sufficient anesthetic plane was induced with 2% vaporized isoflurane with oxygen in a plexiglass induction chamber, mice were placed into a stereotaxic device and fixed in place with ear bars and a gently clamping nose cone. Isoflurane was administered through the nose cone at 0.5-2% combined with oxygen for anesthesia maintenance. Anesthetic depth was monitored regularly by respiratory rate and toe pinches to maintain a stable, deep anesthetic plane. All instruments were autoclaved before surgery to ensure sterility

Lidocaine was administered subcutaneously and topically to the shaved and sterilized scalp. A cocktail of Dexamethasone (4mg/kg), Meloxicam (5mg/kg), and Enrofloxacin (10mg/kg) was administered subcutaneously at the commencement of surgery. To begin, a “teardrop” patch of skin overlying the dorsal surface of the skull was removed with surgical scissors. Care was taken to remove skin up to, but not over, the neck and temporal muscles. The skull was cleaned with H_20_2 followed by Enrofloxacin applied directly to the skull and left in place for 5 minutes. The perimeter of the skull was scored in a crosshatched pattern to facilitate bonding of the Metabond later. Skull pitch was leveled by ensuring bregma and lambda were in the same horizontal plane. Skull roll was leveled by moving to AP -2.1mm from bregma and

A

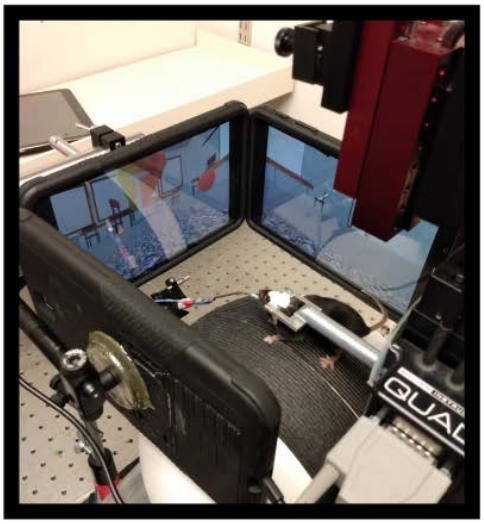

B

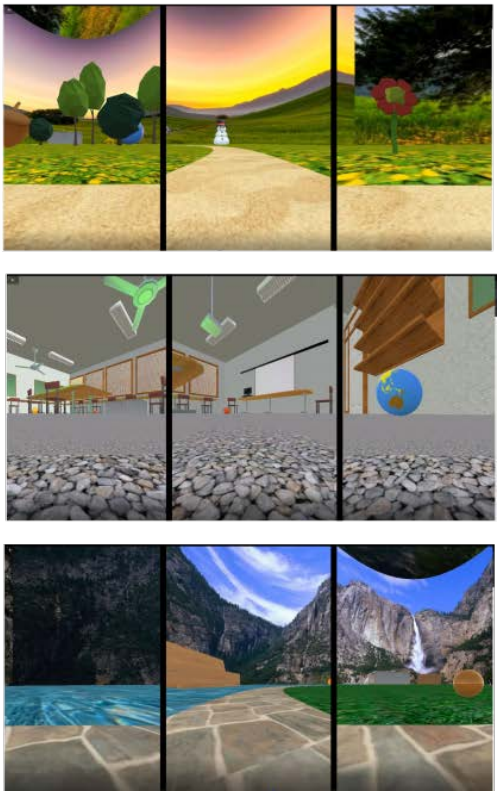

Figure 2.1. Head-fixed visual virtual reality behavioral apparatus. A) 4 views of a mouse in the head-fixation device atop the styrofoam ball, running for reward in virtual reality. B) Mouse perspective view of the 3 visual VR tracks used in this experiment. Top left was always used as VR2. Top right and bottom were used as VR1 and VR3, counterbalanced within Sham and HPC lesion groups. Black vertical lines indicate the boundaries between the tablet displays.

measuring depth +- 2.0mm laterally. It was adjusted until those two points were within 100um in depth.

The skull was then dried. Then Metabond was applied to the perimeter and undersides of the headbars. Then the headbars were gently lowered down to the skull surface. A second layer of metabond was applied over the first and on top of the point where the headbars touch the skull. The remaining exposed skull is then cleaned and dried again with sterile cotton-tipped applicators. Kwik-Cast sealant was then applied over the exposed skull. Once cured, a final layer of Metabond was applied over the Kwik-Cast, completely protecting the skull from exposure to outside air and debris. The same drug cocktail used at surgery onset was administered subcutaneously for 5 days thereafter as an analgesic and antibiotic. Ibuprofen (2mg/mL) and amoxicillin (.5mg/mL) were mixed in the mice's drinking water during this 5-day recovery period, too.

Visual Virtual Reality Behavior

After reaching 80% *ad libitum* body weight following the headbar surgery, mice were trained to run for sweetened, condensed milk reward in a head-fixed visual virtual reality (VR) apparatus (Fig 2.1A). Starting on Day 1, mice were required to run on a styrofoam ball (80cm circumference) on an axle connected to a 16-bit rotary encoder. The signal from the rotary encoder was used to update the position on an annular virtual reality track (see Figure 2.1B). The track was displayed on three tablets positioned directly in front of and 90° to the left and right of the mouse's head. A replaceable cohesive bandage was wrapped along the meridian for added grip. At multiple fixed locations on the track, rewards were administered through a lick tube positioned just in front of the mouse's mouth. The timestamps of rewards, licks, and position changes were all synchronized to the same system clock. Mice initially received 10 rewards per lap, administered at pseudorandomly assigned locations, and were quickly trained to run for only 2 rewards per lap in all VR environments thereafter. All VR environments shared a ~10cm-long grey 'cave' cue that appeared simultaneously on all three tablets at the same position on the annular track, which is the position where all training sessions started. Aside from the recording day, all sessions on the ball were 30 minutes long, and mice were only exposed to one environment per day.

Mice were trained until they ran at least 60 laps (2 laps/min) per session for three consecutive training days. After meeting this criterion in the first environment, VR1 (Old), they were given *ad libitum* access to food for at least 3 days prior to HPC Lesion

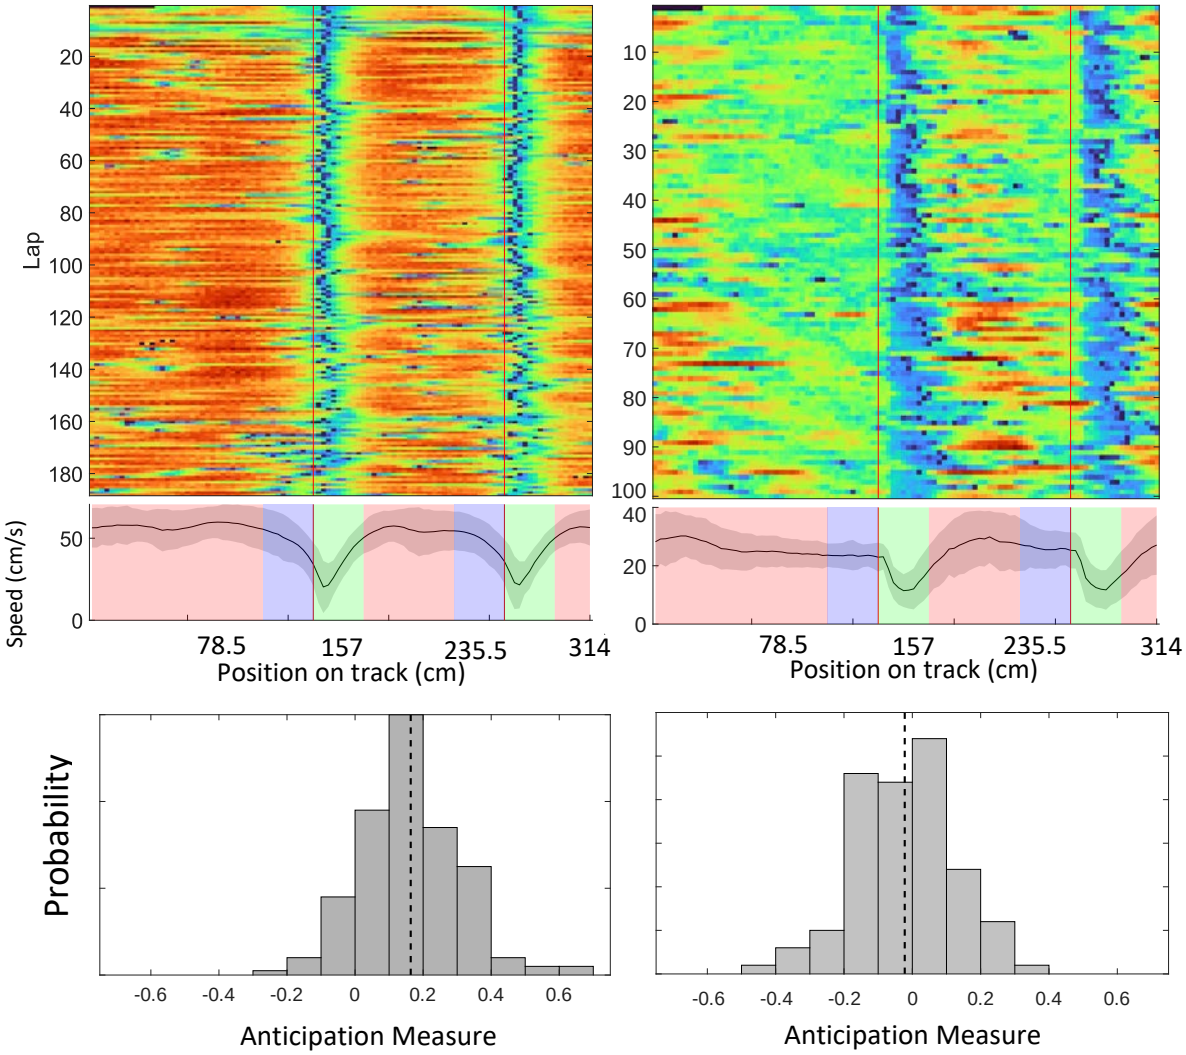

Figure 2.2. Example behavior on head-fixed VR task . A) Top heat maps show the relative speed as a function of lap (y-axis) and position within lap (x-axis). Bottom line plots show the speed as a function of within-lap position average of the entire session. Thin red vertical lines correspond to the reward sites. Red, blue, and green shaded sections indicate the 'outside reward' area (B_O), a 'pre-reward' area (B_R), and a 'post-reward' area, respectively. Darker shading is the standard deviation. B) Distribution of anticipation slowing measures. This measure is calculated for each lap by subtracting speed in the blue regions from speed in the red regions and dividing by speed in the red. The left column shows a well-trained mouse slowing down in anticipation of a reward and the distribution of the anticipatory slowing measure is shifted above 0. The mouse in the right column is consistently slowing down *after* the reward is delivered and not in advance of the reward. The corresponding distributions of the anticipation measure is therefore centered at ~0.

or Sham surgery. Once recovered from surgery, mice were returned to 80% *ad libitum* body weight and returned to the training regimen but in a novel VR track, VR2 (Interim). Once again, mice were trained to the criterion of at least 60 laps per session for three

consecutive training days. Once met, craniotomy and durotomy were performed and recordings were made the following day. On recording day, mice were placed on the ball as usual. Silicon probes were slowly lowered into the cortex and allowed to sit for at least 45 minutes upon reaching the target depth. Recording commenced and after a brief delay (1-2 minutes) either VR1 (Old) or VR3 (New) was loaded onto the tablets. The VR presentation order was counterbalanced within HPC Lesion and Sham groups. Once the VR loaded, mice were free to begin running for reward. After ~60 laps or ~30 minutes, whichever came first, the VR tablets shut off for 1-2 minutes, and then the second VR environment was loaded. Following another ~60 laps or ~30 minutes, the VR tablets were immediately shut off and urethane was administered. Mice were euthanized and transcardially perfused with phosphate-buffered saline and paraformaldehyde upon completion of recording to confirm HPC lesion (or lack thereof in shams) and recording sites histologically.

Anticipatory slowing was used as a behavioral measure from which to infer memory on this task. In general, mice slow down to consume rewards in this kind of behavioral protocol, and if they anticipate an upcoming reward, they will slow down before the reward is administered (Fig 2.2). To quantify this behavioral phenomenon, a mouse's speed leading up to the reward was compared to its speed elsewhere in the VR. More specifically, the track was divided into 100 bins (3.14cm each). The 5 bins before each reward point were considered the pre-reward bins (B_R). The 5 bins after each reward point were ignored because mice slow down, often to a complete stop, to consume the reward. The remaining 80 bins were considered outside the reward areas (B_O). For each lap, this anticipatory slowing measure was calculated. It is calculated by

subtracting the mean speed within B_R from the mean speed in B_O and dividing the result by B_O. For a single session, a distribution of anticipation measures was generated with one observation per lap, and from this distribution the mean value was assumed to capture the degree to which the mouse slowed down in anticipation for the whole session.

Hippocampal (and Sham) Lesion Surgery

Mice were anesthetized with vaporized isoflurane gas (2 % for induction, .5-2% for maintenance). Mice were then placed in ear bars on a stereotax as above. Enrofloxacin (10mg/kg), meloxicam (5mg/kg), and dexamethasone (4mg/kg) were administered via subcutaneous injection. Dental cement overlying the Kwik-Cast sealant from the initial headbar surgery was drilled away with a dental drill. Once all overlying cement was removed, the sealant was removed with forceps. The skull was disinfected with a sterile cotton tipped applicator dipped in 3% H_2O_2. The same skull leveling procedure that was used in headbar surgeries was used here. A fine-point felt-tipped marker was then used to mark positions where the 4 burr holes would be made at the following locations relative to bregma: AP:-2.3mm, ML: ±1.7mm, AP: -3.2mm, ML: ±3.0mm. Shallow divots were then made by pressing the running drill gently into those marked locations. 2 additional divots were made over the cerebellum for ground and reference pins for later recording. The skull was then soaked in lidocaine with epinephrine for 5 minutes to reduce any potential pain and bleeding while drilling the burr holes.

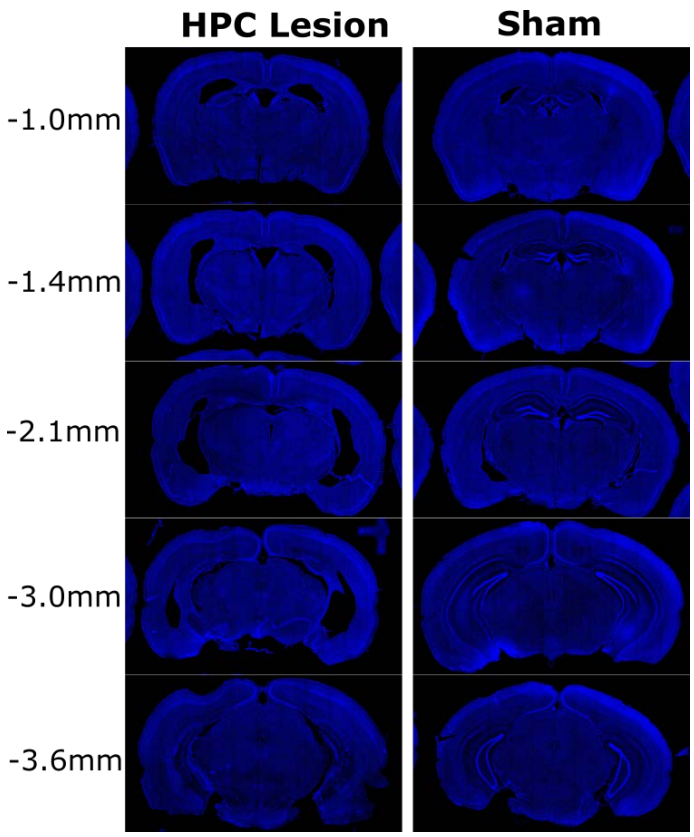

Figure 2.3. Representative HPC Lesion and Sham Coronal Slice Series. DAPI-stained sections from one HPC (left column) and one Sham (right column) lesion mouse. Numbers of the left of each row are the distance from bregma on the anterior-posterior axis. Notice the large ventricles and lack of clear cell layers in the HPC lesion brain relative to the sham. These features are indicative of a successful NMDA lesion. mm, millimeters

While the lidocaine/epinephrine mixture was sitting atop the skull, a 33G beveled metal needle was affixed to the end of a 10uL syringe, front-filled with distilled water. The syringe is pumped such that water is expelled from the needle tip, ensuring that there is a continuous volume of distilled water from the end of the plunger to the end of the needle. 150nL of air was then drawn up through the tip. The needle tip was then inserted into a droplet of sterile artificial cerebrospinal fluid for Sham or 15mg/mL n-methyl-d-aspartate dissolved in 0.9% sterile saline for HPC Lesions, and 5uL of the respective solution was drawn up into the syringe. The coordinates of bregma were then re-recorded using the syringe needle tip affixed to the stereotax as the new marker. Drilling and injection at each site was done one-at-a-time by bringing the syringe needle tip to the desired location, ensuring that the needle was centered in pre-drilled divot, moving the needle away, continuing drilling the burr hole carefully until the skull was removed and the brain underneath exposed. In rare cases, when there was bleeding, any blood was irrigated

with 0.9% sterile saline quickly. The needle was then hovered above the skull at the injected site and a small volume of solution was expelled to ensure the needle was not clogged. Assuming no clog was present, the needle was lowered at a rate of ~10um/s until it reached the target depth (-1.75mm for septal site, -3.45mm for temporal site) and then left to rest in place for 2 minutes. The needle was then withdrawn 0.05mm to form a pocket into which the solution could be injected. The position of the air bubble between injection solution and distilled water was recorded. Injections were then initiated at a rate of 100nL/min for 2 minutes (200nL per site total). The position of the air bubble is checked again to ensure it moved a distant equivalent to 200nL in the syringe barrel. Following injection, the needle was left in place for an additional 2 minutes to allow for diffusion of the solution and then withdrawn at a rate of ~10um/s. During the entire process of lowering, injecting, and withdrawing the needle, the burr hole area is kept moistened with 0.9% saline to ensure that dried tissue and blood did not adhere to the needle. Immediately upon withdrawal of the needle from the brain, solution was expelled from the needle to check again for clogs. If the solution was visibly expelled from the needle tip before and after lowering the needle and the air bubble moved during the injection, it was considered successful. Immediately after each injection, any residual saline in the burr hole was absorbed and bone wax was melted into the hole to seal it using a low-temperature cauterizer. This process of drilling, injecting, and sealing was then repeated for the remaining injection sites one-at-a-time.

Once all injections were made, two additional burr holes were made in the cerebellum for ground and reference gold pins. The male ends of the pins were angled such that once implanted, the female ends pointed away from the skull, thus not

interfering with lowering the probes later. The pins were gently placed in the burr holes so that the male ends rested on the surface of the cerebellum. Once in place, Metabond was applied to the pins to bond them to the skull and the metabond remaining from the previous surgery. At this time, a tiny bit of Metabond was also applied over the bonewax sealing up the injection holes. After the metabond cured, a cotton tipped applicator dipped in H_2O_2 was rubbed against the skull to clean it once again. A dry applicator was then applied to the skull to completely dry it. Kwik-Cast sealant was then applied over the remaining exposed skull. After that cured, a final thin layer of Metabond was laid over the Kwik-Cast such that it was totally encased in Metabond. This ensures the skull was sealed off from outside air which is important in preventing infection.

The same analgesic and antibiotic regimen that was administered after headbar surgery was used again here. In addition, the anticonvulsant Diazepam (5mg/kg; I.P.) was given upon surgery completion and every day for 5 consecutive days thereafter. This prevented seizures that sometimes occur following excitotoxic lesions of the hippocampus. Both HPC Lesion and Sham groups were given the Diazepam regimen.

Results

HPC Lesion Reduces Reward Anticipation while Learning New VR Environment

HPC lesions are well known for causing impairments to spatial learning and memory in freely moving rodents (see O'Keefe & Nadel, 1978 and Bird & Burgess, 2008 for reviews). Because our mice were trained to run in a new environment, VR2 (Interim), after they received HPC or Sham lesions, we were able to test whether or not the HPC

lesions in this head-fixed behavioral protocol would result in similar impairments in spatial memory.

To assess spatial memory in this task, we capitalized on the fact that mice slow down in anticipation of upcoming reward (Fig 2.2). Indeed, in VR1 (Old), before any HPC lesions or Sham surgery was performed, all mice showed anticipatory slowing on at least 2 of the first 3 consecutive days with greater than 60 laps (Fig 2.4A, white bars). Using this behavioral criterion of 2 out of 3 days with a positive anticipatory slowing, there is no difference between groups ($p>0.05$, Fisher's Exact test). After the Sham or HPC Lesion surgery, while mice began running in VR2 (Interim), this was no longer true. There were significantly fewer mice meeting this criterion in the HPC Lesion group (2/6) compared to the Sham group (8/8) ($p = 0.015$, Fisher's Exact test).

Also, performance across the first three consecutive days of running greater than 60 laps in each VR was averaged for each mouse to yield one value for each mouse for each VR environment. A two-way repeated measures ANOVA revealed that while there was no significant interaction (($F(1,12)=3.53$, $p>0.085$) between the two factors, HPC status ($F(1,12)=7.00$, $p=0.021$) and VR identity ($F(1,12)=8.98$, $p=0.011$) both had main effects. Post-hoc multiple comparisons test showed that anticipation was only significantly different between Sham and HPC Lesion mice in the VR2 (Interim) environment ($p=8.1 \times 10^{-3}$, Bonferroni correction; Fig 2.4B) and not in the VR1 (Old) environment ($p>.99$, Bonferroni correction). The lack of a difference between groups in VR1 was expected; this data was gathered before the HPC Lesion or Sham surgery. This lack of difference indicates that there were no intrinsic group differences from the start of the experiment. The results from these comparisons make two things clear: 1)

Sham mice remember the reward location and 2) show more anticipation than HPC Lesion mice.

HPC Lesion Reduces Reward Anticipation in Old VR, too

The above measures of anticipation were all taken from sessions when the mice were running consistently in the same VR environments for several days, so they do not tell us about what happens during the very first exposure to a new VR environment. So, the next part of the experiment asks the following questions: how does anticipation differ when mice are exposed to a novel VR environment for the first time? Do we see anticipation in an old environment they have not been exposed to since before the HPC Lesion or Sham surgery?

On recording day, mice were exposed to both of these circumstances. A two-way repeated measures ANOVA on anticipatory slowing during recording day revealed that HPC status ($F(1,10)=7.435$, $p=0.021$) showed a main effect. There was no main effect of VR identity ($F(1,10)=0.387$, $p=0.548$) and no interaction between the two factors ($F(1,10)=0.489$, $p=0.5$). A post-hoc multiple comparisons test show anticipation was significantly different only in the VR3 (New) track ($p=0.022$, Bonferroni correction; Fig. 2.5C) and not in VR1 (Old) ($p=0.063$, Bonferroni correction). Importantly, two-way repeated measures ANOVA on mean moving speeds showed no main effect of VR identity($F(1,10)=0.304$,$p=0.594$), nor HPC status ($F(1,10)=0.627$,$p=0.447$), nor interaction ($F(1,10)=1.607$,$p=0.234$; Fig 2.5D).

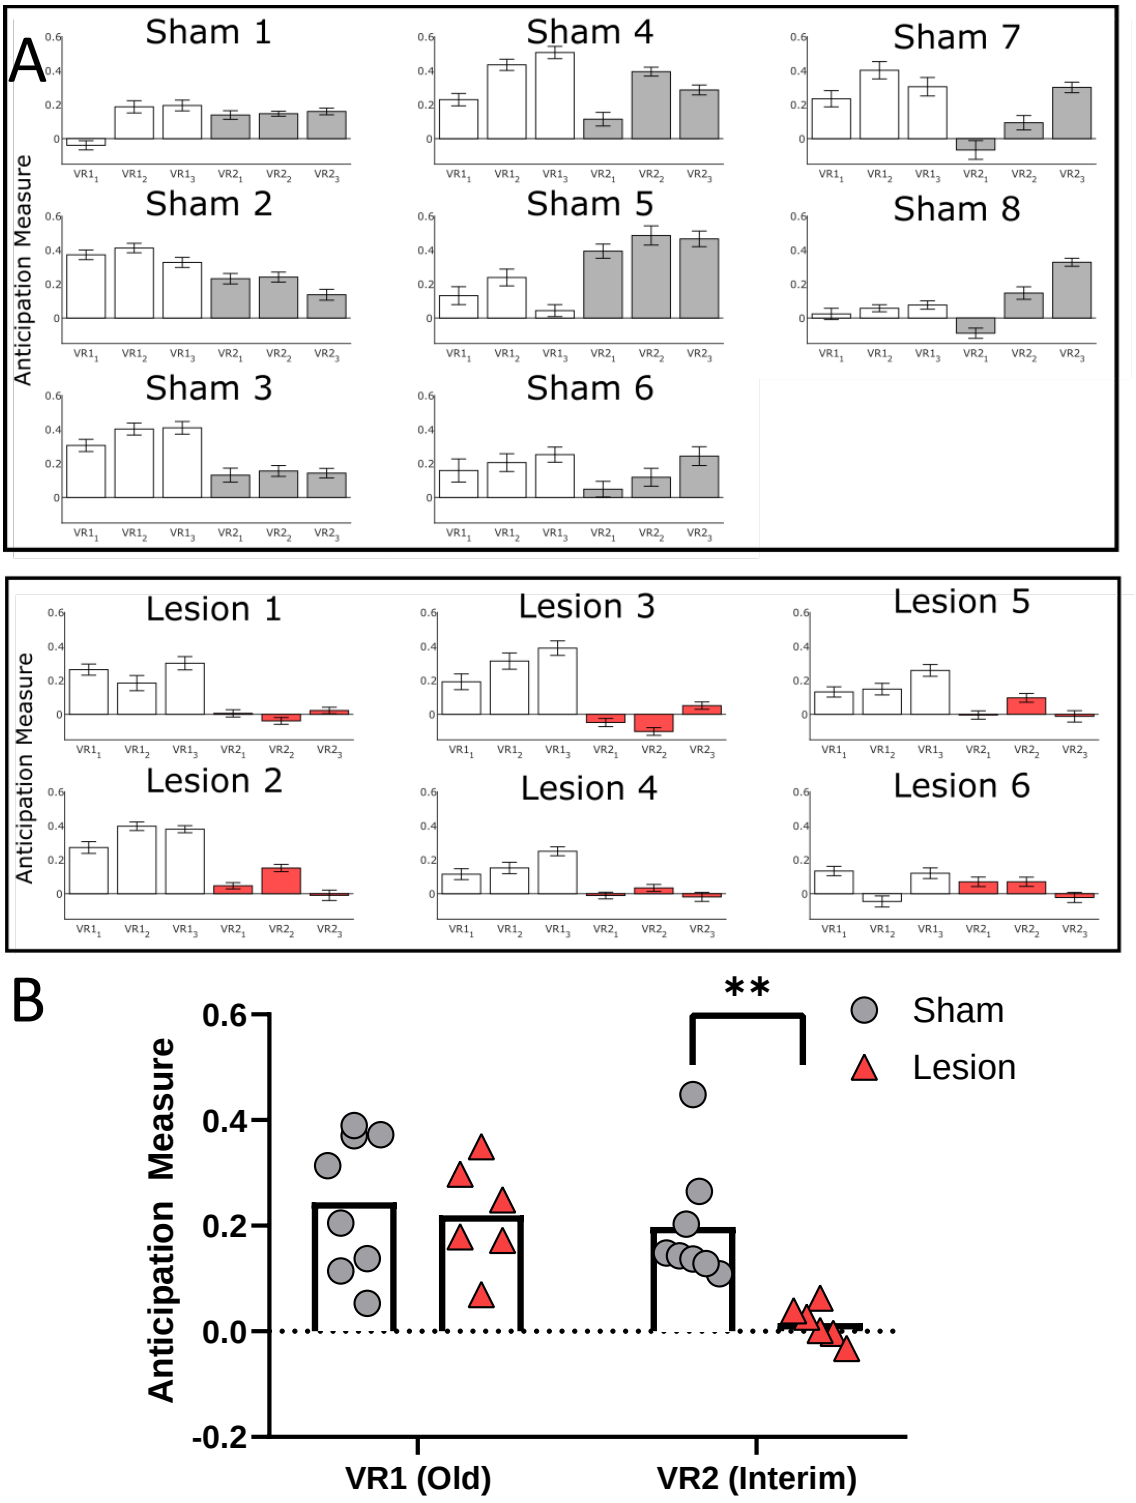

Figure 2.4. Anticipatory slowing during learning in VR1 and VR2. HPC lesion mice show less anticipation while learning VR2 track. A) Mean anticipation on each of the first 3 training days with greater than 60 laps in VR1 (white bars, pre-op) and VR2 (colored bars, post-op). Top box, Sham mice. Bottom box, HPC Lesion. Bars show mean anticipation ± 95% CI. B) Per mouse means (circles, triangles) and group averages (bars) of anticipation in VR1 and VR2. Shams showed more anticipation than HPC Lesion mice in VR2 only ($p=8.3x10^{-3}$, Bonferroni correction)

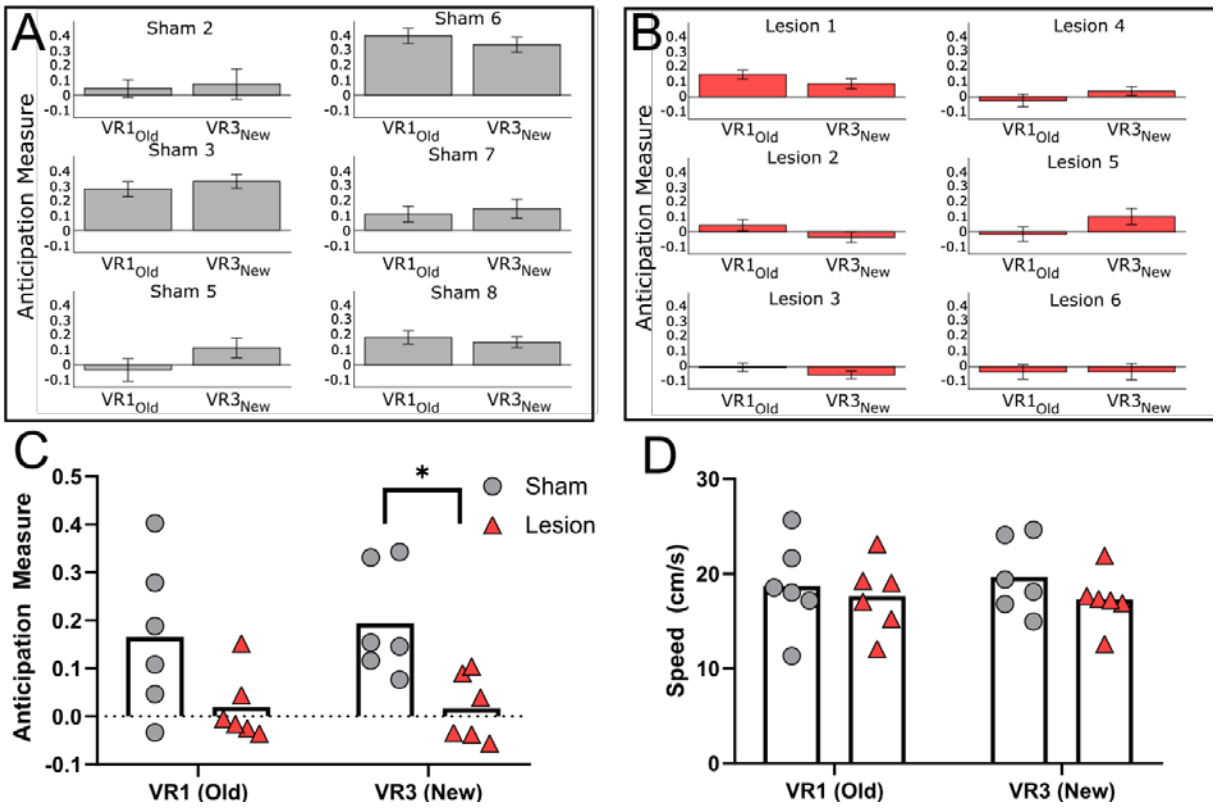

Figure 2.5. Anticipatory slowing during VR1 and VR3 during recording. HPC lesion mice show less anticipatory slowing on recording day. A) Mean anticipation for each Sham mice in the Old (left bars) and New (right bars) VR tracks. B) Same as A but for HPC Lesion mice. Bars show mean anticipation ± 95% CI. C) Group and condition averages (bars) of anticipation in VR1 and VR3 on recording day with individual measure overlayed (circles and triangles). There was a main effect of HPC status (p=0.021), no effect of VR identity (p=0.548) and no interaction(p=0.5). There was also a difference between groups within the VR3 track (p=0.022, Bonferroni correction). D) There were no differences in average movement speed across groups or VRs.

Discussion

Unexpectedly, HPC lesions reduced anticipatory slowing compared to controls when both VR tracks were considered together (Fig 2.5). However, the multiple comparisons test showed that anticipation was significantly reduced in the VR3 (New) track and there was no difference in VR1 (Old). It is likely that this was due to the large variance of anticipatory slowing in VR1 in the Sham group. It is unclear why there was so much variance in the Sham group in VR1 (Old) and relatively less in the VR3 (New).

Intuitively, one would expect the opposite: there might be some savings from the prior learning in VR1 that would lead to *more* anticipatory slowing in that environment.

The reduced anticipation in VR3 (New) in HPC Lesion mice could have been expected, because HPC lesions generally induce anterograde amnesia, and HPC lesions have routinely been shown to impair spatial memory (O'Keefe & Nadel, 1978). It was unclear, however, what would happen in the old VR1, because older memories can become independent of HPC through systems consolidation with time and/or repeated experience (Lehmann et al., 2009) although this is not true for all kinds of HPC-related memory (Ocampo et al., 2017). It appears that, in this task, the mean number of sessions in VR1 that occurred before HPC lesion (20.5 ± 4.32sd, n=6), was insufficient to produce such a consolidation phenomenon such that the task, at least in the familiar VR, is no longer HPC-dependent. This is unsurprising considering that even in the Sham mice, there was little to no savings when running in the VR1 environment. Based only on this study, it is unclear whether or not additional exposures and/or time before HPC lesion would make the memory resistant to HPC lesion.

Using this anticipatory slowing measure while mice were beginning to run consistently on VR1 (Old) and VR2 (Interim), it is clear that the Sham group showed more anticipation than the HPC Lesion group and actually remembered the locations of rewards in VR. The Sham group even appeared to rapidly learn the reward locations in a new VR during their first exposure to show significant anticipatory slowing (Fig 2.5C, Sham VR3).

It is also important to note that HPC Lesion mice might not necessarily have a spatial memory impairment. Because this task does not require the mice to perform an action at a specific place, running at a constant speed until reward is administered is a viable strategy to consistently receive reward. The HPC Lesion group may have simply adopted this strategy rather than forget the location of the rewards. Given the long list of previous experiments demonstrating spatial memory impairments following HPC lesion, it is probable that the reduction in anticipatory slowing in the HPC Lesion group is the result of a spatial memory impairment, but this cannot be concluded with certainty. An additional experiment, perhaps with the mice making an explicit action (e.g., licking, lever pressing, beam breaking with snout) at a specific place on the track, would have to be performed. In that case, if the HPC Lesion mice were still performing the action, but at erroneous locations, it would provide more convincing evidence that spatial memory is indeed impaired in a head-fixed 1D visuospatial task such as this one.

CHAPTER 3:

Neocortical Dynamics with and without a Hippocampus

If NC has a local version of the “index code”, it should share the same properties as the HPC “index code”. Namely, it should use a sparse coding scheme, and it should be capable of generating orthogonal representations to similar experiences to reduce interference. So far, it seems that there are areas of NC that abide by these principals, specifically, the superficial layers of associational NC regions. We also know, however, that bottom-up input from the sensory system dictates responses in the NC to a certain degree. We do not know the precise nature of how these bottom-up (i.e., ascending from sensory systems) and top-down (i.e., descending from HPC) signals are functionally integrated in NC to generate the response properties that we observe: a sparse, orthogonal code in some NC areas under some conditions but a more invariant, sensory-stimuli-driven code in others. The review and experiment described in this section aim to shed some light on this issue.

Methods

This experiment was carried out as described in the previous chapter. Only the details about the craniotomy procedure and recording will be described here.

Craniotomy and Durotomy

After reaching the behavioral criterion of at least 60 laps on 3 consecutive days of training in the VR2 (Interim) environment, mice underwent a final surgical procedure.

In this procedure, two circular craniotomies were made over the M2 and RSC regions of the dorsal NC.

To begin, anesthesia was induced with vaporized isoflurane running through a small plexiglass induction chamber at 2% of the gas composition. Once a deep and stable anesthetic plane was reached and the mouse was placed on the stereotax. The thin layer of Metabond overlying the Kwik-Cast sealant was drilled away, and the sealant was extracted with forceps. The skull was briefly cleaned with a cotton-tipped applicator dipped in H_20_2 and rubbed vigorously over the exposed skull. With a fine needle point attached to a stereotax arm, the bregma landmark was located and the approximate centers of each craniotomy were marked with a fine, felt-tipped marker – AP: 2.5mm, ML 1.0mm for M2 and AP: - 3.0mm, ML:0.75mm. A shallow outline of the craniotomies was made first. The M2 craniotomy was made in a rectangular shape, longer in the ML axis than AP axis. The RSC craniotomy was also made in rectangle, but the long axis was parallel with the AP axis of the skull. The RSC craniotomy extended over the sagittal sinus to allow for probe insertion close to the midline. All craniotomies were made over the mice's right hemispheres.

After the outlines were made, a lidocaine with epinephrine solution kept on ice was applied to the brain surface and allowed to rest in place for a minimum of 5 minutes. This minimizes bleeding during the drilling of the craniotomies. One cleared up and dry, drilling proceeds by going around the craniotomy with the running drill once, clearing any residual bone dust with a compressed air cannister, applying cold saline to the craniotomy area, and dipping the drill bit into the cold saline. Repeating this procedure reduces the likelihood of swelling to due heat and vibration of the drill. The

process is continued gradually until the bone at the center of the craniotomy can be removed with little force using a pair of fine-tipped forceps. After the bone is removed, the area is irrigated with saline until it is clear of debris.

At this point, the dura mater was removed slowly and carefully with a hooked 33G needle, a pair of fine-tipped forceps, and in some cases a saline solution containing diluted collagenase enzyme. In cases in which the dura was difficult to identify, the collagenase was used to gently perforate to allow for easier manipulation with forceps. Once the dura water completely removed, the sites were once again irrigated with saline and a triple antibiotic ointment was applied over them. Over the ointment was placed a layer of Kwik-Cast sealant. The same analgesic and antibiotic injectable mixture was administered following this surgery as it was the others.

Silicon Probe Recordings

The following day, the mice were brought into the behavior room, fixed to the headfork as usual, and placed on the styrofoam ball where the behavioral protocol is carried out. Before starting any of the VR display tablets, the Kwik-cast sealant was removed, and the craniotomies are irrigated again with saline. Once, any residual ointment and/or debris has been cleared, two 4-shank 128-channel Si probes coated in DiI (Fig 2.1) inserted into a custom-made aluminum holder so that they remain at a fixed distance from and orientation with respect to one another. The probe tips are manually moved to the brain surface at the specific AP and ML positions relative to bregma (Fig 3.1B) and then lowered manually until they puncture through the brain surface, which usually occurs at ~300um. They are then lowered to their final depths with an

automated micromanipulator at a speed of 1um/s. At this point, sterile mineral oil is applied generously over the brain and skull surface to act as a barrier between the brain and air and to avoid the brain surface dehydrating. If the unit yield, as estimated by activity in the high pass-filtered LFP, does not look high, the probe may be lowered at speeds of 1um/s until the yield improves. Once a final depth has been reached, the probe is allowed to settle for an additional 45 minutes before recording commences.

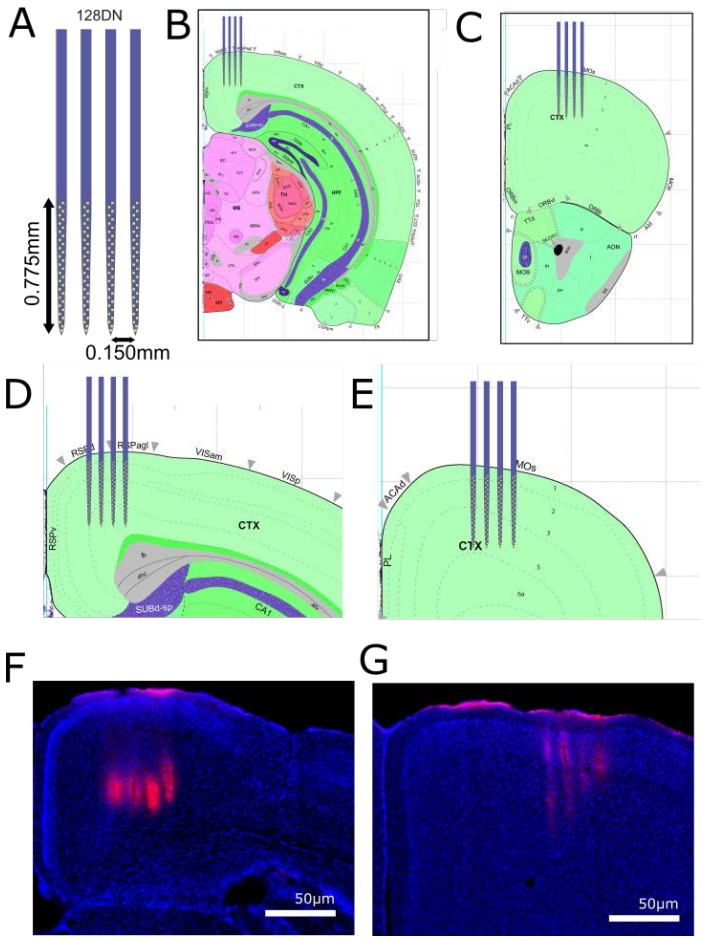

Figure 3.1. Probe dimensions and recording locations. A) Electrode pads span 0.775mm vertically and horizontal spacing between shank tips is 0.15mm. From lab website of Sotiris Masmanidis. B & C) Coronal half-sections showing placement of probes in RSC and M2 during recording. Adapted from the Allen Mouse Brain Atlas D & E) Magnified views of probes in B & C. F & G) Probe tracks shown in brain sections corresponding to those from the atlas. DiI in red; DAPI in blue.

Electrophysiological data from the 256 channels on the probes is sampled at 30kHz and digitized at 16 bits on the headstage amplifier before being transmitted to the acquisition system. Digital synchronization signals are sent to the acquisition system every time the VR-controlling routine updates the mouse's position and every time a reward is administered.

As described in the previous chapter, either the VR1 or VR3 environment (presentation order counterbalanced within groups) is loaded on the tablets, and the mouse can run at its own pace for reward. After ~60 laps or ~30 minutes, the tablets turn black, pause there for 1-2 minutes, and then the other VR environment is loaded, and the mouse is allowed to run freely again for reward. After that, the mouse is injected with urethane to induce anesthesia, the recording continues for several hours, and then the mouse is euthanized. The probes are cleaned with distilled water, 70% ethanol, and a Trypsin protease solution and stored for reuse later.

Electrophysiological Data Processing

The raw data for each probe were processed using Kilosort, an automated spike sorting Matlab package (Steinmetz et al., 2021). Automatically identified spike clusters were then manually curated. Noise clusters were identified by unusual waveform shapes (e.g., square components, multiple peaks, etc.) and high spike rates within the absolute refractory period (2ms). Clusters that have clearly overlapping principal components and compatible cross-correlograms (i.e., few spikes of either cells fall into the refractory period of the other cell) were merged. Clusters that have clearly separate principal components were split. Clusters that had abrupt cutoffs in their waveform amplitude distributions were labelled as noise because, some portion of their spikes are merging with the noise floor.

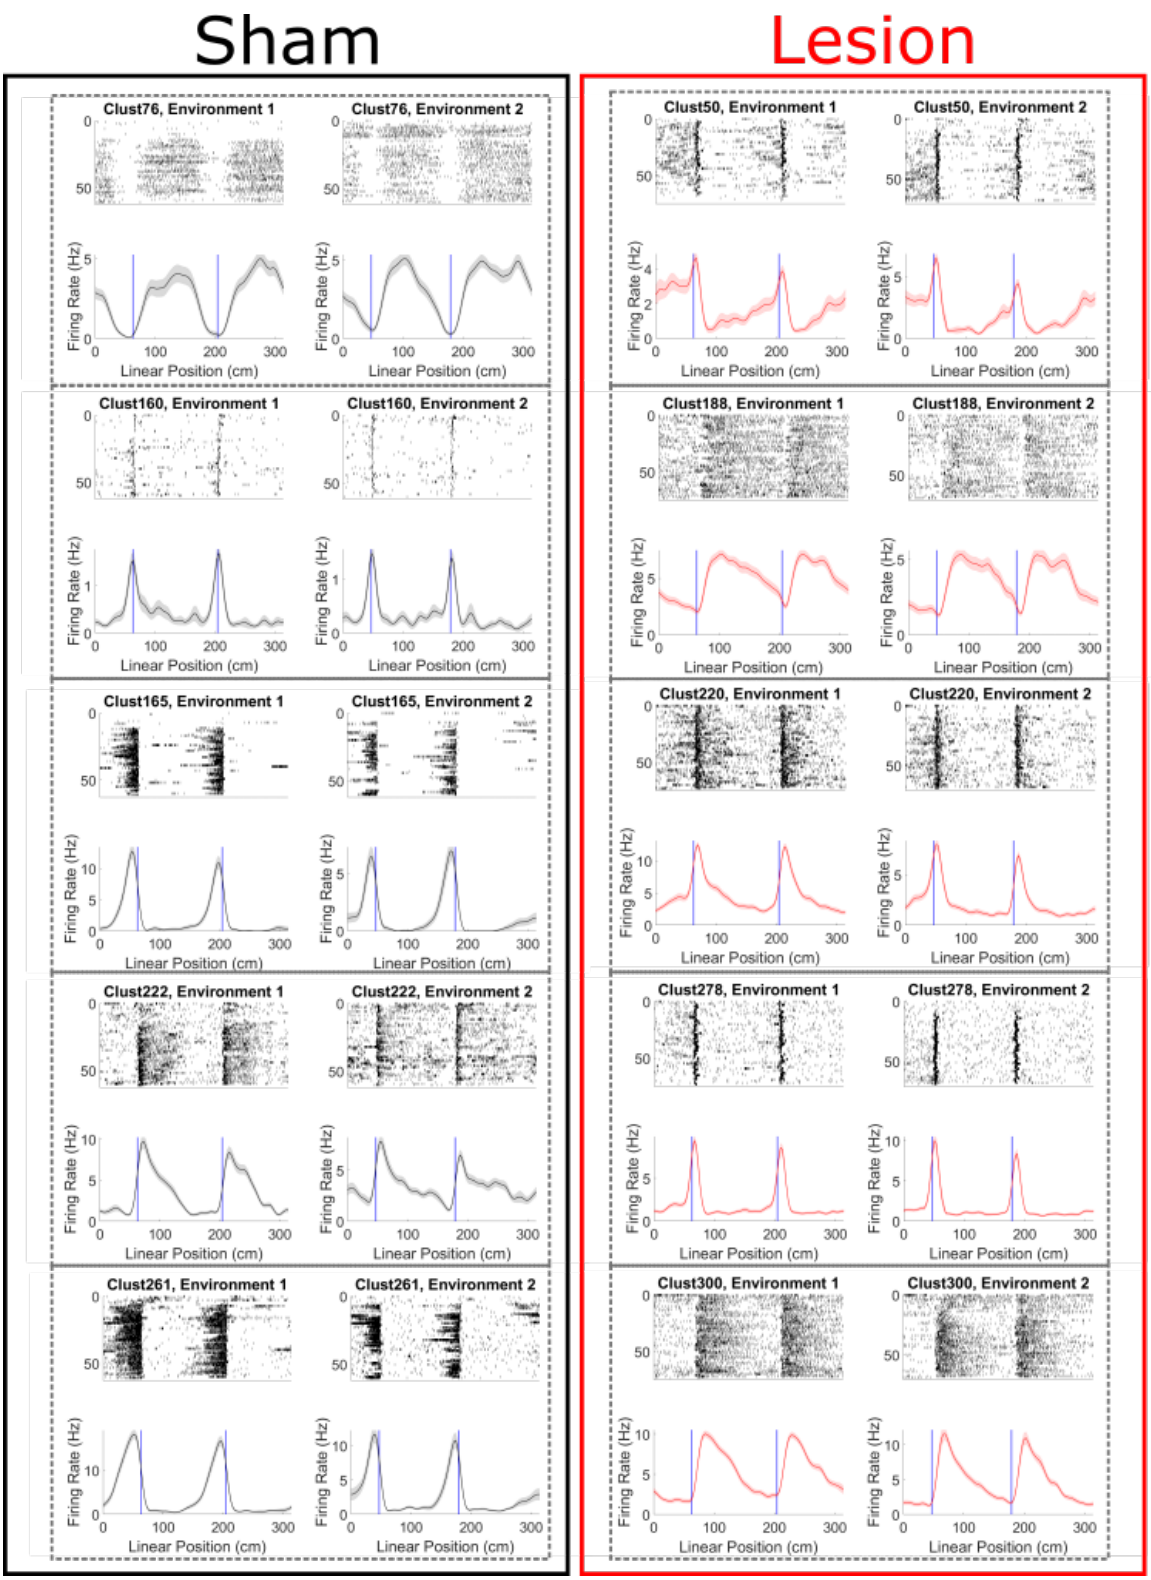

Figure 3.2. Example tuning curves for M2 neurons in Sham and Lesion mice. Each grey box displays the activity of one neuron in two VR environments. Top blots in each box are raster plots of spikes as a function of lap (y-axis) and location in VR (x-axis). Bottom plots are mean firing rates of the neurons. Neurons from sham mice are in black, lesion in red. Vertical blue lines indicate reward location.

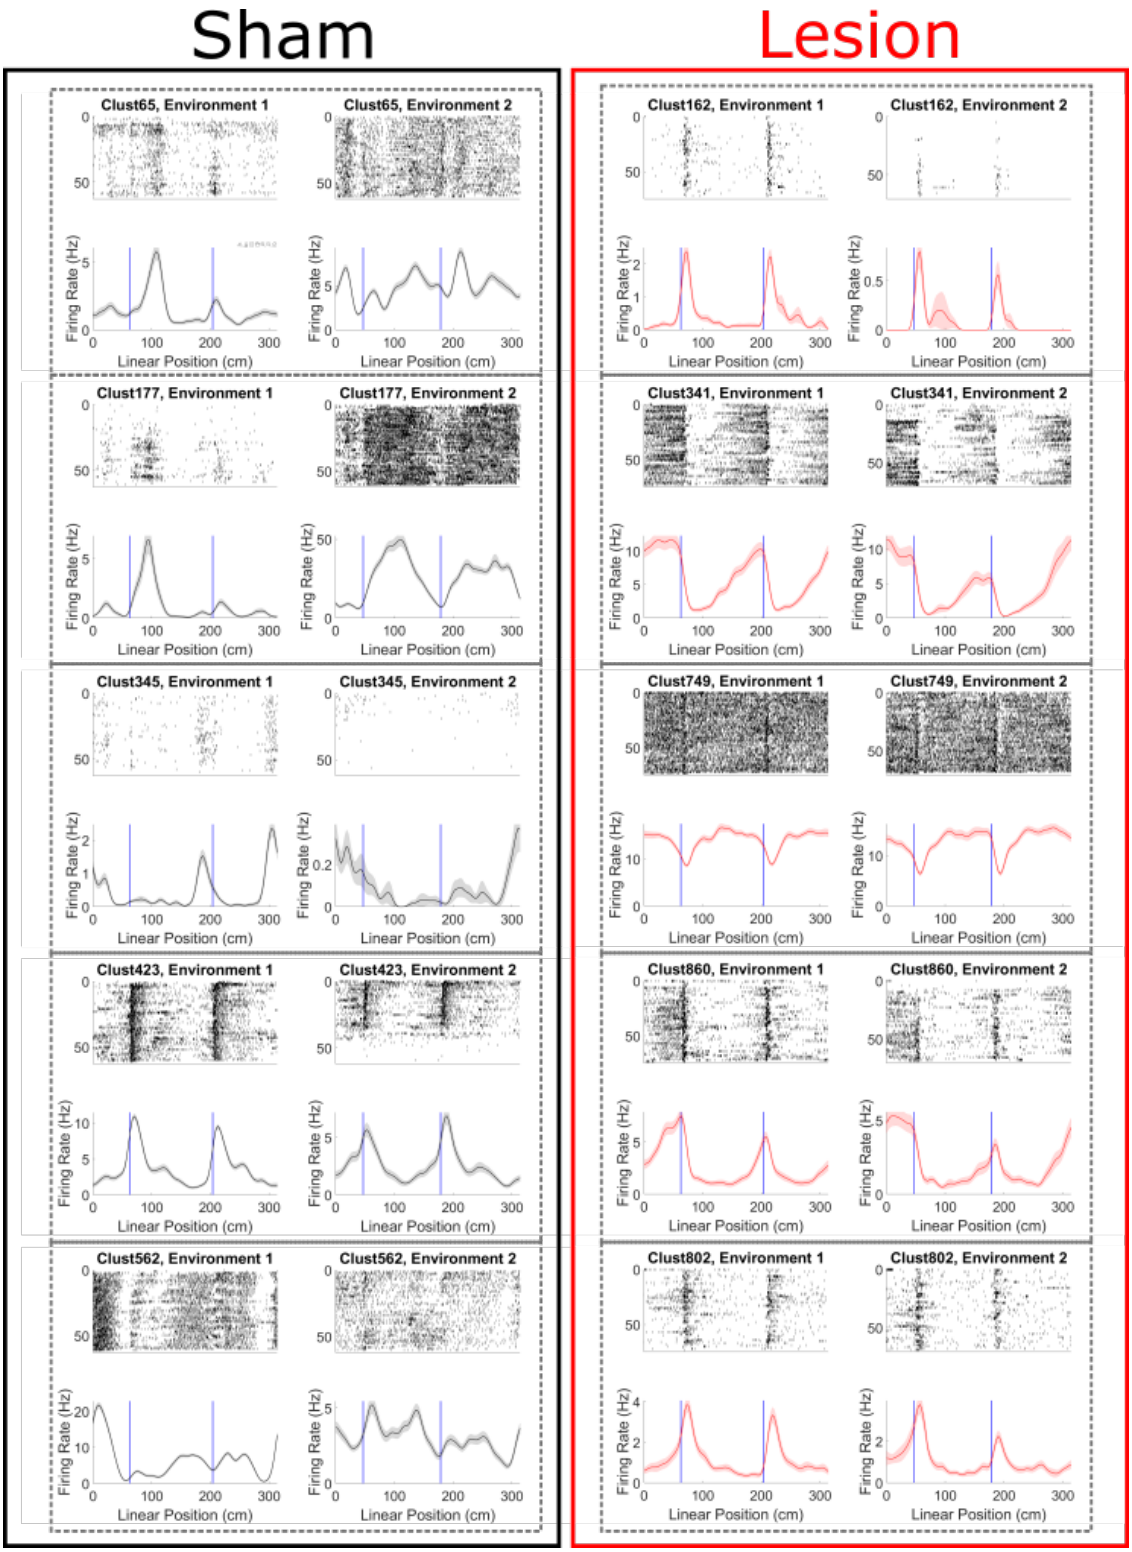

Figure 3.3. Example tuning curves for RSC neurons in Sham and Lesion mice. Each grey box displays the activity of one neuron in two VR environments. Top blots in each box are raster plots of spikes as a function of lap (y-axis) and location in VR (x-axis). Bottom plots are mean firing rates of the neurons. Neurons from sham mice are in black, lesion in red. Vertical blue lines indicate reward location.

Results

Robust activity in M2 around the reward sites

By simply looking at the firing rates of individual M2 neurons with respect to position in the VR environments, it is clear that many were active around the reward sites in both the HPC Lesion and Sham groups (Fig 3.2). While there are clearly some neurons RSC active around the reward sites (Fig 3.3, Fig 3.4A), it appears there are far fewer than those in M2.

To quantitatively test this observation, population activity (i.e., total spikes per second per cell) was binned spatially along the length of the track. Centered at each reward site, a mean-normalized 51-bin segment (+-80cm from reward) of population activity was extracted. This reward-centered mean-normalized population activity trace was averaged across all reward sites in all VR environments for each mouse, such that for each mouse there was one trace for M2 and for RSC (n = 6 each for Sham and Lesion, Fig 3.4B). 3-way repeated measures ANOVA was performed with brain region, HPC status, and position relative to reward considered as factors. The only significant effect was an interaction between position and brain region ($F(50,500)=6.5$, $p<10^{-4}$).

More Inhibitory Cells in RSC than M2

To do assess whether there were differences in the distributions of cell types across regions and conditions, features of the mean waveforms of each putative unit were calculated, including peak-to-trough durations and repolarization times (Fig 3.5

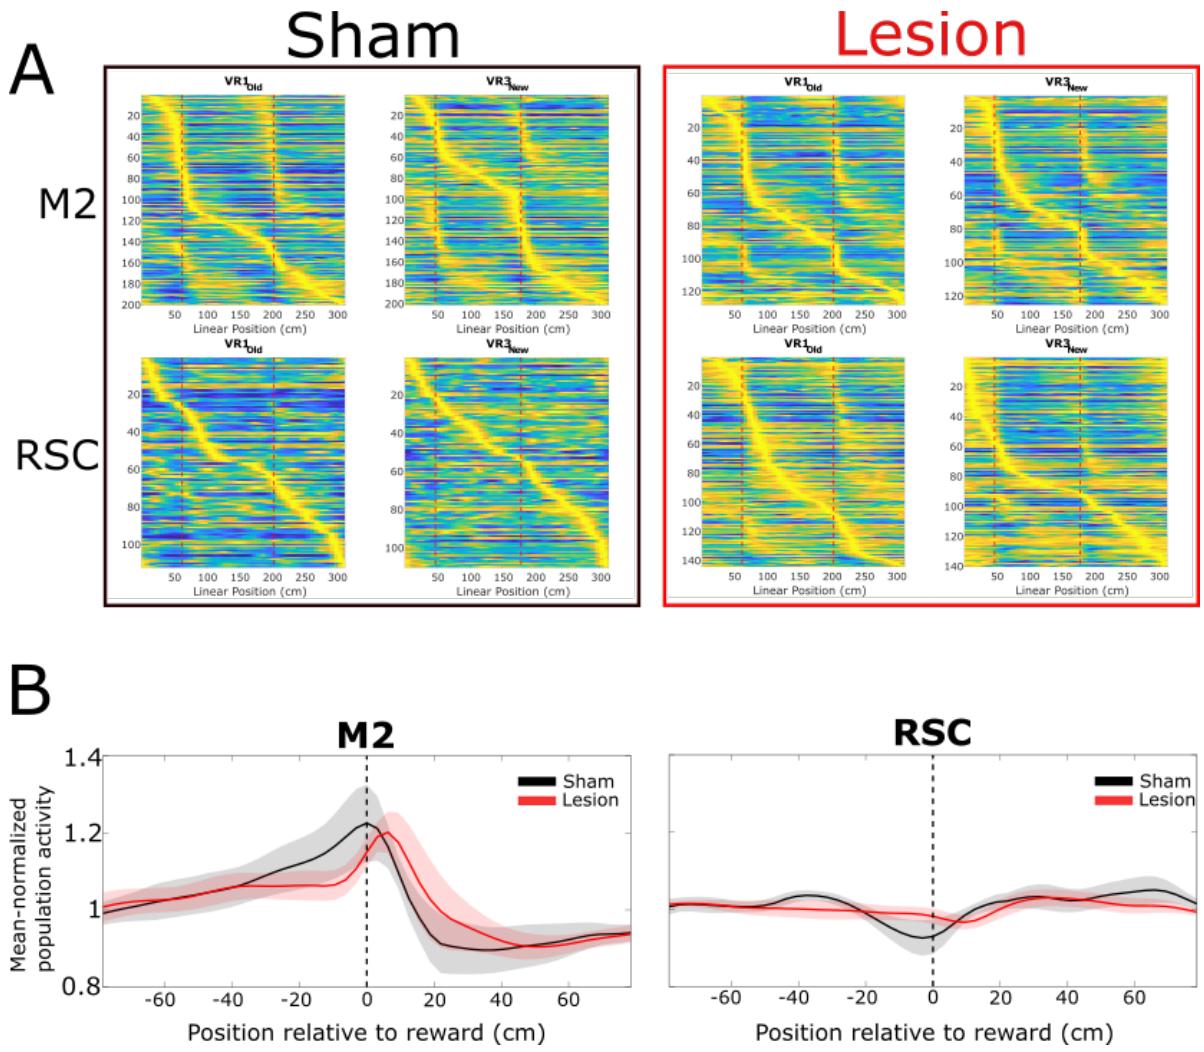

Figure 3.4 M2 and RSC response around reward sites. A) Example population firing rate maps from one mouse in the Sham (black) and Lesion (red) groups. It is clear that in M2 (top plots) there is concentrated activity around the rewards (red dotted lines); many M2 cells even seem to have two 'fields' around the rewards. Also, in the VR1 environment, many cells are active *before* the reward site, while in VR3, many are active *after* the reward. In RSC (bottom plots), there is little to no response around the reward sites, especially in the Sham mouse. B) Mean-normalized spiking activity across the entire ensemble in M2 (left) and RSC (right). For each mouse, the total activity in bins at and around the reward site was normalized to the mean population spike rate (spikes per second per cell). M2 in both Sham and Lesion groups shows activity around the rewards, although Shams show a peak in activity early than Lesions. In RSC, there is no change in activity around the reward site in either group.

A,B). K-means clustering was applied to the derived features space to detect distinct clusters in the mean waveform features with a maximum of 3 clusters. Waveforms of inhibitory neurons have shorter short peak-to-trough durations and repolarization times (Fig 3.5 B,C). A three-way repeated measures ANOVA was conducted using HPC

status, brain region, and cell type as factors with the proportions from each mouse as observations (n=12; 6 Sham, 6 Lesion; Fig 3.5D, E, Table S1). Expectedly, there was a main effect of cell type on their relative proportion ($F(1,10)=167$, $p<1.5x10^{-7}$). There was no main effect of HPC condition on the proportion cell types ($F(1,10)=8.8x10^{-3}$, $p=0.92$) nor was there a main effect of region ($F(1,10)=3.5$, $p=0.09$). There was, however, an interaction between region and cell type ($F(1,10)=9$, $p=0.01$) with a noticeably greater proportion of inhibitory cells in RSC (0.25 +- 0.11sd) than M2 (0.11 +- 0.06sd). No other interactions were present.

Silent Cells in the Superficial M2 Cortex

To determine if there were differences in firing rates of neurons across the depth of the cortex, units were divided based on what half of the recording shank on which they had the maximum amplitude spikes. Spike rates were then calculated across the entire behavioral epoch, including exposures to both VR environments. Spikes rates were calculated separately for cumulative epochs in which the mice ran greater than and less than 2cm/s. These were considered “running” and “motionless” time periods. There was a noticeable population of exceptionally low-firing neurons in “superficial” M2 of HPC lesioned mice, regardless of the speed at which mice were moving (Fig 3.6 top box).

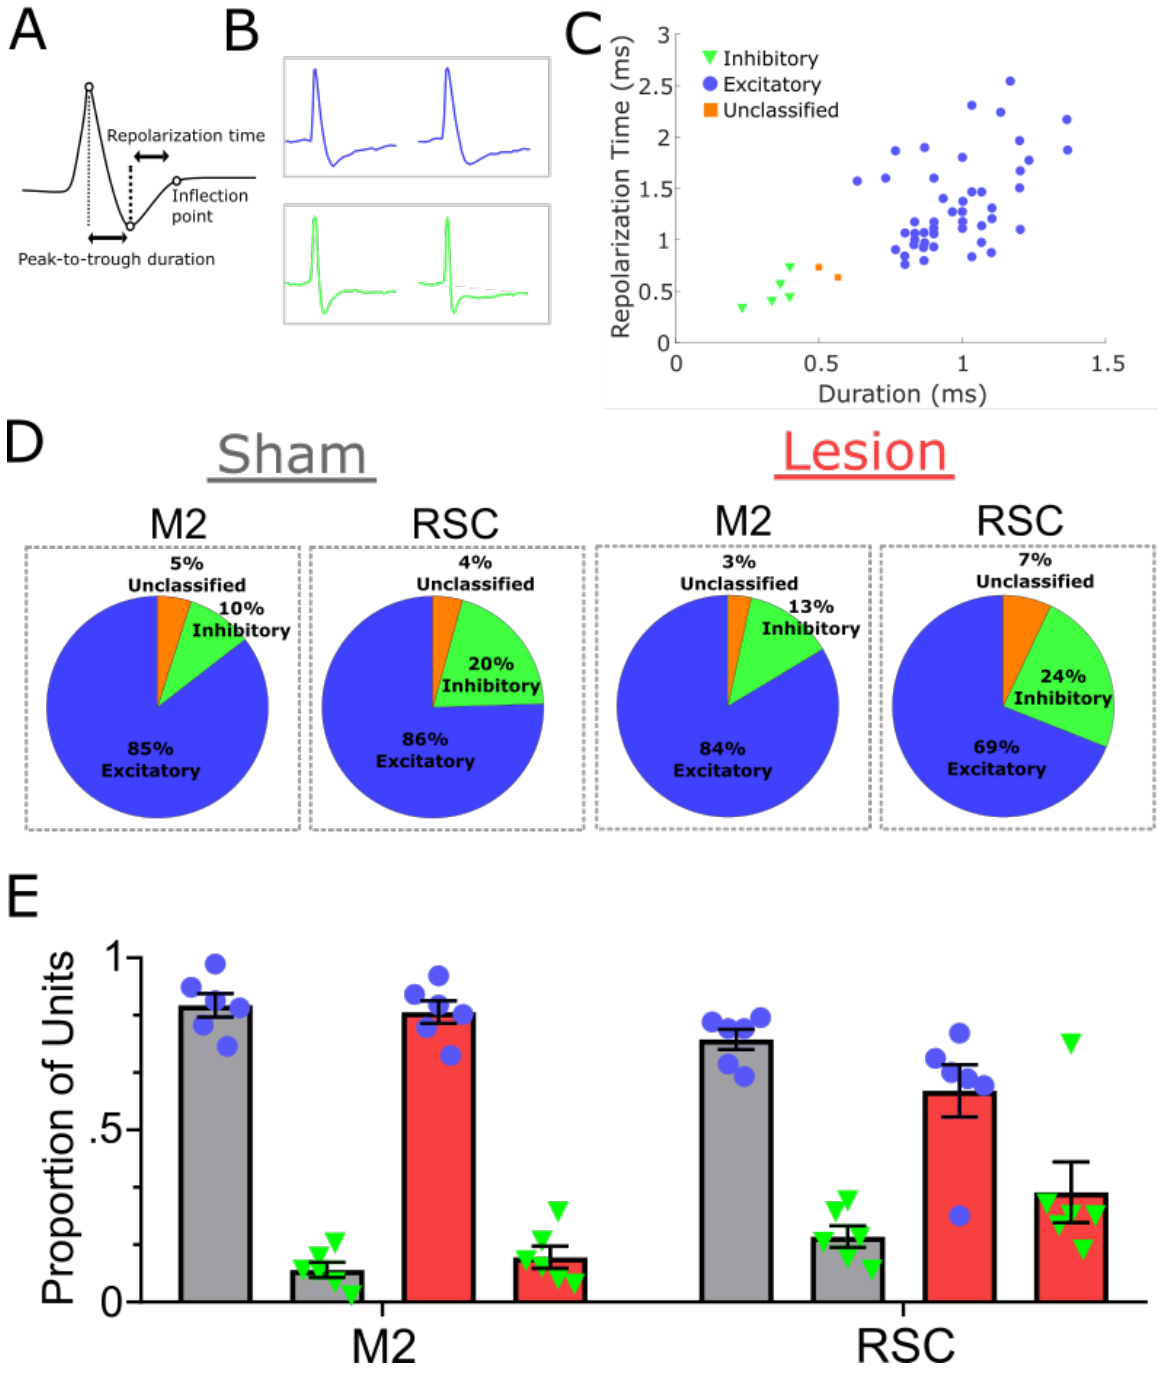

Figure 3.5 Classification and quantification of different cell types. A) Diagram of spike waveform with features used for classification labeled: peak-to-trough and repolarization times. B) Example mean waveforms from two excitatory (blue) and inhibitory (green) neurons. C) Example classification from one mouse (Sham) in one brain region (RSC) demonstrating the difference between clusters of waveforms in different classes. D) Grand total percentages of each neuron class across all mice within a condition and brain region. E) Mean percentages of each neuron class for each condition (gray, sham; red, lesion; M2 left 4 bars, RSC right 4 bars). Blue circles are proportions of excitatory neurons per condition/region. Green triangles are proportions of inhibitory neurons.

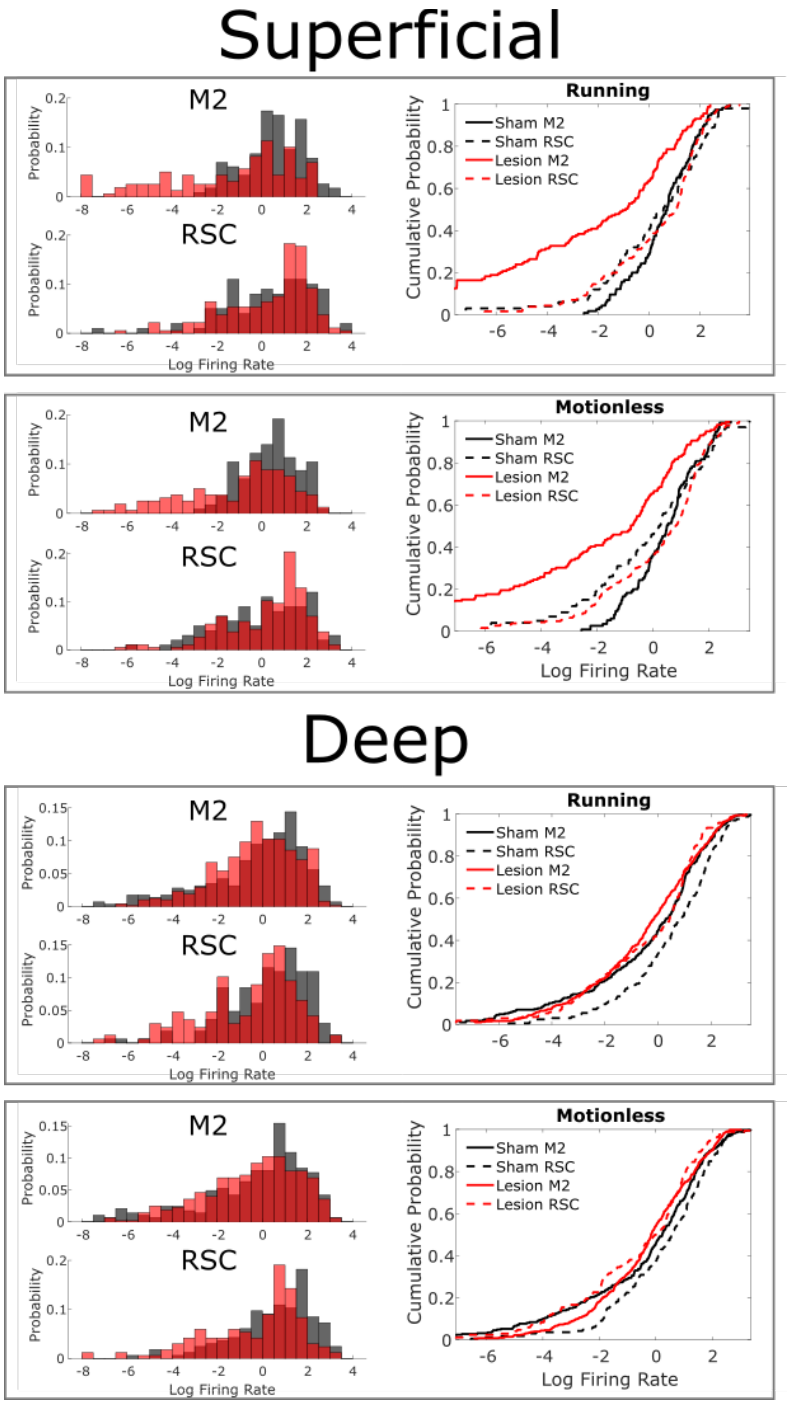

Figure 3.6. Firing rate distributions for neurons at different depths. The top two boxes show distributions from neurons on the top half of the probes ("superficial" cortex). The bottom are from the bottom half of the probes ("deep cortex). Histograms on left show the probabilities for log firing rates for M2 and RSC. Top groups in each box are firing rates calculated while the mice are moving (>2cm/s). Bottom groups in each box are from still periods (<2cm/s)

Spatial Information M2 and RSC is Lower in Lesion Mice

Other studies have found that HPC lesions reduce the spatial selectivity of neurons in NC. To test if this was the case in the two areas of interest in this experiment, M2 and RSC, the spatial information score for every excitatory neuron was calculated, separately for each VR environment (Fig 3.7). However, because the reward areas so strongly modulated neuronal activity, especially in M2, we decided to remove this region from the spatial information calculation. To do this, 10 bins (~30cm) on either side of both reward regions (and the reward bin itself) were removed from the mean firing rate and occupancy distribution vectors, which are used in the spatial information calculation. Spatial information was computed based on the remaining firing and occupancy distribution vectors after reward area removal. Using this spatial information metric, the total distribution of spatial information scores from neurons in Sham mice was skewed more towards higher values than the distribution from neurons in Lesion mice, regardless of VR environment (i.e., Old or New). This is consistent with other findings showing that spatial selectivity is impaired following HPC lesion (Mao et al., 2018; Esteves et al., 2021).

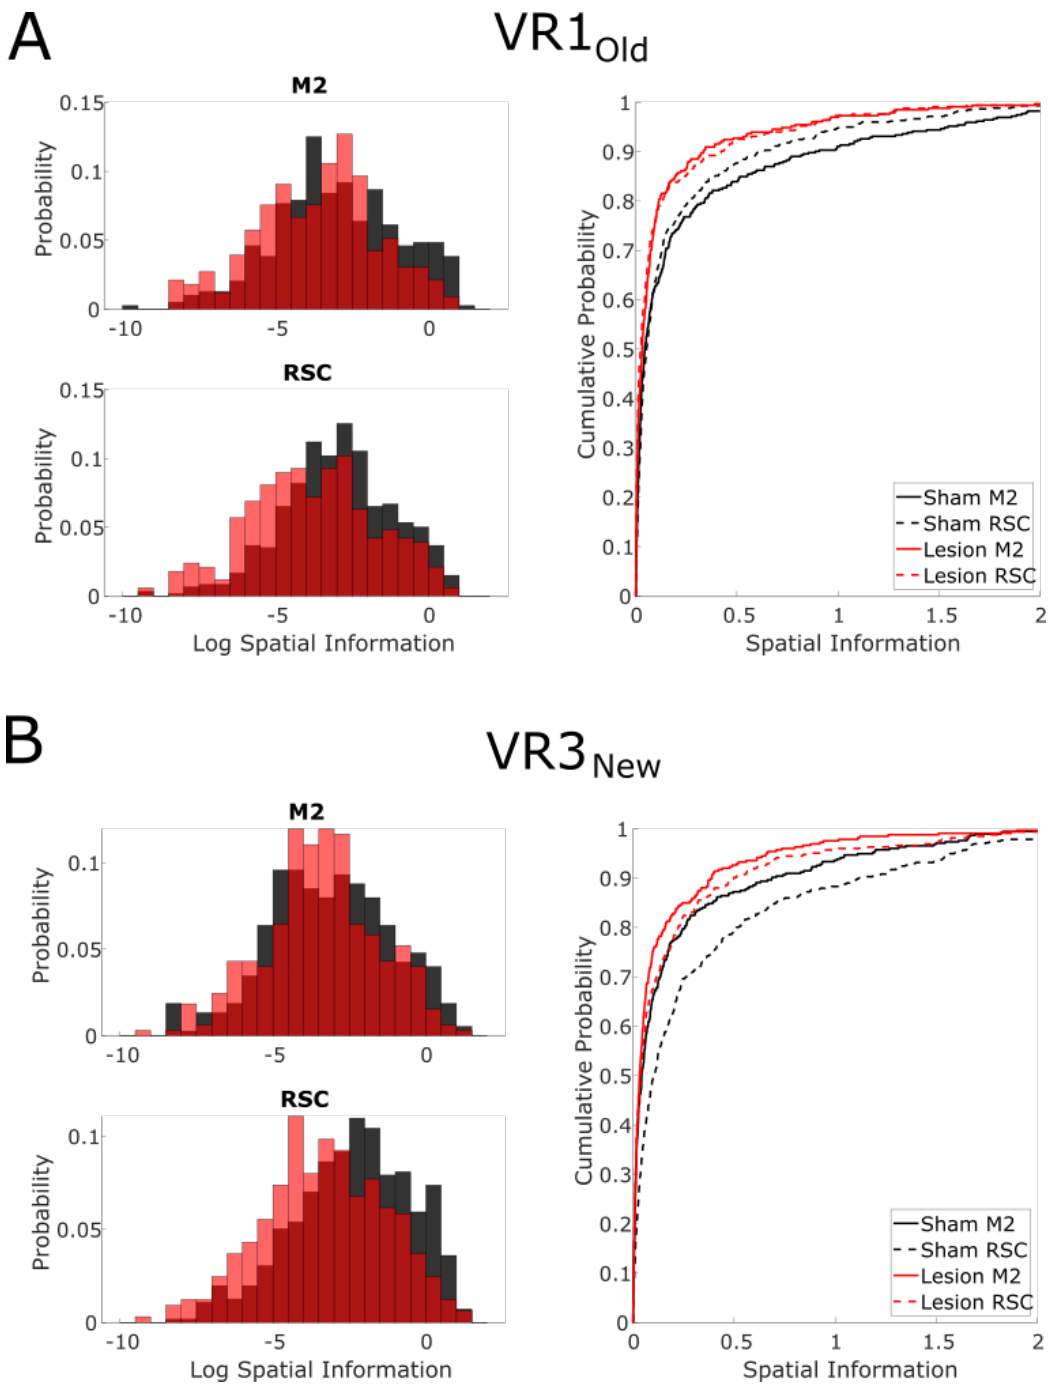

Figure 3.7. Spatial information distributions in old and new VR. When calculating the spatial information score, mean rate vectors were calculated for each neuron. Then, the 10 bins before and after each reward (+- ~30cm) were excluded. Spatial information was then calculated based on the mean rate vector. Cells were split between bottom (deep) and top (sup) half of channels. Only excitatory cells used here.

Firing rate changes are greater across VR environments in NC of HPC Lesion mice

Looking at the firing rate profiles of individual neurons, there were noticeable differences in maximum firing rates across VR environments (Fig. 3.8 A-D). To quantify the magnitude of firing rate change between VR environments, the lap-wise mean firing rate in the 20 laps flanking the VR environment change were examined (Fig 3.8 E). These 20 laps were used to minimize firing rate differences potentially due to drifting electrodes gaining/losing neurons. The mean firing rate from the 10 laps before (FR1) was subtracted from the 10 laps after (FR2) the VR change and that result divided by the sum of FR1 + FR2 to provide a measure of firing rate change between VRs varying from -1 to 1. Pooling neurons across brain regions and individual mice (n=819 from shams, 1275 from HPC lesions), neurons from lesioned mice showed greater firing rate differences across VR environments ($p<0.01$, Wilcoxon rank-sum test, Fig 3.8 F). This was unexpected, because it was anticipated that an intact HPC would bias NC towards remapping between the VR environments, which would yield greater firing rate changes in the sham group. However, as explained in detail in the conclusions section below, an intact HPC may not be remapping across VR environments because essential self-motion cues to update the path integrator and signal that the mouse is in a new spatial environment are missing in this head-fixed behavioral protocol.

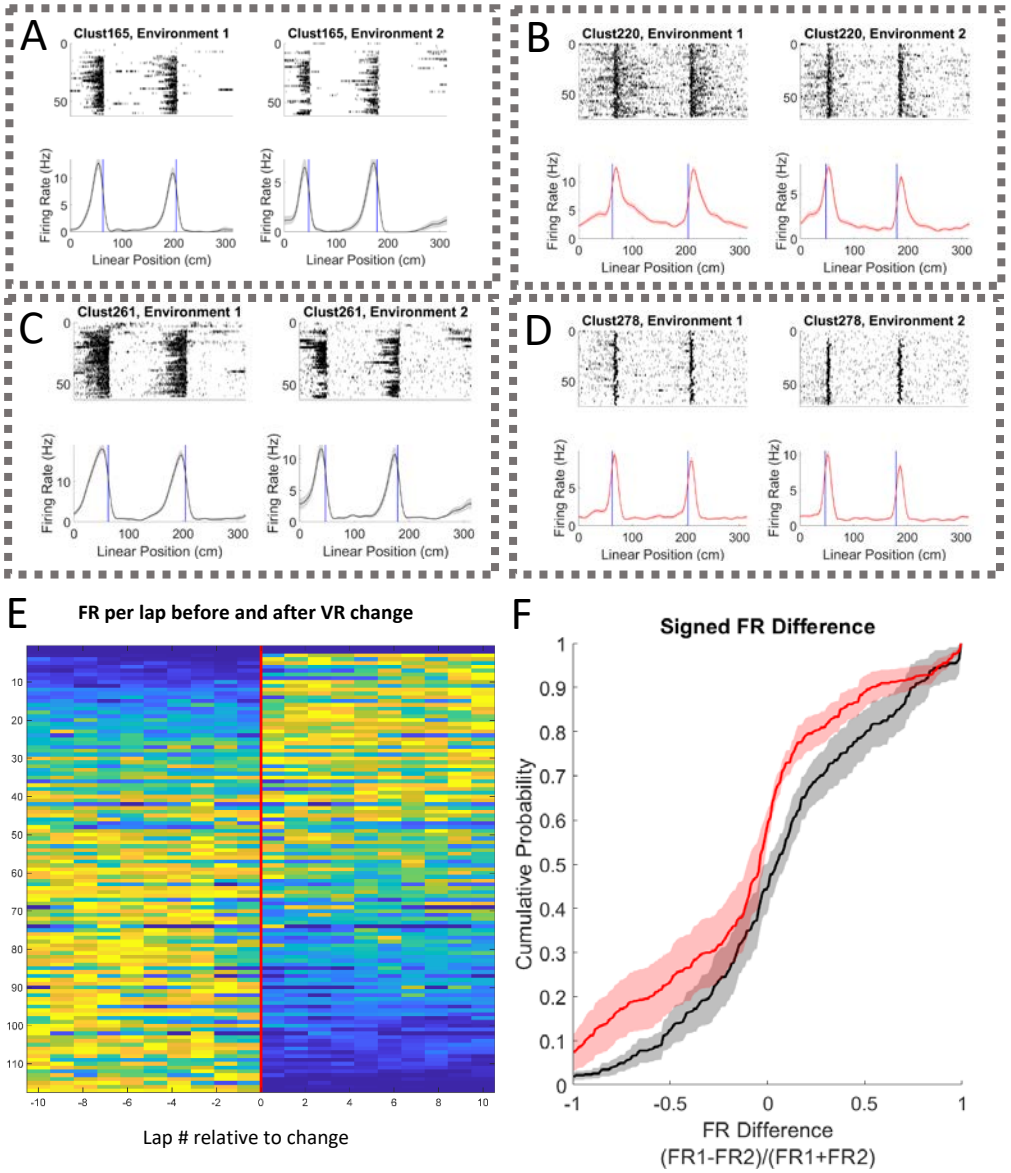

Figure 3.8. Firing rate changes across VR environments. A-D) Example firing rate profiles of 4 neurons. Top plots in each box show spike rasters for a single neuron in two VR environments. Bottom plots show mean firing rate as a function of position with reward sites marked by blue lines. Note that the maximum firing rate of these neurons changes across VR environments. A-D correspond to neurons from conditions-NC areas: M2-Sham, M2-Lesion, RSC-Sham, RSC-Lesion, respectively. E) Normalized lap-wise firing rate in the 10 laps before and after VR environment change from ~120 neurons in one HPC lesioned mouse. The vertical red line indicates VR environment change. F) Cumulative distributions of the signed firing rate difference score. Lesions show larger firing rate changes across VR environments ($p < 0.01$, Wilcoxon rank-sum test).

CHAPTER 4:

Conclusions and Future Directions

The objective of the experiments herein was to address a prediction derived from the of the hippocampal memory indexing theory (Teyler & DiScenna, 1986; McNaughton, 2010). Namely, we investigated whether there would be an “index code” in NC while mice navigate through a 1-dimensional virtual environment. Such an “index code” in NC is theoretically necessary for NC to “decompress” the information it receives from the major HPC output pathways, subiculum and EC, which use dense coding schemes to “compress” information to efficiently transmit it widely across NC. We also asked whether HPC lesions affect the NC “index code”, if there was a difference in index-like properties of cells across cortical depths, and what the nature of the NC “index code” in different regions (i.e., what behavioral or sensory variables affect the code).

4.1 Firing Rate Changes in NC Across VR Environments

In accordance with previous studies (Mao et al., 2017; Esteves et al., 2021), there are unique orthogonal neuronal ensembles active in NC while mice navigate along a 1-dimensional track (Fig 3.4A), which is what is expected from an “index code”. This is similar to what is observed in HPC: at each different location in an environment, a different set of neurons are active (O’Keefe & Dostrovsky, 1971; Wilson & McNaughton, 1993). However, at least in this visual VR environment, NC responses seem to be driven by behavioral variables (reward, anticipation, licking), too. Although HPC activity was not recorded in this experiment, it would be expected not to respond so robustly to

such input and instead depend heavily on path integration computations derived from self-motion signals like it does in freely moving protocols (McNaughton et al., 2006).

In this study, responses of many NC neurons seemed to be anchored to the reward locations in both the old and new VR environments. Furthermore, many of these reward-area-responsive neurons changed firing rate across VR environments even though they still responded at the same position relative to the reward. This may be a NC instantiation of so-called ‘rate-remapping’ seen in HPC, which occurs when external cues in the environment are changed, but an animal’s position remains constant (Leutgeb et al., 2005). This effect appears to be even greater in NC of HPC lesioned relative to sham mice. Theoretically, this could be due to HPC in sham mice not remapping between VR environments. This is a strong possibility given that the brain’s path integrator may not reset in this experimental protocol because the mouse stays head-fixed in place while visual VR environments are switched. Additional HPC recordings in this experimental protocol are necessary to determine whether or not this is true.

In a similar experiment, Saleem et al. (2018) found that V1 and CA1 neurons respond fairly sparsely in a similar visual VR behavioral task, that is, they often have only one field along a 100cm VR track. Furthermore, while V1 activity is sparse and orthogonal along the length of the track and can be used to predict the mouse’s position, CA1 activity shows a more uniform distribution across the track’s length and can be used for more accurate position decoding. Interestingly, V1 neurons also show a preference for visually stimuli portions of the track, but the peak firing within those similar portions is of a different magnitude, indicating some more global positional

information is present, potentially coming from HPC. This is similar to the peak firing rate changes seen in the VR task in this dissertation, however, these seem to occur across VR environments, not across rewards in the same environment as in Saleem et al. (2018). Factors including the behavioral shaping protocol and length of VR track, which were different between these two experiments, may account for the observed differences in the distributions of spatial fields.

4.2 M2 Ramping Activity Before Rewards is Reduced with HPC Lesion

The experiment in this dissertation shows that many M2 neurons are highly active around reward areas of the visual virtual environments. Further analyses are needed to determine whether these neurons are responsive to the behavioral actions around the rewards (e.g., slowing down, licking in anticipation or to consume reward after delivery, accelerating after reward consumption), the presentation of the reward itself (clicking of the solenoid, odor/taste/sight of reward, etc.), or anticipation/expectation of the reward. Overall M2 activity in sham mice ramps up before rewards, which is consistent with many earlier findings that M2 shows early choice-related activity in memory tasks (Barthas & Kwan, 2017). Interestingly, this ramping activity in anticipation of reward is reduced in HPC lesioned mice, which corresponds with the observed deficit in anticipatory slowing.

Additional analyses of these data will address whether the M2 ramping activity is simply a reflection of behavioral changes around the reward (e.g., anticipatory licking, acceleration/deceleration, or speed around rewarded areas). For example, by comparing activity in sections of the track before and after the reward, the degree of

speed modulation can be assessed because there will be points at which speed is the same but position relative to the reward is different. Such an approach would allow for conclusions about reward anticipation versus speed modulation to be made. Similar procedures will be carried out to assess the degree to which M2 activity is affected by licking, acceleration, deceleration, etc.

4.3 Future Directions and Analyses

Because this study is the first to combine electrophysiological recording of NC ensembles and HPC lesions, there are many broad questions that can potentially be answered with these data. For example, in addition to reductions in spatially selectivity of NC neurons (Mao et al., 2018; Esteves et al., 2021), perhaps HPC lesions affect baseline NC dynamics such as firing rate, LFP power and/or coherence, pairwise cross-correlations of neurons within or between NC regions. These questions will be addressed with further analysis of data presented here. Several questions and the planned approaches to answering them are outlined below.

Prevalent models of memory assert that HPC and NC interact to allow for memory consolidation and retrieval (Marr, 1971; McNaughton, 1983; Teyler & DiScenna, 1986; Buzsaki, 1989). Many of these models point to bursts of activity in HPC, known as sharp wave-ripples, which coincide with NC “up-states” (Battaglia et al., 2004), as instances in which HPC and NC interact to reactivate, and thus consolidate, memory traces. Although some work shows that HPC ripples can lead and lag NC “up-states” (Karimi Abadchi et al., 2020), it is unclear if and how HPC organizes NC activity in any way. Data from the present experiment can begin to answer this question.

Because HPC is lesioned in the experimental group and intact in the shams, NC activity in “offline” conditions such as quiet rest or light anesthesia can be compared between groups to determine the effect of HPC lesion on NC up-states. These up-states, for example, may become less diverse and therefore show less hierarchical clustering, reflecting replay of only already-consolidated memory traces in NC rather than interleaving with new memory traces corresponding to recent memories that still depend on HPC for reactivation. LFP coherence across layers or regions in the slow oscillation (0.5-4Hz), theta (4-8Hz), or spindle (8-20Hz) frequencies may be disrupted without HPC to organize the spatially broadly distributed memory trace in NC. These possibilities will be tested with data from the present experiment.

There is another fundamental question about the hippocampal memory indexing theory that remains unanswered: does HPC actually orchestrate the reactivation of NC activity patterns? Experimentally this could be asked in this form: “does removing or inactivating HPC prevent the reinstatement of appropriate spatiotemporal patterns widely across NC?” An experiment would have to meet the following criteria to address that question: 1) elicits memory retrieval/reactivation 2) remove or inactivate HPC at the appropriate time 3) record broadly across multiple different NC areas. The experiment described herein can address this question because it meets these criteria. HPC, being lesioned in the experimental groups, is of course impaired in this experiment. The neural recordings took place in M2 and RSC which are two associational NC areas that would be expected to require a coordinating force to reactivate coherently, at least for relatively new, unconsolidated memories. The criterion that this experiment does not quite meet is the requirement of memory retrieval/reactivation. Because the behavioral

protocol did not explicitly tax spatial memory, for instance by requiring an operant response to receive reward, the mice did not need to retrieve a memory. To receive reward, they could have simply run and licked all along the track until the reward was distributed without any attention paid to the VR environment nor their specific location within it. Furthermore, even though it does appear that Sham mice did remember the reward locations based on their anticipatory slowing, there is no discrete time or position when memory retrieval would likely occur. This makes it difficult to determine when to expect memory retrieval occur. Perhaps memory retrieval is occurring all along the track.

One study that addressed the role of HPC in the reactivation of NC patterns is Cowansage et al. (2014). By using a *c-fos*-based genetic tagging system to selectively tag and activate cells engaged during fear conditioning, they showed that reactivation of this specific set of cells was sufficient to reinstate the fear memory, even when HPC was pharmacologically inactivated. Furthermore, they went on to show that the same ensembles in two amygdalar nuclei, the basal and central, which are downstream of RSC, were activated when RSC was artificially stimulated. The results make a powerful, convincing argument for the necessity of HPC in informing NC activity. However, there are a few aspects of the indexing theory that this study does not address. Firstly, the indexing theory predicts that memory retrieval involves a spatially broad reactivation of neuronal patterns across most, if not all, of NC. This study only observes a reactivation of ensembles in RSC and the amygdala, a subcortical structure. Also, the time course of the activity is insufficient to rigorously test the theory. The index theory predicts, specifically, that NC areas will be reactivated in a specific spatiotemporal order. To do

that, one would need to record activity with superior time resolution, ideally using single-unit electrophysiology like that used here.

The experimental protocol used in this dissertation ended with administering urethane anesthesia immediately after running in the second VR environment. Because urethane induces cyclic alternations between REM and SWS-like brain states, it can be used as a model for natural sleep (Clement et al., 2008), which is when the spontaneous reactivation thought to support memory consolidation occurs. Spontaneous and evoked patterns of activity have actually already been shown to appear repeatedly under urethane anesthesia (Luczak et al., 2009; Bermudez Contreras et al., 2013) although it has never been shown that patterns active during waking are then reactivated while subsequently under urethane anesthesia. Data from this experiment will address that issue. Ultimately, this data may provide an avenue for testing this untested prediction of the hippocampal memory indexing theory: that the HPC *coordinates* reactivation of precise spatiotemporal sequences distributed across NC. Assuming reactivation between M2 and RSC (e.g., correlations between inter-area pairs of neurons or sequences involving neurons from both areas) does occur under urethane anesthesia, this experiment can examine if and how HPC is involved in it. This would be ideal for answering that longstanding prediction of the hippocampal memory indexing theory.

REFERENCES

Ackermann, S., & Rasch, B. (2014). Differential effects of non-REM and REM sleep on memory consolidation. *Current neurology and neuroscience reports*, *14*(2), 430.

Aghajan, Z. M., Acharya, L., Moore, J. J., Cushman, J. D., Vuong, C., & Mehta, M. R. (2015). Impaired spatial selectivity and intact phase precession in two-dimensional virtual reality. *Nature neuroscience*, *18*(1), 121-128.

Alexander, A. S., & Nitz, D. A. (2015). Retrosplenial cortex maps the conjunction of internal and external spaces. *Nature neuroscience*, *18*(8), 1143-1151.

Alexander, A. S., & Nitz, D. A. (2017). Spatially periodic activation patterns of retrosplenial cortex encode route sub-spaces and distance traveled. *Current Biology*, *27*(11), 1551-1560.

Amaral, D. G., Ishizuka, N., & Claiborne, B. (1990). Chapter Neurons, numbers and the hippocampal network. *Progress in brain research*, *83*, 1-11.

Amzica, F., & Steriade, M. (1995). Short-and long-range neuronal synchronization of the slow (< 1 Hz) cortical oscillation. *Journal of neurophysiology*, *73*(1), 20-38.

Amzica, F., & Steriade, M. (1997). The K-complex: its slow (< 1-Hz) rhythmicity and relation to delta waves. *Neurology*, *49*(4), 952-959.

Andersen, P., Sundberg, S. H., Sveen, O., & Wigström, H. (1977). Specific long-lasting potentiation of synaptic transmission in hippocampal slices. *Nature*, *266*(5604), 736-737.

Andrillon, T., Nir, Y., Staba, R. J., Ferrarelli, F., Cirelli, C., Tononi, G., & Fried, I. (2011). Sleep spindles in humans: insights from intracranial EEG and unit recordings. *Journal of Neuroscience*, *31*(49), 17821-17834.

Aronov, D., & Tank, D. W. (2014). Engagement of neural circuits underlying 2D spatial navigation in a rodent virtual reality system. *Neuron*, *84*(2), 442-456.

Aserinsky, E., & Kleitman, N. (1953). Regularly occurring periods of eye motility, and concomitant phenomena, during sleep. *Science*, *118*(3062), 273-274.

Aubrey, J., Smith, C., Tweed, S., & Nader, R. (1999). Cognitive and motor procedural tasks are dissociated in REM and stage two sleep. *Sleep Res. Online*, *2*(Suppl 1), 220.

Axmacher, N., Elger, C. E., & Fell, J. (2008). Ripples in the medial temporal lobe are relevant for human memory consolidation. *Brain*, *131*(7), 1806-1817.

Barnes, C. A., McNaughton, B. L., Mizumori, S. J., Leonard, B. W., & Lin, L. H. (1990). Chapter Comparison of spatial and temporal characteristics of neuronal activity in sequential stages of hippocampal processing. *Progress in brain research*, *83*, 287-300.

Barrett, T. R., & Ekstrand, B. R. (1972). Effect of sleep on memory: III. Controlling for time-of-day effects. *Journal of experimental psychology*, *96*(2), 321.

Barthas, F., & Kwan, A. C. (2017). Secondary motor cortex: where 'sensory'meets 'motor'in the rodent frontal cortex. *Trends in neurosciences*, *40*(3), 181-193.

Battaglia, F. P., Sutherland, G. R., & McNaughton, B. L. (2004). Hippocampal sharp wave bursts coincide with neocortical "up-state" transitions. *Learning & Memory*, *11*(6), 697-704.

Bean, C. H. (1912). The curve of forgetting. Press of the New era printing Company.

Berendse, H. W., Graaf, Y. G. D., & Groenewegen, H. J. (1992). Topographical organization and relationship with ventral striatal compartments of prefrontal

corticostriatal projections in the rat. *Journal of Comparative Neurology*, *316*(3), 314-347.

Berger H. (1929). Uber das Elektrenkephalogramm des Menschen. *Archiv Psychiat Nervenkr*; 87, 527–570

Bird, C. M., & Burgess, N. (2008). The hippocampus and memory: insights from spatial processing. *Nature Reviews Neuroscience*, *9*(3), 182-194.

Burke, S. N., Chawla, M. K., Penner, M. R., Crowell, B. E., Worley, P. F., Barnes, C. A., & McNaughton, B. L. (2005). Differential encoding of behavior and spatial context in deep and superficial layers of the neocortex. *Neuron*, *45*(5), 667-674.

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

Buzsáki, G. (1996). The hippocampo-neocortical dialogue. *Cerebral cortex*, *6*(2), 81-92.

Buzsáki, G. (2015). Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus*, *25*(10), 1073-1188.

Buzsáki, G., & Tingley, D. (2018). Space and time: The hippocampus as a sequence generator. *Trends in cognitive sciences*, *22*(10), 853-869.

Cain, D. P., Humpartzoomian, R., & Boon, F. (2006). Retrosplenial cortex lesions impair water maze strategies learning or spatial place learning depending on prior experience of the rat. *Behavioural brain research*, *170*(2), 316-325.

Carskadon, M. A., & Dement, W. C. (2005). Normal human sleep: an overview. *Principles and practice of sleep medicine*, *4*, 13-23.

Caton, R. (1875). Electrical Currents of the Brain. *The Journal of Nervous and Mental Disease*, *2*(4), 610.

Chawla, M. K., Guzowski, J. F., Ramirez-Amaya, V., Lipa, P., Hoffman, K. L., Marriott, L. K., ... & Barnes, C. A. (2005). Sparse, environmentally selective expression of Arc RNA in the upper blade of the rodent fascia dentata by brief spatial experience. *Hippocampus*, *15*(5), 579-586.

Chen, G., King, J. A., Lu, Y., Cacucci, F., & Burgess, N. (2018). Spatial cell firing during virtual navigation of open arenas by head-restrained mice. *Elife*, *7*, e34789.

Chen, G., Lu, Y., King, J. A., Cacucci, F., & Burgess, N. (2019). Differential influences of environment and self-motion on place and grid cell firing. *Nature communications*, *10*(1), 1-11.

Chen, L. L., Lin, L. H., Green, E. J., Barnes, C. A., & McNaughton, B. L. (1994). Head-direction cells in the rat posterior cortex. *Experimental brain research*, *101*(1), 8-23.

Cho, J., & Sharp, P. E. (2001). Head direction, place, and movement correlates for cells in the rat retrosplenial cortex. *Behavioral neuroscience*, *115*(1), 3.

Cirelli, C., & Tononi, G. (2008). Is sleep essential?. *PLoS biology*, *6*(8), e216.

Clemens, Z., Fabo, D., & Halasz, P. (2005). Overnight verbal memory retention correlates with the number of sleep spindles. *Neuroscience*, *132*(2), 529-535.

Clemens, Z., Fabó, D., & Halász, P. (2006). Twenty-four hours retention of visuospatial memory correlates with the number of parietal sleep spindles. *Neuroscience letters*, *403*(1-2), 52-56.

Clement, E. A., Richard, A., Thwaites, M., Ailon, J., Peters, S., & Dickson, C. T. (2008). Cyclic and sleep-like spontaneous alternations of brain state under urethane anaesthesia. *PloS one*, *3*(4), e2004.

Contreras, D., Destexhe, A., Sejnowski, T. J., & Steriade, M. (1997). Spatiotemporal patterns of spindle oscillations in cortex and thalamus. *Journal of Neuroscience*, *17*(3), 1179-1196.

Contreras, E. J. B., Schjetnan, A. G. P., Muhammad, A., Bartho, P., McNaughton, B. L., Kolb, B., ... & Luczak, A. (2013). Formation and reverberation of sequential neural activity patterns evoked by sensory stimulation are enhanced during cortical desynchronization. *Neuron*, *79*(3), 555-566.

Cooper, B. G., & Mizumori, S. J. (2001). Temporary inactivation of the retrosplenial cortex causes a transient reorganization of spatial coding in the hippocampus. *Journal of Neuroscience*, *21*(11), 3986-4001.

Cowansage, K. K., Shuman, T., Dillingham, B. C., Chang, A., Golshani, P., & Mayford, M. (2014). Direct reactivation of a coherent neocortical memory of context. *Neuron*, *84*(2), 432-441.

Cowey, A., & Bozˆek, T. (1974). Contralateral 'neglect'after unilateral dorsomedial prefrontal lesions in rats. *Brain Research*, *72*(1), 53-63.

Crick, F., & Mitchison, G. (1983). The function of dream sleep. *Nature*, *304*(5922), 111-114.

Crick, F., & Mitchison, G. (1986). REM sleep and neural nets. *The Journal of Mind and Behavior*, 229-249.

Crowne, D. P., & Pathria, M. N. (1982). Some attentional effects of unilateral frontal lesions in the rat. *Behavioural Brain Research*, *6*(1), 25-39.

Csicsvari, J., Hirase, H., Czurko, A., Mamiya, A., & Buzsáki, G. (1999). Fast network oscillations in the hippocampal CA1 region of the behaving rat. *J Neurosci*, *19*(RC20), 1-4.

Czajkowski, R., Jayaprakash, B., Wiltgen, B., Rogerson, T., Guzman-Karlsson, M. C., Barth, A. L., ... & Silva, A. J. (2014). Encoding and storage of spatial information in the retrosplenial cortex. *Proceedings of the National Academy of Sciences*, *111*(23), 8661-8666.

De Kock, C. P. J., Bruno, R. M., Spors, H., & Sakmann, B. (2007). Layer-and cell-type-specific suprathreshold stimulus representation in rat primary somatosensory cortex. *The Journal of physiology*, *581*(1), 139-154.

de Souza, R. T. F., Gerhardt, G. J. L., Schönwald, S. V., Rybarczyk-Filho, J. L., & Lemke, N. (2016). Synchronization and propagation of global sleep spindles. *PloS one*, *11*(3), e0151369.

Dement, W., & Kleitman, N. (1957). Cyclic variations in EEG during sleep and their relation to eye movements, body motility, and dreaming. *Electroencephalography and clinical neurophysiology*, *9*(4), 673-690.

Deshmukh, S. S., & Knierim, J. J. (2013). Influence of local objects on hippocampal representations: Landmark vectors and memory. *Hippocampus*, *23*(4), 253-267.

Diekelmann, S., & Born, J. (2010). The memory function of sleep. *Nature Reviews Neuroscience*, *11*(2), 114.

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.

Donoghue, J. P., & Wise, S. P. (1982). The motor cortex of the rat: cytoarchitecture and microstimulation mapping. *Journal of Comparative Neurology*, *212*(1), 76-88.

Ebbinghaus, H. (1913). Memory: A Contribution to Experimental Psychology (HA Ruger & CE Bussenius, Trans.). New York: Teachers College.(Original work published 1885), 39.

Ego-Stengel, V., & Wilson, M. A. (2010). Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. *Hippocampus*, *20*(1), 1-10.

Ekstrand, B. R. (1967). Effect of Sleep on Memory. Journal of Experimental Psychology, 75(1), 64–72.

Erlich, J. C., Bialek, M., & Brody, C. D. (2011). A cortical substrate for memory-guided orienting in the rat. *Neuron*, *72*(2), 330-343.

Erlich, J. C., Brunton, B. W., Duan, C. A., Hanks, T. D., & Brody, C. D. (2015). Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. *Elife*, *4*, e05457.

Eschenko, O., Mölle, M., Born, J., & Sara, S. J. (2006). Elevated sleep spindle density after learning or after retrieval in rats. *Journal of Neuroscience*, *26*(50), 12914-12920.

Eschenko, O., Ramadan, W., Mölle, M., Born, J., & Sara, S. J. (2008). Sustained increase in hippocampal sharp-wave ripple activity during slow-wave sleep after learning. *Learning & Memory*, *15*(4), 222-228.

Esteves, I. M., Chang, H., Neumann, A. R., Sun, J., Mohajerani, M. H., & McNaughton, B. L. (2021). Spatial Information Encoding across Multiple Neocortical Regions Depends on an Intact Hippocampus. *Journal of Neuroscience*, *41*(2), 307-319.

Euston, D. R., & McNaughton, B. L. (2006). Apparent encoding of sequential context in rat medial prefrontal cortex is accounted for by behavioral variability. *Journal of Neuroscience*, *26*(51), 13143-13155.

Finkenbinder, E. O. (1913). The curve of forgetting. *The American Journal of Psychology*, *24*(1), 8-32.

Fogel, S. M., & Smith, C. T. (2006). Learning-dependent changes in sleep spindles and stage 2 sleep. *Journal of sleep research*, *15*(3), 250-255.

Foster, T. C., Castro, C. A., & McNaughton, B. L. (1989). Spatial selectivity of rat hippocampal neurons: dependence on preparedness for movement. *Science*, *244*(4912), 1580-1582.

Foucault, M. (1913). Introduction à la psychologie de la perception. Experiénces sur l'oubli ou sur l'inhibition regressive. Rev. des Cours et Conférences, 21, 444-454.

Fowler, M. J., Sullivan, M. J., & Ekstrand, B. R. (1973). Sleep and memory. *Science*, *179*(4070), 302-304.

Fujisawa, S., Amarasingham, A., Harrison, M. T., & Buzsáki, G. (2008). Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nature neuroscience*, *11*(7), 823-833.

Gabbott, P. L., Warner, T. A., Jays, P. R., Salway, P., & Busby, S. J. (2005). Prefrontal cortex in the rat: projections to subcortical autonomic, motor, and limbic centers. *Journal of Comparative Neurology*, *492*(2), 145-177.

Gais, S., & Born, J. (2004). Declarative memory consolidation: mechanisms acting during human sleep. *Learning & Memory*, *11*(6), 679-685.

Geiller, T., Fattahi, M., Choi, J. S., & Royer, S. (2017). Place cells are more strongly tied to landmarks in deep than in superficial CA1. *Nature communications*, *8*(1), 1-11.

Gemmell, C., Anderson, M., & O'Mara, S. M. (2002). Deep layer prefrontal cortex unit discharge in a cue-controlled open-field environment in the freely-moving rat. *Behavioural brain research*, *133*(1), 1-10.

Gill, P. R., Mizumori, S. J., & Smith, D. M. (2011). Hippocampal episode fields develop with learning. *Hippocampus*, *21*(11), 1240-1249.

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.

Graves, E. A. (1936). The effect of sleep upon retention. *Journal of Experimental Psychology*, *19*(3), 316.

Hall, R. D., & Lindholm, E. P. (1974). Organization of motor and somatosensory neocortex in the albino rat. *Brain research*, *66*(1), 23-38.

Hartmann, E. (1967). *The biology of dreaming* (Vol. 2). Charles C Thomas Pub Ltd.

Harvey, C. D., Collman, F., Dombeck, D. A., & Tank, D. W. (2009). Intracellular dynamics of hippocampal place cells during virtual navigation. *Nature*, *461*(7266), 941-946.

Hebb, D. O. (1949). *The organisation of behaviour: a neuropsychological theory*. New York: Science Editions.

Heine, R. *Zeitschrift für Psychologie und Physiologie der Sinnesorgane* 68, 161–236 (1914).

Hoover, W. B., & Vertes, R. P. (2007). Anatomical analysis of afferent projections to the medial prefrontal cortex in the rat. *Brain Structure and Function*, *212*(2), 149-179.

Huber, R., Ghilardi, M. F., Massimini, M., & Tononi, G. (2004). Local sleep and learning. *Nature*, *430*(6995), 78.

Isomura, Y., Sirota, A., Özen, S., Montgomery, S., Mizuseki, K., Henze, D. A., & Buzsáki, G. (2006). Integration and segregation of activity in entorhinal-hippocampal subregions by neocortical slow oscillations. *Neuron*, *52*(5), 871-882.

Jadhav, S. P., Kemere, C., German, P. W., & Frank, L. M. (2012). Awake hippocampal sharp-wave ripples support spatial memory. *Science*, 1217230.

Jenkins, J. G., & Dallenbach, K. M. (1924). Obliviscence during sleep and waking. *The American Journal of Psychology*, *35*(4), 605-612.

Jeong, M., Kim, Y., Kim, J., Ferrante, D. D., Mitra, P. P., Osten, P., & Kim, D. (2016). Comparative three-dimensional connectome map of motor cortical projections in the mouse brain. *Scientific reports*, *6*(1), 1-14.

Jinno, Shozo, et al. "Neuronal diversity in GABAergic long-range projections from the hippocampus." *Journal of Neuroscience* 27.33 (2007): 8790-8804.

Johnson, L. A., Euston, D. R., Tatsuno, M., & McNaughton, B. L. (2010). Stored-trace reactivation in rat prefrontal cortex is correlated with down-to-up state fluctuation density. *Journal of Neuroscience*, *30*(7), 2650-2661.

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

Jung, M. W., Qin, Y., McNaughton, B. L., & Barnes, C. A. (1998). Firing characteristics of deep layer neurons in prefrontal cortex in rats performing spatial working memory tasks. *Cerebral cortex (New York, NY: 1991)*, *8*(5), 437-450.

Kandel, A., & Buzsáki, G. (1997). Cellular–synaptic generation of sleep spindles, spike-and-wave discharges, and evoked thalamocortical responses in the neocortex of the rat. *Journal of Neuroscience*, *17*(17), 6783-6797.

Kargo, W. J., Szatmary, B., & Nitz, D. A. (2007). Adaptation of prefrontal cortical firing patterns and their fidelity to changes in action–reward contingencies. *Journal of Neuroscience*, *27*(13), 3548-3559.

Karimi Abadchi, J. K., Nazari-Ahangarkolaee, M., Gattas, S., Bermudez-Contreras, E., Luczak, A., McNaughton, B. L., & Mohajerani, M. H. (2020). Spatiotemporal patterns of neocortical activity around hippocampal sharp-wave ripples. *Elife*, *9*, e51972.

Keene, C. S., & Bucci, D. J. (2008). Contributions of the retrosplenial and posterior parietal cortices to cue-specific and contextual fear conditioning. *Behavioral neuroscience*, *122*(1), 89.

Khodagholy, D., Gelinas, J. N., & Buzsáki, G. (2017). Learning-enhanced coupling between ripple oscillations in association cortices and hippocampus. *Science*, *358*(6361), 369-372.

Konorski, J. (1948). Conditioned reflexes and neuron organization.

Kraus, B. J., Robinson II, R. J., White, J. A., Eichenbaum, H., & Hasselmo, M. E. (2013). Hippocampal “time cells”: time versus path integration. *Neuron*, *78*(6), 1090-1101.

Kudrimoti, H. S., Barnes, C. A., & McNaughton, B. L. (1999). Reactivation of hippocampal cell assemblies: effects of behavioral state, experience, and EEG dynamics. *Journal of Neuroscience*, *19*(10), 4090-4101.

Lehmann, H., Sparks, F. T., Spanswick, S. C., Hadikin, C., McDonald, R. J., & Sutherland, R. J. (2009). Making context memories independent of the hippocampus. *Learning & Memory*, *16*(7), 417-420.

Leonard, C. M. (1969). The prefrontal cortex of the rat. I. Cortical projection of the mediodorsal nucleus. II. Efferent connections. *Brain research*, *12*(2), 321-343.

Leutgeb, S., Leutgeb, J. K., Barnes, C. A., Moser, E. I., McNaughton, B. L., & Moser, M. B. (2005). Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science*, *309*(5734), 619-623.

Leutgeb, S., Leutgeb, J. K., Treves, A., Moser, M. B., & Moser, E. I. (2004). Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science*, *305*(5688), 1295-1298.

Levy, W. B., & Steward, O. (1979). Synapses as associative memory elements in the hippocampal formation. *Brain research*, *175*(2), 233-245.

Loomis, A. L., Harvey, E. N., & Hobart, G. (1935a). Potential rhythms of the cerebral cortex during sleep. *Science*.

Loomis, A. L., Harvey, E. N., & Hobart, G. (1935b). Further observations on the potential rhythms of the cerebral cortex during sleep. *Science*, *82*(2122), 198-200.

Luczak, A., Barthó, P., & Harris, K. D. (2009). Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron*, *62*(3), 413-425.

Luczak, A., Barthó, P., Marguet, S. L., Buzsáki, G., & Harris, K. D. (2007). Sequential structure of neocortical spontaneous activity in vivo. *Proceedings of the National Academy of Sciences*, *104*(1), 347-352.

Lynch, G. S., Dunwiddie, T., & Gribkoff, V. (1977). Heterosynaptic depression: a postsynaptic correlate of long-term potentiation. *Nature*, *266*(5604), 737-739.

MacDonald, C. J., Carrow, S., Place, R., & Eichenbaum, H. (2013). Distinct hippocampal time cell sequences represent odor memories in immobilized rats. *Journal of Neuroscience*, *33*(36), 14607-14616.

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

Maingret, N., Girardeau, G., Todorova, R., Goutierre, M., & Zugaro, M. (2016). Hippocampo-cortical coupling mediates memory consolidation during sleep. *Nature neuroscience*, *19*(7), 959.

Mao, D., Kandler, S., McNaughton, B. L., & Bonin, V. (2017). Sparse orthogonal population representation of spatial context in the retrosplenial cortex. *Nature communications*, *8*(1), 1-9.

Mao, D., Neumann, A. R., Sun, J., Bonin, V., Mohajerani, M. H., & McNaughton, B. L. (2018). Hippocampus-dependent emergence of spatial sequence coding in retrosplenial cortex. *Proceedings of the National Academy of Sciences*, *115*(31), 8015-8018.

Marr D. A theory of the cerebellar cortex. J Physiol. 1969;202:437–470.

Marr, D. (1971) Simple memory: a theory of archicortex. *Phil. Trans. R. Soc. Lond B Biol Sci., 262*(841), 23-81.

Marshall, L., Helgadóttir, H., Mölle, M., & Born, J. (2006). Boosting slow oscillations during sleep potentiates memory. *Nature*, *444*(7119), 610.

McCormick, D. A., & Bal, T. (1997). Sleep and arousal: thalamocortical mechanisms. *Annual review of neuroscience*, *20*(1), 185-215.

McNaughton, B. L. (2010). Cortical hierarchies, sleep, and the extraction of knowledge from memory. *Artificial Intelligence*, *174*(2), 205-214.

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

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

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

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

McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., & Moser, M. B. (2006). Path integration and the neural basis of the ‘cognitive map'. *Nature Reviews Neuroscience*, *7*(8), 663-678.

McNaughton, B. L., Douglas, R. M., & Goddard, G. V. (1978). Synaptic enhancement in fascia dentata: cooperativity among coactive afferents. *Brain research*, *157*(2), 277-293.

Mednick, S. C., McDevitt, E. A., Walsh, J. K., Wamsley, E., Paulus, M., Kanady, J. C., & Drummond, S. P. (2013). The critical role of sleep spindles in hippocampal-dependent memory: a pharmacology study. *Journal of Neuroscience*, *33*(10), 4494-4504.

Miyashita, T., & Rockland, K. S. (2007). GABAergic projections from the hippocampus to the retrosplenial cortex in the rat. *European Journal of Neuroscience*, *26*(5), 1193-1204.

Mölle, M., Marshall, L., Gais, S., & Born, J. (2004). Learning increases human electroencephalographic coherence during subsequent slow sleep oscillations. *Proceedings of the National Academy of Sciences of the United States of America*, *101*(38), 13963-13968.

Monaghan, D. T., & Cotman, C. W. (1986). Identification and properties of N-methyl-D-aspartate receptors in rat brain synaptic plasma membranes. *Proceedings of the National Academy of Sciences*, *83*(19), 7532-7536.

Moruzzi, G., & Magoun, H. W. (1949). Brain stem reticular formation and activation of the EEG. *Electroencephalography and clinical neurophysiology*, *1*(1-4), 455-473.

Müller, G. E., & Pilzecker, A. (1900). *Experimentelle beiträge zur lehre vom gedächtniss* (Vol. 1). JA Barth.

Murakami, M., Vicente, M. I., Costa, G. M., & Mainen, Z. F. (2014). Neural antecedents of self-initiated actions in secondary motor cortex. *Nature neuroscience*, *17*(11), 1574-1582.

Nakazawa, K., Sun, L. D., Quirk, M. C., Rondi-Reig, L., Wilson, M. A., & Tonegawa, S. (2003). Hippocampal CA3 NMDA receptors are crucial for memory acquisition of one-time experience. *Neuron*, *38*(2), 305-315.

Navratilova, Z., Hoang, L. T., Schwindel, C. D., Tatsuno, M., & McNaughton, B. L. (2012). Experience-dependent firing rate remapping generates directional selectivity in hippocampal place cells. *Frontiers in neural circuits*, *6*, 6.

Neafsey, E. J., Bold, E. L., Haas, G., Hurley-Gius, K. M., Quirk, G., Sievert, C. F., & Terreberry, R. R. (1986). The organization of the rat motor cortex: a microstimulation mapping study. *Brain research reviews*, *11*(1), 77-96.

Nicolau, M. C., Akaârir, M., Gamundí, A., González, J., & Rial, R. V. (2000). Why we sleep: The evolutionary pathway to the mammalian sleep. Progress in Neurobiology, 62(4), 379–406.

O'Neill, J., Pleydell-Bouverie, B., Dupret, D., & Csicsvari, J. (2010). Play it again: reactivation of waking experience and memory. *Trends in neurosciences*, *33*(5), 220-229.

Ocampo, A. C., Squire, L. R., & Clark, R. E. (2017). Hippocampal area CA1 and remote memory in rats. *Learning & Memory*, *24*(11), 563-568.

Oh, S. W., Harris, J. A., Ng, L., Winslow, B., Cain, N., Mihalas, S., ... & Zeng, H. (2014). A mesoscale connectome of the mouse brain. *Nature*, *508*(7495), 207-214.

O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. *Brain research*.

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

Park, E., Dvorak, D., & Fenton, A. A. (2011). Ensemble place codes in hippocampus: CA1, CA3, and dentate gyrus place cells have multiple place fields in large environments. *PloS one*, *6*(7), e22349.

Pastalkova, E., Itskov, V., Amarasingham, A., & Buzsáki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. *Science*, *321*(5894), 1322-1327.

Patel, J., Schomburg, E. W., Berényi, A., Fujisawa, S., & Buzsáki, G. (2013). Local generation and propagation of ripples along the septotemporal axis of the hippocampus. *Journal of Neuroscience*, *33*(43), 17029-17041.

Paus, T. (2000). Functional anatomy of arousal and attention systems in the human brain. In *Progress in brain research* (Vol. 126, pp. 65-77). Elsevier.

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

Peyrache, A., Khamassi, M., Benchenane, K., Wiener, S. I., & Battaglia, F. P. (2009). Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nature neuroscience*, *12*(7), 919.

Poucet, B. (1997). Searching for spatial unit firing in the prelimbic area of the rat medial prefrontal cortex. *Behavioural brain research*, *84*(1-2), 151-159.

Racine, R. J., Milgram, N. W., & Hafner, S. (1983). Long-term potentiation phenomena in the rat limbic forebrain. *Brain research*, *260*(2), 217-231.

Radossawljewitsch, P. R. (1907). Das Behalten und Vergessen bei Kindern und Erwachsenen nach experimentellen Untersuchungen. E. Meumann (Ed.). Otto Nemnich.

Ramadan, W., Eschenko, O., & Sara, S. J. (2009). Hippocampal sharp wave/ripples during sleep for consolidation of associative memory. *PloS one*, *4*(8), e6697.

Rasch, B., & Born, J. (2013). About sleep's role in memory. *Physiological reviews*, *93*(2), 681-766.

Reep, R. L., & Corwin, J. V. (1999). Topographic organization of the striatal and thalamic connections of rat medial agranular cortex. *Brain research*, *841*(1-2), 43-52.

Reep, R. L., Corwin, J. V., Hashimoto, A., & Watson, R. T. (1987). Efferent connections of the rostral portion of medial agranular cortex in rats. *Brain research bulletin*, *19*(2), 203-221.

Reep, R. L., Goodwin, G. S., & Corwin, J. V. (1990). Topographic organization in the corticocortical connections of medial agranular cortex in rats. *Journal of Comparative Neurology*, *294*(2), 262-280.

Rich, P. D., Liaw, H. P., & Lee, A. K. (2014). Large environments reveal the statistical structure governing hippocampal representations. *Science*, *345*(6198), 814-817.

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.

Rolls, E. T. (1989). Functions of neuronal networks in the hippocampus and neocortex in memory. In *Neural models of plasticity* (pp. 240-265). Academic Press.

Rolls, E. T. (2008). *Memory, attention, and decision-making: A unifying computational neuroscience approach*. Oxford: Oxford University Press.

Sadowski, J. H., Jones, M. W., & Mellor, J. R. (2016). Sharp-wave ripples orchestrate the induction of synaptic plasticity during reactivation of place cell firing patterns in the hippocampus. *Cell reports*, *14*(8), 1916-1929.

Sakata, S., & Harris, K. D. (2009). Laminar structure of spontaneous and sensory-evoked population activity in auditory cortex. *Neuron*, *64*(3), 404-418.

Saleem, A. B., Diamanti, E. M., Fournier, J., Harris, K. D., & Carandini, M. (2018). Coherent encoding of subjective spatial position in visual cortex and hippocampus. *Nature*, *562*(7725), 124-127.

Schabus, M., Gruber, G., Parapatics, S., Sauter, C., Klösch, G., Anderer, P., ... & Zeitlhofer, J. (2004). Sleep spindles and their significance for declarative memory consolidation. *Sleep*, *27*(8), 1479-1485.

Schacter, D. L., Eich, J. E., & Tulving, E. (1978). Richard Semon's theory of memory. *Journal of Verbal Learning and Verbal Behavior*, *17*(6), 721-743.

Scoville, W. B. (1954). The limbic lobe in man. *Journal of neurosurgery*, *11*(1), 64-66.

Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of neurology, neurosurgery, and psychiatry*, *20*(1), 11.

Sejnowski, T. J. (1998). The computational neuroethology of sleep. In *New Neuroethology on the Move: Proceedings of the 26th Gottingen Neurobiology Conference. New York: Georg Theime Verlag Stuttgart*.

Sejnowski, T. J., & Destexhe, A. (2000). Why do we sleep? 1. *Brain research*, *886*(1-2), 208-223.

Semon, R. The nmeme. London: George Allen & Unwin, 1921.

Sesack, S. R., Deutch, A. Y., Roth, R. H., & Bunney, B. S. (1989). Topographical organization of the efferent projections of the medial prefrontal cortex in the rat: an anterograde tract-tracing study with Phaseolus vulgaris leucoagglutinin. *Journal of Comparative Neurology*, *290*(2), 213-242.

Siapas, A. G., & Wilson, M. A. (1998). Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep. *Neuron*, *21*(5), 1123-1128.

Siegel, J. M. (2001). The REM sleep-memory consolidation hypothesis. *Science*, *294*(5544), 1058-1063.

Sirota, A., Csicsvari, J., Buhl, D., & Buzsáki, G. (2003). Communication between neocortex and hippocampus during sleep in rodents. *Proceedings of the National Academy of Sciences*, *100*(4), 2065-2069.

Smith, C., & Macneill, C. (1994). Impaired motor memory for a pursuit rotor task following Stage 2 sleep loss in college students. *Journal of sleep research*, *3*(4), 206-213.

Smith, D. M., Barredo, J., & Mizumori, S. J. (2012). Complimentary roles of the hippocampus and retrosplenial cortex in behavioral context discrimination. *Hippocampus*, *22*(5), 1121-1133.

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

Staresina, B. P., Bergmann, T. O., Bonnefond, M., Van Der Meij, R., Jensen, O., Deuker, L., ... & Fell, J. (2015). Hierarchical nesting of slow oscillations, spindles and ripples in the human hippocampus during sleep. *Nature neuroscience*, *18*(11), 1679.

Steinmetz, N. A., Aydin, C., Lebedeva, A., Okun, M., Pachitariu, M., Bauza, M., ... & Harris, T. D. (2021). Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science*, *372*(6539).

Steriade, M. (1999). Coherent oscillations and short-term plasticity in corticothalamic networks. *Trends in neurosciences*, *22*(8), 337-345.

Steriade, M., Deschenes, M., Domich, L., & Mulle, C. (1985). Abolition of spindle oscillations in thalamic neurons disconnected from nucleus reticularis thalami. *Journal of neurophysiology*, *54*(6), 1473-1497.

Steriade, M., Nunez, A., & Amzica, F. (1993a). A novel slow (< 1 Hz) oscillation of neocortical neurons in vivo: depolarizing and hyperpolarizing components. *Journal of neuroscience*, *13*(8), 3252-3265.

Steriade, M., Nuñez, A., & Amzica, F. (1993b). Intracellular analysis of relations between the slow (< 1 Hz) neocortical oscillation and other sleep rhythms of the electroencephalogram. *Journal of Neuroscience*, *13*(8), 3266-3283.

Stuesse, S. L., & Newman, D. B. (1990). Projections from the medial agranular cortex to brain stem visuomotor centers in rats. *Experimental brain research*, *80*(3), 532-544.

Sul, J. H., Jo, S., Lee, D., & Jung, M. W. (2011). Role of rodent secondary motor cortex in value-based action selection. *Nature neuroscience*, *14*(9), 1202-1208.

Sul, J. H., Kim, H., Huh, N., Lee, D., & Jung, M. W. (2010). Distinct roles of rodent orbitofrontal and medial prefrontal cortex in decision making. *Neuron*, *66*(3), 449-460.

Sutherland, G. R., & McNaughton, B. (2000). Memory trace reactivation in hippocampal and neocortical neuronal ensembles. *Current opinion in neurobiology*, *10*(2), 180-186.

Swanson, L. W. (1981). A direct projection from Ammon's horn to prefrontal cortex in the rat. *Brain research*, *217*(1), 150-154.

Tabuchi, E., Furusawa, A. A., Hori, E., Umeno, K., Ono, T., & Nishijo, H. (2005). Neural correlates to action and rewards in the rat posterior cingulate cortex. *Neuroreport*, *16*(9), 949-953.

Tennant, K. A., Adkins, D. L., Donlan, N. A., Asay, A. L., Thomas, N., Kleim, J. A., & Jones, T. A. (2011). The organization of the forelimb representation of the C57BL/6 mouse motor cortex as defined by intracortical microstimulation and cytoarchitecture. *Cerebral cortex*, *21*(4), 865-876.

Teyler, T. J., & DiScenna, P. (1986). The hippocampal memory indexing theory. *Behavioral neuroscience*, *100*(2), 147.

Teyler, T. J., & Rudy, J. W. (2007). The hippocampal indexing theory and episodic memory: updating the index. *Hippocampus*, *17*(12), 1158-1169.

Treves, A., & Rolls, E. T. (1992). Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. *Hippocampus*, *2*(2), 189-199.

Tseng, K. Y., Chambers, R. A., & Lipska, B. K. (2009). The neonatal ventral hippocampal lesion as a heuristic neurodevelopmental model of schizophrenia. *Behavioural brain research*, *204*(2), 295-305.

Urban, N. N., Henze, D. A., & Barrionuevo, G. (2001). Revisiting the role of the hippocampal mossy fiber synapse. *Hippocampus*, *11*(4), 408-417.

Van Eden, C. G., Lamme, V. A. F., & Uylings, H. B. M. (1992). Heterotopic cortical afferents to the medial prefrontal cortex in the rat. A combined retrograde and anterograde tracer study. *European Journal of Neuroscience*, *4*(1), 77-97.

van Groen, T., & Michael Wyss, J. (1990). Connections of the retrosplenial granular a cortex in the rat. *Journal of Comparative Neurology*, *300*(4), 593-606.

van Groen, T., & Wyss, J. M. (1992). Connections of the retrosplenial dysgranular cortex in the rat. *Journal of Comparative Neurology*, *315*(2), 200-216.

Van Groen, T., & Wyss, J. M. (2003). Connections of the retrosplenial granular b cortex in the rat. *Journal of Comparative Neurology*, *463*(3), 249-263.

Van Ormer, E. B. (1932). Retention after intervals of sleep and waking. *Archives of Psychology*.

Van Ormer, E. B. (1933). Sleep and retention. *Psychological bulletin*, *30*(6), 415.

Vann, S. D., & Aggleton, J. P. (2002). Extensive cytotoxic lesions of the rat retrosplenial cortex reveal consistent deficits on tasks that tax allocentric spatial memory. *Behavioral neuroscience*, *116*(1), 85.

Vedder, L. C., Miller, A. M., Harrison, M. B., & Smith, D. M. (2017). Retrosplenial cortical neurons encode navigational cues, trajectories and reward locations during goal directed navigation. *Cerebral Cortex*, *27*(7), 3713-3723.

Vertes, R. P. (2004). Memory consolidation in sleep: dream or reality. *Neuron*, *44*(1), 135-148.

Vertes, R. P., & Siegel, J. M. (2005). Time for the sleep community to take a critical look at the purported role of sleep in memory processing. *Sleep*, *28*(10), 1228-1229.

Walker, M. P., & Stickgold, R. (2006). Sleep, memory, and plasticity. *Annu. Rev. Psychol.*, *57*, 139-166.

Wallace, M. N., & Palmer, A. R. (2008). Laminar differences in the response properties of cells in the primary auditory cortex. *Experimental brain research*, *184*(2), 179-191.

Wei, Z., Lin, B. J., Chen, T. W., Daie, K., Svoboda, K., & Druckmann, S. (2020). A comparison of neuronal population dynamics measured with calcium imaging and electrophysiology. *PLoS computational biology*, *16*(9), e1008198.

Wilber, A. A., Skelin, I., Wu, W., & McNaughton, B. L. (2017). Laminar organization of encoding and memory reactivation in the parietal cortex. *Neuron*, *95*(6), 1406-1419.

Willshaw, D. J., & Buckingham, J. T. (1990). An assessment of Marr's theory of the hippocampus as a temporary memory store. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, *329*(1253), 205-215.

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

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

Wise, S. P., & Murray, E. A. (2000). Arbitrary associations between antecedents and actions. *Trends in neurosciences*, *23*(6), 271-276.

Yamawaki, N., Radulovic, J., & Shepherd, G. M. (2016). A corticocortical circuit directly links retrosplenial cortex to M2 in the mouse. *Journal of Neuroscience*, *36*(36), 9365-9374.

Yaroush, R., Sullivan, M. J., & Ekstrand, B. R. (1971). Effect of sleep on memory: II. Differential effect of the first and second half of the night. *Journal of experimental psychology*, *88*(3), 361.

Zhang, S., Xu, M., Chang, W. C., Ma, C., Do, J. P. H., Jeong, D., ... & Dan, Y. (2016). Organization of long-range inputs and outputs of frontal cortex for top-down control. *Nature neuroscience*, *19*(12), 1733-1742.

Table S1. Cell type counts per animal and brain region

Mouse	HPC Condition	Region	Total Units	Exc	Inh	Unc
U1	Lesion	M2	141	101	37	3
		RSC	171	121	43	7
U2	Lesion	M2	79	63	14	2
		RSC	35	22	10	3
U4	Sham	M2	70	52	12	6
		RSC	61	40	16	5
U5	Sham	M2	206	176	19	11
		RSC	117	93	22	2
U7	Lesion	M2	117	101	14	2
		RSC	45	30	10	5
U8	Lesion	M2	188	157	19	12
		RSC	131	85	33	13
U9	Sham	M2	55	54	1	0
		RSC	23	19	4	0
U10	Sham	M2	35	32	2	1
		RSC	43	35	4	4
U11	Lesion	M2	179	160	12	7
		RSC	119	93	18	8
U12	Sham	M2	56	49	5	2
		RSC	39	31	5	3
U13	Sham	M2	46	37	6	3
		RSC	68	47	20	1
U14	Lesion	M2	58	55	3	0
		RSC	12	3	9	0
Total	Sham	M2	468	400	45	23
		RSC	351	265	71	15
Total	Lesions	M2	762	637	99	26
		RSC	513	354	123	36
Grand Totals	Both	M2	1230	1037	144	49
		RSC	864	619	194	51

Table S2. Superficial and deep cell counts per animal and brain region

Mouse	HPC Condition	Region	Superficial	Deep
U1	Lesion	M2	28	113
		RSC	119	52
U2	Lesion	M2	2	77
		RSC	0	35
U4	Sham	M2	34	36
		RSC	23	38
U5	Sham	M2	87	119
		RSC	28	89
U7	Lesion	M2	1	116
		RSC	16	29
U8	Lesion	M2	66	122
		RSC	54	77
U9	Sham	M2	0	55
		RSC	10	13
U10	Sham	M2	4	31
		RSC	19	24
U11	Lesion	M2	79	100
		RSC	62	57
U12	Sham	M2	2	54
		RSC	6	33
U13	Sham	M2	1	45
		RSC	53	15
U14	Lesion	M2	13	45
		RSC	9	3
Total	Sham	M2	128	340
		RSC	139	212
Total	Lesions	M2	189	573
		RSC	260	253
Grand Totals	Both	M2	317	913
		RSC	399	465