

Behavioral Correlates of Hippocampal Neural Sequences

Anoopum S. Gupta

CMU-RI-TR-11-36

*Submitted in partial fulfillment of the requirements for
the degree of Doctor of Philosophy in Robotics*

The Robotics Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

September 2011

Thesis Committee

David S. Touretzky (Chair)

Tai Sing Lee

Reid Simmons

George Stetten

A. David Redish, University of Minnesota

Copyright © 2011 by Anoopum S. Gupta. All rights reserved.

This work was supported by NIH National Research Service Award F30 MH-091821,
NSF IGERT DGE-0549352, NIH T32 NS007433, NIM R01 MH-080318,
and the Pennsylvania Tobacco Settlement Fund.

Abstract

Sequences of neural activity representing paths in an environment are expressed in the rodent hippocampus at three distinct time scales, with different hypothesized roles in hippocampal function. As an animal moves through an environment and passes through a series of place fields, place cells activate and deactivate in sequence, at the time scale of the animal's movement (i.e., the behavioral time scale).

Moreover, at each moment in time, as the animal's location in the environment overlaps with the firing fields of many place cells, the active place cells fire in sequence during each cycle of the 4-12 Hz theta oscillation observed in the hippocampal local field potentials (i.e., the theta time scale), such that the neural activity, in general, represents a short path that begins slightly behind the animal and ends slightly ahead of the animal. These sequences have been hypothesized to play a role in the encoding and recall of episodes of behavior.

Sequences of neural activity occurring at the third time scale are observed during both sleep and awake but restful states, when animals are paused and generally inattentive, and are associated with sharp wave ripple complexes (SWRs) observed in the hippocampal local field potentials. During the awake state, these sequences have been shown to begin near the animal's location and extend forward (forward replay) or backward (backward replay), and have been hypothesized to play a role in memory consolidation, path planning, and reinforcement learning.

This thesis uses a novel sequence detection method and a novel behavioral spatial decision task to study the functional significance of theta sequences and SWR sequences. The premise of the thesis is that by investigating the behavioral content represented by these sequences, we may further our understanding of how these sequences contribute to hippocampal function.

The first part of the thesis presents an analysis of SWR sequences or replays, revealing several novel properties of these sequences. In particular it was found that instead of

preferentially representing the more recently experienced parts of the maze, as might be expected for memory consolidation, paths that were not recently experienced were more likely to be replayed. Additionally, paths that were never experienced, including shortcut paths, were observed. These observations suggest that hippocampal replay may play a role in constructing and maintaining a "cognitive map" of the environment.

The second part of the thesis investigates the properties of theta sequences. A recent study found that theta sequences extend further forward at choice points on a maze and suggested that these sequences may be partly under cognitive control. In this part of the thesis I present an analysis of theta sequences showing that there is diversity in theta sequences, with some sequences extending more forward and others beginning further backward. Furthermore, certain components of the environment are preferentially represented by theta sequences, suggesting that theta sequences may reflect the cognitive "chunking" of the animal's environment.

The third part of the thesis describes a computational model of the hippocampus which explores how synaptic learning due to neural activity during navigation (i.e., theta sequences) may enable the hippocampal network to produce forward, backward, and shortcut sequences during awake rest states (i.e., SWR sequences).

Acknowledgements

I have been very fortunate to be surrounded by mentors, colleagues, and friends who have enriched my training and without whom this work would not have been possible. I am grateful to my advisor, David Touretzky, for his endless wisdom and guidance over the last four years. Dr. Touretzky has not only supported the research in this thesis, but has mentored me in all aspects of academic life. I want to thank David Redish for welcoming me into the Redish lab, teaching me the ways of electrophysiology, and for his constant mentorship throughout my training. Dr. Redish has been like a second thesis advisor to me and I am deeply thankful for all the time he has spent helping me grow as a scientist. George Stetten, Tai Sing Lee, and Reid Simmons have each contributed a unique perspective to the work in this thesis, and I am very thankful for their support and advice.

My experience as a graduate student would not have been the same without my interactions with the Redish lab: Matthijs van der Meer, Adam Johnson, Adam Steiner, John Ferguson, Chris Boldt, Adam Vogel, Nate Powell, Andrew Wikenheiser, and Andy Papale. In particular, I want to thank Matt for all the impromptu discussions and for his invaluable help learning how to perform experiments in the Redish lab.

Thanks to Mark Fuhs, Andrew Maurer, Piotr Dollar, Rob Kass, Bruce McNaughton, Michael Hasselmo, the CMU Robotics community, the Pitt/CMU CNBC and MSTP, and friends and colleagues from NS&B 2009 for collaborations and discussions that have enriched my experience as a graduate student.

Thank you to Rachelle, Boz, and my friends for giving me a life outside of research. Finally, thank you to Mom, Dad, Sumi, Dan, Apurve, and Nitika for all their love and guidance.

Contents

| | |
|--|-----------|
| Abstract | 2 |
| Acknowledgements | 4 |
| 1 Introduction | 8 |
| 2 Hippocampal function and physiology | 11 |
| Hippocampal roles in spatial navigation and episodic memory | 11 |
| Representation of the environment by hippocampal place cells | 12 |
| Functional implications of theta phase precession | 15 |
| 3 Hippocampal neural sequences | 20 |
| Hippocampal theta sequences | 20 |
| Hippocampal network states: theta and LIA | 21 |
| Hippocampal replay sequences | 22 |
| 4 Extracting behavioral information from neural sequences | 27 |
| Decoding the animal's location in the environment | 27 |
| New method for decoding paths in the environment | 30 |
| 5 General methods used in electrophysiology experiments | 36 |
| Behavioral task: The 2T maze | 36 |
| Surgery, recording, and histology. | 37 |
| Place fields | 38 |
| Replay sequence specific analyses | 40 |
| Theta sequence specific analyses | 43 |
| 6 Hippocampal replay sequences | 45 |
| 6.1 Summary | 45 |

| | | |
|----------|--|-----------|
| 6.2 | Results | 45 |
| | Forward and backward replay of local and non-local trajectories | 46 |
| | Spatial distributions of forward and backward replay | 50 |
| | Comparison of data with current proposed mechanisms | 50 |
| | Task-dependence of the content of replay | 53 |
| | Never-experienced shortcut sequences | 55 |
| 6.3 | Discussion | 59 |
| 7 | Hippocampal theta sequences | 65 |
| 7.1 | Summary | 65 |
| 7.2 | Results | 65 |
| | Ahead and behind length are anti-correlated and occur during different parts of the theta cycle | 68 |
| | Theta period and gamma cycles increase with path length | 68 |
| | Ahead and behind representations | 71 |
| | Average ahead and behind length over space | 71 |
| | Spatial distribution of represented paths | 73 |
| | Ahead length increases with acceleration, behind length increases with de- celeration | 77 |
| | Lengths show a complex relationship with velocity | 78 |
| 7.3 | Discussion | 80 |
| | Diversity in theta sequences | 83 |
| | Information processing: Theta and Gamma | 85 |
| | Behavioral and Neural Chunking | 85 |
| 8 | Model for backward replay generation | 88 |
| 8.1 | Abstract | 88 |
| 8.2 | Introduction | 88 |

| | | |
|----------|---|------------|
| 8.3 | Theoretical Model: potential contribution of the theta phase gradient and longitudinal projections for learning reverse connections | 90 |
| 8.4 | Synaptic Learning Model | 95 |
| 8.5 | Replay Model | 99 |
| 8.6 | Discussion | 105 |
| 9 | Conclusion | 109 |
| 9.1 | Summary | 109 |
| 9.2 | Future work | 111 |
| | References | 115 |

1 Introduction

The work described in this thesis was performed with the goal of understanding how the spiking of a population of neurons in a particular region of the brain may contribute to the region's cognitive functions. The link between neural activity and cognition can be made in two steps. By observing the spiking of a particular neuron with respect to an animal's behavior (e.g., movement in an environment), one can determine if the activity is informative (predictive) of the behavioral variable (the animal's location). Methods can then be used to decode the behavioral content represented by the neural activity. These methods allow investigators to determine what the brain is representing purely from the neural activity. And since cognitive functions such as memory and planning involve representing information that is distinct from the animal's immediate sensory/behavioral state (i.e., non-local information), such techniques allow us to study how these processes are expressed in the brain.

Spiking neurons are a unique and advantageous window into brain function at multiple levels. For one, the ability to discern the activity of an individual neuron (thought to be the fundamental processing unit of the brain) allows for the extraction of information in systems in which neurons very near to one another in the brain are representing different behavioral or cognitive information. Without the ability to distinguish individual neuron activity, the information would be merged and lost in such systems. Furthermore, it has been shown that behavioral information is sometimes represented at fast time scales that require recording methods with fine temporal precision in order to resolve. The ability to record neural activity at fast time scales additionally provides us with a connection between the behavioral information represented in spike trains and neural plasticity. Spike timing dependent plasticity (STDP) has been shown to occur when the pre and post synaptic cells fire within 80 ms of one another. Recording from populations of neurons at sub-millisecond time scales enables the study of neural sequences to determine how sequences of behavioral information may

be stored in the synaptic connections in the brain.

The work described in this thesis studies the behavioral information represented by a population of spiking neurons in the CA1 region of the rodent hippocampus during the performance of a spatial navigation decision task. The thesis focuses on the hippocampus because it is known to play an important role in episodic memory and spatial navigation, and because it has recently been implicated in higher cognitive functions such as self-projection and imagination. The rodent hippocampus is an extremely informative system because it is relatively easy to access with electrodes, and because the firing of hippocampal pyramidal cells has a clear behavioral correlate: the animal's location in an environment. At any given time as an animal moves through an environment, the population of pyramidal cells in the hippocampus is representing the animal's physical location. The strong correlation between neural activity and the animal's physical location on the maze enables accurate decoding of the animal's location from subsequent neural activity alone. It also allows us to determine (via decoding) if and when non-local information (e.g., a different location or path on the maze) is being represented, as has been found during pausing at a choice point, and during sleep and awake rest states. These studies have suggested that the non-local representations play a role in memory consolidation, decision making, reinforcement learning, and planning. In this thesis I explore the content of non-local representations on a novel task that has several advantages over previous studies, allowing me to test some of the suggested functions of non-local representations. Additionally, I construct a computational model to explore how the known synaptic learning rules in the hippocampus may enable the expression of some of the non-local representations that are observed.

The thesis is organized into eight additional chapters. **Chapter 2** provides an overview of hippocampal function and physiology, while **Chapter 3** focuses on experimental work on neural sequences observed in the hippocampus. **Chapter 4** surveys different decoding methods for characterizing sequences of neural activity, and then describes the decoding method developed here and the advantages it imparts given data constraints. **Chapter 5**

describes the general methods used in the experimental studies in this thesis. **Chapter 6** and **Chapter 7** present original experimental work on neural sequences recorded from the hippocampus as animals are performing a novel behavioral task. **Chapter 6** presents an analysis of neural sequences as animals are awake, but paused and resting at feeder locations on the maze. **Chapter 7** presents an analysis of neural sequences as animals are actively navigating the maze to reach the feeder locations. **Chapter 8** describes a computational model that explores how the backward and shortcut sequences presented in **Chapter 6** may be produced, taking into account the anatomical and physiological constraints of the hippocampal network. Finally, **Chapter 9** summarizes the advances made in this thesis and outlines important future work.

2 Hippocampal function and physiology

Hippocampal roles in spatial navigation and episodic memory

The hippocampus has been long known to play an important role in spatial navigation (O'Keefe and Nadel, 1978; Morris et al., 1982; Redish, 1999) and encoding the memory of autobiographical events (i.e., episodic memory, Scoville and Milner, 1957; Squire, 1992; Eichenbaum et al., 1999; Tulving, 2002). Some of the first evidence of hippocampal involvement in episodic memory came from studies of patients who had lesions of the hippocampus and surrounding neural structures. Studies of the famous patient HM, who had a bilateral hippocampal resection due to intractable epilepsy, revealed a specific deficit in declarative memory (i.e., the acquisition of new fact and event memory). HM's ability to form new memories of autobiographical experiences was lost (anterograde amnesia, Scoville and Milner, 1957). The memory of recent experiences was also impaired (retrograde amnesia), but past experiences and childhood memories were preserved, pointing to a role for the hippocampus in encoding memories, but not in the long term storage of episodic memory (Milner et al., 1968). Numerous studies performed on HM and other patients over a half century have revealed that hippocampal lesions specifically impair episodic memory and that other types of memories, including working memory and procedural memory (e.g., delay conditioning and habit learning) are left intact (Squire and Zola-Morgan, 1991).

Similarly, lesion studies in animals have elucidated the role of the hippocampus in spatial navigation. Morris et al. (1982), using a novel experimental paradigm (the now famous Morris water maze), demonstrated that animals with hippocampal lesions were impaired in their ability to learn a spatial navigation task. The water maze task involved a cylindrical tank of opaque water, with a hidden platform located at a fixed location in the tank, just beneath the water. Rodents were dropped in the tank, and finding the cold water unpleasant, searched for the hidden platform. Cues were present in the room such that normal rats quickly learned the location of the hidden platform and were able to navigate direct routes

to the platform even when initially placed at a novel location in the tank. Hippocampus lesioned animals had no trouble taking direct routes to a visible platform in the tank, but could not learn to take direct routes to a hidden platform, thus demonstrating the impairment in spatial memory. This experiment supported the idea that the hippocampus is involved in the formation and use of cognitive maps (i.e., a mental representation of the environment) to flexibly navigate or take novel paths through space (Tolman, 1948; O'Keefe and Nadel, 1978). Several other experiments have supported a role for the hippocampus in enabling flexible navigation (Tolman et al., 1946; Nadel et al., 1975; Packard and McGaugh, 1996).

An important similarity between episodic memory and spatial navigation is the need to encode and recall sequenced information. In the case of episodic memory, the sequential information is a series of events in time. For spatial navigation, the sequential information is a series of locations in space. An elegant study (Fortin et al., 2002) showed that hippocampal lesions in rats specifically disrupted their ability to perform a task requiring a representation of sequence. Normal and hippocampus lesioned rats were presented with a sequence of odors and were asked to perform a recognition task (choose the odor not presented in the sequence) or a sequential order task (choose the odor presented earlier in the sequence). Normal rats were able to learn both tasks. Hippocampus lesioned rats could perform the recognition task, but not the sequential order task, demonstrating a specific impairment in sequence memory. Thus, several lines of evidence suggest that a basic function of the hippocampus is to process sequential information. As will be discussed, observations at the neurophysiological level (i.e., place cell activity in the rodent hippocampus) provide evidence that sequential information is represented in the hippocampus and in a way that could allow for sequence learning at a synaptic level.

Representation of the environment by hippocampal place cells

The discovery of cells with spatially localized firing fields (place cells) in the rodent hippocampus (O'Keefe and Dostrovsky, 1971), has provided researchers with a way to link

neural activity in the hippocampus with an animal's behavior. This has been a key to understanding how the hippocampus represents physical environments and an animal's experience in that environment. Forty years of research has revealed numerous properties of place cells, some of which will be reviewed here.

Place cells are pyramidal cells located in the CA1 and CA3 subregions of the hippocampal formation and are defined by their tendency to fire when the animal occupies a particular region of the environment: the cell's place field (**Figure 1**). Most place cells have a single place field, but individual place cells can have multiple distinct firing fields in an environment. Place fields are uniformly distributed over the entire environment (in general there is not a propensity for certain parts of the environment to have a higher density of place fields, O'Keefe, 1976; Olton et al., 1978; McNaughton et al., 1983) and each place field has a roughly Gaussian shape in terms of its firing rate over space. A typical rate map is shown in **Figure 1**. Due to these properties, an animal's location in an environment can be accurately decoded from the activity of a population of place cells during a small window of time (a few hundred milliseconds) using statistical techniques (Zhang et al., 1998).

Not every pyramidal cell in CA1 has a place field in every environment. Guzowski et al. (1999) used the immediate early gene *Arc* to label active CA1 neurons as animals explored an environment and determined that approximately 45% of neurons had place fields in a given environment, which is consistent with cellular recording studies (Wilson and McNaughton, 1993; Gothard et al., 1996). This demonstrates the prominence of place cells in the rodent hippocampus and the large proportion of cells involved in representing a given environment.

In the hippocampus, there is no topographical mapping between the anatomical location of place cells in the brain structure and the location of their place fields in the physical environment (Redish et al., 2001; Dombeck et al., 2010). This is an important feature of the hippocampus, as it has been shown that place cells "remap" between environments; when an animal is moved from environment A to environment B, cells with place fields at particular




Figure 1: The rate maps of 80 cells, simultaneously recorded as an animal foraged in a square box environment. Each panel is the average firing rate of a single neuron as a function of location in the square box (red is high, blue is low). Some of the cells had spatially localized firing fields (place cells), some were silent in the environment, and others fired at a high rate throughout the environment (interneurons). (Reprinted from Wilson and McNaughton, 1993, with permission from author and publisher.)

locations in environment A have new, independent place field locations in environment B (Wilson and McNaughton, 1993). When an animal is returned to environment A, place cells remap again to take on their original place field locations (Wilson and McNaughton, 1993). Thus, place cells in the hippocampus form a unique and stable representation of each new environment that an animal encounters.

The spatially selective firing fields of place cells are not derived from a single sensory modality, consistent with the multimodal sensory input that the hippocampus receives from the entorhinal cortex. Place fields persist in darkness (O'Keefe, 1976; Quirk et al., 1990; Markus et al., 1994) and have normal fields even after olfactory, visual, or both visual and auditory inputs are eliminated (Save et al., 1998; Hill and Best, 1981). However, when vision is not impaired, place fields are influenced by local and distal visual cues (O'Keefe, 1976; Olton et al., 1978; Kubie and Ranck, 1983; Muller and Kubie, 1987; Rossier et al., 2000). Zinyuk et al. (2000) used an experimental paradigm to independently alter local and distal cues, showing that the degree to which place fields were influenced by local versus distal cue perturbations depended on the degree to which the cues were used for the navigation task. When the task required the animal to attend to the distal cues, perturbation of the distal cues altered place field locations more than the perturbation of local cues. These studies indicate that place cells in the hippocampus reflect a high-level multimodal representation of the environment. Such a representation could be useful for encoding behavioral experiences in an environment, an important component of both episodic memory and spatial navigation.

Functional implications of theta phase precession

The experiments described above indicate the way in which place cells use a rate code to represent the environment. In a landmark study, O'Keefe and Recce (1993) provided the first evidence that hippocampal place cells may also represent information using a temporal code. The authors observed that there was a special relationship between the firing of

place cells and the 4-12 Hz theta rhythm observed in the hippocampal local field potentials (Vanderwolf, 1969; Buzsáki et al., 1983). As an animal entered a cell's place field, the cell initially fired late in the theta cycle and as the animal passed through the field, the cell fired earlier and earlier in the cycle (**Figure 2**). This phenomenon was termed theta phase precession. It has been shown that information about the phase of spikes can be used to improve the accuracy of position decoding (Jensen and Lisman, 2000) and provides information about an animal's direction of motion (Huxter et al., 2008).

The discovery of phase precession has spawned a wealth of research concerned with the mechanism and function of the phenomenon. Many models have been developed to attempt to explain the mechanism of phase precession. In general, one class of models proposes that phase precession is largely an individual cell phenomenon, with each cell precessing independently of other cells. Another class of models proposes that the phenomenon arises due to network properties and synaptic communication between cells with overlapping place fields. The proposed mechanisms have been nicely reviewed by Maurer and McNaughton (2007). Although phase precession describes the relationship between the firing of a single place cell with the theta oscillation, the phenomenon implies a spike timing relationship between cells with overlapping place fields (Skaggs et al., 1996; Tsodyks et al., 1996; Jensen and Lisman, 1996). Imagine an animal running through a series of ordered place fields A, B, C, D, and E (**Figure 3A**). When the animal is located at the intersection of these place fields, the animal is at the end of place field A, the middle of place field C, and the beginning of place field E. Due to the phase precession phenomenon, place cell A fires first in the theta cycle, C fires in the middle of the theta cycle, and E fires at the end of the cycle. Thus, the order of behavioral experience is preserved and temporally compressed during each 125 ms theta cycle.

Phase precessing neurons imply that hippocampus is not representing the animal's precise location at each moment in time, but is instead representing a short trajectory during each theta cycle, which begins slightly behind the animal and ends slightly ahead of the

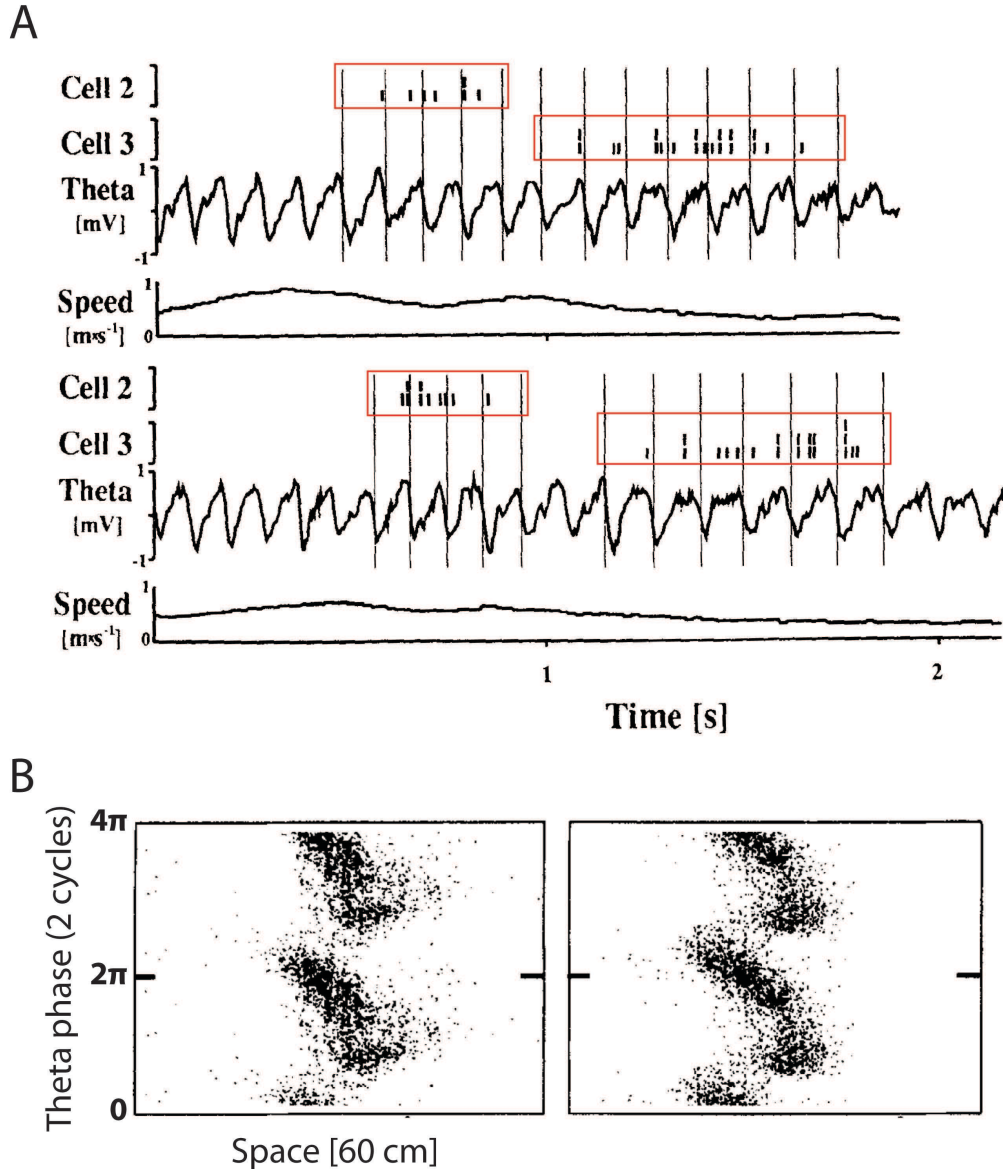


Figure 2: Theta phase precession by hippocampal place cells. **(A)** The relationship between the spiking of two neurons and the theta oscillation. Each set of short vertical lines surrounded by a red box represents the set of spikes emitted by a single cell as an animal runs through that cell's place field. As the animal runs through the field (left to right in time), the spikes fire earlier and earlier in the theta cycle (i.e., they phase precess). (Adapted from O'Keefe and Recce, 1993, with permission from author and publisher) **(B)** Phase precession by two cells (one cell per panel) compiled over many passes through each cell's place field. Each dot in the panel represents a spike occurring with a particular theta phase (y-axis) at a particular location on the maze (x-axis). When the animal enters the place field (left side of panel), spikes occur late in the theta cycle. As the animal passes through the field, eventually exiting on the right side of the panel, spikes occur with earlier phases. (Adapted from Skaggs et al., 1996, with permission from author and publisher.)

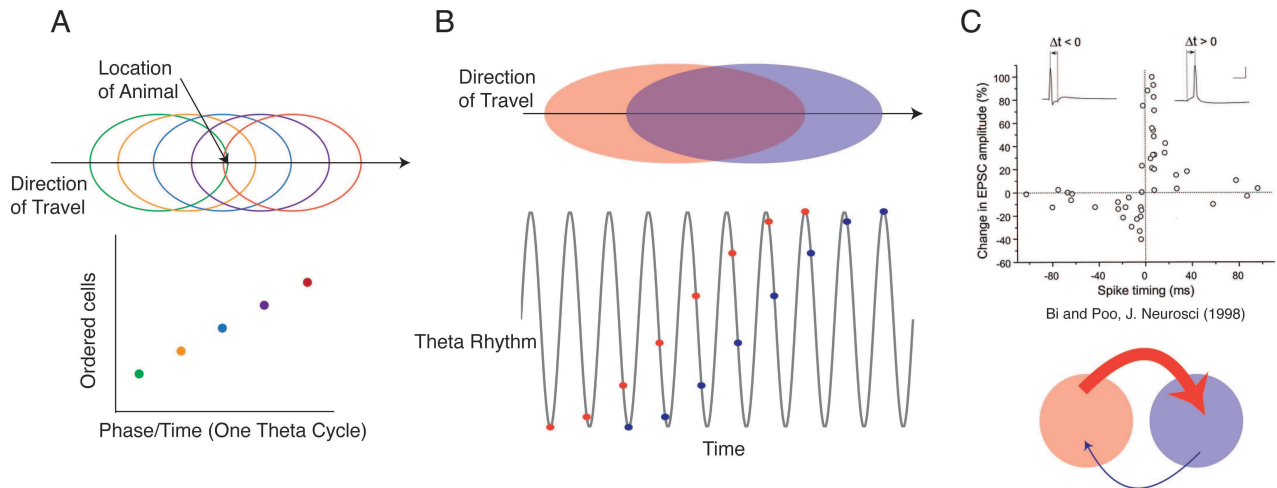


Figure 3: **(A)** Due to phase precession, if an animal runs from left to right through a series of place fields (ellipses in top of panel), the order of place cell firing matches the order of place field visitation during each theta cycle (bottom of panel: colored dots represent the firing of each corresponding place cell. X-axis is time and y-axis is spatially ordered place cell). **(B)** Schematic showing how the phase precession of two cells with overlapping place fields results in the first cell firing before the second in each theta cycle (red dots are spikes corresponding to the red place cell, blue dots correspond to the blue place cell). **(C)** The spike timing dependent plasticity (STDP) learning rule (top of panel, reprinted from Bi and Poo, 1998, with permission from author and publisher). Due to STDP, the connection from the red cell to the blue cell in panel B would be strengthened and the connection from the blue cell to the red cell would be weakened (depicted in bottom of panel).

animal (Foster and Wilson, 2007). The trajectory represented progresses in the same direction that the animal is moving and the order of place cell firing during each theta cycle matches the order of place fields being visited by the animal. This temporal structure is especially interesting given the observation of spike timing dependent plasticity (STDP) in the hippocampus (Bi and Poo, 1998).

STDP is a synaptic learning rule in which the synapse enabling communication from neuron A to neuron B is potentiated (strengthened) if neuron A fires before neuron B and is depressed (weakened) if neuron B fires before neuron A (**Figure 3B,C**). Synaptic potentiation or depression takes place when the two neurons fire within 80 ms of each other (**Figure 3C**, Bi and Poo, 1998). Place cells in hippocampal subregion CA3 have strong recurrent connections (Lorento do Nó, 1934; Swanson et al., 1978), meaning that place cells within this region synapse onto other place cells in this region, and project strongly to place cells in subregion CA1 (Schaffer collaterals, Swanson et al., 1980, 1981). Thus,

the anatomical connections of neurons in the hippocampus combined with the ordered firing of place cells during each theta cycle, may enable synaptically connected hippocampal neurons to encode episodes of behavior in the precise order that it is experienced (Skaggs et al., 1996; Tsodyks et al., 1996; Jensen and Lisman, 1996), which is an important component of both episodic memory and spatial navigation. It should be noted that these spike sequences are occurring during 125 ms theta cycles and are therefore consistent with the time scale of STDP.

3 Hippocampal neural sequences

Hippocampal theta sequences

As described in **Chapter 2**, hippocampal phase precession implies that the hippocampus represents a short trajectory in the environment during each theta cycle (i.e., theta sequences), beginning slightly behind the animal and ending just ahead of the animal. Recent work has begun to explore the properties of theta sequences by analyzing the spiking of neuron pairs with overlapping place fields (Dragoi and Buzsaki, 2006) or by analyzing the entire population of neurons active during each theta cycle (Foster and Wilson, 2007; Johnson and Redish, 2007; Maurer et al., 2011).

Dragoi and Buzsaki (2006) showed that some pairs of neurons with overlapping place fields fired with a relatively constant temporal relationship during each theta cycle. The time offset between the firing of each cell was less than the time it took for the animal to traverse the distance between the centers of the two place fields, indicating that the neural representation of space was temporally compressed with respect to behavior (by 10 to 15 times, Dragoi and Buzsaki, 2006). Furthermore, it was shown that the temporal coordination of pairs of neurons was stronger than would be expected if each cell phase precessed independently, suggesting that synaptic interactions were involved. Foster and Wilson (2007) extended this work by moving beyond cell pairs to analyze the entire neural sequence taking place during each theta cycle. They confirmed the prediction that phase precession results in a neural sequence representing a forward path beginning behind the animal and extending forward. Similar to Dragoi and Buzsaki (2006), they showed that the integrity of the forward ordered sequence was beyond what would be expected if each cell phase precessed independently, further supporting synaptic involvement in the coordination of theta sequences.

Johnson and Redish (2007) were the first to examine the behavioral content represented by neural activity at the theta time scale. Using a Bayesian decoding approach (Brown et al.,

1998; Johnson et al., 2008, discussed in **Chapter 4**), the authors decoded physical locations represented by the neural activity in sliding 20 ms windows. Decoding at this fine time scale (approximately an eighth of a theta cycle), allowed for the visualization of paths represented during theta cycles. Using this technique, Johnson and Redish (2007) observed that when an animal came to the choice point on a modified T-maze, and moved its head back and forth as if considering its options (vicarious trial and error or VTE, Meunzinger, 1938; Tolman, 1939), the paths represented did not end just ahead of the animal as previously observed, but swept considerably ahead of the animal. These "sweeps" represented each path option separately and specifically occurred at the choice point. The authors suggested that sweeps could reflect the mental search of a "cognitive map" (a mental model of the environment, Tolman, 1948; O'Keefe and Nadel, 1978), in order to evaluate possible options before making a decision (Johnson and Redish, 2007; van der Meer and Redish, 2009). Thus, by analyzing the behavioral content expressed by theta sequences, Johnson and Redish (2007) provided neural evidence for the recall of previously experienced sequential information, and the first evidence that theta sequences may be under cognitive influences and play a role in decision making. **Chapter 7** describes original work in which I examine the behavioral content represented by theta sequences using a novel decoding method described in **Chapter 4**. Briefly, I found that the paths represented by these sequences shifted forward or backward in space depending on the animal's acceleration, and they parsed the environment in segments, providing a potential mechanism for cognitive chunking of experience.

Hippocampal network states: theta and LIA

When animals are awake, there are two main network states in the hippocampus, reflected by the local field potentials (Vanderwolf, 1969; Buzsáki et al., 1983). One network state is the "theta" state which occurs while animals are moving around in an environment or are stationary but attentive and engaged in a task. This state is characterized by high power in

the theta frequency band (4-12 Hz) and it is in this state that the theta sequences described above occur. It is generally thought that experience encoding and recall occurs during the theta state, as animals are interacting with the world. The neural activity and the behavioral correlates of theta sequences support this view (Skaggs et al., 1996; Tsodyks et al., 1996; Jensen and Lisman, 1996).

The other main network state consists of a non-rhythmic pattern in the hippocampal local field potentials known as large amplitude irregular activity (LIA, O'Keefe and Nadel, 1978; Buzsáki et al., 1983). This is the main pattern observed while animals are awake and eating, grooming, or resting, and is also the pattern observed during slow-wave sleep (SWS). Neural activity is on average lower than during the theta state, but there are intermittent bursts of high spike rate activity, which are accompanied by a short duration (40-120 ms), high frequency (180-220 Hz) oscillation in the local field potentials, called sharp wave ripples (SWRs, O'Keefe and Nadel, 1978; Buzsáki et al., 1983; Buzsáki, 1989). It has been suggested that the high frequency neural activity during SWRs is well suited for the long term potentiation of synapses (Buzsáki, 1989) and that this activity occurring while animals are inattentive and at rest may play a role in the transfer or consolidation of experiential information from the hippocampus to the cortex (Marr, 1971; Buzsáki, 1989; Sutherland and McNaughton, 2000). The next section describes the analysis of neural activity during SWRs.

Hippocampal replay sequences

If neural activity during inattentive awake states and during sleep reflects a consolidation process involving the transfer of information from hippocampus to cortex, the neural activity would likely represent information from the animal's recent experiences. Wilson and McNaughton (1994) tested this hypothesis by recording from place cells during active behavior and subsequently during slow-wave sleep (SWS). The authors identified pairs of neurons that fired together (within 100 ms) during behavior and found that these neuron pairs prefer-

entially reactivated together during SWRs in SWS. Pairs of neurons that did not fire together during behavior did not reactivate together during SWRs. Skaggs and McNaughton (1996) extended this work by showing that correlated neuron pairs not only reactivate together during SWS, but the order of firing observed during behavior is preserved during sleep. These initial observations have been replicated and refined using different experimental and analytical techniques, demonstrating that reactivation preferentially occurs after experience compared to before (Kudrimoti et al., 1999; Lee and Wilson, 2002), occurs during SWRs in SWS as well as during SWRs in the awake state (Kudrimoti et al., 1999; Nadasdy et al., 1999), and is a property not only of neuron pairs, but involves sequences of neurons activating in concert ("replay" instead of the more general term, "reactivation", Nadasdy et al., 1999; Lee and Wilson, 2002).

Thus, neural sequences representing paths in the environment are observed not only during active behavior (theta sequences), but also when animals are asleep or awake and inattentive, during SWRs (replay sequences). Whereas theta sequences are thought to be involved in the rapid encoding and recall of episodes of behavior (Skaggs et al., 1996; Tsodyks et al., 1996; Jensen and Lisman, 1996), replay sequences are thought to play a role in memory and learning consolidation (Marr, 1971; Buzsáki, 1989; Sutherland and McNaughton, 2000). Studies have directly disrupted SWRs (Girardeau et al., 2009; Ego-Stengel and Wilson, 2010) or, as in Nakashiba et al. (2009), have blocked CA3 output (the site of SWR generation, Ylinen et al., 1995), showing that these manipulations impair performance on hippocampal-dependent tasks. This provides evidence that SWRs and associated neural activity play an important role in learning and memory.

Several studies have begun to look at the behavioral content represented during replay sequences, with the goal of understanding in more detail how this neural activity may contribute to learning and memory. The experiments described above showed that replay sequences matched the order of behavior: when the locations A, B, C, D, and E were visited in that order, the replay sequences observed also represented the ordered locations A, B,




Figure 4: Examples of forward and backward replay sequences recorded from an animal running on a linear track. The y-axis is place cells ordered by the location of their respective place fields along the linear track. Thus, when running along the linear track the animal passes through place cell 1's firing field first and place cell 13's firing field last. The x-axis is time. Each short vertical line represents a spike emitted by a cell. The local field potential during one lap on the track is shown in the top of the figure, and the corresponding velocity of the animal at each moment in time plotted in the bottom of the figure. The dense neural activity in the middle of the figure occurs as the animal is running, and parts of the neural activity occurring before and after the run while paused (boxed in gray), are replotted at a finer time scale on either side of the figure. The activity during the pause before the run (left of figure) is a forward replay since the order of firing is the same as the order of firing during behavior. The activity after the run is a backward replay because the order of firing (reflected by the negative slope) is the reverse of behavior. (Reprinted from Diba and Buzsáki, 2007, with permission from author and publisher.)

C, D, and E. Foster and Wilson (2006) observed that during some SWRs in the awake state, sequences replayed experienced trajectories in the reverse order (E, D, C, B, and then A). These "reverse" or "backward" replay events (**Figure 4**) were observed when animals were paused after running the length of a linear track, and always began at the animal's location and spread backwards.

Foster and Wilson (2006) speculated that these backward replays may be a mechanism to resolve the "credit assignment" problem described in reinforcement learning models (Sutton and Barto, 1998). The credit assignment problem arises when a sequence of actions leads to later reward. The agent needs to determine which of its past actions were responsible for this reward, as those are the actions it will want to repeat. Proper credit assignment may not be clear from a single trial, but can be inferred from repeated trials with varying behavior. Memory is crucial to this process, as the agent must have access to the entire sequence of actions that led to a reward in order to apportion credit appropriately. A

phenomenon such as backward replay, which begins at the animal's location and spreads backward over the trajectory that was just experienced (Foster and Wilson, 2006), could play a role in passing outcome information back through the action sequence to update the value of each action at each location along the path.

Backward replay is a surprising phenomenon because as described earlier, the phase precession mechanism results in forward ordered theta sequences during behavior, which are thought to synaptically encode paths in the order they were experienced (i.e., forward, not backward), due to the STDP learning rule. If experienced paths are encoded asymmetrically in the forward order, how can the paths be replayed in the reverse order? Foster and Wilson (2006) proposed that backward replays may occur due to a non-synaptic mechanism, specifically a decaying activity trace: cells whose place fields were the most recently experienced are the most excitable and therefore firing occurs in approximately the reverse order of experience. Davidson et al. (2009) examined backward replay as animals ran back and forth on a long track and found that backward replay events could begin at locations away from the animal's position. As will be described in **Chapter 6**, I observed backward replays over non-local and non-recently experienced paths. These observations are incompatible with a decaying activity trace mechanism. **Chapter 8** presents potential mechanisms that could enable the hippocampal network to encode experiences in a reverse order and produce backward replay events of non-recent and non-local paths.

Diba and Buzsáki (2007) replicated the findings of Foster and Wilson (2006), also on a linear track maze. They observed that in addition to replaying the last experience in a backward order, the hippocampus replays the path that is about to be experienced (and has been experienced previously by the animal) in a forward order, while the animal is awake and just prior to movement (**Figure 4**). Forward replays, like backward replays, were observed to begin at the animal's position on the maze and extend down the track. The authors suggested that forward replay events, since they traverse the upcoming path, may play a role in planning future behaviors. The studies of Foster and Wilson (2006) and Diba

and Buzsáki (2007), in addition to highlighting the robust nature of forward and backward neural sequences during awake SWRs, also demonstrate how analyzing the behavioral content of neural sequences can help generate hypotheses about the role of the sequences in cognitive functions such as learning, memory, and planning. In **Chapter 6** I present a detailed analysis of the behavioral content represented by forward replay events during the awake state, in part showing that they are not restricted to beginning at the animal's position on a maze, and suggest that forward and backward replays may play a role in learning a cognitive map of the environment.

4 Extracting behavioral information from neural sequences

Decoding the animal's location in the environment

In order to analyze the behavioral content represented by hippocampal activity sequences, we need analytical methods that can translate place cell spike trains into paths in the environment. The problem of decoding paths represented by neural sequences is a relatively new undertaking, but several studies have shown that an animal's location in the environment can be accurately decoded from the population activity of a group of place cells (Wilson and McNaughton, 1993; Zhang et al., 1998; Brown et al., 1998). It is intuitive that it should be possible to decode an animal's location from the activity of many place cells: if three cells fire in a small window of time and they each have a place field located in the top right corner of an environment, we should be able to say that the neural activity is representing the top right corner. As will be discussed, neural data is inherently noisy and current recording technologies are limited, which raises challenges for decoding neural activity. There are two main techniques that have been used to decode an animal's location in an environment. One is template matching (Wilson and McNaughton, 1993) and the other is Bayesian decoding (Zhang et al., 1998; Brown et al., 1998).

Template matching

The first step of the template matching method is to parse the environment into a number of equally sized spatial bins and then compute the average firing rate (over the entire recording session) of each neuron in each spatial bin (**Equation 1**; $f(\mathbf{x})$ is the average firing rate of a neuron at each spatial location, $S(\mathbf{x})$ is the total number of spikes that the neuron fired in each spatial bin, $P(\mathbf{x})$ is the number of position samples recorded in each spatial bin, and dt is the position sampling period in seconds.)

$$f(\mathbf{x}) = \frac{S(\mathbf{x})}{P(\mathbf{x})dt} \quad (1)$$

This procedure results in an $N \times X$ matrix (space matrix) of firing rates, where N is the number of cells and X is the number of spatial bins in the environment. Thus, the matrix contains the average firing rate of each neuron at each location in the environment. The next step is to bin the neural activity recorded over the session into small time bins (on the order of a few hundred milliseconds), in order to create an $N \times T$ matrix (time matrix) of firing rates, where N is once again the number of cells and T is the number of time bins. Each column of this matrix represents the instantaneous firing rate of the population of place cells. To decode the animal's location from neural activity at a moment in time, a column in the time matrix (representing the instantaneous firing rates at a moment in time) is compared with each column in the space matrix. The vector comparison, which can be the Euclidean distance or dot product, results in a similarity score for the time vector with each space vector. The location of the spatial bin corresponding to the space vector with the highest score is the decoded location. Wilson and McNaughton (1993) performed template matching on datasets obtained from three rats, with 34, 24, and 40 place cells recorded from each rat. Using 1 s time bins to construct the time matrix, they were able to decode the animal's location with average errors of 2 to 5 cm. Using smaller time bins increased the average error and the authors estimated that in order to decode locations with a finer temporal precision (i.e., every 100 ms), with the same average error obtained using 1 s time bins, 380 place cells would be required. Thus, template matching can accurately decode an animal's location in the environment, but has poor temporal and spatial resolution.

Bayesian decoding

In the context of template matching, the space matrix was introduced as a collection of N -length vectors, with each vector containing the average firing rate of each neuron at a particular location in the environment. However, by transposing the space matrix, it can also be viewed as a collection of X -length vectors, with each vector containing the average firing rate of a single neuron at each location in the environment. In this interpretation, each vector is a neuron's place field or tuning curve. The Bayesian decoding technique (Zhang

et al., 1998; Brown et al., 1998) converts the N -collection of tuning curves, $f_N(\mathbf{x})$, into the conditional probability distribution $P(\mathbf{n}|\mathbf{x})$, which is the probability of observing the vector of spike counts, \mathbf{n} , at each location in the environment (**Equation 2**; \mathbf{n} is an N -length vector of spike counts observed in a time window, \mathbf{x} is each spatial bin in the environment, and τ is the size of the time window). This conversion assumes that the firing of each neuron is independent (allowing for the product of each neuron's probability distribution) and assumes that the firing of place cells follows Poisson statistics (Tuckwell, 1988) with the λ parameter equal to the mean firing rate of each neuron at each location in the environment, multiplied by τ ($\lambda_i(\mathbf{x}) = \tau f_i(\mathbf{x})$). The entire expression inside the product is the Poisson probability mass function and represents the probability that a neuron would fire n spikes at each location in the environment.

$$P(\mathbf{n}|\mathbf{x}) = \prod_{i=1}^N P(n_i|\mathbf{x}) = \prod_{i=1}^N \frac{(\tau f_i(\mathbf{x}))^{n_i}}{n_i!} \exp(-\tau f_i(\mathbf{x})) \quad (2)$$

Once the conditional probability distribution, $P(\mathbf{n}|\mathbf{x})$, has been calculated, Bayes' rule (**Equation 3**) can be used to calculate the likelihood of each location in the environment given the spike count vector ($P(\mathbf{x}|\mathbf{n})$). Taking the mode of the likelihood distribution provides the decoded location in the environment at a moment in time.

$$P(\mathbf{x}|\mathbf{n})P(\mathbf{n}) = P(\mathbf{n}|\mathbf{x})P(\mathbf{x}) \quad (3)$$

$$P(\mathbf{x}|\mathbf{n}) = P(\mathbf{x}) \left(\prod_{i=1}^N f_i(\mathbf{x})^{n_i} \right) \exp(-\tau \sum_{i=1}^N f_i(\mathbf{x})) \quad (4)$$

The normalizing constant, $P(\mathbf{n})$, has been dropped from **Equation 4** because it doesn't depend on \mathbf{x} and therefore does not affect the mode of $P(\mathbf{x}|\mathbf{n})$. $P(\mathbf{x})$ is the prior probability distribution over \mathbf{x} , and the decision to use a non-uniform prior distribution depends on the application. Non-uniform priors improve the decoding of an animal's location when the prior

is set to the likelihood distribution calculated at the previous time step ($P_t(\mathbf{x}) = P_{t-1}(\mathbf{x}|\mathbf{n})$). However, if the neural activity is representing locations that begin a new sequence or jump away from the animal's location, as is the case during replay events, using a non-uniform prior can lead to inaccurately decoded paths.

Zhang et al. (1998) showed that the Bayesian decoding approach both with and without a prior, was more accurate compared to template matching for decoding an animal's location. However, the accuracy of Bayesian methods rapidly decreased when the time window was reduced beyond 500 ms, similar to the template matching method, indicating the poor temporal precision of these algorithms.

New method for decoding paths in the environment

Going from decoding single locations to decoding paths represented by neural sequences is a non-trivial step. First, there are additional pieces of information that need to be extracted from the data. The beginning and end of the sequence need to be identified. In the case of theta sequences, neural sequences generally occur during each theta cycle and the beginning and end of the theta cycle can be used as a guide for determining the ends of the sequence. However, replay sequences have variable lengths and begin at seemingly random times during pauses, thus it can be a challenge to accurately identify the beginning and end of the sequence. Furthermore, it must be determined if the neural activity represents a coherent path and not just random locations.

The decoding of neural sequences is also very sensitive to limited data. Because the paths represented by neural sequences are 10-20 times sped up compared to behavior, activity must be analyzed in small time windows in order to decode the path. Theta sequences last on average about 125 ms and replay sequences last 60-500 ms. Thus, time windows need to be small and in the 10-20 ms range in order to decode the path. Using time windows in the 500-1000 ms range, as required for accurate template or Bayesian decoding, would blur the path and completely miss its temporal structure.

Given the challenges presented above, I will now describe a new method I developed to decode the behavioral content represented by neural sequences while animals ran laps on a modified T-maze with a left and a right loop (**Figure 5**). The algorithm required the maze to be linearized, such that the animal's location on the maze was described by a single scalar value (Schmitzer-Torbert and Redish, 2004). This linearization and the algorithm described below were performed separately for the left and right loops of the maze and separately for detecting forward and backward sequences. The input to the algorithm was the neural activity of each place cell over time and the location of the place field centers of each place cell (cells could have multiple place fields).

This algorithm was used to detect the beginning and end of replay sequences, determine if the path was forward or backward, and decode the spatial path represented. Results are presented in **Chapter 6**. The same algorithm was used to decode theta sequences (**Chapter 7**), except that sequences were restricted to the time period between the peaks of each theta cycle and only forward paths were detected.

Sequence detection.

Stepping through each recording session in 10 ms steps, the sequence score of place cell activity within a flexible time window (40 ms to 600 ms, to accommodate longer sequences) was calculated as follows. Using the place field centers and spike times, each spike in the flexible time window was pairwise compared with other “local” spikes in the same time window. A “local” spike is a spike that occurred within 80 ms of its neighbor and whose place field center was separated by less than 80 cm along the linearized path (this 80 cm restriction was subsequently removed during significance testing (see below) to ensure that all spikes contributed equally to the scoring of the sequence). Thus, within each flexible time window, a “local region” was defined around each spike and the score contributed by each spike came from comparing it with the other spikes within its local region. If the place field center corresponding to spike A was traversed before the place field center for spike B (positive value in ΔC matrix, see below) and spike A occurred before spike B (positive value

in ΔT matrix), the forward score was +1 and the backward score was -1. Within the flexible time window, all pairwise comparisons were summed to determine the cumulative score of the forward or backward sequence.

Basic sequence scoring algorithm:

Given:

1. Spike times: T vector: [nSpikes x 1]
2. Spike place field centers: C vector: [nSpikes x 1]

Compute:

1. Time difference for each spike pair: ΔT matrix: [nSpikes x nSpikes]
2. Field center distance for each spike pair: ΔC matrix: [nSpikes x nSpikes]
3. Element-wise multiply ΔT and ΔC to create ΔTC matrix: [nSpikes x nSpikes]
4. Binarize ΔTC . For each element in ΔTC : if element is positive, replace with '+1'; if element is negative, replace with '-1'.
5. Sum over all elements in ΔTC to get the sequence score.
 - A positive sequence score indicates that the sequence is 'forward' ordered. A negative score indicates that the sequence is 'backward'.

The scoring presented above requires that each spike is associated with a single place field. For cells with multiple place fields, it is unclear which place field is being represented by a spike. In order to address this issue, the algorithm assigned each spike from a multi-field cell to the place field which best fit with the rest of the firing sequence. This was performed by maximizing the forward score (for forward replays) and maximizing the backward score (for backward replays) over each place field available to a spike. Always picking the place field that optimizes the replay score could result in artificially high sequence scores. In order to provide a fair and statistically justified comparison, this same score maximization

procedure was performed for the control sequences produced in the bootstrapping procedure described below.

The time window was allowed to keep increasing until there were five consecutive time window steps (50 ms) without an increase in the score or the window reached 600 ms. Once the beginning and end of the sequence were detected, along with the place fields contributing positively to the sequence score, the spatial coverage of replays was defined: the coverage of each replay event was represented as an X -element binary vector (where X is the number of spatial pixels in the environment), with a '1' at each position where a contributing place field existed.

Sequence significance testing.

The preceding algorithm resulted in a series of time windows (start and stop times for each replay), place field center-labeled spikes, and scores for each putative forward or backward sequence occurring on left and right laps during a single session. These sequences were then analyzed to identify significant replay events using two independent bootstrapping procedures. The first method preserved the identity of the cell to which each spike belonged, and randomly shuffled the spike times for each event 300 times. Each shuffled event had its sequence score recomputed, once again assigning spikes from cells with multiple place fields to the place field center that maximized the sequence score. This ensured that there was no bias resulting from assigning spikes to fields that maximized the sequence score. The second bootstrapping procedure preserved the spike trains, but randomly reassigned each spike train to a different cell's firing field(s) (Davidson et al., 2009). Each event was shuffled 300 times, once again assigning spikes from cells with multiple place fields to the place field center that maximized the sequence score. If the unshuffled replay score was greater than 95% of both independent sets of shuffled sequences, the replay was deemed significant and included for analysis. These methods ensured that only significant replay events were included in the analysis. Because sequences were initially identified for left and right laps separately, identical pairs (i.e. replay pairs occurring over

the central stem) were merged. This method provides an effective and statistically justified means of identifying sequences directly from the recorded spike trains.

Comparison of our method with Bayesian decoding.

This method for identifying sequences was used instead of a Bayesian decoding approach for two reasons. First, it can identify sequence structure even when a small number of neurons are co-active by considering neural activity over a relatively large time window, without sacrificing spatial precision. Bayesian decoding relies on the activity of several cells during a small time window (on the order of 10-20 ms for SWR replay events) in order to accurately and precisely decode the spatial locations being represented. The decoded locations over time are then analyzed to determine if a forward or backward behavioral sequence is present, which relies on the quality of the decoding. As mentioned above, large time windows are often needed to accurately decode locations using the Bayesian approach. The method described here operates in the opposite order. Pair-wise comparisons of spikes separated by less than 80 ms are used to identify behaviorally meaningful sequential structure in the pattern of place cell activity. Once the sequence has been identified, then the trajectory corresponding to the entire sequence is decoded by combining the score-maximizing place field from each spike contributing positively to the sequence score. Since whether a spike contributes positively to the sequence score depends on its relationship with all spikes occurring within 80 ms of it, data over a relatively large time window is considered, which enables accurate decoding even with sparse data. Since this large time window is used to detect sequential structure, not spatial location, the precision of spatial decoding is not affected.

Second, the method described here effectively handles the multi-field place cells present in complex mazes such as the 2T maze. When a multi-field cell fires, the Bayesian decoding approach increases the decoded probability distribution over all the cell's fields. This is problematic when many cells have multiple place fields and only one or two cells are active in a 15-20 ms time window, and results in inaccurate spatial decoding and subsequent

degradation of sequence identification. The sequence detection method described here assigns each spike emitted by a multi-field cell to the place field that fits best within the sequential structure of the replay (as described above). This produces an improvement in trajectory decoding.

5 General methods used in electrophysiology experiments

In this chapter, I discuss the experimental and analysis methods used to study hippocampal replay sequences (**Chapter 6**) and theta sequences (**Chapter 7**). The majority of these methods were also described in Gupta et al. (2010).

Behavioral task: The 2T maze

Animals were trained on the 2T-Maze (**Figure 5**). The task consisted of two T intersections, with return rails after the second turn, making it a lap based task in which the environment was experienced in one direction. Food was delivered at 2 sites on each return rail contingent on the animal's choice at the second turn. On left laps the first reward site delivered banana-flavored pellets (5TUL-banana, TestDiet, Richmond IN) and on right laps the first reward site delivered fruit-flavored pellets (5TUL-fruit). On both left and right laps the second reward site delivered white unflavored pellets (5TUL). Each training and recording session lasted 40 min. Training on the task was performed in two phases. During phase one, rats were trained to run laps on one side of the maze, while the other side was blocked. After running at least 40 laps on two consecutive days, phase two of training began. Blocks were removed and on a given day, rats had to run all left laps (L), all right laps (R), or alternating left and right laps (A) in order to receive reward. After consistently getting 80% of the laps correct on all three tasks, rats were implanted with hyperdrives. During initial training, rats attempting to run backwards on the maze were blocked by the experimenter. While there were occasions in which rats faced backwards during their early experiences, rats rarely ran backwards on the track. After recovery from surgery, phase two of training resumed until the rats regained proficiency and tetrodes were in the cell layer. At this point, training ended and a six-day sequence of recording sessions began. During the six recording sessions, the task contingency changed approximately midway through the session (mean: 18.07 min \pm 1.13 min (SD)). There were six recording sessions to allow for all possible pairings of the




Figure 5: **(A)** The two-choice T maze. The maze had two possible physical configurations, the second indicated by dotted green lines. Noteworthy locations on the maze are labeled as follows: maze start (MS), turn 1 (T1), turn 2 (T2), feeder 1 (F1), and feeder 2 (F2). In addition the task entailed three reward contingencies reflecting a decision made at the second choice point (T2): animals were trained to turn left at the final choice, right at the final choice, or to alternate on a lap-by-lap basis. During recording days, the contingency was changed approximately midway through the task. **(B)** Correctness of final choice (T2) aligned to the start of the session. If the task was alternation, the first lap was always deemed correct, so chance performance over all tasks is $2/3=66\%$. **(C)** Correctness of final choice (T2) aligned to the time of contingency switch. Chance of a pre-switch correct behavior being a correct behavior after the switch is $1/6=16\%$. (From Gupta et al., 2010.)

three tasks (L-R, R-L, L-A, A-L, R-A, A-R). The order in which each rat experienced the six potential pairings was randomized across animals. Animals learned the contingency quickly at the onset of the recording session and again after the contingency switch (**Figure 5B,C**).

Surgery, recording, and histology

Subjects.

Six male Fisher-Brown-Norway hybrid rats (Harlan, Indianapolis, IN, age 7-10 months at time of implantation) were maintained on a synchronous day/night cycle. Rats were food deprived to no less than 80% of their body weight during behavioral training and water was available *ad libitum* in the home cage at all times. All procedures were in accordance with National Institutes of Health guidelines for animal care and were approved by the Institutional Animal Care and Use Committee at the University of Minnesota.

Surgery, recording, and histology.

After pre-training to proficiency (19-24 days), three rats were implanted with a single

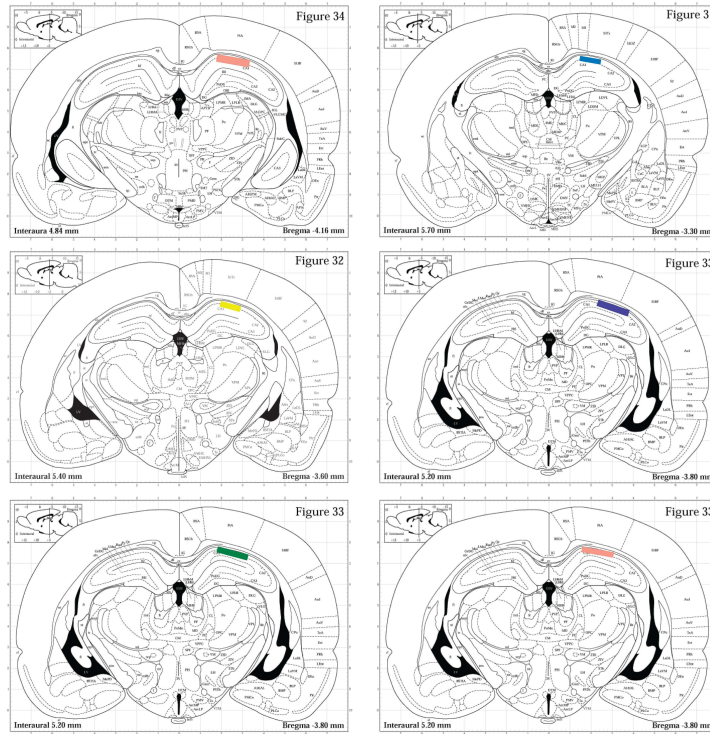
bundle 12-tetrode, 2-reference microdrive (Neuro-Hyperdrive; Kopf, Tujunga, CA) directed toward the CA1 hippocampal (HC) subfield (3.8 mm posterior and 2.5 mm right-lateral from bregma). Another three rats were implanted with a double bundle 12-tetrode, 2-reference microdrive directed toward CA1 and ventral striatum (CA1 targets 3.8 mm posterior and 2.5 mm right-lateral from bregma, only HC data analyzed here). The remaining details of the surgery were as presented previously (Jackson et al., 2006). Tetrodes and references were slowly advanced toward the pyramidal cell layer over approximately two weeks after surgery. For HC-only implants, one reference was lowered to the HC fissure and one was left in corpus callosum or a quiet region of cortex to be used as a superficial reference. For the dual-structure implants, the HC reference was placed in the fissure and the ventral striatum reference was placed near the corpus callosum; spike data was referenced against the HC reference and LFP data against the ventral striatal reference. Neural activity was recorded and spikes were sorted into putative cells as presented previously (Jackson et al., 2006). All units were examined to have a reasonable Isolation distance and L-ratio (Schmitzer-Torbert et al., 2005), and no spikes during the two second refractory period. The median Isolation distance and L-ratio were 21 and 0.09, respectively. Only well isolated units were included in the analyses. Recordings from an example tetrode are shown in **Figure 6**.

After task performance, rats were overdosed on Nembutal (150 mg/kg, i.p.) and perfused intracardially with formalin. After 24 h in formalin, brains were transferred to a 30% sucrose-formalin solution, sliced, and stained with cresyl violet using standard procedures. All HC recording locations were verified to lie in the CA1 region of the dorsal HC.

Place fields

Cells that fired more than 15000 spikes over the 40 minute session (6.25 Hz) were excluded in the analysis to filter putative inhibitory interneurons. Additionally, cells with fewer than 100 spikes (approximately 0.04 Hz) were excluded. Sessions with > 40 cells were considered for analysis (31 out of the 36 recording sessions). A total of 2046 cells were isolated from

A



B

Tetrode 4 taken from session R153-2008-09-15

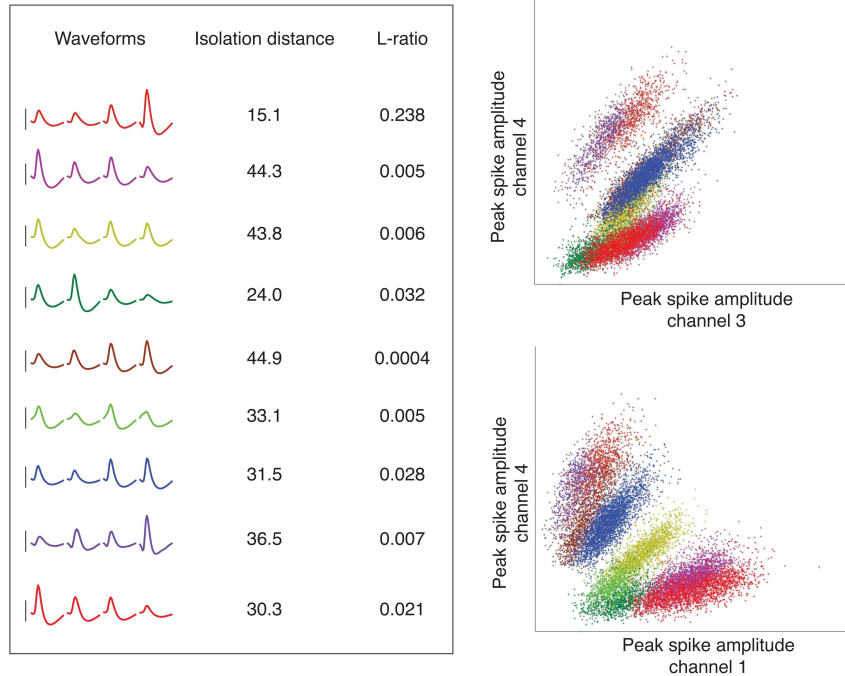


Figure 6: (A) Schematics showing the position of implanted electrodes directed towards hippocampus for each of the six rats. **(B)** Example tetrode recordings from one recording session. Each spike cluster is shown with different colors on two projections (right). One millisecond waveforms for each of the four tetrode channels are plotted for each separable cluster along with the cluster's Isolation distance and L-ratio (left). Scale bar next to each cluster's waveforms is 100 μ V.

31 recording sessions. 1881 cells satisfied the number of spikes criteria and 1470 cells contained at least one place field on the maze. For each session separately, position along the maze was linearized separately for left and right laps such that the rat's position along a lap could be described by a single scalar value (Schmitzer-Torbert and Redish, 2004). Place fields were then identified as contiguous linear pixels (one linear pixel is approximately 3.5 cm along the linearized maze) with average activity $> 5\%$ of the maximum rate observed over the session for any cell at any pixel (cells could have more than one place field, although place fields separated by only a single pixel were merged). Low time speeds were not excluded in place field identification. In this way we identified 2729 place fields from the 1470 place cells. This allowed for the determination of place field centers along left and right laps. The centers were ordered from maze start (MS) to MS in the direction the rats traveled around the maze.

Replay sequence specific analyses

Place field centers obtained as described above and spike times were used to detect both replay sequences and theta sequences (**Chapter 6** and **Chapter 7**). For replay sequences, the algorithm was performed 4 times separately for each session (left lap centers, forward sequences; left lap centers, backward sequences; right lap centers, forward sequences; right lap centers, backward sequences).

Classification of same-side and opposite-side replays.

Once sequences were decoded as described in **Chapter 4**, replay sequences were identified as occurring on one side of the maze ($> 10\%$ coverage on the non-central stem of one side and $< 5\%$ coverage on the other) or else were labeled as central-stem replays. If a replay occurred over the right portion of the maze while the rat was at a right-side reward site, it was labeled a same-side replay. If the rat was at a right-side reward site and the replay covered a trajectory on the left side of the maze, it was labeled an opposite-side replay. Same-side and opposite-side replays were defined analogously when the rat was

sitting at a reward site on the left side of the maze. The vector representation of each replay was used to create the spatial distribution plots in **Figure 10** (page 51).

Simulating scenarios to test the relationship of replay with experience.

Three scenarios based on three proposals (1. replays reflect the most recent experience; 2. replays reflect accumulated experience within a session; 3. replays are independent of experience) were simulated and compared with our data. As described above, each observed replay was classified as occurring on the left, right, or central stem portion of the maze. This information was then used to identify same-side and opposite-side replays based on the rat's location at the time of the replay. In order to generate a comparable dataset that was representative of the three scenarios, the actual replay times and rat locations were kept constant and the left/right classification was changed for each individual replay (central stem replays were left intact). For scenario 1, replays were restricted to cover only the parts of the maze that were traversed over the last lap (i.e. all replays occurring while the animal is paused at a reward location on the left side of the maze after performing a left lap were assigned to the left side of the maze). For scenario 2, replays stochastically preferred portions of the maze that were most visited at the time of the replay (i.e. regardless of the rat's paused location on the maze, the probability of a replay being assigned to the left side of the maze was equal to the proportion of accumulated experience on the left side of the maze at the time of the replay). For scenario 3, each non-central stem replay was randomly assigned to be on the left or right, independent of recent or accumulated experience during the session. From these new assignments, same-side and opposite-side classifications were made and time since last experience, laps since last experience, and proportion of accumulated experience on the side of the replay were calculated for each scenario. Each of these distributions was calculated 100 times and compared to the actual data using Kolmogorov-Smirnov tests. The mean p-value and mean distribution was reported for each comparison (see **Figure 11A-C** and **Figure 12A** on pages 54 and 56).

Additionally, we performed 1-step Bayesian decoding (Zhang et al., 1998) with a uniform

spatial prior and a 20 ms sliding time window on the identified replay times. Replays were classified as occurring on the side of the maze with the largest cumulative probability and the results from **Figure 11A-C** and **Figure 12A** were replicated (**Figure 13** on page 57). See **Chapter 4** for a discussion of my decoding method compared with Bayesian decoding.

Shortcut detection and analysis.

A shortcut sequence on the two-choice maze can be viewed as a combination of one forward and one backward replay that are temporally and spatially adjacent on the top or bottom of the maze. There are 32 possible combinations of forward and backward replays on the maze. There are 4 types of forward replays: top left (TL), top right (TR), bottom left (BL), bottom right (BR). Each of the forward replays can be paired with the 4 types of backward replay (TL, TR, BL, BR), to make $4 \times 4 = 16$ combinations. Since order matters, there are $16 \times 2 = 32$ possible combinations. Out of these combinations, 4 give rise to shortcut sequences (bwdTR-fwdTL, bwdTL-fwdTR, fwdBR-bwdBL, fwdBL-bwdBR), 8 give rise to a coherent trajectory that does not result in a shortcut path, and 20 do not create a coherent trajectory at all (see **Figure 15**). Thus the theoretical probabilities of shortcut, non-shortcut, and disjoint combinations are $4/32$, $8/32$, and $20/32$, respectively. These probabilities assume that forward and backward replays are evenly distributed on the top and bottom of the maze. We used a bootstrapping procedure to determine the actual probabilities of each combination from the distributions of replays observed. This procedure randomly paired each forward replay with 100 backward replays and each backward replay with 100 forward replays. The probability of shortcut, non-shortcut, and disjoint combinations were then calculated from the number of occurrences of the 32 possible combinations.

A simple algorithm was then used to identify the actual numbers of shortcut, non-shortcut, and disjoint combinations present in the data. The algorithm first identified replays that occurred primarily in each corner of the maze (TL, TR, BL, BR). Using this set of replays, the algorithm identified all pairs of forward and backward replays whose mid-times were separated by less than 150 ms (the mean duration of the replays). Other methods of

identification produced similar results. The algorithm identified 16 shortcut sequences, 3 non-shortcut sequences, and 0 disjoint pairings. To test the likelihood that 19/19 pairings would be contiguous sequences based on the replay distributions in our data, we calculated the probability that the random alignment of forward and backward replays could give rise to 0/19 disjoint sequences (binomial test, see **Table 1** on page 59). Out of the 32 combinations, 12 give rise to sequences (back-to-back replays in which the second replay picks up where the first left off), 4 of which are shortcut sequences. To test whether the shortcut sequences resulted from two back-to-back replays in which the second replay picked up where the first one left off, we calculated the probability under this theory that 16/19 sequence pairings would result in shortcuts (binomial test, see **Table 1**).

Theta sequence specific analyses

For theta sequences, the sequence detection algorithm (**Chapter 4**) was performed 2 times separately for each session (left lap centers, forward sequences; right lap centers, forward sequences). Since theta sequences occur during each theta cycle, the beginning and end of the sequences were confined to within a single theta cycle.

Path, Pro, and Ret length.

As described above, the sequence detection algorithm resulted in start and end times (and start and end theta phases), place field center-labeled spikes, and scores for each sequence occurring on left and right laps during a single session. For theta sequence analyses, there were three important quantities: prospective (pro) length (distance from the animal's location to the end of the represented path in space), retrospective (ret) length (distance from the beginning of the spatial path to the animal's location), and overall path length (length of the represented path). The represented path in space was determined from the place field center-labeled spikes. The mean place field center location of the first two spikes in the sequence was the start location of the path and the mean place field center location of the last two spikes in the sequence was the end location of the path.

Results were qualitatively identical when only the first and last spike or the first and last three spikes were used to determine the start and end location. These locations defined the spatial path represented by each theta sequence. Path length was computed by taking the distance between the start and end location, Pro length was computed by taking the distance between the animal's location and the end location, and Ret length was computed by taking the distance between the start location and the animal's location.

Bayesian decoding. The main theta sequences results were replicated using a Bayesian decoding approach (Zhang et al., 1998) to identify the path represented during each theta cycle. Decoding was performed on each theta cycle using a 40 ms sliding time window (shifting by 10 ms at each step). Decoding was restricted to an approximately 190 centimeter region centered at the animal's location (half a lap). Thus, decoding over each theta cycle resulted in an $n\text{TimeBins} \times n\text{SpaceBins}$ matrix of probabilities, with each row containing the probability distribution over space for a particular time bin.

To detect the beginning of the spatial path represented, the probability distributions over space for the first half of the theta cycle (matrix size $n\text{TimeBins}/2 \times n\text{SpaceBins}$) were summed to create a single probability distribution over space (a $n\text{SpaceBins}$ -element vector). The 5% tail of this distribution (tail on the end of the distribution that is behind the animal) was located, providing the beginning location of the spatial path. Similarly, the end of the represented path was determined by summing the probability distributions over space for the second half of the theta cycle and locating the 5% tail of that distribution (tail on the end of the distribution that is ahead of the animal). The Bayesian decoding results presented in **Figure 18** (page 70), **Figure 20** (page 74), and **Figure 23** (page 81) were robust to other time windows (20 and 30 ms), other tail sizes (10%, 15%, and 20%), and were similar when using the first and last time bin for determining the ends of the spatial path or when summing the spatial distributions over the entire theta cycle rather than the first and second half separately.

6 Hippocampal replay sequences

This chapter contains a reproduction of original work published by my colleagues and me: *Hippocampal replay is not a simple function of experience* (Gupta et al., 2010). In this paper, we reported several novel findings and suggested an additional way that replay may contribute to hippocampal function (see commentary by Derdikman and Moser, 2010).

6.1 Summary

Replay of behavioral sequences in the hippocampus during sharp wave ripple complexes (SWRs) provides a potential mechanism for memory consolidation and the learning of knowledge structures. Current hypotheses imply that replay should straightforwardly reflect recent experience. However, we find these hypotheses to be incompatible with the content of replay on a task with two distinct behavioral sequences (A and B). We observed forward and backward replay of B even when rats had been performing A for >10 minutes. Furthermore, replay of non-local sequence B occurred more often when B was infrequently experienced. Neither forward nor backward sequences preferentially represented highly-experienced trajectories within a session. Additionally, we observed the construction of never-experienced novel-path sequences. These observations challenge the idea that sequence activation during SWRs is a simple replay of recent experience. Instead, replay reflected all physically available trajectories within the environment, suggesting a potential role in active learning and maintenance of the cognitive map.

6.2 Results

To investigate the relationship between experience and the content of hippocampal replay, we recorded neural ensembles from the CA1 region of hippocampus in 6 rats trained on a spatial decision task with two possible paths (**Figure 5A**), allowing for the separation of re-

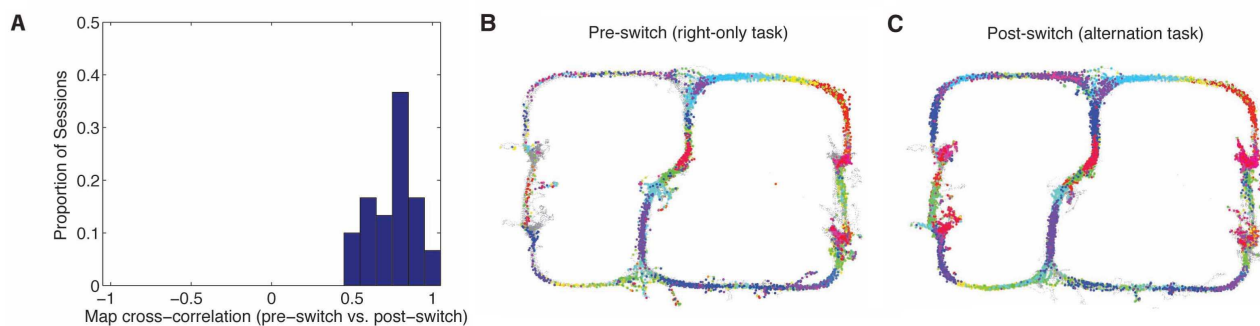


Figure 7: Place fields did not change between contingencies. Each session was split at the time of the switch in contingencies, and place fields were determined for each half-session. Similarity was measured as the cross-correlation between the two place field distributions (as in Barnes et al., 1997). **(A)** Histogram showing the similarity between the pre and post-switch spatial tuning curves for each recording session. Cross-correlation was consistently high, comparable to using the same map between half-sessions (compare Barnes et al. 1997). **(B)** Place fields from an example pre-switch half-session. **(C)** Place fields from the corresponding post-switch half-session. Each cell's place fields are plotted in the same color between the two half-sessions. The place fields remain in the same locations after the switch, even though the task changes from right-only to alternation. (From Gupta et al., 2010.)

cent and non-recent experiences. The maze consisted of two choices with a daily-changing contingency: the rat was rewarded for turning left (L), turning right (R), or alternating left and right turns (A) at the final choice point. During recording sessions, the reward contingency switched approximately midway through the session. Thus results were analyzed in half-sessions which could be L, R, or A. This means that within a session, rats sampled both loops of the maze, but how long ago each loop was experienced and the total amount of experience on each loop of the maze could differ depending on the contingency. Place cells did not "remap" after the switch (i.e. place cells maintained the same place fields after the switch, **Figure 7**). The vast majority of experience on this track occurred in one direction (**Figure 8**).

Forward and backward replay of local and non-local trajectories

We analyzed sequences (replays) of place cell activity during sharp wave-ripple (SWR) events as rats paused at the reward locations. A sequence detection algorithm was applied to quantify the degree to which spiking activity during SWRs represented a coherent se-




Figure 8: Rats primarily ran forward on the maze. Direction of travel was measured from sequential position samples using the Janabi-Sharifi algorithm (Janabi-Sharifi et al., 2000; Masimore et al., 2005), providing a velocity vector $\langle dx, dy \rangle$ for each position sample. Spatial points were gridded on a 64×48 grid. For each grid point, the average velocity vector for all position samples within 5 pixels (1.7 cm) was calculated, giving an average velocity vector for each grid point $\langle \widetilde{dx}, \widetilde{dy} \rangle$. Then for each grid point, the proportion of position samples with direction of motion velocities in the opposite quadrant of the mean direction ($\text{sign}(dx) = -\text{sign}(\widetilde{dx})$ & $\text{sign}(dy) = -\text{sign}(\widetilde{dy})$) was measured. Only grid samples with > 5 position samples in range were included in the analysis, but every session (including all training and test sessions) for each animal was included. Plots show the proportion opposite for all grid points from all sessions. A small random jitter was added to the x and y coordinates for the grid points for display only. For display, points were sorted so that the brightest colors (highest proportion opposite) were on top for display. Most areas include a vast oversampling of forward relative to backward directions. (From Gupta et al., 2010.)

quence, identifying 1719 replays across 31 sessions from 6 rats (see Methods, **Chapter 5**). The median number of active cells during these events was 10 with a median “sequence score” of 90 (see Methods, **Chapter 5**). In general, replays varied in their starting location and direction and were identified as occurring on either the left, right, or central stem of the maze. In line with previous reports (Foster and Wilson, 2006; Diba and Buzsáki, 2007; Karlsson and Frank, 2009) we observed replays occurring over spatially proximal portions of the environment (rat sitting on the same side of the maze as the replayed trajectory, “same-side replay”, $n = 869$, 450 forward, 419 backward, **Figure 9A**). However, we also observed frequent replays of spatially distal portions of the environment (rat on opposite side of the maze as the replay, “opposite-side replay”, $n = 534$, 307 forward, 227 backward, **Figure 9B**). The median number of active cells (MAC) and the median sequence score (MSS) were similar across all groups (same-side, forward: MAC = 8 and MSS = 65; same-side, backward: MAC = 10 and MSS = 85; opposite-side, forward: MAC = 10 and MSS = 86; opposite-side, backward: MAC = 11 and MSS = 96). These data show that forward and backward replay robustly occurred over both loops in the maze when the animal paused at a reward site on one side of the maze. This is consistent with recent studies showing re-activation of experiences from previous environments (Jackson et al., 2006; Gelbard-Sagiv et al., 2008), replay of forward sequences from previous environments (Karlsson and Frank, 2009), and backward replay of trajectories that don’t begin at the animal’s location on a linear track (Davidson et al., 2009). However, unlike backward replay on linear tracks, this result demonstrates that backward replay robustly occurred in an environment experienced in one direction, over trajectories that had never been experienced in the direction of replay. While rats would sometimes face backwards during early experiences, for each rat there were backward trajectories that were replayed but never experienced and the vast majority (>96%) of all experience was in the forward direction. Yet, forward and backward replays were observed with similar proportion.

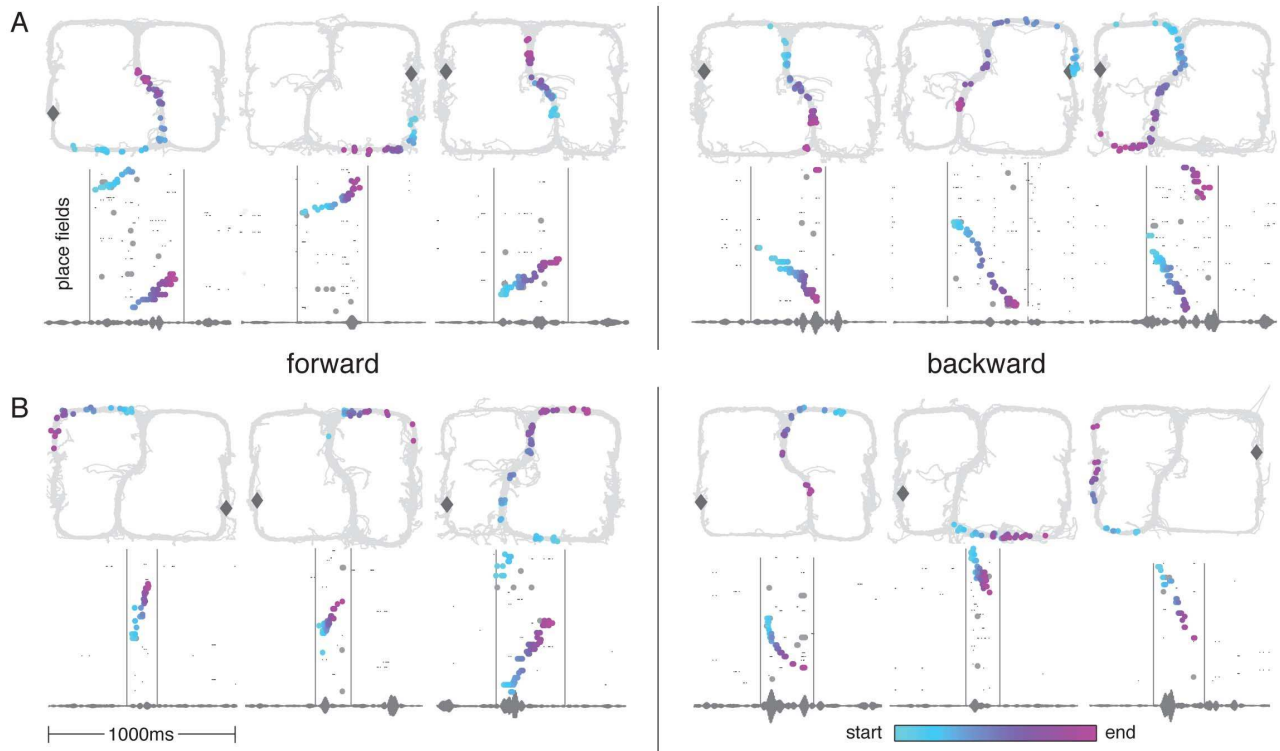


Figure 9: Examples of forward and backward replays. Row **(A)**: Examples of same-side and central stem replays (3 forward, 3 backward). Row **(B)**: Examples of opposite-side replays (3 forward, 3 backward). Gray diamond indicates the rat's location at the time of the replay. On the bottom panel of each subfigure, spikes are plotted by ordered place field (spatial firing field) center (along either a left or right loop of the maze) over a one second period (see Methods, **Chapter 5**). LFPs filtered between 180 Hz and 220 Hz are plotted at the bottom of the panel. Colored points indicate spikes that contribute positively to the sequence score of the replay according to the automatic sequence detection algorithm. The color of the spike indicates its relative time within the replay (light blue = early, light purple = late). Gray points are spikes that do not contribute positively to the score. For cells with multiple place fields, small black points are plotted at every place field center belonging to the cell (colored points occupy the place field center that contributes maximally to the score). Each colored point from the bottom panel is plotted on the 2D maze in the top panel at the location of its 2D place field center. Note that forward replays could begin near the rat's location (row **(A)**, left) or on the opposite side of the maze (row **(B)**, left). Similarly, backward replays could begin near the rat's location (row **(A)**, right) or on the opposite side of the maze (row **(B)**, right). Backward replays occurred over parts of the environment that were rarely or never experienced in the reverse direction (**Figure 8** and **Figure 14**). (From Gupta et al., 2010.)

Spatial distributions of forward and backward replays

Given that replays could start at any location on the maze, we investigated whether certain portions of the maze were preferentially represented by forward or backward replays. While same-side replays occurred as expected, with forward replays ahead of the animal and backward replays behind it, this pattern was reversed for opposite-side replays (**Figure 10**). Here, forward replays preferentially occurred on the segment leading up to reward sites, and backward replays similarly covered trajectories ending near reward sites. Thus, distant/non-local replays not influenced by the animal's current location did not uniformly represent the environment, but preferred certain trajectories. A potential concern is that if the distribution of place fields is non-uniform, more replays will be detected along segments of the track that have more place fields. To test this possibility we compared the spatial distributions of forward and backward replays with the spatial distribution of place fields and found that they were significantly different ($p < 10^{-71}$, forward vs place fields; $p < 10^{-176}$, backward vs place fields; Kolmogorov-Smirnov tests). As shown by the place field distribution in **Figure 10A,B**, the trajectory preferences for forward and backward replay on the top and bottom of the maze were not reflected by the spatial distribution of place fields. Backward replays frequently included the central stem, whereas forward replays did so less often. Overall spatial distributions of forward and backward replays were also significantly different ($p < 10^{-197}$, Kolmogorov-Smirnov test). Since trajectories around a given loop were equally experienced, but not equally replayed, this demonstrates a dissociation between the two, which is inconsistent with the hypothesis that the content of replay reflects the amount of previous experience.

Comparison of data with current proposed mechanisms

Current proposals for the mechanism of replay suggest that it should reflect recent experience (Foster and Wilson, 2006; Jackson et al., 2006) or total experience (Buzsáki, 1989;

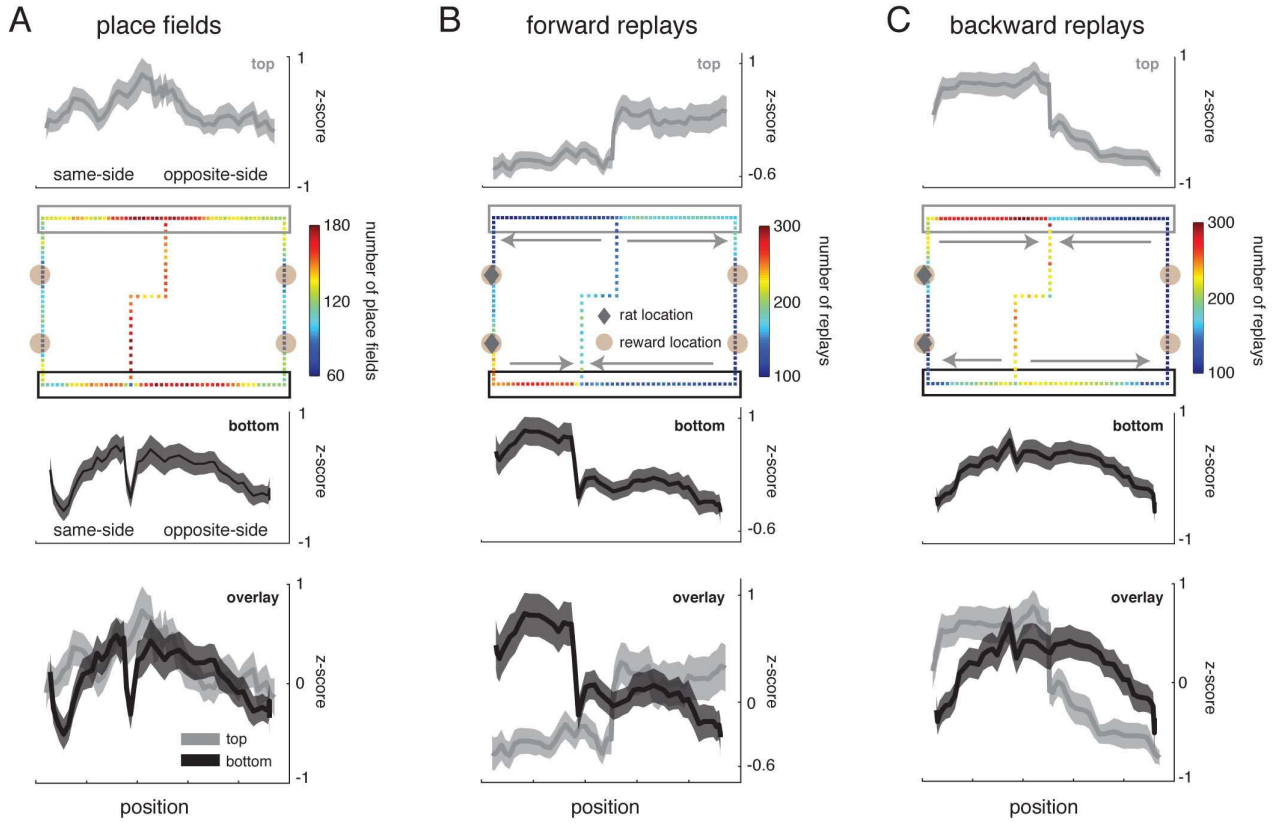


Figure 10: Spatial distribution of forward and backward replays. **(A)** Place field spatial distribution. **(B)** Forward replay spatial distribution. **(C)** Backward replay spatial distribution. Spatial distributions over the entire environment are shown in the second row. The first and third rows display the spatial distribution over the top and bottom of the maze (indicated by gray and black boxes in row 2), respectively. Spatial distributions over the top and bottom of the maze are overlaid for comparison in the fourth row. All mazes were flipped and aligned such that the animal's location was always on the left side of the maze (indicated by gray diamonds) at the time of the replay. Therefore, the spatial distribution on the left side of the maze reflects same-side replays and the distribution on the right reflects opposite-side replays. The pixel color indicates the total number of replays that represented that particular location in the environment (see colorbar in image). Errorbars in rows 1, 3, and 4 are SEMs over 31 sessions. Note that forward and backward replays preferentially represented certain portions of the maze, which could not be explained by the place field distribution. Overall distributions were significantly different ($p < 10^{-71}$, forward vs place fields; $p < 10^{-176}$, backward vs place fields; $p < 10^{-197}$, forward vs backward; Kolmogorov-Smirnov tests). (From Gupta et al., 2010.)

Wilson and McNaughton, 1994; Redish and Touretzky, 1998; Redish, 1999; Sutherland and McNaughton, 2000). To test these ideas, we examined the distributions of time and number of experiences that had elapsed between replaying a sequence and the most recent experience of that sequence. Since same-side replays involve a recent behavioral experience, we only considered the 534 opposite-side replays for this analysis. Histograms showing the amount of time elapsed since the last experience and the number of laps that had been run since the last experience for each replayed trajectory are shown in **Figure 11A,B**. The median time since the last experience was 337 s and the median number of laps since the last experience was 9 laps (366 s, 11 laps for forward replays and 298 s, 7 laps for backward replays). A substantial number of forward and backward trajectories were replayed more than 10 minutes or 15 laps after they were last experienced.

We also asked whether replayed trajectories favored the side of the maze that had been the most experienced within a session. For each non-central stem replay (1403 replays), the total number of laps the rat had run on the same side of the maze as the replayed trajectory was divided by the total number of laps that the animal had run on either side of the maze at the time of replay. A histogram of this proportion (the proportion of the rat's total experience that had occurred over the replayed trajectory) is shown in **Figure 11C**. The symmetry of this histogram indicates that poorly and extensively experienced portions of the maze were both replayed with a similar frequency (median of distribution = 0.5 for all replays). Contrary to the predictions of current proposals, replay did not preferentially represent trajectories that had been experienced most often within a session. This also held true for forward replays only and for backward replays only (not shown).

To directly test the recent experience and total experience hypotheses, we compared the data in **Figure 11A-C** with the distributions that would be expected under three scenarios of replay content generation: (1, blue) replays preferentially include the most recent experience, (2, red) replays preferentially include trajectories that are highly experienced within a session, and (3, green) replays are independent of experience (all portions of the

environment have an equal chance of being replayed, independent of recent or total experience). For each scenario, we used the actual rat behavior and replay times, but reassigned the content of each replay depending on the scenario (see Methods, **Chapter 5**). For example, under Scenario 1 (most recent experience), replay was assumed to reflect the most recently experienced lap. Thus, the number of laps since the last experience (**Figure 11B**) was always 1, but the time since the last experience (**Figure 11A**) depended on the time the animal had been paused at the reward sites.

Our observations are not compatible with the hypothesis that replay reflects the most recent experience, nor with the hypothesis that it reflects the accumulated experience within a session. Instead, our data is most similar to what would be expected if replays were independent of experience within a session. (Data vs accumulated experience: laps ($p < 10^{-12}$), time ($p < 10^{-10}$), proportion ($p < 10^{-32}$); Data vs most recent experience: laps ($p < 10^{-54}$) time ($p < 10^{-31}$), proportion ($p < 10^{-34}$); Data vs experience independent: laps ($p < 0.05$), time ($p < 0.05$), proportion ($p < 0.05$), Kolmogorov-Smirnov tests, see Methods, **Chapter 5**.) Since the number of left and right laps that each animal experienced over its lifetime was counterbalanced, a scenario in which replay equally reflects all experiences over the animal's lifetime would look like the experience-independent scenario. Thus, this result is not compatible with the recent or accumulated experience within a session scenarios, but may be explained by a scenario in which replay reflects the animal's total experience over its lifetime. However, the preferential replay of certain trajectories (described above) and the pattern of replays (below) cannot be explained by the experience-independent scenario either.

Task-dependence of the content of replay

Since the task contained alternation (A) as well as left-only (L) and right-only (R) half-sessions, we asked whether the proportion of opposite-side replays was different when the animals were alternating compared to when they were performing laps on only one




Figure 11: Replay content is incompatible with scenarios based on recent experience or accumulated experience within a session. **(A)** Histogram showing the time elapsed since the last behavioral experience over each replayed trajectory. **(B)** Histogram showing the number of laps elapsed since the last behavioral experience over each replayed trajectory. **(C)** Histogram showing the proportion of total behavioral experience on the same side of the maze as each replayed trajectory at the time of the replay. The actual data (gray bars) is compared against three scenarios, where replay coverage is determined by recent experience (blue line), by accumulated experience (red line), or independent of experience (green line). The curves representing each scenario were constructed based on the animals' actual behavior and replay times (see Methods, **Chapter 5**). Thus, the shapes of the curves reflect the finite session length and the behavioral contingencies. For example, the time-since-last-experience curve for the experience independent scenario (panel **A**) slopes downward due to the decreased probability that a replay event and the last experience over the replayed trajectory will be separated by 2000 s in a 2400 s recording session. Similarly in panel **C**, the experience independent curve has three peaks, reflecting the fact that experience was evenly distributed between left and right laps (during alternation) or was primarily on one side of the maze (during left or right-only contingencies). Note that in panel **C**, only the experience-independent scenario yields a good match to these data, because the other two scenarios fail to account for the peak in replays on the poorly-experienced side (black arrow). These results were replicated using a Bayesian decoding approach (**Figure 13**). (From Gupta et al., 2010.)

side of the maze. During alternation half-sessions, experiences on the opposite side of the maze occurred more recently in the past and would occur sooner in the future compared with opposite-side experiences during left-only or right-only half-sessions. The median time since the last experience over the replayed trajectory and the number of laps since the last experience for opposite-side replays during alternation was 110 s and 1 lap, while the median time and laps since the last experience for opposite-side replays during left-only or right-only half-sessions was 382 s and 11 laps. These distributions were significantly different (time since last experience, alternation vs left or right-only: $p < 10^{-4}$; laps since last experience, alternation vs left or right-only: $p < 10^{-7}$, Mann-Whitney tests). The *recent experience* and *total accumulated experience within a session* scenarios predict an increased number of opposite-side replays during alternation sessions and the *experience independent* scenario predicts an equal proportion of opposite-side replays during alternation and left-only and right-only sessions. Contrary to all three of these scenarios, a larger proportion of opposite-side replays occurred during left-only or right-only half-sessions compared to the alternation half-sessions (**Figure 12A, B**, $p < 0.002$, Mann-Whitney test). This result suggests that replays may serve a role in preserving the representation of non-recent trajectories. The results shown in **Figure 11** and **Figure 12** were replicated using a Bayesian 1-step decoding method with a 20 ms sliding time window (Zhang et al. (1998), see **Figure 13** and Methods, **Chapter 5**).

Never-experienced shortcut sequences

19 shortcut sequences were observed over 8 sessions from 3 rats. These sequences traversed a straight path on the top (R152: 1 occurrence; R153: 11 occurrences over 3 sessions) or bottom (R153: 1 occurrence; R158: 6 occurrences over 3 sessions) of the maze between the reward sites. Trajectories spanning the bottom of the maze were very infrequently experienced (R153: 4 times out of approximately 1780 laps; R158: 12 times out of approximately 1550 laps), never rewarded, and at the time of recording a minimum of three

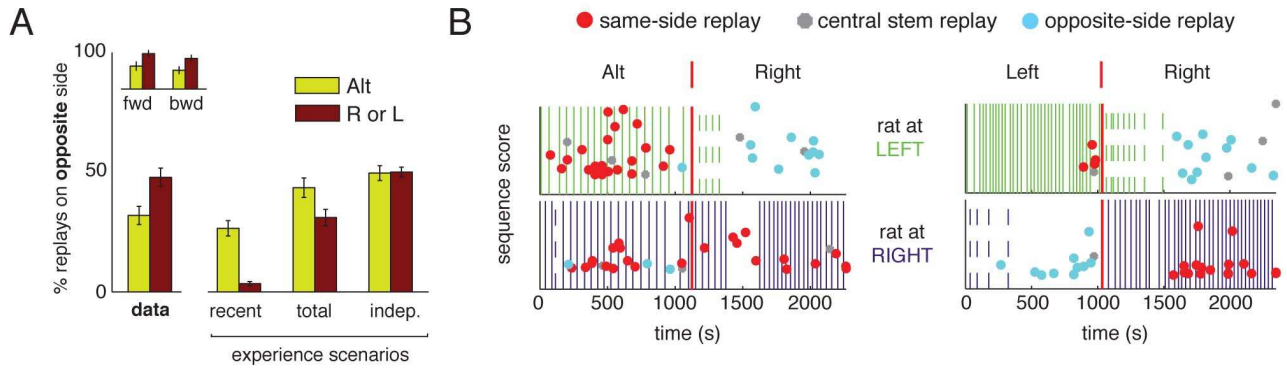


Figure 12: Influence of the behavioral contingency on the content of replay. **(A)** Proportion of opposite-side replays on alternation (A) vs right (R) or left-only (L) contingencies. The “Data” group is shown with all replays combined, as well as with forward and backward replays separately. Whereas the scenarios (described in text) all predicted a higher or equal proportion of opposite-side replays during alternation half-sessions, the data contained a significantly higher proportion of opposite-side replays during right or left-only half-sessions. [Errorbars are SEMs over half-sessions] This result was replicated using a Bayesian decoding approach (**Figure 13**). **(B)** Examples of same-side and opposite-side replays on L and R contingencies versus an A contingency (2 sessions shown). For each session, the top panel shows left lap experiences (vertical green lines), which were followed by pauses at left reward locations, and replays (points) that occurred over the left side of the maze, while the animal was paused at either left or right reward locations. The bottom panel shows right lap experiences (vertical purple lines) and replays (points) that occurred over the right side of the maze. Dashed lines indicate error laps (e.g. the animal performs a right lap when only left laps are rewarded). Opposite-side replays are indicated by blue points, same-side replays by red points, and central stem replays by small gray points. Vertical red lines mark the contingency switch. X axis is time, Y axis is replay sequence score (see Experimental Procedures). Note the large number of blue points (opposite-side replays) during L or R contingencies compared to A. Thus, there was a bias for replaying non-local trajectories when they were infrequently experienced. This is in contrast to all three scenarios for replay content generation, including the experience-independent scenario. (From Gupta et al., 2010.)

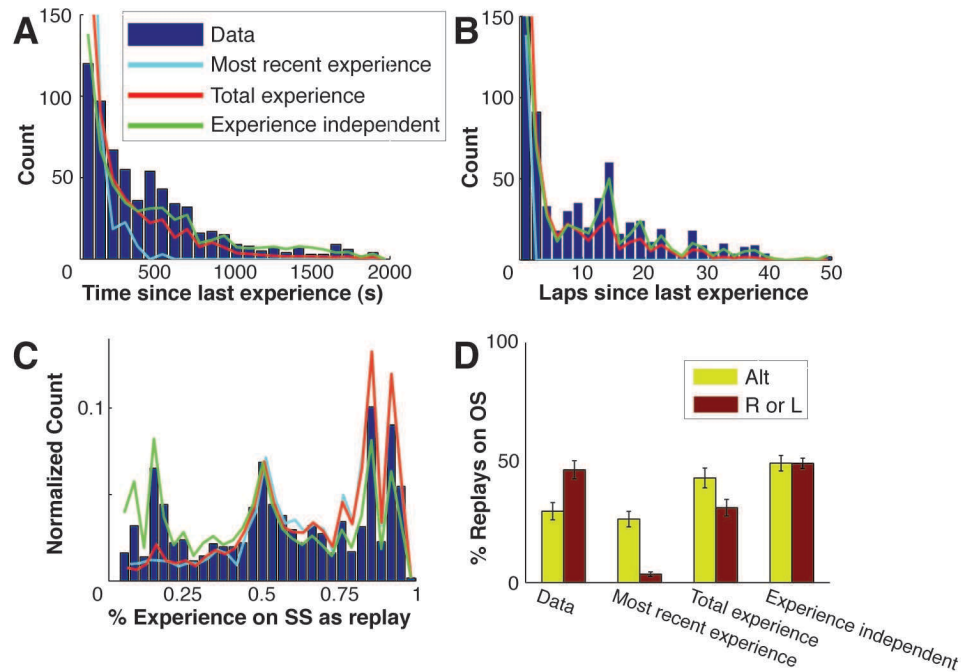


Figure 13: Replay content versus scenarios using Bayesian decoding. (A-C) Comparison of opposite-side replays (classified using Bayesian decoding) with the same three scenarios presented in **Figure 11**. (D) Proportion of opposite-side replays on alternation (A) vs right (R) or left-only (L) contingencies determined from Bayesian decoding of replay sequences. Compare with **Figure 12**. (From Gupta et al., 2010.)

days had passed since the animal had experienced the trajectory. Trajectories spanning the top of the maze were never experienced by one rat (R153) and only experienced 6 times by another (R152). All experience on the maze, including all training sessions, was analyzed (**Figure 14**). Therefore, rat R153 constructed at least 11 never-experienced shortcut sequences across the top of the maze over 3 separate recording sessions (**Figure 15A**). Each of these novel trajectories began close to the animal's location on one side of the maze and ended close to the reward location on the opposite side. Some of the trajectories appear to dip down slightly at the choice point. This is likely to be a consequence of the assignment of place field centers and can be explained by the non-uniform sampling of the region around the choice point (lots of sampling below, but not above) and by the observation that the animals usually cut the corner at the choice point, creating a sampling bias that resulted in downwardly displaced place field centers at the choice point.

In order to control for the possibility that back-to-back forward and backward replays

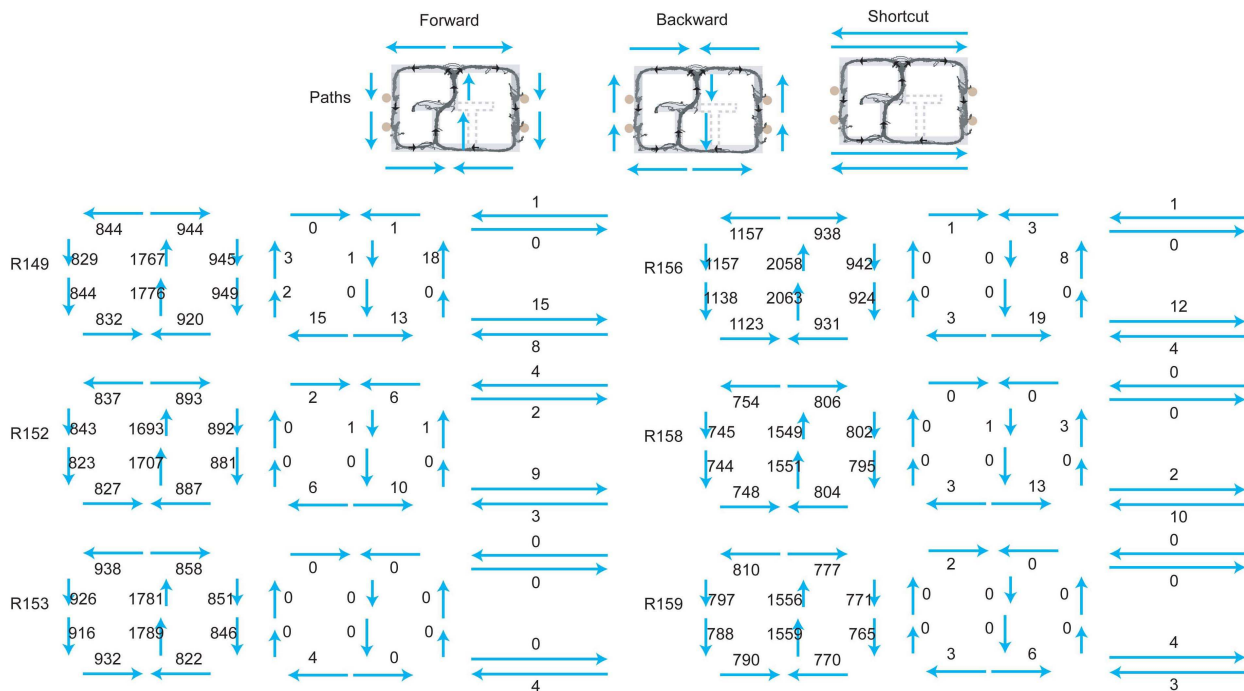


Figure 14: Physical paths taken by each rat. Each maze from each day was divided into nine boxes and sequential sampling of each pair of boxes was measured. Thus, if a rat was at the top right corner of the maze, if the next box entered was the center portion of the return rail on the right side, this would contribute one count to a forward path from top-right to center-right. However, if the next box entered was the choice point (top-center), then this would contribute one count to a backward path from top-right to top-center. Total number of sampled paths were counted over all training and test sessions. The vast majority of paths were forward. Note that some of the backward replay sequences observed occurred over paths never completely experienced in the backward direction (implying that they are truly backward replays, see **Figure 9**). Also note that shortcut replays occurred that were never completely experienced in either forward or backward directions (implying that they are truly novel sequences, see **Figure 15**). There are multiple reasons why the number of trajectories identified around each lap are not identical: tracking errors, clipping of the first few seconds of some sessions (due to tracking errors as the animal is first placed on the maze), removal of the animal at the end of the session at different points on the maze, and possible errors by the detection algorithm. Nevertheless, it is clear that the vast majority of trajectories were forward and that there are trajectories that were replayed but never experienced in the direction of replay. (From Gupta et al., 2010.)

| Sequence type | expected (E) probability | bootstrap probability | observed (O) probability | significance of O-E difference (binomial test) |
|---|--------------------------|-----------------------|--------------------------|--|
| Disjoint sequences | 20/32 = 62.5% | 64% | 0/19 = 0% | 8×10^{-9} |
| Non-shortcut sequences | 8/32 = 25% | 24% | 3/19 = 15% | 0.15 |
| Shortcut sequences | 4/32 = 12.5% | 12% | 16/19 = 84% | 2×10^{-12} |
| Shortcut sequences under back-to-back replay theory | 4/12 = 33.3% | 33% | 16/19 = 84% | 7×10^{-6} |

Table 1: Comparison of shortcuts observed with expected proportions of shortcuts given random pairings of forward and backward replays. (From Gupta et al., 2010.)

aligned to give rise to what appeared to be a shortcut trajectory, we measured the likelihood that the random alignment of forward and backward replays would account for the number of shortcuts actually observed. The chance of this happening was very small ($p < 10^{-8}$, binomial test, see **Table 1**). Similarly, the possibility that the novel trajectories were actually two replays in which the second replay picked up where the first one left off was also unlikely to account for the shortcut trajectories ($p < 10^{-5}$, binomial test, see **Table 1**). Given all possible combinations of forward and backward replays, we calculated the distribution of disjoint, non-shortcut, and shortcut sequences that would occur if forward and backward replays were randomly paired. Both the expected distribution (based on a uniform distribution of forward and backward replays) and the bootstrapped distribution (based on the actual distribution of replays) differed from the observed distribution (**Figure 15B, C**). This strongly suggests that the observed sequences were novel sequence constructions.

6.3 Discussion

By considering the content of what is replayed during a task that manipulates the recency and frequency of experiences, our data speaks to two main issues: the first concerning the relationship between replay and experience within a session and the second concerning the content of what is replayed. We found that within a session the hippocampus replayed both frequently and infrequently experienced trajectories in both forward and backward orders. Both forward and backward replay robustly occurred over trajectories not experienced in

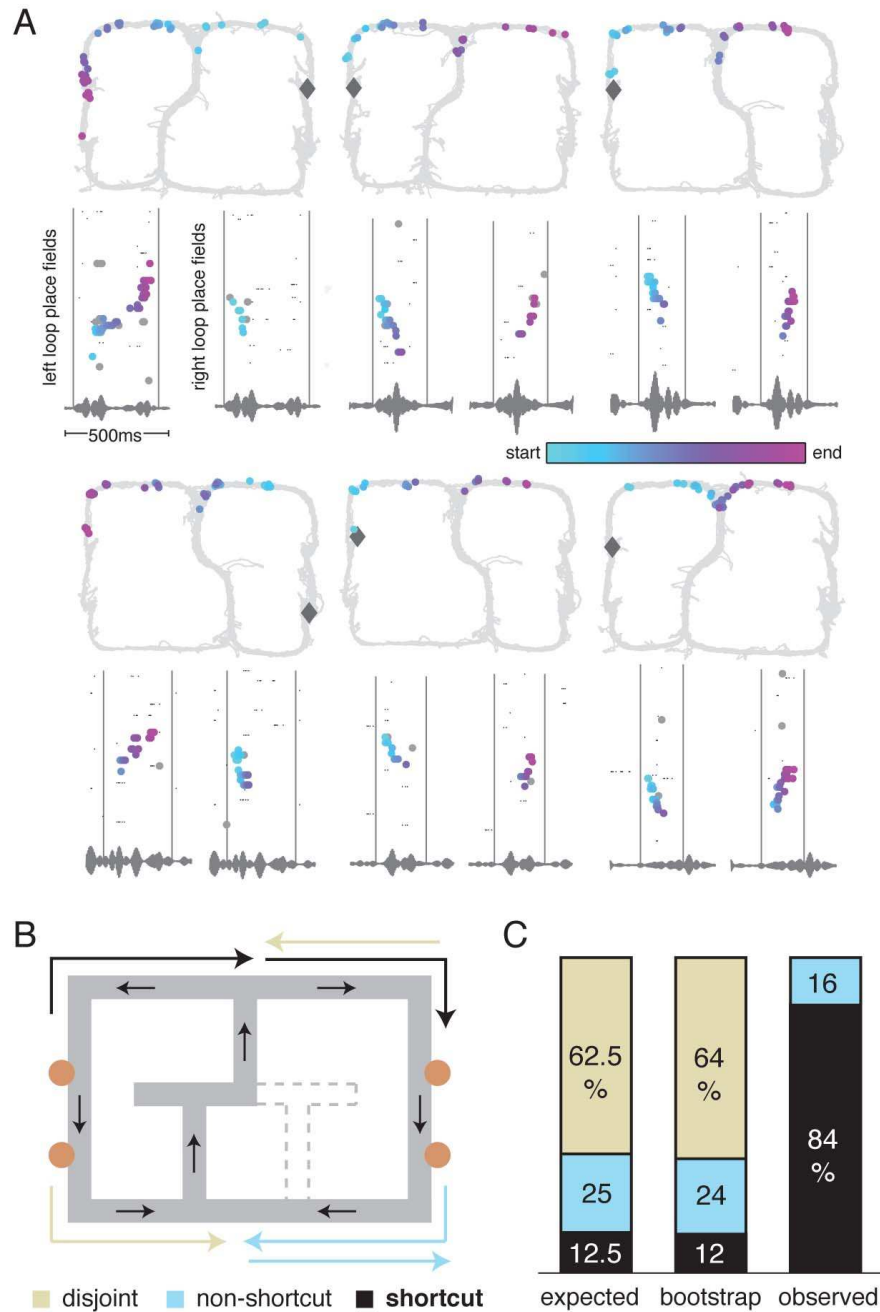


Figure 15: Construction of novel shortcuts. **(A)** Examples of novel trajectories. On the bottom panels, spikes are plotted by ordered place field center for both left and right loops over the same 0.5 s period. The gray vertical lines mark the beginning and end of the shortcut sequence and capture the exact same period of time on both left and right loop raster plots (as can also be seen in the repeated LFP trace). The temporally color coded spikes (as described in the **Figure 9** caption) are plotted on the 2D maze (top panels) to visualize the shortcut trajectories spanning the top of the maze. **(B)** Examples of a shortcut trajectory (black), a disjoint trajectory (beige), and a non-shortcut sequence (light blue). **(C)** Expected, bootstrapped, and observed distributions of disjoint, non-shortcut, and shortcut trajectories (see Methods, **Chapter 5**). This analysis shows that the observed shortcuts were extremely unlikely to arise from chance alignments of forward and backward replays, supporting the notion that rats can mentally construct spatially coherent, but never-experienced paths (see **Figure 14** for paths experienced by each rat). (From Gupta et al., 2010.)

more than 10 minutes or 15 laps, in an environment with the vast majority of experience in one direction. Furthermore, during left-only and right-only half-sessions, trajectories along the non-recent (opposite-side) loop were replayed with a similar frequency to trajectories on the recent (same-side) loop. This observation was in contrast to alternation half-sessions in which opposite-side loops were replayed less frequently. These observations indicate that current proposals for potential mechanisms of replay that rely on recency or frequency of experience are inadequate.

In theories of learning and consolidation, the information replayed is the information learned (Marr, 1971; Buzsáki, 1989; Wilson and McNaughton, 1994; Redish and Touretzky, 1998; Redish, 1999; Sutherland and McNaughton, 2000; Foster and Wilson, 2006). In our data, trajectories in the environment leading toward reward locations and/or locations where the animal paused after a lap were preferentially replayed. This result could not be accounted for by the distribution of place fields on the task and suggests that certain trajectories were actively replayed. Proposals suggesting that replay is a simple function of experience cannot explain the preferential replay of particular trajectories.

We also found that the content of replay changed depending on the behavioral task at hand: during alternation half-sessions a lower proportion of replays occurred along the opposite-side loop compared to left-only and right-only half-sessions, even though the opposite-side loop had been traversed more recently in the past and would occur sooner in the future during the alternation half-sessions. This finding is consistent with the prediction that the content of replay should depend on the animal's behavior, but is the opposite pattern to that predicted by the recency and frequency proposals (compare simulations from Redish and Touretzky, 1998 and the discussion in Foster and Wilson, 2006). While previous studies (Wilson and McNaughton, 1994; Jackson et al., 2006; O'Neill et al., 2006; Karlsson and Frank, 2009) have found increasing replay with experience, these studies did not compare tasks with different behavioral requirements in which parts of the environment were experienced more than others. Therefore, the increase in replay with experience seen in

these earlier studies may be due to general experience in the environment rather than the experience of particular trajectories.

Previous interpretations of the role of replay have arisen from the concept that it provides a mechanism by which recent experiences are written out from hippocampus to cortex (Marr, 1971; Squire et al., 1984; Squire, 1987; Buzsáki, 1989; Cohen and Eichenbaum, 1993; Alvarez and Squire, 1994; Hasselmo and Bower, 1993; Redish and Touretzky, 1998). Correlational studies have found interactions between hippocampal replay and cortical learning (Hoffmann and McNaughton, 2002; Euston et al., 2007; Ji and Wilson, 2007) as well as transfers of dependence (Squire, 1987; Maviel et al., 2004) and learning changes from disruption of SWRs (Girardeau et al., 2009; Jackson et al., 2009; Ego-Stengel and Wilson, 2010). While our data does not preclude the possibility that information is transferred between hippocampus and cortex during SWRs, it suggests that the information available for transfer is more reflective of the entire set of navigationally available paths rather than the specific experiences themselves.

The results presented above suggest that sequence generation during SWRs is likely to be involved in learning and maintaining a representation of the environment. The forward and backward replay of all trajectories in the environment (not just the recent and well-experienced ones) could be a mechanism to establish and reinforce connections between nearby locations in the representation. During behavior, the combination of phase precession and spike timing dependent plasticity (STDP) has been proposed to enable the storage of sequences in the order experienced by the animal (asymmetric connections are learned) (Levy, 1996; Skaggs et al., 1996; Mehta and McNaughton, 1996). While this mechanism would produce representations of the forward paths traversed by the animal, it would not produce backward connections, nor would it produce novel (never-experienced) connections. Our finding of backward replay in an environment experienced primarily in one direction, over trajectories never experienced in the direction of replay, suggests that backward replay could be an important mechanism for learning a navigationally-complete representa-

tion of the environment: a representation that reflects not only the trajectories experienced by the animal, but also the reverse trajectories available in the environment. This interpretation is consistent with the observation that reactivation and backward replay events are more frequent in novel environments (Foster and Wilson, 2006; Cheng and Frank, 2008), which could be a mechanism to rapidly acquire a complete representation.

Our finding that novel sequences never experienced by the animal are also played out during awake SWRs may also reflect mechanisms for learning the navigationally-complete representation of the environment through the internal exploration of potential shortcuts (Samsonovich and Ascoli, 2005). Although they were rare, the likelihood that these novel sequences could be accounted for by chance alignments of forward and backward replays was very low — the distributions observed were significantly different from those expected by a bootstrap pairing of replays. The presence of shortcut sequences offers strong support that the hippocampal network contains a navigationally-complete representation (a cognitive map) of the environment (Tolman, 1948; O’Keefe and Nadel, 1978). These sequences further support the idea that backward replay is not simply an experience replayed in the reverse order, but rather reflects a constructed sequence representing a trajectory in the environment. The properties of sequence “play” described here suggest that hippocampal reactivation during SWRs could be an important part of the mechanism that learns this cognitive map.

One problem with maintaining a stable representation of the environment is the issue of catastrophic interference: when a network simultaneously learns and encodes many sequences, sequences that aren’t being rehearsed tend to degrade as other activated sequences interfere with them (O’Reilly and McClelland, 1994; McClelland et al., 1995). Preferentially replaying non-local trajectories when they aren’t being visited (as was the case for left-only and right-only half-sessions but not for alternation half-sessions) could be a mechanism to prevent the representation of the non-recently experienced trajectories from degrading (Pomerleau, 1991). Reactivation of memories from previous tasks within a day

(Jackson et al., 2006; Karlsson and Frank, 2009) may serve a similar purpose.

The hippocampal formation has been shown to be important for the imagination of novel situations (Hassabis et al., 2007) and has been implicated in self-projection (Gelbard-Sagiv et al., 2008), the ability to consciously explore the world from different perspectives (Buckner and Carroll, 2007). Our data that sequential place cell activation during awake states reflects forward, backward, and novel sequences spanning the environment, with a pattern more consistent with maintaining a representation of the environment than with replaying recent experiences, supports the hypothesis that the hippocampus may provide a potential substrate for self-projection-like processes. The fact that these properties occur during the awake state, while the animal is paused but still engaged in the behavioral task, allows the intriguing speculation that “replay” could contribute to the animal’s real-time representation of the world, providing access to information spanning the cognitive map, thereby supporting flexible and goal-driven behaviors.

7 Hippocampal theta sequences

This chapter contains original work submitted by my colleagues and me: *Hippocampal theta sequences: dynamic modulation and the segmentation of spatial experience* (Gupta et al., in revision). This work demonstrates a relationship between animal behavior and hippocampal spatial representations and provides a potential mechanism for the cognitive chunking of experience.

7.1 Summary

In freely-moving rodents, neural activity during each theta cycle represents a path beginning just behind the animal and ending just ahead of the animal. We report that, under certain conditions, as animal velocity increased, theta sequences represented longer paths in the environment. As animals accelerated, paths included more space ahead of the animal and during deceleration, paths included more space behind the animal. Furthermore, as path length increased, there were linear increases in both theta period and number of gamma cycles per theta cycle. Finally, we observed that neural activity represented the environment in segments bounded by physical maze landmarks. These observations suggest a potential neural mechanism for the cognitive chunking of experience.

7.2 Results

To investigate the relationship between behavior and the paths represented by theta sequences, we recorded neural ensembles from the CA1 region of hippocampus in six rats trained on a spatial decision task (see Methods, **Chapter 5**, **Figure 16A**). The beginning and end of each theta cycle was identified and the neural activity during each cycle was analyzed. Using the sequence decoding algorithm described in **Chapter 5**, the paths represented during each theta cycle were determined, and only those sequences with statistically

significant structure and three or more active cells were included for analysis (analyses were also performed using a Bayesian decoding approach Zhang et al., 1998). Furthermore, theta cycles with a mean sharp wave ripple (SWR) power greater than one standard deviation above the mean were excluded from analysis to avoid the inclusion of replay sequences. Out of 618,408 theta cycles over 31 recording sessions, 168,770 cycles had three or more active cells, SWR power less than one standard deviation above the mean, and theta power greater than one standard deviation below the mean. Of these cycles, 33,397 (19.8%) contained statistically significant neural sequences (see Methods, **Chapter 5**). **Figure 16C** displays the distribution of animal occupancy (time spent) over the linearized maze for all theta cycles (dotted line) and only theta cycles with significant sequences (solid line). As expected, animals spent the most time at the feeder locations (dotted line) and a smaller proportion of theta cycles were significant at the feeder locations compared with other locations on the maze (contrast dotted and solid lines).

For each theta sequence, the detection algorithm determined the start and end location of the represented path as well as the starting and ending theta phase of spikes in the sequence. From the start and end location of the path, as well as the animal's location, we calculated three key distances: the distance from the animal's location to the end of the path (the "ahead" length), the distance from the start of the path to the animal's location (the "behind" length), and the total length of the path ("path" length).

We examined the relationship between animal velocity and path length by creating a 2D histogram of velocity as a function of path length. As shown in **Figure 16B**, there was a low-velocity and a high-velocity population of sequences. The low-velocity sequences were preferentially observed when animals were at the feeder locations and the choice point (**Figure 16E**), whereas high-velocity sequences were evenly distributed across the maze (**Figure 16D**). For the remainder of the analyses presented here, we focus on high-velocity theta sequences (29,351 sequences), although all results are unchanged by the inclusion of low-velocity sequences. We come back to the relationship between velocity and the length

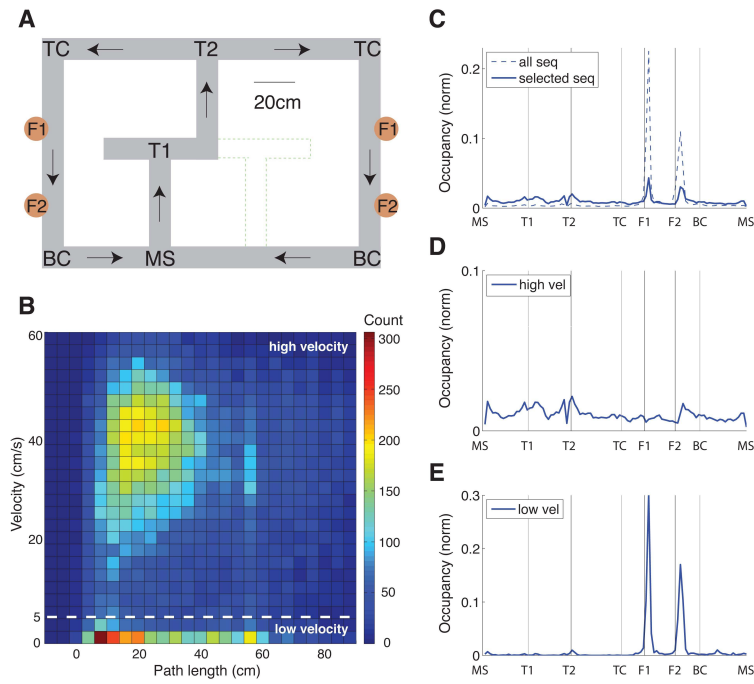


Figure 16: (A) The two-choice T maze. The maze had two possible physical configurations, the second indicated by dotted lines. Noteworthy landmarks on the maze include the maze start (MS), turn 1 (T1), turn 2 (T2), top corner (TC), feeder 1 (F1), feeder 2 (F2), and bottom corner (BC). Arrows indicate the direction of maze traversal. (B) Theta sequence histogram of animal velocity versus path length. The color of each pixel indicates the number of theta sequences with a specified path length (x-axis) recorded while an animal was running at a particular velocity (y-axis). The horizontal dotted line separates low velocity theta sequences from high velocity ones. Average velocity as a function of path length for high-velocity sequences is plotted in white with thickness indicating the SEMs. Short and long theta sequences were divided based on the peak of the average velocity curve (vertical dotted line). (C) Overall and selected sequence count over the linearized maze. Overall sequence count over the maze (dotted line) is a reflection of the amount of time animals spent at each location on the maze. Selected sequence count (solid line) indicates the number of sequences at each location on the maze that passed the inclusion criteria. (D) High-velocity sequence count over the linearized maze. The number of high-velocity sequences (sequences above horizontal line in panel B) at each location on the maze. (E) Low-velocity sequence count over the linearized maze.

of space represented later in the Results section.

Ahead and behind length are anti-correlated and occur during different parts of the theta cycle

To determine the relationships between path, ahead, and behind length, we plotted each mean length with respect to the others. As shown in **Figure 17E**, as path length increased, both ahead and behind length increased, as would be expected since the path length of each theta sequence is the sum of its ahead and behind length. On the other hand, as behind length increased, ahead length decreased. Thus, ahead length and behind length were anti-correlated ($r = -0.24$, $p < 0.01$): spatial paths that extended further forward did not begin as far back and vice versa.

Next, we determined the mean starting and ending phase of each spike sequence as a function of path, ahead, and behind length. As path length increased, the spike sequence began earlier in the theta cycle and ended later in the theta cycle (**Figure 18A**, left column). As ahead length increased, the spike sequence ended later in the theta cycle, but the starting phase of firing was relatively constant (**Figure 18A**, middle column). The opposite was true for increasing behind length: as behind length increased, the spike sequence began earlier in the theta cycle and the ending phase was relatively constant. Thus, as the represented path extended further ahead of the animal, spikes extended later in the theta cycle, and as the path represented began increasingly behind the animal, spikes began earlier in the theta cycle.

Theta period and gamma cycles increase with path length

As path length increased, we observed that theta period and the number of gamma cycles per theta period also increased (**Figure 18B,C**). Using the relationship between theta period and velocity (**Figure 17D**) and the observed relationship between velocity and path length

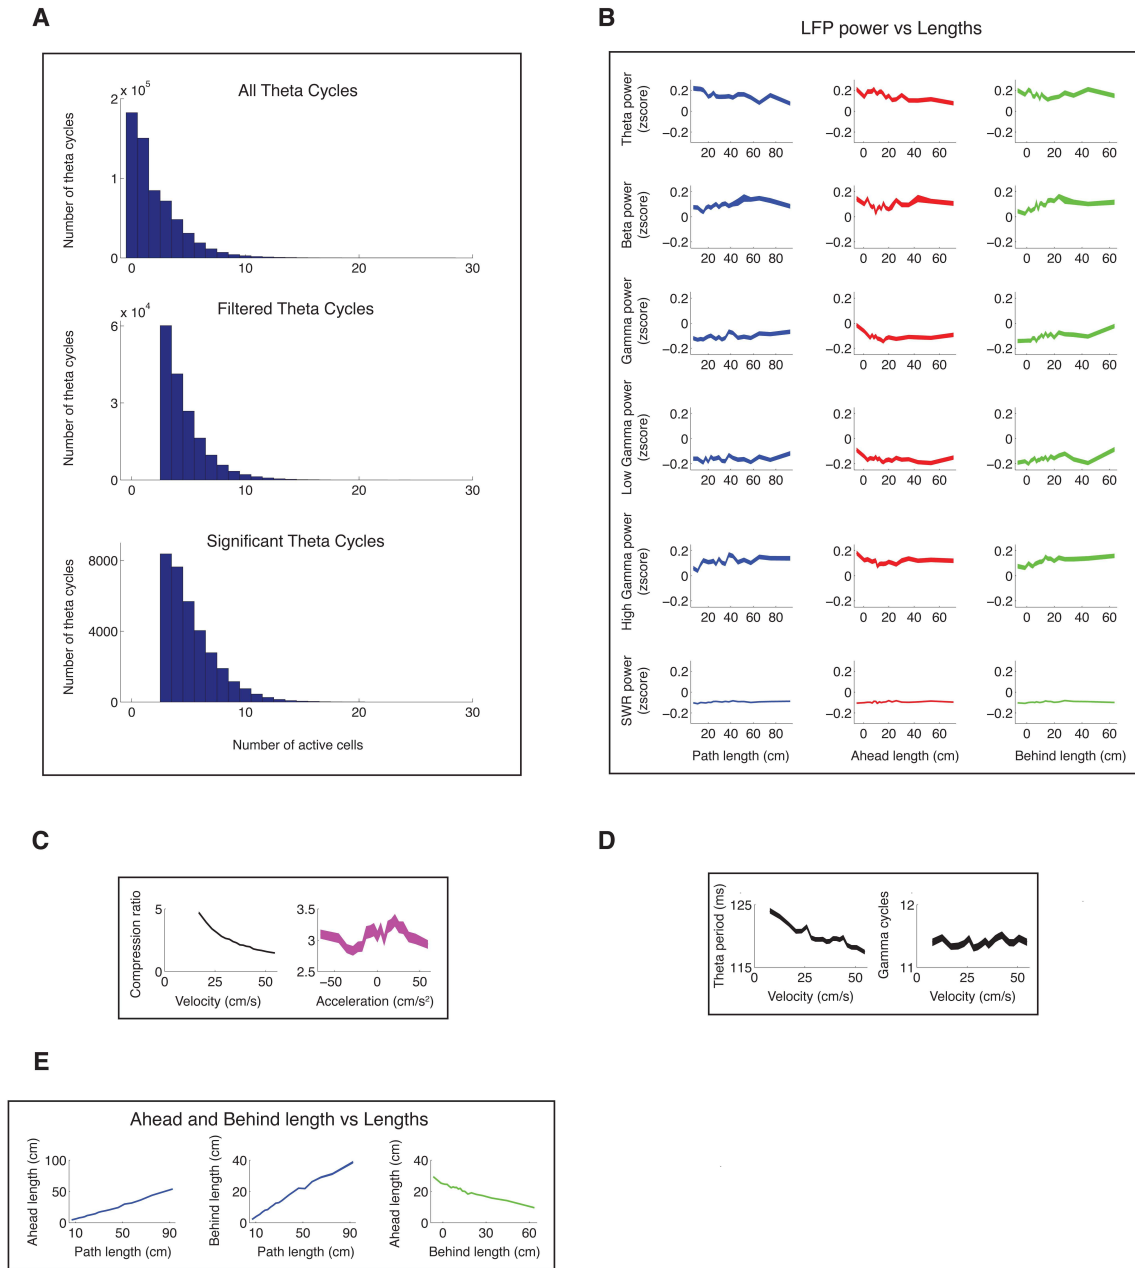


Figure 17: (A) Histograms of the number of theta cycles over the number of active cells per theta cycle. Top panel includes all theta cycles over 31 recording sessions (618,408 cycles). Middle panel includes theta cycles with three or more active cells, mean sharp wave ripple power less than one standard deviation above the mean, and theta power greater than one standard deviation below the mean (168,770 cycles). Bottom panel includes theta cycles whose neural activity contains significant sequential structure as described in **Chapter 5** (33,397 cycles). **(B)** Mean LFP power (z-scored by session) as a function of path, ahead, and behind length. **(C)** Compression ratio (length of spatial path represented divided by distance traveled by the animal) as a function of velocity and acceleration. For clarity, compression ratio versus velocity is shown for velocities greater than approximately 20 cm/s. For slower velocities, as velocity decreases, the compression ratio increases exponentially. **(D)** Theta period and number of gamma cycles as a function of velocity. **(E)** Left to right: ahead length as a function of path length, behind length as a function of path length, and ahead length as a function of behind length. Note the inverse relationship between ahead length and behind length.

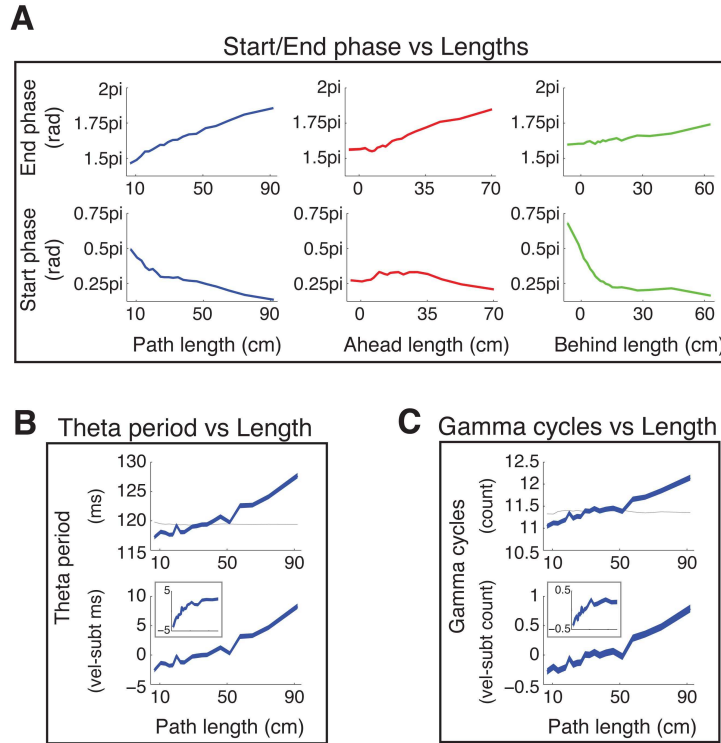


Figure 18: (A) Starting and ending theta phase as a function of path, ahead, and behind length. Starting phase refers to the theta phase of the first spikes in each theta sequence. Similarly, ending phase refers to the theta phase of the last spikes in each sequence. **(B)** Theta period as a function of path length. The blue curve in the top panel displays average theta period (y-axis) with increasing path length (x-axis). Errorbars are SEMs. The gray curve demonstrates the relationship between theta period and path length if the relationship was purely a secondary affect of the relationship between theta period and velocity and path length and velocity. The bottom panel shows the subtraction of the gray and blue curves in the top panel, which removes the velocity component of the relationship. **(C)** Gamma cycles as a function of path length. See panel **B** description, with the number of gamma cycles replacing theta period. Figure insets show results replicated using Bayesian decoding (see Methods, **Chapter 5**).

(**Figure 16B**), we determined the relationship between theta period and path length if theta period were purely a function of velocity (**Figure 18B**, top, gray curve) and subtracted it from the observed relationship (**Figure 18B**, bottom). As shown in this figure, when controlled for velocity, theta period increased linearly with increasing path length. Similarly, the number of gamma cycles increased linearly with increasing path length (**Figure 18C**). In summary, longer theta periods, more gamma cycles per theta cycle, and spikes occurring earlier and ending later in the theta cycle were associated with longer paths represented by theta sequences.

Ahead and behind representations

As described above, we observed that the represented path sometimes extended well ahead of the animal, and at other times began farther behind the animal.

Johnson and Redish (2007) observed that when animals arrived at a choice point, the CA3 region of hippocampus represented spatial paths extending substantially ahead of the animal. These representations began near the animal's location and traversed each path the animal had the option to take, one at a time. Power in the theta band was strong during these representations. It is possible that the spatial paths they observed in CA3 are analogous to the longer, ahead representations we report here from CA1 recordings. Examples of ahead and behind representations are shown in **Figure 19**. Ahead sequences were observed at the choice point as well as at other specific locations on the maze.

Average ahead and behind length over space

In order to determine whether theta sequences preferentially represented space ahead or behind the animal at certain locations on the maze, the average ahead length and the average behind length were plotted as a function of linearized location on the maze. As shown in **Figure 20A,C**, we observed clear peaks and troughs in average ahead length

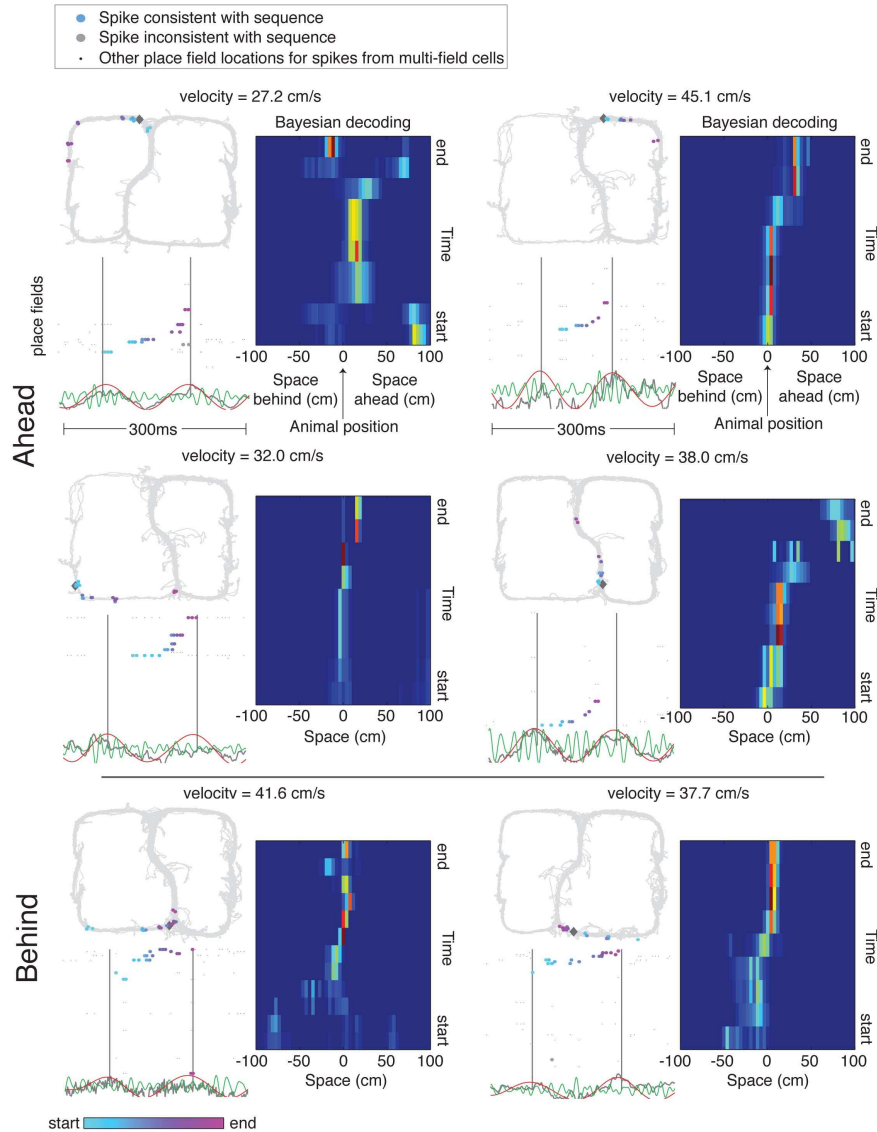


Figure 19: Examples of ahead and behind sequences. **Top row:** ahead sequences while animals were located at the choice point. **Middle row:** ahead sequences while animals were at other locations on the maze. **Bottom row:** behind sequences on the maze. Each example is represented by a left and right panel. **Left panel:** the gray diamond indicates the rat's location at the time of the sequence. The animal's average velocity over the theta cycle is displayed at the top of each subfigure. On the bottom panel of each subfigure, spikes are plotted by ordered place field center (along either a left or right loop of the maze) over a 300 ms period (see Experimental Procedures). LFPs filtered between 6 Hz and 12 Hz (red), 40 Hz and 100 Hz (green), and unfiltered (gray) are plotted at the bottom of the panel. Colored points indicate spikes that contribute positively to the sequence score according to the automatic sequence detection algorithm. The color of the spike indicates its relative time within the sequence (light blue = early, light purple = late). Gray points are spikes that do not contribute positively to the score. For cells with multiple place fields, small black points are plotted at every place field center belonging to the cell (colored points occupy the place field center that contributes maximally to the score). Each colored point from the bottom panel is plotted on the 2D maze in the top panel at the location of its 2D place field center. **Right panel:** the Bayesian decoded spatial probability distributions computed over the theta sequence (see Methods, **Chapter 5**).

over space and average behind length over space. Furthermore, average ahead length and average behind length were anti-correlated over space ($r = -0.5$, $p < 0.01$), with ahead length peaks aligning with behind length troughs, and vice versa. This observation is consistent with, but not a necessary result of the inverse relationship between ahead and behind length presented above (**Figure 17E**). Thus, at certain locations on the maze, the hippocampus represented more space ahead of the animal and less space behind, and at other locations on the maze, the hippocampus represented more space behind the animal, and less space ahead.

The peaks of ahead length and the peaks of behind length had a relationship with respect to physical landmarks on the maze (i.e., turns and feeders). Ahead length peaks were generally observed just after landmarks (e.g., MS, T2, F1, BC) and behind length peaks were observed just before landmarks. Thus, as animals entered a segment of the maze, defined by two physical landmarks, the spatial representation was more ahead of the animal and as animals exited the segment, the representation was more behind the animal.

Spatial distribution of represented paths

The alternating pattern of ahead and behind representations resulted in the preferential representation of segments of the maze lying between landmarks, as reflected in the spatial distribution of represented paths (**Figure 20B**). The distribution demonstrated that paths represented by theta sequences preferentially occurred between landmarks on the maze, instead of across them. The distribution shown was normalized such that the paths represented at each location on the maze contributed equally to the distribution (occupancy normalization). This ensured that factors such as the animal consistently spending more time at a particular location on the maze did not misleadingly influence the distribution around that point.

The observed distribution was compared to a distribution created by a shuffling procedure. Keeping the animal location of each theta cycle constant, we shuffled the de-

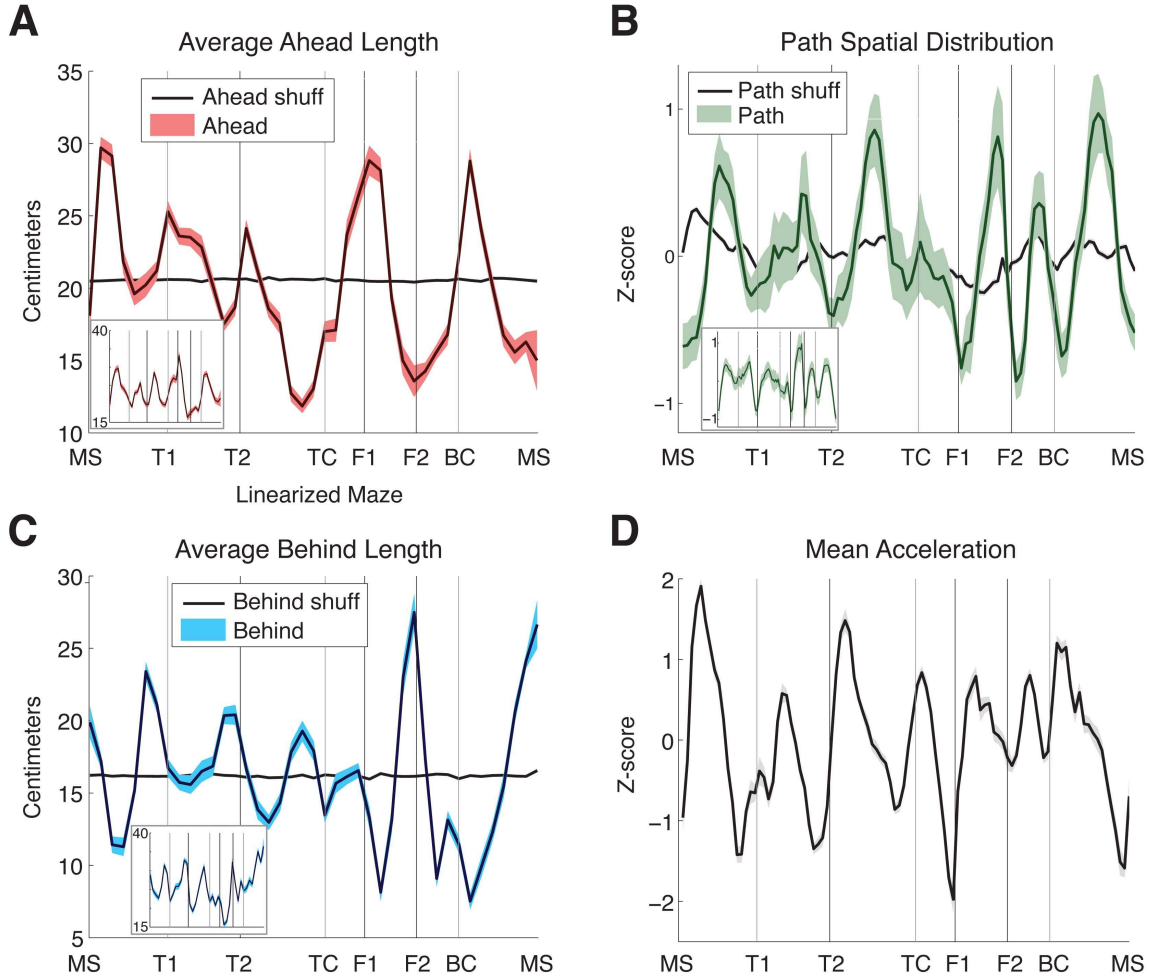


Figure 20: **(A)** Ahead length as a function of location on maze. Peaks indicate locations at which the represented path extended farther ahead of the animal, on average. Troughs indicate locations where the represented path ended closer to the animal's location. Figure insets in panels **A**, **B**, and **C** show results replicated using Bayesian decoding (see Methods, **Chapter 5**). **(B)** Density of represented paths as a function of location on maze. Peaks indicate regions of the maze that were overrepresented by theta sequences. Y axis is average Z-score (Z-score computed over each session). **(C)** Behind length as a function of location on maze. Peaks indicate locations at which the represented path began farther behind the animal, on average. Troughs indicate locations where the represented path began closer to the animal's location. **(D)** Acceleration as a function of location on maze. Y axis is average Z-score (Z-score computed over each session). Errorbars are SEMs. As a control, panels **A-C** include a curve produced by randomly reassigning theta sequences with different animal locations. This shuffling procedure was performed 100 times and the average shuffled distribution (black curve) is plotted alongside the real data.

coded paths and randomly assigned them to the animal locations. This was performed 100 times and the average spatial distribution was plotted along with the observed distribution (**Figure 20B**). As would be expected from the occupancy normalization, the shuffled distribution was relatively flat, demonstrating that the observed spatial distribution of paths was a product of specific sequences occurring at particular locations on the maze. The shuffling procedure was also performed for the average ahead and behind length over space plot (**Figure 20A,C**), and as expected the shuffled distribution was flat.

We compared the spatial distribution of represented paths (and ahead and behind length over space) with the place field density, place field length, the animal's velocity, and the animal's acceleration. The place field density, place field length, and average velocity over space did not match the shape of the path distribution or ahead and behind length (**Figure 21**). Place field length, plotted by location on the maze, showed some peaks in between landmarks, however the overall shape and magnitude was not consistent with the ahead, behind, and spatial path distributions. For example, peaks in the ahead and behind distributions between the maze start (MS) and turn 1 (T1) had peak to trough distances of approximately 12 cm (18 to 30 cm for ahead and 11 to 23 for behind) (**Figure 20A,C**), whereas the peak to trough distance of the place field length distribution was approximately half that distance (6 cm; 37 to 43 cm) (**Figure 21B**). Furthermore, large peaks in the ahead and behind distribution were seen between the feeders (F1 and F2) in the absence of a peak in the place field length distribution.

However, average acceleration over space (**Figure 20D**) shared striking similarities with the path distribution. For each peak in the path distribution, there was a corresponding peak in acceleration, which was shifted slightly to the left. The peaks and troughs of acceleration were actually aligned with peaks in ahead length and peaks in behind length, respectively (**Figure 20A,C**) and ahead length was positively and behind length was negatively correlated with acceleration over space (ahead: $r = 0.41$, $p = 0.016$; behind: $r = -0.74$, $p < 0.01$). This observation predicts a relationship between acceleration and ahead and behind length

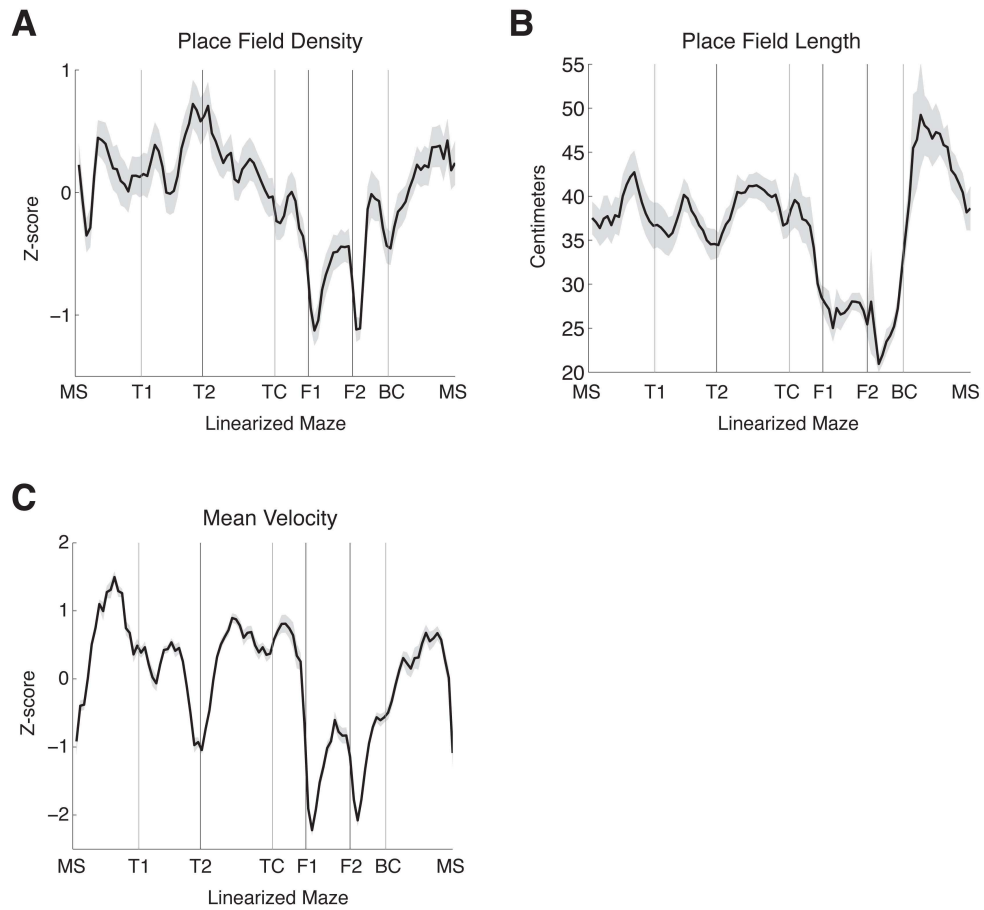


Figure 21: **(A)** Place field density as a function of location on linearized maze. The height of the curve reflects the number of place fields overlapping a particular location in the environment. Y-axis is z-scored place field distribution by session, x-axis is linearized maze. **(B)** Mean place field length as a function of location. Y-axis is centimeters averaged by session, x-axis is linearized maze. **(C)** Animal velocity as a function of location on linearized maze. Y-axis is z-scored velocity by session, x-axis is linearized maze.

(see below).

In summary, the spatial distribution of paths represented by theta sequences reflected an over-representation of certain segments or "chunks" of the environment that were bounded by physical landmarks on the maze.

Ahead length increases with acceleration, behind length increases with deceleration

Next, we investigated the relationship between animal acceleration and path, ahead, and behind length. We observed that as ahead length increased, acceleration also increased, and, conversely, as behind length increased, acceleration decreased (**Figure 22A**, top row). In contrast, acceleration remained constant or decreased slightly with increasing path length.

These relationships were strong for shorter ahead and behind lengths and dropped off for longer lengths. This was also seen in the relationship between velocity and path, ahead, and behind length (**Figure 22B**, top row). Thus, we divided sequences into a set of shorter sequences and a set of longer sequences, based on the peaks of the *velocity* versus path, ahead, and behind length curves (marked by vertical lines in **Figure 22A,B**, top row). The peaks were located at approximately 25 cm (path) and 20 cm (ahead and behind). Acceleration continued to increase (with ahead length) and decrease (with behind length) slightly past these drop-off points, indicating the relative strength of the relationship between acceleration and ahead and behind length compared to velocity.

Ahead length and behind length for the shorter sequences displayed a linear relationship with acceleration (**Figure 22A**, middle row). Ahead length doubled in distance over the range of acceleration. Conversely, behind length decreased by approximately 50% with increasing acceleration. Whereas changes in ahead and behind length were clear for the shorter sequences, ahead and behind length were relatively constant over acceleration for the longer sequences (**Figure 22A**, bottom row), demonstrating that the path length of shorter and longer theta sequences depend on behavior in different ways. These relation-

ships were robust, as the exclusion of the spike with place field center farthest ahead or farthest behind the animal did not alter the relationships (**Figure 23A,B**). Furthermore, the use of a Bayesian decoding approach to determine the represented path yielded qualitatively similar results (**Figure 23C**).

In summary, there was a strong relationship between an animal's acceleration and the spatial path represented that was present for most, but not all theta sequences. As animals accelerated on the maze, the path represented extended farther ahead of the animal and began just behind the animal. In contrast, while decelerating, the path represented began farther behind the animal and ended just ahead of the animal. On the other hand, path length was relatively constant with acceleration. This demonstrates that the represented path included more space ahead of the animal (and less space behind) during acceleration and included more space behind (and less space ahead) during deceleration.

Lengths show a complex relationship with velocity

To determine the relationship between animal velocity and lengths, we plotted changes in velocity as a function of path, ahead, and behind length. As shown in **Figure 22B**, top row, as each of the three lengths increased, there was initially an increase in velocity and subsequently a drop off.

As described above, we used the drop-off points to divide path, ahead, and behind length into shorter and longer sequences. For the shorter theta sequences, path, ahead, and behind length all increased with increasing velocity (**Figure 22B**, middle row). However, for the longer sequences, path, ahead, and behind length all decreased with increasing velocity (**Figure 22B**, bottom row). As with the relationship with acceleration, the relationship of these lengths with respect to velocity were robust to spike exclusion and a Bayesian decoding analysis (**Figure 23C**).

This observation and the presence of low and high-velocity theta sequences presented earlier (**Figure 16B**) indicate that there was not a simple relationship between velocity and

S = Shorter sequences

L = Longer sequences

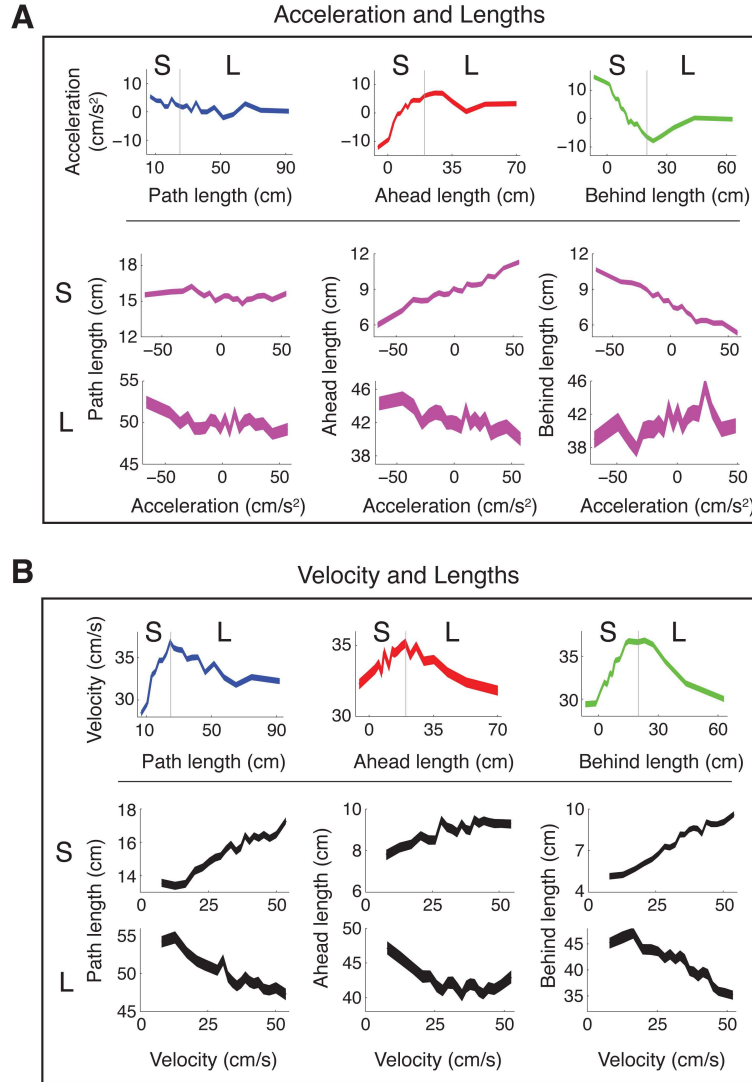


Figure 22: (A) Relationship between acceleration and path, ahead, and behind length. The top row shows acceleration as a function of path, ahead, and behind length. Vertical lines mark the peak of each velocity curve in panel **B**, which was used to divide sequences into short (S) and long (L). Below each panel of the top row, short and long sequences are plotted separately, with path, ahead, and behind length now as a function of acceleration. **(B)** Relationship between velocity and path, ahead, and behind length. As in panel **A**, the top row shows velocity as a function of path, ahead, and behind length. Vertical lines mark the peak of each curve and divide short (S) and long (L) sequences. Below each panel of the top row, short and long sequences are plotted separately, with path, ahead, and behind length as a function of velocity.

the paths represented during theta sequences. Some sequences representing both short and long paths occurred while animals were stationary or moving at low velocities in the environment as observed by Johnson and Redish (2007). Other sequences occurred as animals were moving at faster velocities. These high-velocity sequences could be divided into a set of short-path sequences which increased in length with increasing animal velocity (Maurer et al., 2011), and a set of long-path sequences that decreased or remained constant in length with increasing velocity. It should be noted that theta, beta, gamma, and SWR power were relatively constant with increasing path, ahead, and behind length (**Figure 17B**), thus the LFP profile did not change as the represented path increased in length. Furthermore, consistent with Maurer et al. (2011), we observed an exponential decrease in the sequence compression ratio (path length/distance traveled) with increasing velocity (**Figure 17C**).

7.3 Discussion

Neural activity during hippocampal theta provides a potential mechanism for encoding current experience and may reflect the use of recalled information from past experience to assist behavior. It has been previously shown that neural activity during each theta cycle is sequential and represents a short path that begins just behind the animal and ends just in front of the animal (Skaggs et al., 1996; Maurer and McNaughton, 2007; Foster and Wilson, 2007). It has also been demonstrated that the path represented increases in length with animal velocity (Maurer et al., 2011) and extends further ahead of the animal during pauses at a choice point (Johnson and Redish, 2007). The results presented here support these observations and show that there is considerable diversity in the paths represented during hippocampal theta, which is sometimes, but not always related to animal velocity, acceleration, and physical landmarks in the environment.

We found coherent theta sequences when animals were paused or moving at very low

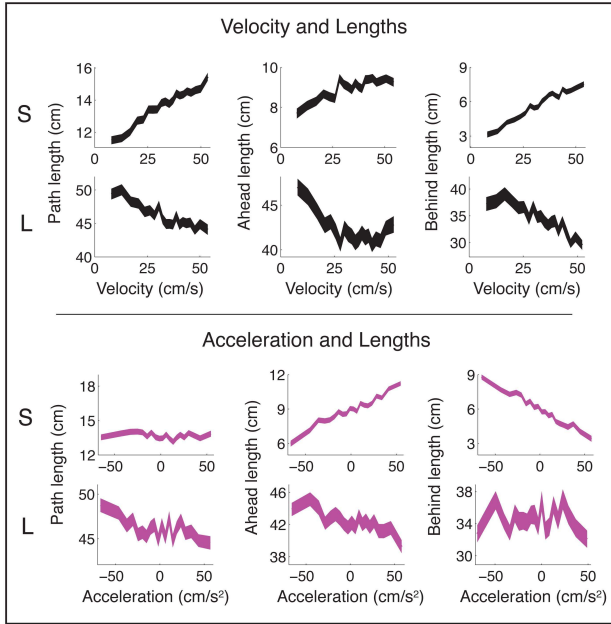
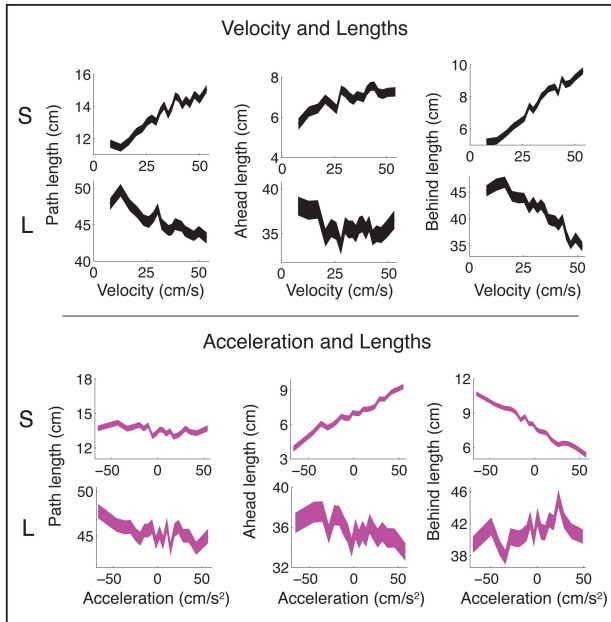
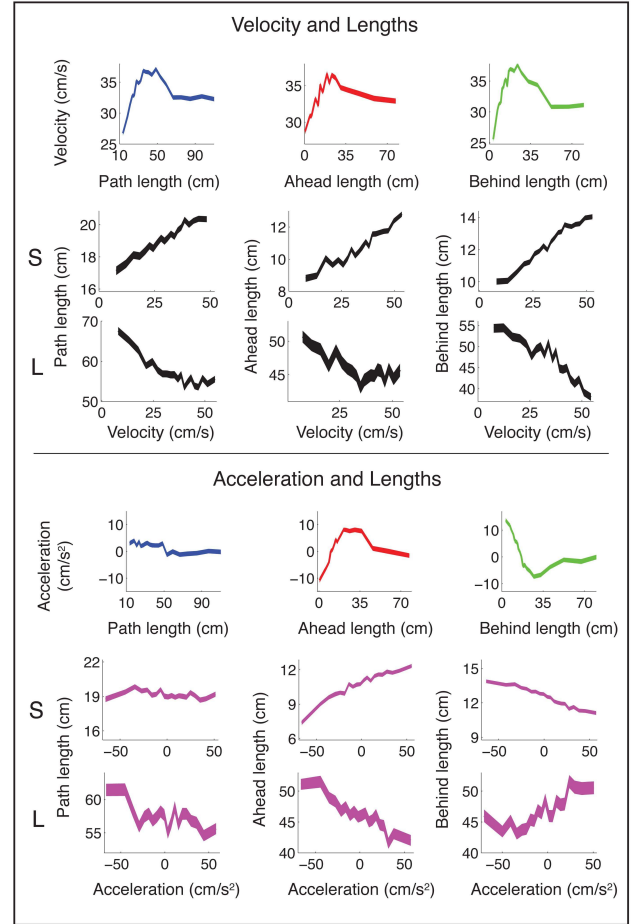
A**B****C**

Figure 23: (A) Relationship between velocity and acceleration and path, ahead, and behind length, with the *last spike* in the theta sequence excluded. **(B)** Relationship between velocity and acceleration and path, ahead, and behind length, with the *first spike* in the theta sequence excluded. Results are qualitatively similar to those presented in **Figure 22**. **(C)** Analyses in **Figure 22** performed using Bayesian decoding (Zhang et al., 1998) to detect the represented path during each theta cycle. See Methods, **Chapter 5** for details.

velocities, which is consistent with the sequences observed by Johnson and Redish (2007) in CA3. Consistent with the recent report by Maurer et al. (2011), we found that the length of the spatial path represented during high-velocity theta sequences increased with animal velocity. Interestingly, we found that more space was represented ahead of the animal during acceleration and more space was represented behind the animal during deceleration. These relationships between animal behavior and represented path held true for many, but not all high-velocity theta sequences. Furthermore, the length of the path represented was related to the length of the theta cycle, the number of gamma cycles occurring during each theta cycle, and the starting and ending theta phase of the spike sequence. Finally, we observed that theta sequences represented the maze in segments, which corresponded with the segmented nature of animal behavior over the maze.

Sequence representations in the hippocampus have been extensively studied in the theta phase precession literature. Phase precession refers to the observation that as an animal passes through a cell's place field, the cell initially fires late in the theta cycle and gradually fires earlier and earlier in the theta cycle as the animal traverses the place field (O'Keefe and Recce, 1993). At the population level, phase precession implies that compressed sequences of place cell activity emerge during each theta cycle (Skaggs et al., 1996; Tsodyks et al., 1996; Jensen and Lisman, 1996; Dragoi and Buzsaki, 2006; Foster and Wilson, 2007; Maurer et al., 2011). A number of different models have been proposed to account for phase precession: these include dual-oscillator interference models at the single cell (O'Keefe and Recce, 1993) and local network level (Yamaguchi et al., 2007), as well as network-level proposals (Tsodyks et al., 1996; Jensen and Lisman, 1996; Hasselmo and Eichenbaum, 2005). Currently, these mechanisms cannot account for the diversity in theta sequences reported here. Future evolutions of phase precession models will be required to explain how sequences can flexibly represent more space ahead, represent more space behind, and vary in path length depending on the animal's behavior.

A number of studies have investigated whether hippocampal activity is reflective of the

past and/or the future. Some of these studies, like the one here, have focused on the spatial path represented during the fast time scale neural activity observed in each theta cycle (Dragoi and Buzsaki, 2006; Johnson and Redish, 2007; Foster and Wilson, 2007; Maurer et al., 2011). Other studies have analyzed hippocampal activity with respect to the animal's past and future behaviors and have shown that the activity of individual place cells or place cell ensembles (at relatively slow time scales) is sometimes more reflective of past behaviors and is sometimes more reflective of future behaviors (Wood et al., 2000; Frank et al., 2000; Ferbinteanu and Shapiro, 2003; Ji and Wilson, 2008). Muller and Kubie (1989) made the argument that hippocampal activity represents the animal's immediate future, due to the fact that place fields were on average tighter when spikes were shifted to precede the animal's location by 120 ms. Thus, there is evidence that the hippocampus represents more than just the current state of the animal during behavior and does so at multiple time scales (Yartsev, 2008). We add to this body of work by showing that at the theta time scale, the representation of future space and past space is dissociable and dependent on the animal's acceleration and landmarks in the environment. It should be noted that while it has been suggested that the "true place field" is different from the observed place field (Jensen and Lisman, 1996; Lisman and Redish, 2009), we do not make any assumptions about the true place field and define the point of representation of each spike as the center of the place field (sequence detection) or the entire place field (Bayesian decoding).

Diversity in theta sequences

We suggest that some of the theta sequences observed here reflect a representation of the animal's immediate demands during navigation. As an animal increases its speed and traverses a larger segment of the environment during a theta cycle, the neural sequence correspondingly represents a longer path in space (Maurer et al., 2011). When an animal is accelerating toward a location up ahead, the paths represented are shifted forward in space, possibly in preparation or anticipation of reaching a desired location. As an animal

decelerates upon approach to the goal or subgoal, the paths represented shift backward in space, potentially in review of the experience. This indicates that sequence representations during theta are flexible, sometimes reflective of a "look ahead" function and at other times reflective of "look behind", a result further supported by the observation that ahead and behind length were anti-correlated in time and location on the maze. It is possible that the ahead representations observed here reflect a predictive recall process to aid in navigation (Lisman and Redish, 2009; Maurer et al., 2011) while the behind representations reflect an encoding process to store what was just experienced (Buzsáki, 2005; Hasselmo, 2005). This type of recall process could be cued by sensory input from upcoming landmarks, which were previously associated with locations in space, as modeled in Hasselmo (2009). Furthermore, the observation that ahead sequences extend later into the theta cycle and behind sequences begin earlier in the cycle is consistent with the proposal that encoding and recall occur during different parts of the theta cycle (Hasselmo et al., 2002).

In addition to the theta sequences described above that increased in length with velocity and shifted with acceleration, we observed sequences that appeared to be less related to the animal's immediate behavior. These include sequences representing a variety of path lengths that were observed while animals were paused in the environment. They also include longer path length sequences observed during movement that did not increase in length with velocity or shift forward and backward with acceleration. These sequences may reflect a more cognitive process, as seen by Johnson and Redish (2007). Whereas Johnson and Redish (2007) observed long, ahead representations while animals were paused at the choice point, we observed that these sequences occurred at other locations on the maze, did not only occur when animals were paused, and at times represented more space ahead of the animal and at other times represented more space behind the animal. This indicates that theta sequences are not always a simple function of animal behavior and can flexibly access more remote information by extending farther ahead of the animal and by beginning farther behind the animal.

Information processing: Theta and Gamma

It has been proposed that the information represented during each theta cycle is discretized into parts by each cycle of gamma (Lisman and Idiart, 1995; Lisman, 2005). In this view, the information capacity of a given cycle of theta would be determined by the number of gamma cycles that occur over the course of the theta cycle. We observed a linear increase in theta period and the number of gamma cycles per theta cycle as the length of the represented path increased. Furthermore, as the path length increased, spikes began earlier and ended later in the theta cycle. In support of Lisman and Idiart (1995) and Lisman (2005), it is possible that the increase in gamma cycles per theta cycle allows for an increase in the number of information processing steps, thus contributing to longer represented paths.

Behavioral and Neural Chunking

"Chunking" refers to a phenomenon in which sequential information is grouped into intuitive, information-rich units. This recoding is thought to enable the handling of more information at once (Miller, 1956) and to enable planning through subgoals (Newell, 1990; Newell and Simon, 1972). From a navigation perspective, it has been previously shown that individual ventral striatal neurons, which receive hippocampal projections, have receptive fields that represent navigationally-relevant segments in an environment (Mulder et al., 2004). In contrast, hippocampal place cells do not represent environments in discrete segments (**Figure 21**). As discussed below, we show for the first time that *sequences* of place cell activity preferentially represent particular, navigationally-relevant chunks of the environment.

From a behavioral perspective, chunking was evident on the 2T task from the rat's acceleration over the linearized maze. There was a clear and smooth increase and decrease in acceleration over segments of the maze bounded by physical landmarks. This was not surprising, because landmarks on the maze were either turns or feeders, and it would be expected for animals to slow down as they approached a turn and speed up afterwards.

However, it should be noted that we observed an increase in acceleration prior to arrival at the top corner of the maze, thus the behavioral chunking did not always perfectly align with the physical landmarks.

Chunking was also evident from the neural representation of space. The spatial distribution of paths represented corresponded peak for peak with acceleration. Even the acceleration peak that occurred over the top corner of the maze was matched with a peak in the spatial distribution of paths. Thus, chunking was present both in the behavior of the animal and in the neural representation of space.

It should be noted that although the spatial path distribution and mean acceleration plots (**Figure 20B,D**) had corresponding peaks, the corresponding peaks did not have similar size characteristics. In particular, although there was a clear acceleration peak that occurred just past the top corner (TC) of the maze, there was only a small peak in the spatial path distribution at the TC location on the maze. This can be explained by the lack of a clear peak in ahead length (**Figure 20A**) at the TC location. Instead, ahead length continued to increase from before the TC until the first feeder location (F1). On the other hand, the trough in behind length was intact at the TC (**Figure 20C**). This difference was reflected in the degree of spatial correlation between behind length and acceleration ($r = -0.74$) compared with ahead length and acceleration ($r = 0.41$). Furthermore, average ahead length and behind length over space changed by 10-15 cm between landmarks on several instances (**Figure 20A,C**), whereas average ahead and behind length as a function of acceleration changed by about 6 cm over the full range of acceleration (**Figure 22A**, middle row).

Based on these observations, we suggest that the chunking observed in the neural representation does not imply a causal relationship between acceleration and ahead and behind length. Instead, we suggest that the chunking present in behavior and the chunking in the neural representation are parallel processes and the correlation between the two processes may result in the observed relationship between acceleration and ahead and behind length.

While theta sequences preferentially represented segments of the environment, such segmentation was not seen by replay sequences occurring during sharp wave ripples, while animals were awake and paused at feeder locations on the 2T maze (Gupta et al., 2010). Davidson et al. (2009) found that replay sequences can represent long paths in space and can occur across multiple ripple events, leading to the suggestion that replay events consist of chains of subsequences. Thus, it is possible that theta sequences preferentially encode particular segments of experience, which are then concatenated during extended, multi-ripple replay events.

Segmentation of the environment by hippocampal theta sequences has implications for the encoding of experience and for how the hippocampus represents space. Since neural activity during theta represents experience as it occurs and operates at the time scale required for STDP, it is thought that synaptic learning during each theta cycle is the basis for storing experience. Thus, theta sequences representing the environment in segments may result in the encoding of experience in chunks. Storing experience this way could be a coding scheme that the brain uses to organize vast amounts of information in a way that is efficient and behaviorally useful to animals.

8 Model for backward replay generation

8.1 Abstract

The hippocampus is thought to consolidate information from past experiences by replaying sequences of neural activity during rest states. Navigation paths are replayed both in the order they were experienced (forward replay) and in reverse order (backward replay). It was recently shown that backward replay can occur over spatially and temporally remote trajectories, and that novel paths that were never traversed are also represented. To explore how known hippocampal physiology may enable the expression of these replay events, I built a computational model of hippocampal place cells in the dorsal third of CA3. Explicitly building in theta phase precession, the theta phase gradient, inter theta cycle learning, and hippocampal longitudinal projections into the simulations enabled the learning of both forward and backward associations using an STDP learning rule. Simulations showed that under certain conditions the network produced forward, backward, and novel path replays after traversal of a path in only one direction. If the hippocampus learns forward and backward associations this way, this may provide a basis for creating a cognitive map of an animal's environment, allowing it to flexibly navigate and plan novel paths.

8.2 Introduction

As a rat travels a path $A \rightarrow B \rightarrow C$, the place cells representing locations along that path become active in sequential order as the animal passes through each place field (the time scale of behavior). However, sequential firing at the behavior time scale is not sufficient for encoding the sequence: if each place cell fired within its place field, but fired independently at the time scale for spike timing dependent plasticity (STDP, Levy and Steward, 1983; Bi and Poo, 1998), in some instances the neuron representing location A would fire before B and in others B would fire before A. Despite the fact that hippocampal pyramidal cells

show STDP, the resulting synaptic connections would be relatively symmetric. However, individual place cells phase precess (O'Keefe and Recce, 1993; Skaggs et al., 1996; Maurer and McNaughton, 2007) and the order of firing is preserved within each theta cycle such that the neuron representing location A reliably fires before the neuron representing location B. This precisely coordinated firing, in conjunction with an STDP learning rule, enables the hippocampal network to develop asymmetric neural connections representing experiences in the order they were encountered (Levy and Steward, 1983; Skaggs et al., 1996; Blum and Abbott, 1996; Mehta et al., 1997).

However, as discussed in **Chapter 3**, during sharp wave ripples, hippocampal replay sequences represent paths both in the order that they were experienced (forward replay) and, surprisingly, in the reverse order of experience (backward replay). Backward replay has now been shown to occur over trajectories remote in time (see **Chapter 6**), which is difficult to explain using the proposal that backward replay results from a decaying trace of recently active neurons. In addition, sequences can be played that were never traversed by the animal (shortcut replay, **Chapter 6**), which cannot be explained by such a mechanism. These results imply that reverse connections are present in the hippocampal network and raises the question of how such connections are formed given that STDP learning rules applied to forward order sequences can only produce forward connections. A network that learns both forward and reverse connections could provide the basis for a cognitive map (Tolman, 1948; O'Keefe and Nadel, 1978; Muller et al., 1991), which could be used to plan novel behaviors and support flexible spatial navigation.

In this work, we explore how four aspects of hippocampal physiology and anatomy may contribute to the establishment of reverse connections and the generation of backward and shortcut replay sequences. The four components studied include hippocampal theta phase precession, neural activity occurring during adjacent theta cycles (inter-theta cycle activity), anatomical projections across the longitudinal or septo-temporal (ST) axis of the hippocampus, and the hippocampal theta phase gradient across the longitudinal axis (Lubenov and

Siapas, 2009). The way in which longitudinal projections and the theta phase gradient may contribute to learning reverse connections will be introduced in the next section.

As described above, hippocampal phase precession organizes the timing of spikes within each theta cycle such that forward sequences are observed and forward connections can be established according to STDP learning rules. However, it is unknown how neural activity in adjacent theta cycles may contribute to synaptic learning in the hippocampal network. If the neural sequences observed in two consecutive theta cycles are $A \rightarrow B \rightarrow C \rightarrow D \rightarrow E$ and $A \rightarrow B \rightarrow C \rightarrow D \rightarrow E$ (Figure 24), forward connections (e.g., $A \rightarrow B$) would be strengthened from the activity in a single theta cycle, as has been previously proposed. However, it is also possible that backward connections (e.g., $E \rightarrow A$) would be strengthened when considering the activity across theta cycles (inter-theta cycle learning). Although there is some evidence that synaptic learning occurs preferentially during certain phases of the theta cycle (Hasselmo et al., 2002; Hyman et al., 2003), it is unclear whether synaptic learning can only occur as a result of activity taking place within a single theta cycle.

In the following sections we present (1) a theoretical model of how the theta phase gradient could influence learning of excitatory connections between place cells, (2) computer simulations investigating the affects of phase precession, inter-theta cycle learning, and the theta phase gradient on the learning of forward and backward connections along the long axis, and (3) a computational model of a network of hippocampal place cells that exhibits forward, backward, and shortcut replay of learned sequences under certain conditions.

8.3 Theoretical Model: potential contribution of the theta phase gradient and longitudinal projections for learning reverse connections

The hippocampal circuit has typically been modeled as a two dimensional structure, with the third anatomical dimension (the septotemporal or long axis) not widely incorporated in modeling work to date. However, anatomical studies indicate that the hippocampus, in terms of

its physical connections, isn't a purely laminar structure, as projections in many hippocampal subregions extend substantially across the long axis (i.e., longitudinal projections, Amaral and Witter, 1989; Ishizuka et al., 1990, 1995; Wittner et al., 2007). Recurrent connections in hippocampal subregion CA3 extend up to two thirds of the way along the septotemporal axis (Swanson et al., 1978; Amaral and Witter, 1989; Ishizuka et al., 1990; van Strien et al., 2009; Bast et al., 2009; Adhikari et al., 2010). In fact, it has been shown that the density of projections is strongest 1 mm or more septally or temporally from the level at which the projections originated (Ishizuka et al., 1990). Experimental studies have demonstrated that neurons residing at different levels along the septotemporal axis have different place field sizes (Jung et al., 1994; Kjelstrup et al., 2008; Royer et al., 2010), that there are differences in functional connectivity with other brain structures along the septotemporal axis (Risold and Swanson, 1996; Dolorfo and Amaral, 1998; Cenquizca and Swanson, 2007), and that there is a phase gradient in hippocampal theta across the long axis (Lubenov and Siapas, 2009). These observations underlie the importance of understanding how the three dimensional spatio-temporal structure of the hippocampus contributes to its function. Here we describe how the hippocampal theta phase gradient may contribute to the encoding of experiences in a forward and reverse order in at least the dorsal third of the hippocampus.

Lubenov and Siapas (2009) observed that the hippocampal theta rhythm behaves as a traveling wave, propagating from dorsal to ventral (septal to temporal) hippocampus. They showed that along at least the dorsal third of the hippocampus, the phase gradient is approximately $21^{\circ}/\text{mm}$, which results in a phase difference of up to 124° . Furthermore, they recorded from neurons at different locations along the long axis and observed that there is a corresponding temporal shift in neural firing along this axis. They concluded that instead of the hippocampus representing a single location, it represents a trajectory at a moment in time, where neurons at each cross-section (level) of the hippocampus perpendicular to the long axis represent a location along the trajectory.

This implies that during a single theta cycle, the activity of neurons representing the

spatial trajectory $A \rightarrow B \rightarrow C \rightarrow D \rightarrow E$ has the spatiotemporal structure shown in **Figure 24**. As observed in **Figure 24**, the order of neural firing within a level of the hippocampus matches the order in which the trajectory was experienced (green arrows), with the consequence of learning the forward sequence via the STDP rule. However, due to the delay in activity across the long axis, it is apparent that some neurons representing locations later in the trajectory (e.g. E_3) become active prior to some neurons representing locations earlier in the trajectory (e.g. D_6). Since the recurrent connections in dorsal CA3 span up to two thirds of the long axis, this could enable the encoding of reverse connections via the STDP rule. Furthermore, the phase gradient combined with learning across theta cycles (inter-theta cycle learning, discussed earlier) could enable the learning of reverse connections from a more temporally located neuron to a more septally located neuron (e.g. D_6 to A_3).

Hippocampal plasticity studies have shown that STDP maximally occurs when the presynaptic cell spikes 0 to 20 ms before the postsynaptic cell (Bi and Poo, 1998). If two neurons representing locations A and B fire 20 ms apart (representing the extreme case), a neuron representing location A should be delayed by 20 to 40 ms in order for a reverse $B \rightarrow A$ connection to be established. This delay can be realized by selecting the neuron representing location A to be located more ventrally than the neuron representing location B. Assuming a theta frequency of 8 Hz, a 20 – 40 ms time delay corresponds to a $58 - 115^\circ$ theta phase shift. This phase shift, representing the extreme case, easily falls within the observed phase shift of up to 124° (Lubenov and Siapas, 2009). Of course, lower magnitude phase shifts could also enable reverse connections to be established. If neurons representing locations A and B fired with an offset of 5 ms, a delay in A of 10 ms (corresponding to a theta phase shift of 28.8°) would cause a neuron at another level of the long axis representing location A to fire 5 ms after the neuron representing location B, enabling the establishment of a $B \rightarrow A$ connection.

If reverse ordered firing is what establishes backward connections, why hasn't reverse ordered firing been observed in electrophysiology experiments? As shown in **Figure 24**




Figure 24: The proposed mechanism for encoding both forward and backward sequences. The figure depicts the activity of neurons within two theta cycles with time on the x-axis and location along the septotemporal (long) axis of the hippocampus on the y-axis. Each circle represents the firing of a single neuron. The letter (A-F) within the circle indicates the ordered location along a trajectory that the neuron is representing. The subscripted number indicates the depth of the neuron along the long axis. As shown in the figure, due to the temporal shift in firing along the long axis, it is possible for a neuron representing location 'E' to fire before a neuron representing location 'D' (red arrow from E₃ to D₆). Green arrows indicate spike timing that gives rise to forward connections.

the reverse ordered firing relationship occurs between neurons located at different positions along the long axis of the hippocampus. Since typical recordings from hippocampal place cells occur over a narrow range of the long axis of the hippocampus, reverse ordered firing relationships would likely not be seen, even if they existed. Even though backward firing is not seen during theta cycles, while animals are moving in the environment, as previously discussed forward and backward sequences are observed during sharp wave ripple complexes (SWRs). This observation could result if during SWRs neural activity was synchronized across the long axis such that at a brief moment in time only one location was being represented across the axis (as opposed to a trajectory, **Figure 25**). The next set of active neurons would then be driven by input from the previously active set of neurons. Assuming a refractory period after neural firing, the replay would propagate in one direction. Whether a replay event was a forward or backward sequence would be probabilistic and depend on the set of neurons activated by the initial population burst and the subsequent set of neurons that fire in response to the initial burst. Global inhibition would ensure that a single location was being represented and stochastic activation would allow the same initial




Figure 25: Activity spread during replay events. The figure depicts the activity of neurons during the first two time steps of a hypothetical replay event with spatial location along a trajectory on the horizontal axis and depth along the septotemporal (long) axis on the vertical axis. Each circle represents a single neuron. The letter (A-E) within the circle indicates the location of the neuron's place field center along the trajectory. The subscripted number indicates the position of the neuron along the long axis. Dark magenta circles indicate neurons that are firing during the time step and light magenta indicates neurons that are receiving input and will fire at the next time step. Light blue indicates neurons that are not activated due to global inhibition in the network and light yellow indicates neurons that are in a refractory state. In this schematic, we show that if spatial input causes a single location to be represented across the long axis at a moment in time, activity can spread either forward or backward through the network. In this example, activity spreads backward through the network.

population burst to sometimes result in a forward replay and sometimes a backward replay.

There is evidence for spatial coherence of SWRs across the long axis of the hippocampus (Chrobak and Buzsáki, 1996; Csicsvari et al., 2000). The spatial coherence increases as SWR power increases (Csicsvari et al., 2000) and during electrophysiology experiments, replay events are not observed during every SWR. Thus, it is possible that when spatial coherence is not achieved (i.e. during low power SWRs), coherent replay events might not be reliably initiated in the network.

In summary, we have described how the anatomical connections in the hippocampus, combined with the shifted neural activity arising from the theta phase gradient, can support the encoding of backward sequences, in addition to forward ones, using the STDP learning rule.

8.4 Synaptic Learning Model

In order to visualize the 3D anatomical structure of the forward and backward connections that may be learned in hippocampal region CA3, we simulated a network of place cells residing across 10 cross-sections (levels) of the dorsal third of the long axis, using cell counts and activity levels derived from anatomical and electrophysiological studies. We focused on CA3 because its recurrent projections provide a potential scaffold for sequence storage and because the region is thought to be involved in generating the SWRs associated with replay events (Buzsáki, 1989). The rodent hippocampus contains approximately 250,000 CA3 pyramidal cells (West et al., 1991; Rapp and Gallagher, 1996). Assuming a uniform distribution of cells across the long axis, we estimated that there would be a total of 83,333 pyramidal cells in the dorsal third of CA3. Approximately forty percent of place cells are active in a single environment (Guzowski et al., 1999), therefore we estimated that there would be 33,333 neurons in the dorsal third of the hippocampus with place fields in a given environment. Furthermore, we calculated from our 2T maze data (Gupta et al., 2010) that each neuron had fields covering 10% of the environment on average, meaning that approximately 3,333 neurons in the dorsal third of the hippocampus would be active at a particular location in a given environment and 333 neurons active at a particular location in each of the 10 long axis levels. Thus, we simulated a network with a total of 33,333 neurons, each with a single Gaussian place field determining the probability of firing, which were identically sized and uniformly distributed along a linear track. Because we were modeling the dorsal third of the hippocampus only, we ignored the increase in place field size along the septotemporal axis (Jung et al., 1994; Kjelstrup et al., 2008; Royer et al., 2010).

For a single neuron residing at each of the 10 levels, we simulated 40 unidirectional paths through the neuron's place field and considered the learning of synaptic connections between that neuron and neurons with overlapping place fields residing in each level. Simulations were performed separately for eight different parameter sets: with and without the

theta phase gradient, with and without inter theta cycle learning, and with and without theta phase precession.

The order of cell firing for neurons in a particular cross-section was determined once per theta cycle (5 cycles total, although more cycles produce similar results) during the place field pass by considering the neurons with place fields overlapping with the animal's location. In the case that phase precession was implemented, the firing of these cells was ordered such that neurons with place fields behind the animal fired before neurons with place fields ahead of the animal (each cell could only fire once per theta cycle). For simulations without phase precession, the order of neuron firing was random within each theta cycle. In order to increase the learning rate, the probability of firing obtained from the Gaussian tuning curves was scaled such that on average 90% of the neurons with fields overlapping with the animal's location fired during each theta cycle (although results were qualitatively similar for lower percentages) and the timing of spikes was uniformly spaced across the theta period. Although this is a simplification of the activity seen during each theta cycle, these measures ensured that for neurons within a level, when phase precession was implemented, the order of behavioral experience was preserved in the compressed sequence of firing during each theta cycle (Dragoi and Buzsaki, 2006), and over the course of the place field pass neurons would gradually fire earlier and earlier in the theta cycle (i.e., phase precess, O'Keefe and Recce, 1993; Skaggs et al., 1996). When implementing the theta phase gradient reported by Lubenov and Siapas (2009), a 120 degree phase shift (corresponding to a $120 \text{ deg} \times (125 \text{ ms} / 360 \text{ deg}) = 41.667 \text{ ms}$ delay in neural firing, assuming an 8 Hz theta frequency) was imposed between the septal-most level and the temporal-most level, as shown schematically in **Figure 24**.

Each CA3 pyramidal cell forms efferent synapses with approximately 1,242 other CA3 pyramidal cells residing in the dorsal third of the hippocampus, and these synapses can be approximated as being uniformly distributed throughout the long axis (Wittner et al., 2007). However, in order to calculate the average connection strengths between a particular

neuron and neurons with place fields ahead and behind the animal at different levels of the long axis, we implemented a fully connected network.

A 33,333 element vector describing the efferent synaptic strengths between a neuron at one level of the long axis and all the other neurons at all levels was initially set to zero and was updated during each unidirectional place field pass according to the relative spike timing of neuron pairs (see equations below for the update rule and **Figure 26C** for a graphical representation). This update rule closely resembles the STDP curve presented in Bi and Poo (1998). To study the effect of our proposed learning rule, we selected one neuron per level and trained its weights. In the equations below, dt is the time difference between the spikes emitted from two neurons, and dW is the change in the synaptic weight between the neuron pair.

For $dt \geq 0$:

$$dW = \exp(-dt/20)$$

For $dt < 0$:

$$dW = \frac{-\exp(dt/30)}{2}$$

For $|dt| > 80$:

$$dW = 0$$

The learned forward and backward connection strengths from a single neuron at each level of the long axis to neurons with overlapping place fields are represented in **Figure 26**, for the simulation in which the theta phase gradient, inter-theta cycle learning, and phase precession were all implemented. As shown in the figure, forward and backward connections are established between neurons at all levels of the long axis of the hippocampus and the pattern of connections established depends on the neuron's location along the long axis.

In order to compare how the connections learned depend on the theta phase gradient,

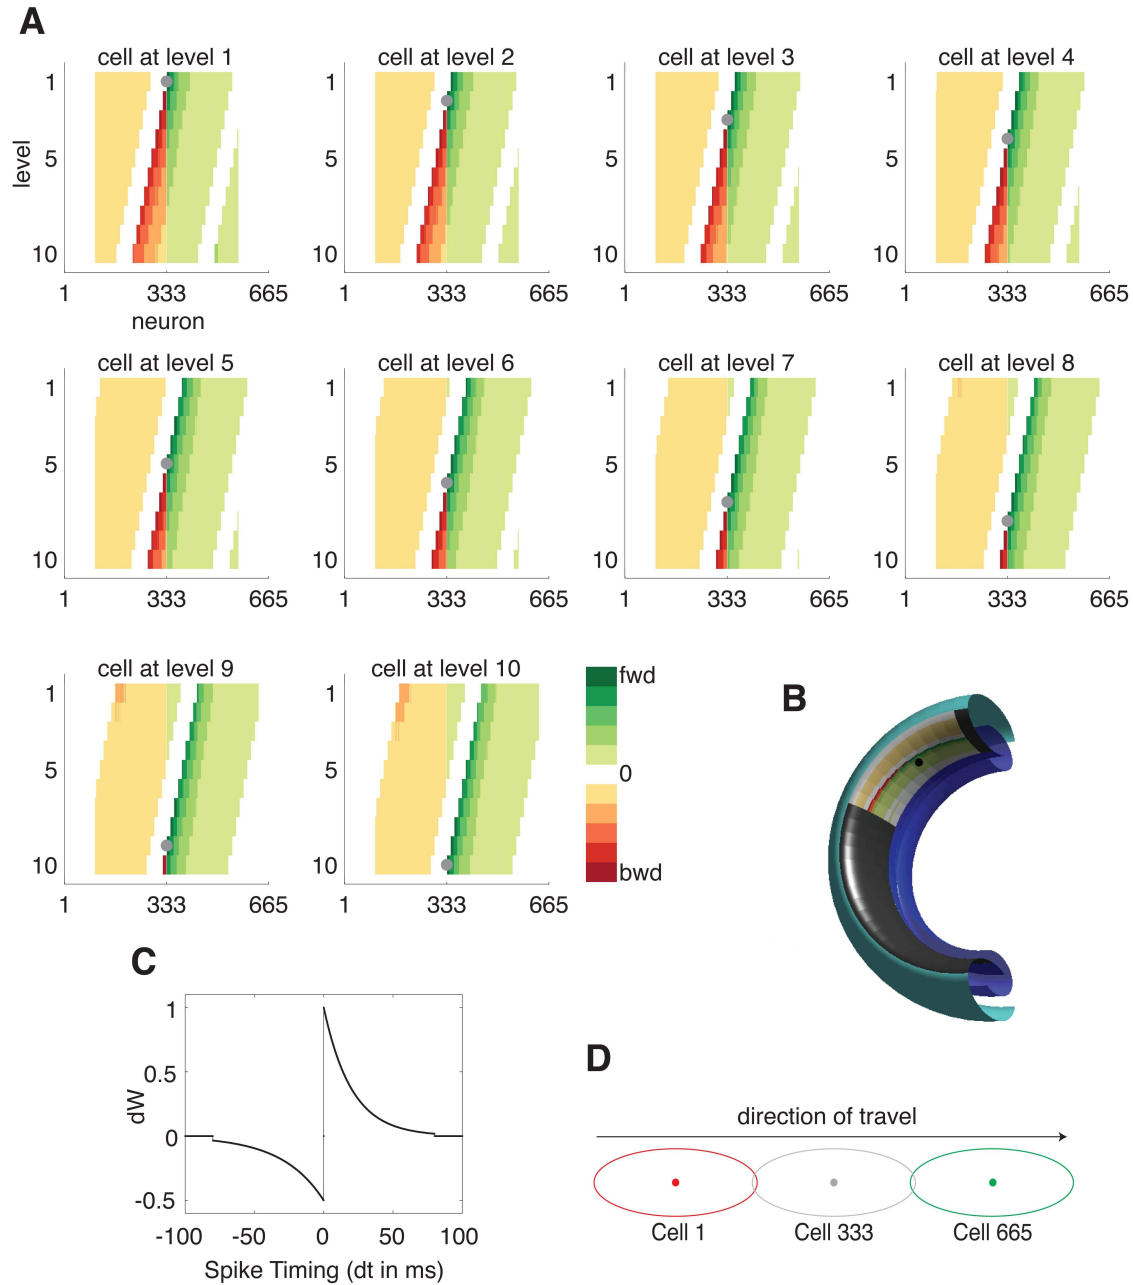


Figure 26: Connection strengths between a single neuron at each level and all other neurons. (A) The y-axis is 10 septo-temporal levels and the x-axis is cells ordered according to place field center. As depicted in panel D, cell 1 is a neuron with a place field behind and barely overlapping with the considered cell (shown in gray) and cell 665 is a neuron with a place field ahead and barely overlapping with the considered cell. The considered cell is always a cell at position 333, but resides at different levels along the long axis in each of the panels in A (shown by gray dots). Color on these panels indicates the forward or backward connection strength between the considered cell and all other cells. The color map ranges from dark green (strong forward connection) to white (zero connection strength) to dark red (strong backward connection). As required from the definition of forward and backward connections, forward connections are made with neurons to the right of cell 333 (shown by the gray dots) and backward connections are made with neurons to the left of cell 333. (B) The connection strengths for a cell residing at level 5 (fifth plot in panel A) were projected onto a 3D surface representing the rodent hippocampus. In this figure, the dentate gyrus is blue, CA3 is dark gray, and CA1 is green. (C) The STDP learning rule (equation in text) that was used to determine the synaptic strengths plotted in panel A.

inter-theta cycle learning, and phase precession, the connections between a single neuron residing at level 5 of the long axis were plotted for each combination of these phenomena (**Figure 27**). In all but one case (the no phase gradient, no inter-theta cycle learning, and phase precession case), both forward and backward connections were observed. In the four simulations involving no phase precession, connections were established relatively uniformly from the neuron at level 5 and all other neurons, independent of place field location or anatomic location across the long axis. Thus, forward and backward connections were created, but there was no observable structure to the connections. In the two phase gradient and phase precession simulations, forward connections were established that were proportional to the distance between place fields. Furthermore, strong backward connections were established between neurons with nearby place fields. For the simulation including inter-cycle learning, additional, weaker, backward connections were also established. For the last simulation, with no phase gradient, inter-theta cycle learning, and phase precession, weaker backward connections were established with cells with place fields some distance away. These simulations indicate that the three phenomena simulated strongly affect the pattern of forward and backward connections established during behavior. In the next section, we explore which combinations of these phenomena allow forward, backward, and shortcut replays to be produced by the network.

8.5 Replay Model

Given that most of the simulated networks described above learned both forward and backward connections, we asked if the connections created in these simulations could allow for the forward, backward, and shortcut sequences that we had observed on the 2T maze (Gupta et al., 2010). To address this question, we built a second, smaller network with a two dimensional grid of place cells. Similar to the previous network, this network had place cells residing across 10 levels of the dorsal third of the long axis of the hippocampus. Each level

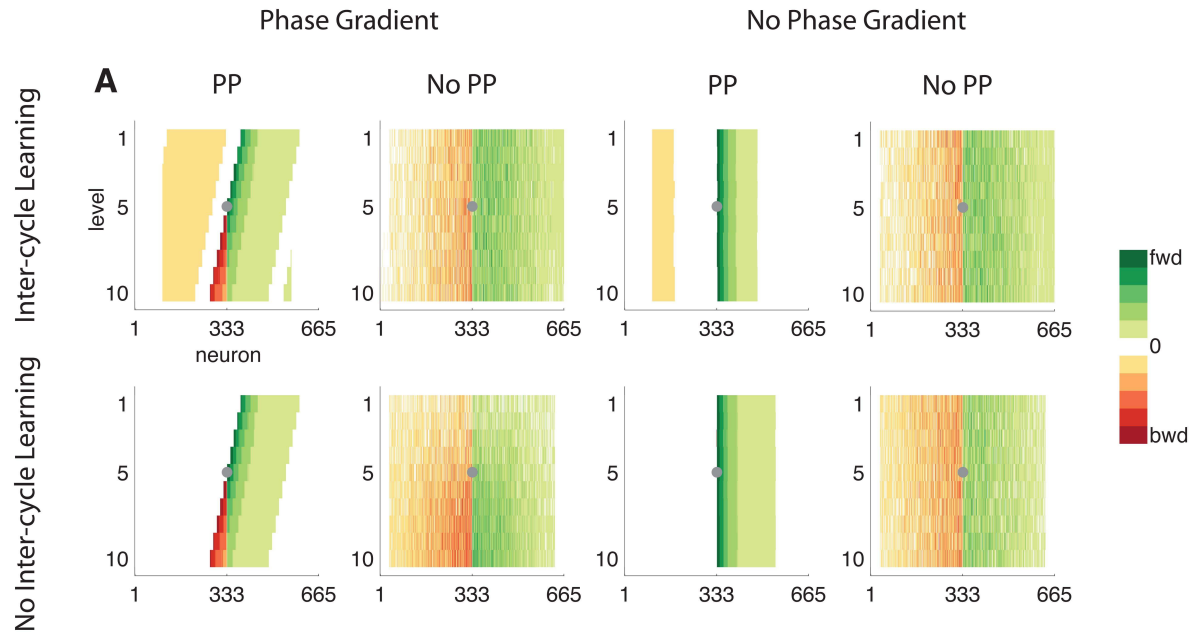


Figure 27: Connection strengths between a single neuron at level 5 and all other neurons for each simulation. See caption for **Figure 26A**. "PP" stands for phase precession.

contained 289 place cells, with place field centers distributed over a 17x17 point grid spanning a 2D square environment. Thus, there were a total of 2890 place cells in the network (10 at each location on the 17x17 grid). We implemented a fully-connected network (every neuron projected to every other neuron), due to the relatively small number of neurons simulated here compared with the rodent hippocampus. Each place field center was randomly jittered so that no two cells had place field centers at the exact same location in the environment. For simplicity and since we modeled the dorsal third of the hippocampus and not more ventral regions where place field sizes increase substantially (Jung et al., 1994; Kjelstrup et al., 2008; Royer et al., 2010), we assumed a uniform spatial scale across levels: each place cell, regardless of level, was randomly assigned a place field covering between 4 and 9% of the environment. However, our results were quantitatively similar when we imposed a gradient of place field sizes from 2 to 9% of the environment along the dorsal third of the hippocampus. This gradient was chosen based on the Kjelstrup et al. (2008) data showing that place field sizes in intermediate hippocampus were on average approximately

four times the size of place fields in dorsal hippocampus.

To simulate neural activity during behavior, the average figure eight trajectory taken by rats on the 2T maze (see **Figure 28G**) was superimposed on the grid of place cells. We simulated a constant running speed through the figure eight trajectory. The figure eight trajectory took a total of 30 seconds and was repeated 40 times in the simulation (similar to a typical rodent recording session on the 2T maze, Gupta et al., 2010). Once per theta cycle during behavior (every 125 ms), the neural activity during the theta cycle was determined as described below.

1) The potential set of active neurons consisted of those cells whose place fields overlapped with the animal's location at the start of the theta cycle. Each neuron fired stochastically according to its Gaussian tuning curve, and as in the previous section, the probability of firing was scaled such that on average 90% of the neurons with fields overlapping with the animal's location fired during each theta cycle. Noise was added to the inter spike intervals (ISI) such that both the set of active neurons and their spike timing could differ for each lap of the simulation. The minimum ISI was set by finding the location in the environment with the largest number of overlapping place fields. By assuming each cell spikes and evenly distributing the spikes over a 125 ms theta cycle, we determined the minimum ISI as the timing between adjacent spikes. The maximum ISI was variable and depended on the actual number of active cells during a given theta cycle of the simulation. For example, if 10 cells were active during a theta cycle, the maximum ISI would be 12.5 ms. Noise was added such that the actual ISI was uniformly distributed between the minimum and maximum ISI.

2) For simulations including the phase precession phenomenon, the order of neuron spiking was determined by projecting the place field center of each firing neuron onto the vector representing the animal's direction of travel, with the cell whose place field center projected furthest back on the vector firing first. This computation, which may not be the one performed by the brain, nevertheless produces the forward sequenced firing patterns and theta phase precession observed during electrophysiology experiments. For simulations

not including phase precession, the spike order was randomly determined.

3) As in the previous section, for simulations involving the theta phase gradient and accompanying delay in neural activity across the long axis of hippocampus, a 120 degree phase shift (corresponding to a 41.667 ms shift in neural firing, assuming an 8 Hz theta frequency) was imposed between the septal most level and the temporal most level. The phase shift was set to zero for simulations lacking the phase gradient.

The $n \times n$ ($n = 10 \times 17 \times 17$) synaptic strength matrix was initially set to zero and was updated during each theta cycle using the same STDP learning rule used in the previous section (see **Figure 26C**). This learning procedure was performed separately for each binary combination of the following phenomena (i.e., with and without): the theta phase gradient, anatomical longitudinal connections, inter-theta cycle learning, and phase precession.

After simulating 40 laps of the figure eight trajectory in a single direction of travel, we visualized forward and backward connections between sets of neurons representing locations in the environment for one combination of the hippocampal phenomena: (+) theta phase gradient, (+) anatomical longitudinal connections, (+) inter-theta cycle learning, and (+) phase precession. Each location in the environment (each grid point) was represented by 10 neurons per level and the average connection strength between each grid point and all other grid points (based on the synaptic strengths of their neural connections) was computed. By comparing the direction of the connection with the direction of the traveled trajectory, the connection was determined to either be a forward connection (green) or backward connection (red). For this particular combination of phenomena, the subset of neural connections that occur within a level (intralevel connections) of hippocampus should be primarily forward (**Figure 28A,B**) and the subset of connections occurring between levels (interlevel connections) should be both forward and backward (**Figure 28C,D,E,F**). As shown in **Figure 28**, this was indeed the case.

As described above, during learning the neural activity was solely a function of spatial input from an extrahippocampal source (such as MEC, Fyhn et al. (2004)). The synaptic

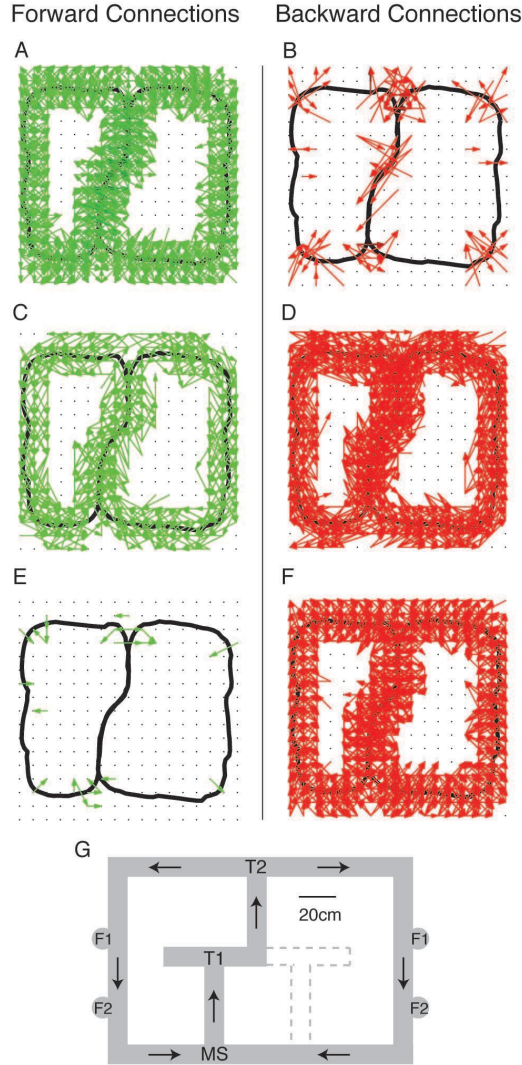


Figure 28: 2T maze, forward and backward synaptic strengths. Grid points, each representing the approximate place field center (due to jittering) of 10 neurons positioned at 10 different locations along the long axis, are plotted in black. Green arrows, beginning at one grid point and ending at another grid point, indicate that a forward connection was learned between a pair of neurons with place field centers at the two grid points. Red arrows, similarly, represent backward connections. In order to visualize connections being learned within a level and across levels, we plotted intralevel connections (connections learned between neurons residing at the same level along the long axis) in panels A and B and interlevel connections (connections learned between neurons residing at different levels along the long axis) in panels C, D, E, and F. Panels C and D show the interlevel projections from neurons in levels 9 and 10 to neurons in levels 1 and 2, while panels E and F show projections from neurons in levels 1 and 2 to neurons residing in levels 9 and 10. These panels indicate that forward and backward connections are not equal between all level pairs as there are few strong forward connections from neurons in levels 1 and 2 to neurons in levels 9 and 10 (E), while there are many strong forward connections from neurons in levels 9 and 10 to neurons in levels 1 and 2 (C). As can be seen by these plots and as predicted by the theoretical model for the (+) theta phase gradient, (+) anatomical longitudinal connections, (+) inter-theta cycle learning, and (+) phase precession combination of phenomena, primarily forward connections are learned between neurons at the same level (A and B), whereas both forward and backward connections are learned between neurons at different levels (C, D, E, and F). For clarity, connection strengths one standard deviation or more below the mean are omitted from these plots. Panel G displays the 2T maze from Gupta et al. (2010). Arrows indicate the direction of travel around the maze. MS is maze start, T1 is turn 1, T2 is turn 2, F1 is feeder 1, and F2 is feeder 2.

connections between neurons within the network did not contribute to the activity of individual neurons.

In the simulation of replay, the starting location of each replay event was determined by spatial input from an extrahippocampal source, which caused neurons with place fields at that location to fire. Subsequent neural activity was determined by synaptic connections learned during behavior plus a fixed inhibitory network (as described above, learning was different depending on the set of hippocampal phenomena included in the simulation). To implement the inhibitory network, negative connection strength weights (representing the inhibitory interneuron network) were imposed between neurons whose place fields did not overlap (i.e., their learned synaptic weights were zero). The purpose of the inhibitory network was to impart attractor-like properties to the model to support the representation of a single location at each moment in time (Wilson and McNaughton, 1993; Samsonovich and McNaughton, 1997; Zhang, 1996; Redish, 1999). The negative connection strength parameter was tuned such that the network was forced to represent a single location at each moment in time during the replay event (value = -0.5).

The simulation of each replay event was performed in discrete cycles. Below is a summary of the first cycle (subsequent cycles only involve step 2):

- 1) Spatial input specifying a random location in the environment (location A) causes some neurons with place fields overlapping with location A to spike. To impose high variability in the set of neurons that fire, each neuron receiving spatial input had an equal probability of firing, and roughly 5% of neurons fired during each cycle of the replay event. The equal probability of firing is consistent with the hyper-excitability of neurons during SWRs.

- 2) The spiking neurons send input (according to the learned synaptic strengths and the inhibitory network) to another population of neurons, representing the next location (location B). A refractory property assures that neurons that fired during the previous cycle cannot reactivate too quickly. This enables the replay event to move away from the previously represented location and form a coherent path. As in step 1, on average 5% of the neurons

fired and each neuron receiving input fired with equal probability.

Step 2 was repeated 15 times for a single replay event, in order to generate a sequence of activity representing a path in the environment. Replay events were generated for each simulation (with different sets of hippocampal phenomena included) as described in the Synaptic Learning Model section. For each simulation, replay events were analyzed to determine if (1) coherent paths were represented, (2) if forward, backward, and shortcut paths were expressed, and (3) if the same paths were represented at different levels of the long axis (i.e., synchrony along the long axis).

For the (+) theta phase gradient, (+) anatomical longitudinal connections, (+) inter-theta cycle learning, and (+) phase precession simulation, coherent forward, backward, and shortcut sequences were observed. Furthermore, these sequences were synchronous across the long axis of the hippocampus (**Figure 29**). Interestingly, removing the theta phase gradient did not affect these observations: forward, backward, and shortcut sequences were present and synchronous across the long axis. Simulations without the longitudinal connections showed two differences: synchrony across the long axis was lost and backward and shortcut sequences were rarely observed. When inter-theta cycle learning was turned off, only forward replays were observed. Thus, inter-theta cycle learning proved to be critical for the generation of backward and shortcut sequences. Finally, for simulations lacking phase precession (active cells fired at random times within the theta cycle), coherent paths were not observed. Replays began at a particular location and spread in all directions rather than in a single direction.

8.6 Discussion

We have shown through simulation that the hippocampal phenomena studied here allow for the learning of forward and backward connections and the generation of forward, backward, and shortcut replay events. Furthermore, through the simulation of different combinations of

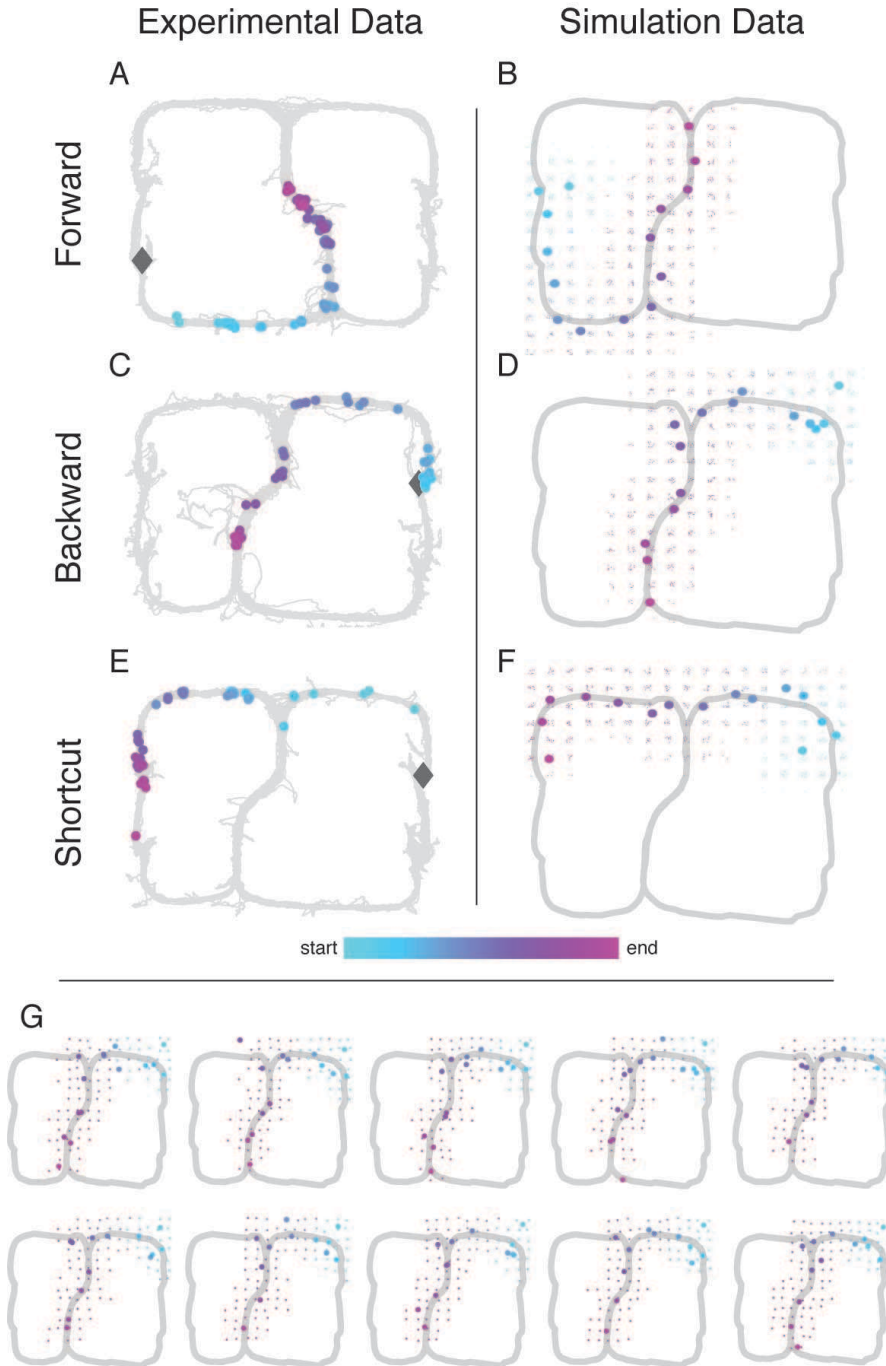


Figure 29: Forward, backward, and shortcut sequences. The left column shows actual replay events recorded on the 2T maze (Gupta et al., 2010) and the right column shows simulated replay events. Left column: colored points are spikes plotted at the firing neurons' place field centers. Right column: two types of points are plotted. The small points are spikes plotted at the simulated neurons' place field centers. Each large point is the center of mass of the spiking neurons' place field centers for each cycle of the replay event. For all panels, the color of the point indicates the relative time within the replay event (light blue = early, light purple = late). Panel G takes the simulated backward replay event from panel D and plots the activity of neurons at each level of the septotemporal axis separately (top row, left to right is levels 1-5; bottom row, left to right is levels 6-10). As seen in this panel, clear backward replays are observed at each level separately, matching with experimental observations.

hippocampal phenomena, potential contributions of each phenomena to replay were identified. For example, inter-theta cycle learning was important for generating backward and shortcut sequences, longitudinal connections were necessary for replay synchrony across the long axis, and phase precession was required for generating sequences representing coherent paths in the environment.

For the learning of backward associations between hippocampal place cells, we proposed two potential mechanisms. The first mechanism was based on the hippocampal theta phase gradient and the second mechanism was based on inter-theta cycle learning. In the Theoretical model section, we explained in detail how the phase gradient could enable the learning of backward connections and subsequently showed in the Synaptic learning model that the phase gradient and inter-theta cycle mechanisms independently learned forward and backward connections (**Figure 27**). However, in the simulation of replay, we found that only inter-theta cycle learning (not the theta phase gradient) was required for the generation of backward and shortcut replay events.

Thus, we have presented a potential network mechanism for backward replay. It was previously proposed that backward replay would result from a decaying remnant signal in recently active neurons due to changes in the intrinsic properties of place cells (Foster and Wilson, 2006) or the short-term modification of synaptic connections (Buzsáki, 1989; Diba and Buzsáki, 2007). In brief, the proposal was that if an animal runs the path $A \rightarrow B \rightarrow C$ and then pauses, because a neuron representing C fired most recently it is the most excitable and therefore fires first, followed by B and then A. However, the recent observations that backward replay can occur over spatially and temporally remote trajectories (Davidson et al., 2009; Gupta et al., 2010) imply that such a mechanism is insufficient. In the models simulated here, backward associations were stored in synaptic connections between neurons, thereby allowing remote backward sequences to be represented during replay events. In addition, associations stored in this way enabled the expression of single replay events composed of both forward and backward sequences, as observed in recent experimental

work (Davidson et al., 2009; Gupta et al., 2010).

The details of synaptic learning during each theta cycle are currently unknown. It has been proposed that the encoding and recall of experience occurs during separate phases of each theta cycle, mediated in part by the neurotransmitter acetylcholine (Hasselmo et al., 2002; Hasselmo, 2006). Thus, it is possible that synaptic learning is susceptible only to neural activity occurring within a single theta cycle. If this is the case, inter-theta cycle learning would not be possible and the backward connections required to produce backward and shortcut replay events in the model described here would not be created. Future experimental work investigating synaptic learning between hippocampal place cells during behavior will help determine whether the creation of backward associations via inter-theta cycle learning is possible.

In summary, we built a model of hippocampal CA3, which incorporated several aspects of hippocampal physiology and anatomy: the theta phase gradient, anatomical longitudinal connections, inter-theta cycle learning, and theta phase precession. We showed that these phenomena (in particular, inter-theta cycle learning) enabled the synaptic learning of forward and backward connections and the replay of local and non-local forward, backward, and shortcut sequences. Thus, the hippocampus contains hardware and software that could support the construction and mental navigation of a cognitive map of the environment.

9 Conclusion

9.1 Summary

In this thesis, I have studied hippocampal neural activity during awake rest and awake behavior, revealing two new potential contributions of hippocampal sequences to cognitive function.

First, the spatial paths represented by neural sequences during awake rest (replay sequences) reflected the learning, maintenance, and presence of a cognitive map of the environment. Although only a subset of possible paths in the environment were experienced by animals, I found that replay sequences represented all possible paths in the environment: paths in the opposite direction than experience and shortcut paths that were never traversed. Furthermore, parts of the environment that were less frequently experienced were more frequently replayed. These findings demonstrate that replay does not simply reflect the consolidation of experience and suggests a role for replay in maintaining and traversing an internal representation or cognitive map of the environment. The concept of cognitive maps is of particular interest in cognitive science because the mechanisms by which the brain constructs internal representations of the world may provide valuable insights into how we learn, understand, and interact with our environment. These maps may also support higher order cognitive functions such as planning, self-projection, and creativity.

Second, during active behavior the spatial paths represented by hippocampal theta sequences suggested the involvement of these sequences in the cognitive chunking of experience. Previous work has shown that the represented path rigidly begins just behind the animal and ends just ahead of the animal, implying that spatial experience is uniformly represented. I found that the path represented was flexible and depended on the animal's behavior and the structure of the environment. When animals were accelerating, the represented path shifted ahead of the animal, potentially to aid in planning and in anticipation of an upcoming location. Conversely, when animals were slowing down, the path shifted

behind the animal, potentially in review of what was just experienced. The observed shift in represented path relative to the animal's location, allowed for the possibility that theta sequences non-uniformly represented the environment. I found that this was in fact the case, with sequences over-representing segments of the environment between important landmarks on the maze. Thus, theta sequences offer a potential mechanism for the cognitive chunking of experience: a process by which information is recoded into computationally efficient and useful units, allowing for more information to be processed at once. This reorganization of information is thought to enhance brain function (for example, by increasing memory capacity and improving planning efficiency). Thus, as with the concept of cognitive maps, the neural mechanisms underlying cognitive chunking is an exciting topic in cognitive science.

In the last part of the thesis, I simulated a network of hippocampal place cells to explore how known hippocampal physiology could support the learning of forward and backward synaptic connections. Such connections could allow the hippocampus to produce the forward, backward, and shortcut replays that I found on the 2T task, and may be the basis of cognitive maps in the hippocampus. The simulations revealed that inter-theta cycle learning was a key component for learning backward connections, and when included in the simulation, it allowed the network to produce the three types of replay sequences observed experimentally. While it hasn't been shown that synaptic learning occurs across theta cycles, our simulations show that if it were to occur, the hippocampus could learn backward connections and produce backward replays. Future experimental work could test this intriguing possibility.

More generally, the work in this thesis demonstrates that information represented by the hippocampus reflects a role not only in memory, but also in the manipulation of experiential information in potential support of higher order cognitive functions (i.e., cognitive maps and cognitive chunking). Furthermore, this work shows how decoding the information represented by populations of spiking neurons in behaving animals can reveal insights about the

contribution of brain structures to cognition.

9.2 Future work

Replay sequences

The work in this thesis and other recent observations raise questions about the functional significance of replay sequences and the behavioral factors that influence the content that is represented. Important topics of future research on replay sequences are outlined below.

1. *The representation of novel paths.* As discussed above, I found that novel shortcut paths were represented by replay sequences. These were relatively rare occurrences and this finding needs to be replicated in novel experimental paradigms to determine if this is a general property of hippocampal replay (Dragoi and Tonegawa, 2011). This is important as it has implications for the role of replay in more generalized learning (e.g., creating cognitive maps) as opposed to strictly memory consolidation.

2. *What is backward replay?* Work in this thesis has shown that backward replay occurs over paths in an environment that have only been experienced in the forward direction. However, it is unclear whether backward replay is the representation of an experienced path in the reverse direction, or the representation of a new path in the environment that is physically available to the animal. Understanding what backward replay is will help us understand the role it plays in hippocampal function. Experiments that simultaneously record replay events in the hippocampus and a separate structure with direction-specific representations (e.g., head direction cells), could help answer this question. Brandon et al. (2011) observed that head direction cells in the postsubiculum do not show replay, but it is possible that examining the behavioral content of replay firing patterns in structures such as parietal cortex (Qin et al., 1997; Nitz, 2006), may provide insight into the identity of backward replay.

3. *Construction versus traversal of a map.* If the representation of novel paths is a general property of replay, what purpose is it serving? Given that replay sequences occur

at the time scale required for spike timing dependent plasticity, these replays could play a role in constructing a cognitive map of the environment. On the other hand, novel path replay may reflect the traversal of an already existing map (Samsonovich and Ascoli, 2005), potentially to explore available paths that are useful to the animal. Another possibility is that replay is traversing a map that exists in one structure (e.g., hippocampus) in order to store the information in another structure, such as cortex.

4. *Awake versus sleep replay.* Backward replay has been observed during the awake state, but not during sleep sharp wave ripples. Early evidence has suggested that only forward sequences are present during sleep (Skaggs and McNaughton, 1996). It is important to continue to investigate differences in the content of replay during awake and sleep states (O'Neill et al., 2006; Karlsson and Frank, 2009), as that would suggest there are functional differences. It is possible that sleep replay is a more literal reflection of experience, serving mainly a memory consolidation function, whereas awake replay represents all possible paths, reflecting a general learning or search function.

5. *Sensory influence on awake replay.* It has been shown that sleep and particularly awake replay, both preferentially represent the animal's location in the environment (Jarosiewicz and Skaggs, 2004; O'Neill et al., 2006). During the awake state, animals have access to sensory input which could influence the content of replay, thereby contributing to the preferential replay of spatial information at the animal's location (O'Neill et al., 2006). This thesis and other work have demonstrated that the paths represented during awake replay are not always tied to the animal's location in the environment, but it is possible that visual or other sensory cues from remote locations contribute to the representations. Future studies that manipulate the animal's sensory environment by presenting various cues when animals are awake but restful, could test whether sensory input influences the content of replay.

6. *Relationship with behavior.* How does the content of replay influence the animal's behavior? It has been shown that disrupting sharp wave ripples during sleep affects perfor-

mance on spatial tasks (Girardeau et al., 2009; Ego-Stengel and Wilson, 2010). However, it is unclear if and how the specific content of replay affects the animal's future behavior. For example, is it necessary for the hippocampus to represent novel shortcut paths in an environment in order for the animal to realize the path is available?

Theta sequences

The study of neural sequences at a population level during each theta cycle is relatively new and there are important questions to address in future work, with the goal of understanding the role that these sequences play in experience encoding and recall. Specific topics of interest are outlined below.

1. *Changes in represented path during learning.* When animals are first placed in a novel environment, it is possible that the hippocampus is primarily encoding experience, rather than performing both encoding and recall functions. As animals gain experience in the environment, animals may begin to recall information from past experiences to aid in navigation. It will be important to examine how the path represented during theta sequences evolves with learning in a novel environment. One hypothesis is that the amount of space represented ahead of the animal increases with learning, potentially reflective of increased information recall.

2. *Evolution of neural chunking.* Does neural chunking arise immediately in a novel environment or does it gradually commence over time? If neural chunking by theta sequences is a reflection of cognitive chunking, studying how neural chunking arises may help us understand the mechanisms underlying cognitive chunking.

3. *Behavioral versus neural chunking.* In the 2T maze experiment behavioral and neural chunking occurred in parallel, potentially due to the physically segmented structure of the maze, with turns and stopping points. An important next step is to develop experimental paradigms in which there are sensory or task related segments, instead of physical segments. Such experiments could help reveal whether behavioral and neural chunking are

separable and therefore independent processes. On tasks in which behavioral and neural chunking is not separable, looking at the evolution of each type of chunking could provide insight into a possible causal relationship between these two processes.

4. *Attentional and sensory influences on theta sequences.* I found that when animals entered a segment of the maze, initially the representation was more ahead of the animal and towards the end of the segment, the representation was more behind the animal. One speculation is that upon entering the segment, the animal's attention is further ahead of the animal, which shifts the representation forward, potentially via sensory cues. This observation is consistent with the finding of forward representations during vicarious trial and error (VTE) events, as animals deliberate options at a choice point (Johnson and Redish, 2007). Experiments that manipulate the animal's attention and sensory environment can help us understand the factors influencing the content of theta sequences.

5. *Short versus long sequences.* Investigating the relationship between the path represented by theta sequences and the animal's acceleration and velocity revealed that some sequences, which generally represented shorter paths in the environment, were influenced by the animal's behavior. However, other sequences representing longer paths either did not have a strong relationship with behavior, or had the opposite relationship with behavior compared to the shorter sequences. This raises the question as to whether there are multiple classes of theta sequences. Future studies could look for other differences between the shorter and longer sequences. One question of particular interest is whether the identification of type 1 and type 2 theta cycles (Kramis et al., 1975) can distinguish between the shorter and longer sequences.

References

- Adhikari, A., Topiwala, M. A., and Gordon, J. A. (2010). Synchronized activity between the ventral hippocampus and the medial prefrontal cortex during anxiety. *Neuron*, 65(2):257–269.
- Alvarez, P. and Squire, L. R. (1994). Memory consolidation and the medial temporal lobe: A simple network model. *Proceedings of the National Academy of Sciences, USA*, 91:7041–7045.
- Amaral, D. G. and Witter, M. P. (1989). The three-dimensional organization of the hippocampal formation: A review of anatomical data. *Neuroscience*, 31(3):571–591.
- Barnes, C. A., Suster, M. S., Shen, J., and McNaughton, B. L. (1997). Multistability of cognitive maps in the hippocampus of old rats. *Nature*, 388(6639):272–275.
- Bast, T., Wilson, I. A., Witter, M. P., and Morris, R. G. M. (2009). From rapid place learning to behavioral performance: a key role for the intermediate hippocampus. *PLoS Biol*, 7(4):e1000089.
- Bi, G. and Poo, M. (1998). Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neuroscience*, 18(10464-10472).
- Blum, K. I. and Abbott, L. F. (1996). A model of spatial map formation in the hippocampus of the rat. *Neural Computation*, 8(1):85–93.
- Brandon, M. P., Bogaard, A. R., Andrews, C. M., and Hasselmo, M. E. (2011). Head direction cells in the postsubiculum do not show replay of prior waking sequences during sleep. *Hippocampus*.
- Brown, E. N., Frank, L. M., Tang, D., Quirk, M. C., and Wilson, M. A. (1998). A statistical paradigm for neural spike train decoding applied to position prediction from ensemble

- firing patterns of rat hippocampal place cells. *Journal of Neuroscience*, 18(18):7411–7425.
- Buckner, R. L. and Carroll, D. C. (2007). Self-projection and the brain. *Trends in Cognitive Sciences*, 11(2):49–57.
- Buzsáki, G. (1989). Two-stage model of memory trace formation: A role for “noisy” brain states. *Neuroscience*, 31(3):551–570.
- Buzsáki, G. (2005). Theta rhythm of navigation: link between path integration and landmark navigation, episodic and semantic memory. *Hippocampus*, 15(7):827–40.
- Buzsáki, G., Leung, L. W., and Vanderwolf, C. H. (1983). Cellular bases of hippocampal EEG in the behaving rat. *Brain Research*, 287(2):139–171.
- Cenquizca, L. A. and Swanson, L. W. (2007). Spatial organization of direct hippocampal field ca1 axonal projections to the rest of the cerebral cortex. *Brain Res Rev*, 56(1):1–26.
- Cheng, S. and Frank, L. M. (2008). New experiences enhance coordinated neural activity in the hippocampus. *Neuron*, 57:303–313.
- Chrobak, J. J. and Buzsáki, G. (1996). High-frequency oscillations in the output networks of the hippocampal-entorhinal axis of the freely behaving rat. *Journal of Neuroscience*, 16(9):3056–3066.
- Cohen, N. J. and Eichenbaum, H. (1993). *Memory, Amnesia, and the Hippocampal System*. MIT Press, Cambridge, MA.
- Csicsvari, J., Hirase, H., Mamiya, A., and Buzsáki, G. (2000). Ensemble patterns of hippocampal CA3-CA1 neurons during sharp wave-associated population events. *Neuron*, 28:585–594.

- Davidson, T. J., Kloosterman, F., and Wilson, M. A. (2009). Hippocampal replay of extended experience. *Neuron*, 63(4):497–507.
- Derdikman, D. and Moser, M.-B. (2010). A dual role for hippocampal replay. *Neuron*, 65(5):582–4.
- Diba, K. and Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience*, 10:1241–1242.
- Dolorfo, C. L. and Amaral, D. G. (1998). Entorhinal cortex of the rat: topographic organization of the cells of origin of the perforant path projection to the dentate gyrus. *J Comp Neurol*, 398(1):25–48.
- Dombeck, D. A., Harvey, C. D., Tian, L., Looger, L. L., and Tank, D. W. (2010). Functional imaging of hippocampal place cells at cellular resolution during virtual navigation. *Nat Neurosci*, 13(11):1433–40.
- Dragoi, G. and Buzsaki, G. (2006). Temporal encoding of place sequences by hippocampal cell assemblies. *Neuron*, 50(1):145–157.
- Dragoi, G. and Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature*, 469(7330):397–401.
- Ego-Stengel, V. and Wilson, M. A. (2010). Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. *Hippocampus*, 20(1):1–10.
- Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M., and Tanila, H. (1999). The hippocampus, memory, and place cells: Is it spatial memory or a memory space? *Neuron*, 23:209–226.
- Euston, D. R., Tatsuno, M., and McNaughton, B. L. (2007). Fast-forward playback of recent memory sequences in prefrontal cortex during sleep. *Science*, 318(5853):1147–1150.

- Ferbinteanu, J. and Shapiro, M. L. (2003). Prospective and retrospective memory coding in the hippocampus. *Neuron*, 40(6):1227–1239.
- Fortin, N. J., Agster, K. L., and Eichenbaum, H. B. (2002). Critical role of the hippocampus in memory for sequences of events. *Nature Neuroscience*, 5(5):458–462.
- Foster, D. J. and Wilson, M. A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, 440(7084):680–683.
- Foster, D. J. and Wilson, M. A. (2007). Hippocampal theta sequences. *Hippocampus*, 17(11):1093–1099.
- Frank, L. M., Brown, E. N., and Wilson, M. (2000). Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron*, 27:169–178.
- Fyhn, M., Molden, S., Witter, M. P., Moser, E. I., and Moser, M.-B. (2004). Spatial representation in the entorhinal cortex. *Science*, 305(5688):1258–1264.
- Gelbard-Sagiv, H., Mukamel, R., Harel, M., Malach, R., and Fried, I. (2008). Internally generated reactivation of single neurons in human hippocampus during free recall. *Science*, 322(5898):96–101.
- Girardeau, G., Benchenane, K., Wiener, S. I., Buzsaki, G., and Zugaro, M. B. (2009). Selective suppression of hippocampal ripples impairs spatial memory. *Nat Neurosci*, 12(10):1222–1223.
- Gothard, K. M., Skaggs, W. E., and McNaughton, B. L. (1996). Dynamics of mismatch correction in the hippocampal ensemble code for space: Interaction between path integration and environmental cues. *Journal of Neuroscience*, 16(24):8027–8040.
- Gupta, A. S., van der Meer, M. A. A., Touretzky, D. S., and Redish, A. D. Hippocampal theta sequences: dynamic modulation and the segmentation of spatial experience. *Submitted*.

- Gupta, A. S., van der Meer, M. A. A., Touretzky, D. S., and Redish, A. D. (2010). Hippocampal replay is not a simple function of experience. *Neuron*, 65(5):695–705.
- Guzowski, J. F., McNaughton, B. L., Barnes, C. A., and Worley, P. F. (1999). Environment-specific expression of the immediate-early gene *ARC* in hippocampal neuronal ensembles. *Nature Neuroscience*, 2(12):1120–1124.
- Hassabis, D., Kumaran, D., Vann, S. D., and Maguire, E. A. (2007). Patients with hippocampal amnesia cannot imagine new experiences. *PNAS*, 104:1726–1731.
- Hasselmo, M. E. (2005). The role of hippocampal regions CA3 and CA1 in matching entorhinal input with retrieval of associations between objects and context: Theoretical comment on Lee et al. (2005). *Behavioral Neuroscience*, 119(1):342–345.
- Hasselmo, M. E. (2006). The role of acetylcholine in learning and memory. *Curr Opin Neurobiol*, 16(6):710–5.
- Hasselmo, M. E. (2009). A model of episodic memory: mental time travel along encoded trajectories using grid cells. *Neurobiol Learn Mem*, 92(4):559–73.
- Hasselmo, M. E., Bodelon, C., and Wyble, B. P. (2002). A proposed function for hippocampal theta rhythm: separate phases of encoding and retrieval enhance reversal of prior learning. *Neural Comput*, 14(4):793–817.
- Hasselmo, M. E. and Bower, J. M. (1993). Acetylcholine and memory. *Trends in Neurosciences*, 16(6):218–222.
- Hasselmo, M. E. and Eichenbaum, H. (2005). Hippocampal mechanisms for the context-dependent retrieval of episodes. *Neural Netw*, 18(9):1172–90.
- Hill, A. J. and Best, P. J. (1981). Effects of deafness and blindness on the spatial correlates of hippocampal unit activity in the rat. *Experimental neurology*, 74:204–217.

- Hoffmann, K. L. and McNaughton, B. L. (2002). Coordinated Reactivation of Distributed Memory Traces in Primate Neocortex. *Science*, 297(5589):2070–2073.
- Huxter, J. R., Senior, T. J., Allen, K., and Csicsvari, J. (2008). Theta phase-specific codes for two-dimensional position, trajectory and heading in the hippocampus. *Nature Neuroscience*, 11:587–594.
- Hyman, J. M., Wyble, B. P., Goyal, V., Rossi, C. A., and Hasselmo, M. E. (2003). Stimulation in hippocampal region ca1 in behaving rats yields long-term potentiation when delivered to the peak of theta and long-term depression when delivered to the trough. *J Neurosci*, 23(37):11725–31.
- Ishizuka, N., Cowan, W. M., and Amaral, D. G. (1995). A quantitative analysis of the dendritic organization of pyramidal cells in the rat hippocampus. *J Comp Neurol*, 362(1):17–45.
- Ishizuka, N., Weber, J., and Amaral, D. G. (1990). Organization of intrahippocampal projections originating from ca3 pyramidal cells in the rat. *J Comp Neurol*, 295(4):580–623.
- Jackson, J. C., Bos, J. J., Donga, A. B., Lankelma, J. V., and Pennartz, C. M. A. (2009). Method of investigating the influence of hippocampal sharp wave-associated ripples on information processing in extrahippocampal structures. *Society for Neuroscience Abstracts*.
- Jackson, J. C., Johnson, A., and Redish, A. D. (2006). Hippocampal sharp waves and reactivation during awake states depend on repeated sequential experience. *Journal of Neuroscience*, 26:12415–12426.
- Janabi-Sharifi, F., Hayward, V., and Chen, C.-S. J. (2000). Discrete-time adaptive windowing for velocity estimation. *IEEE Transactions on Control Systems Technology*, 8(6):1003–1009.

- Jarosiewicz, B. and Skaggs, W. E. (2004). Hippocampal place cells are not controlled by visual input during the small irregular activity state in the rat. *J Neurosci*, 24(21):5070–7.
- Jensen, O. and Lisman, J. E. (1996). Hippocampal CA3 region predicts memory sequences: accounting for the phase precession of place cells. *Learning and Memory*, 3(2-3):279–287.
- Jensen, O. and Lisman, J. E. (2000). Position reconstruction from an ensemble of hippocampal place cells: contribution of theta phase encoding. *Journal of Neurophysiology*, 83(5):2602–2609.
- Ji, D. and Wilson, M. A. (2007). Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nat Neurosci*, 10(1):100–107.
- Ji, D. and Wilson, M. A. (2008). Firing Rate Dynamics in the Hippocampus Induced by Trajectory Learning. *J. Neurosci.*, 28(18):4679–4689.
- Johnson, A., Jackson, J., and Redish, A. D. (2008). Measuring distributed properties of neural representations beyond the decoding of local variables — implications for cognition. In Hölscher, C. and Munk, M. H. J., editors, *Mechanisms of information processing in the Brain: Encoding of information in neural populations and networks*, pages 95–119. Cambridge University Press.
- Johnson, A. and Redish, A. D. (2007). Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *Journal of Neuroscience*, 27(45):12176–12189.
- Jung, M. W., Wiener, S. I., and McNaughton, B. L. (1994). Comparison of spatial firing characteristics of the dorsal and ventral hippocampus of the rat. *Journal of Neuroscience*, 14(12):7347–7356.
- Karlsson, M. P. and Frank, L. M. (2009). Awake replay of remote experiences in the hippocampus. *Nat Neurosci*, 12(7):913–918.

- Kjelstrup, K. B., Solstad, T., Brun, V. H., Hafting, T., Leutgeb, S., Witter, M. P., Moser, E. I., and Moser, M.-B. (2008). Finite scale of spatial representation in the hippocampus. *Science*, 321(5885):140–143.
- Kramis, R., Vanderwolf, C. H., and Bland, B. H. (1975). Two types of hippocampal rhythmical slow activity in both the rabbit and the rat: relations to behavior and effects of atropine, diethyl ether, urethane, and pentobarbital. *Exp Neurol*, 49(1 Pt 1):58–85.
- Kubie, J. L. and Ranck, Jr., J. B. (1983). Sensory-behavioral correlates in individual hippocampus neurons in three situations: Space and context. In Seifert, W., editor, *Neurobiology of the Hippocampus*, pages 433–447. Academic Press, New York.
- Kudrimoti, H. S., Barnes, C. A., and McNaughton, B. L. (1999). Reactivation of hippocampal cell assemblies: Effects of behavioral state, experience, and EEG dynamics. *Journal of Neuroscience*, 19(10):4090–4101.
- Lee, A. K. and Wilson, M. A. (2002). Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron*, 36:1183–1194.
- Levy, W. B. (1996). A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus*, 6(6):579–591.
- Levy, W. B. and Steward, O. (1983). Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Neuroscience*, 8(4):791–7.
- Lisman, J. (2005). The theta/gamma discrete phase code occurring during the hippocampal phase precession may be a more general brain coding scheme. *Hippocampus*, 15(7):913–922.
- Lisman, J. and Idiart, M. A. (1995). Storage of 7 +/- 2 short-term memories in oscillatory subcycles. *Science*, 267(5203):1512–1515.

- Lisman, J. and Redish, A. D. (2009). Prediction, sequences and the hippocampus. *Philosophical Transactions of the Royal Society B Biological Sciences*, 364:1193–1201.
- Lorento do Nó, R. (1934). Studies on the structure of the cerebral cortex II. Continuation of the study of the ammonic system. *J. Psychol. Neurol.*, 46:113–177.
- Lubenov, E. V. and Siapas, A. G. (2009). Hippocampal theta oscillations are travelling waves. *Nature*, 459(7246):534–539.
- Markus, E. J., Barnes, C. A., McNaughton, B. L., Gladden, V. L., and Skaggs, W. E. (1994). Spatial information content and reliability of hippocampal CA1 neurons: Effects of visual input. *Hippocampus*, 4:410–421.
- Marr, D. (1971). Simple memory: A theory of archicortex. *Philosophical Transactions of the Royal Society of London*, 262(841):23–81.
- Masimore, B., Schmitzer-Torbert, N. C., Kakalios, J., and Redish, A. D. (2005). Transient striatal [gamma] local field potentials signal movement initiation in rats. *NeuroReport*, 16(18):2021–2024.
- Maurer, A. P., Burke, S. N., Lipa, P., Skaggs, W. E., and Barnes, C. A. (2011). Greater running speeds result in altered hippocampal phase sequence dynamics. *Hippocampus*.
- Maurer, A. P. and McNaughton, B. L. (2007). Network and intrinsic cellular mechanisms underlying theta phase precession of hippocampal neurons. *Trends in Neurosciences*, 30(7):325–333.
- Maviel, T., Durkin, T. P., Menzaghi, F., and Bontempi, B. (2004). Sites of neocortical reorganization critical for remote spatial memory. *Science*, 305(5680):96–99.
- McClelland, J. L., McNaughton, B. L., and O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes

- and failures of connectionist models of learning and memory. *Psychological Review*, 102(3):419–457.
- McNaughton, B. L., Barnes, C. A., and O'Keefe, J. (1983). The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Experimental Brain Research*, 52:41–49.
- Mehta, M. R., Barnes, C. A., and McNaughton, B. L. (1997). Experience-dependent, asymmetric expansion of hippocampal place fields. *Proceedings of the National Academy of Sciences, USA*, 94:8918–8921.
- Mehta, M. R. and McNaughton, B. L. (1996). Rapid changes in hippocampal population code during behavior: A case for Hebbian learning *in vivo*. Presented at CNS*96, (the fifth annual *Computation in Neural Systems* meeting) .
- Meunzinger, K. F. (1938). Vicarious trial and error at a point of choice. I. a general survey of its relation to learning efficiency. *J. Genet Psychol*, 53:75–86.
- Miller, G. (1956). The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review*, 63(2):81–97.
- Milner, B., Corkin, S., and Teuber, H. (1968). Further analysis of the hippocampal amnesia syndrome: 14-year follow-up study of H. M. *Neuropsychologia*, 6:215–234.
- Morris, R. G. M., Garrud, P., Rawlins, J. N. P., and O'Keefe, J. (1982). Place navigation impaired in rats with hippocampal lesions. *Nature*, 297:681–683.
- Mulder, A. B., Tabuchi, E., and Wiener, S. I. (2004). Neurons in hippocampal afferent zones of rat striatum parse routes into multi-patch segments during maze navigation. *European Journal of Neuroscience*, 19:1923–1932.
- Muller, R. U. and Kubie, J. L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *Journal of Neuroscience*, 7:1951–1968.

- Muller, R. U. and Kubie, J. L. (1989). The firing of hippocampal place cells predicts the future position of freely moving rats. *Journal of Neuroscience*, 9(12):4101–4110.
- Muller, R. U., Kubie, J. L., Bostock, E. M., Taube, J. S., and Quirk, G. J. (1991). Spatial firing correlates of neurons in the hippocampal formation of freely moving rats. In Paillard, J., editor, *Brain and Space*, pages 296–333. Oxford University Press, New York.
- Nadasdy, Z., Hirase, H., Czurko, A., Csicsvari, J., and Buzsaki, G. (1999). Replay and time compression of recurring spike sequences in the hippocampus. *J Neurosci*, 19(21):9497–9507.
- Nadel, L., O'Keefe, J., and Black, A. (1975). Slam on the brakes: A critique of Altman, Brunner, and Bayer's response-inhibition model of hippocampal function. *Behavioral Biology*, 14:151–162.
- Nakashiba, T., Buhl, D. L., McHugh, T. J., and Tonegawa, S. (2009). Hippocampal *ca3* output is crucial for ripple-associated reactivation and consolidation of memory. *Neuron*, 62(6):781–787.
- Newell, A. (1990). *Unified Theories of Cognition*. Harvard.
- Newell, A. and Simon, H. (1972). *Human Problem Solving*. Prentice-Hall, Englewood Cliffs, NJ.
- Nitz, D. A. (2006). Tracking route progression in the posterior parietal cortex. *Neuron*, 49(5):747–56.
- O'Keefe, J. (1976). Place units in the hippocampus of the freely moving rat. *Experimental Neurology*, 51:78–109.
- O'Keefe, J. and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely moving rat. *Brain Research*, 34:171–175.

- O'Keefe, J. and Nadel, L. (1978). *The Hippocampus as a Cognitive Map*. Clarendon Press, Oxford.
- O'Keefe, J. and Recce, M. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, 3:317–330.
- Olton, D. S., Branch, M., and Best, P. J. (1978). Spatial correlates of hippocampal unit activity. *Experimental Neurology*, 58:387–409.
- O'Neill, J., Senior, T., and Csicsvari, J. (2006). Place-selective firing of ca1 pyramidal cells during sharp wave/ripple network patterns in exploratory behavior. *Neuron*, 49:143–155.
- O'Reilly, R. C. and McClelland, J. L. (1994). Hippocampal conjunctive encoding, storage, and recall: Avoiding a trade-off. *Hippocampus*, 4(6):661–682.
- Packard, M. G. and McGaugh, J. L. (1996). Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. *Neurobiology of Learning and Memory*, 65:65–72.
- Pomerleau, D. A. (1991). Efficient training of artificial neural networks for autonomous navigation. *Neural Computation*, 3(1):88–97.
- Qin, Y. L., McNaughton, B. L., Skaggs, W. E., and Barnes, C. A. (1997). Memory reprocessing in corticocortical and hippocampocortical neuronal ensembles. *Philos Trans R Soc Lond B Biol Sci*, 352(1360):1525–33.
- Quirk, G. J., Muller, R. U., and Kubie, J. L. (1990). The firing of hippocampal place cells in the dark depends on the rat's recent experience. *Journal of Neuroscience*, 10(6):2008–2017.
- Rapp, P. R. and Gallagher, M. (1996). Preserved neuron number in the hippocampus of aged rats with spatial learning deficits. *Proc Natl Acad Sci U S A*, 93(18):9926–9930.

- Redish, A. D. (1999). *Beyond the Cognitive Map: From Place Cells to Episodic Memory*. MIT Press, Cambridge MA.
- Redish, A. D., Battaglia, F. O., Chawla, M. K., Ekstrom, A. D., Gerard, J. L., Lipa, P., Rosenzweig, E. S., Worley, P. F., Guzowski, J. F., McNaughton, B. L., and Barnes, C. A. (2001). Independence of firing correlates of anatomically proximate hippocampal pyramidal cells. *Journal of Neuroscience*, 21(RC134):1–6.
- Redish, A. D. and Touretzky, D. S. (1998). The role of the hippocampus in solving the Morris water maze. *Neural Computation*, 10(1):73–111.
- Risold, P. Y. and Swanson, L. W. (1996). Structural evidence for functional domains in the rat hippocampus. *Science*, 272(5267):1484–1486.
- Rossier, J., Kaminsky, Y., Schenk, F., and Bures, J. (2000). The place preference task: A new tool for studying the relation between behavior and place cell activity in rats. *Behavioral Neuroscience*, 114(2):273–284.
- Royer, S., Sirota, A., Patel, J., and Buzsaki, G. (2010). Distinct representations and theta dynamics in dorsal and ventral hippocampus. *J Neurosci*, 30(5):1777–1787.
- Samsonovich, A. V. and Ascoli, G. A. (2005). A simple neural network model of the hippocampus suggesting its pathfinding role in episodic memory retrieval. *Learn. Mem.*, 12(2):193–208.
- Samsonovich, A. V. and McNaughton, B. L. (1997). Path integration and cognitive mapping in a continuous attractor neural network model. *Journal of Neuroscience*, 17(15):5900–5920.
- Save, E., Cressant, A., Thinus-Blanc, C., and Poucet, B. (1998). Spatial firing of hippocampal place cells in blind rats. *Journal of Neuroscience*, 18(5):1818–1826.

- Schmitzer-Torbert, N. and Redish, A. D. (2004). Neuronal activity in the rodent dorsal striatum in sequential navigation: separation of spatial and reward responses on the multiple t task. *J Neurophysiol*, 91(5):2259–2272.
- Schmitzer-Torbert, N. C., Jackson, J., Henze, D., Harris, K., and Redish, A. D. (2005). Quantitative measures of cluster quality for use in extracellular recordings. *Neuroscience*, 131(1):1–11.
- Scoville, W. B. and Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of Neurology, Neurosurgery, and Psychiatry*, 20:11–21.
- Skaggs, W. E. and McNaughton, B. L. (1996). Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science*, 271:1870–1873.
- Skaggs, W. E., McNaughton, B. L., Wilson, M. A., and Barnes, C. A. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*, 6(2):149–173.
- Squire, L. R. (1987). *Memory and Brain*. Oxford University Press, New York.
- Squire, L. R. (1992). Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychology Review*, 99(2):195–231.
- Squire, L. R., Cohen, N. J., and Nadel, L. (1984). The medial temporal region and memory consolidation: A new hypothesis. In Weingartner, H. and Parker, E. S., editors, *Memory Consolidation: Psychobiology of Cognition*. Erlbaum, Hillsdale NJ.
- Squire, L. R. and Zola-Morgan, S. M. (1991). The medial temporal lobe memory system. *Science*, 253:1380–1386.
- Sutherland, G. R. and McNaughton, B. L. (2000). Memory trace reactivation in hippocampal and neocortical neuronal ensembles. *Current Opinion in Neurobiology*, 10(2):180–6.

- Sutton, R. S. and Barto, A. G. (1998). *Reinforcement Learning: An introduction*. MIT Press, Cambridge MA.
- Swanson, L. W., Sawchenko, P. E., and Cowan, W. M. (1980). Evidence that the commissural, associational and septal projections of the regio inferior of the hippocampus arise from the same neurons. *Brain Res*, 197(1):207–212.
- Swanson, L. W., Sawchenko, P. E., and Cowan, W. M. (1981). Evidence for collateral projections by neurons in ammon's horn, the dentate gyrus, and the subiculum: a multiple retrograde labeling study in the rat. *J Neurosci*, 1(5):548–559.
- Swanson, L. W., Wyss, J. M., and Cowan, W. M. (1978). An autoradiographic study of the organization of intrahippocampal association pathways in the rat. *J Comp Neurol*, 181(4):681–715.
- Tolman, E. C. (1939). Prediction of vicarious trial and error by means of the schematic sowbug. *Psychological Review*, 46:318–336.
- Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, 55:189–208.
- Tolman, E. C., Ritchie, B. F., and Kalish, D. (1946). Studies in spatial learning. I. Orientation and the short-cut. *Journal of Experimental Psychology*, 36:13–24.
- Tsodyks, M. V., Skaggs, W. E., Sejnowski, T. J., and McNaughton, B. L. (1996). Population dynamics and theta rhythm phase precession of hippocampal place cell firing: A spiking neuron model. *Hippocampus*, 6(3):271–280.
- Tuckwell, H. C. (1988). *Introduction to Theoretical Neurobiology: Volume 2, Nonlinear and Stochastic Theories*. Cambridge University Press.
- Tulving, E. (2002). Episodic memory: From mind to brain. *Annual Review of Psychology*, 53(1):1–25.

- van der Meer, M. A. A. and Redish, A. D. (2009). Covert expectation-of-reward in rat ventral striatum at decision points. *Frontiers in Integrative Neuroscience*, 3(1):1–15.
- van Strien, N. M., Cappaert, N. L. M., and Witter, M. P. (2009). The anatomy of memory: an interactive overview of the parahippocampal-hippocampal network. *Nat Rev Neurosci*, 10(4):272–282.
- Vanderwolf, C. H. (1969). Hippocampal electrical activity and voluntary movement in the rat. *Electroencephalogr Clin Neurophysiol*, 26(4):407–418.
- West, M. J., Slomianka, L., and Gundersen, H. J. (1991). Unbiased stereological estimation of the total number of neurons in the subdivisions of the rat hippocampus using the optical fractionator. *Anat Rec*, 231(4):482–497.
- Wilson, M. A. and McNaughton, B. L. (1993). Dynamics of the hippocampal ensemble code for space. *Science*, 261:1055–1058.
- Wilson, M. A. and McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, 265:676–679.
- Wittner, L., Henze, D. A., Zaborszky, L., and Buzsaki, G. (2007). Three-dimensional reconstruction of the axon arbor of a ca3 pyramidal cell recorded and filled in vivo. *Brain Struct Funct*, 212(1):75–83.
- Wood, E. R., Dudchenko, P. A., Robitsek, R. J., and Eichenbaum, H. (2000). Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron*, 27:623–633.
- Yamaguchi, Y., Sato, N., Wagatsuma, H., Wu, Z., Molter, C., and Aota, Y. (2007). A unified view of theta-phase coding in the entorhinal-hippocampal system. *Curr Opin Neurobiol*, 17(2):197–204.

- Yartsev, M. M. (2008). Dissociating the effects of past and future on neural encoding of sequences in the hippocampus. *J Neurosci*, 28(34):8383–4.
- Ylinen, A., Bragin, A., Nadasdy, Z., Jando, G., Szabo, I., Sik, A., and Buzsáki, G. (1995). Sharp wave-associated high-frequency oscillation (200 Hz) in the intact hippocampus: Network and intracellular mechanisms. *Journal of Neuroscience*, 15(1):30–46.
- Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. *Journal of Neuroscience*, 16(6):2112–2126.
- Zhang, K., Ginzburg, I., McNaughton, B. L., and Sejnowski, T. J. (1998). Interpreting neuronal population activity by reconstruction: Unified framework with application to hippocampal place cells. *Journal of Neurophysiology*, 79:1017–1044.
- Zinyuk, L., Kubik, S., Kaminsky, Y., Fenton, A. A., and Bures, J. (2000). Understanding hippocampal activity by using purposeful behavior: place navigation induces place cell discharge in both task-relevant and task-irrelevant spatial reference frames. *Proceedings of the National Academy of Sciences, USA*, 97(7):3771–3776.