

HIPPOCAMPAL NEURAL FIRING DYNAMICS DURING SLEEP

BY

ANDRES D. GROSMARK

A Dissertation submitted to the
Graduate School-Newark
Rutgers, The State University of New Jersey

In partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Behavioral and Neural Sciences

Written under the direction of

Dr. György Buzsáki

And approved by

Newark, New Jersey

May, 2014

© 2014

Andres Daniel Grosmark

ALL RIGHTS RESERVED

ABSTRACT OF THE DISSERTATION

Hippocampal Neural Firing Dynamics During Sleep

By ANDRES D. GROSM0ARK

Dissertation Director:

Dr. György Buzsáki

The focus of my thesis was the assessment of the presence and dynamics of non-specific, putatively homeostatic, as well as specific memory-related, modifications in the structure of hippocampal firing patterns during sleep. As a step toward understanding sleep function, I developed an appropriate, open-source Matlab-based application for the visualization, annotation and detailed sleep scoring of long-time scale electrophysiological data relevant to brain state dynamics.

In my study of putatively homeostatic changes in excitability I found that the overall firing rates of hippocampal CA1 neurons decreased across sleep concurrent with an increased recruitment of neuronal spiking to brief ‘ripple’ episodes, resulting in a net increase in neural synchrony. Unexpectedly, within non-REM episodes overall firing rates were found to gradually increase together with a decrease in recruitment of spiking to ripples. The rate increase within non-REM rate episodes was counteracted by a larger rapid decrease of discharge frequency during the interleaving REM episodes. Both the decreasing firing rates and the increasing synchrony during the course of sleep were correlated with the power of theta activity during REM episodes. These findings suggest a prominent role of REM sleep in sleep-related neuronal plasticity.

Lastly, in order to gauge the interaction between non-specific and memory-specific contributions to sleep-related firing dynamics, I recorded hours of sleep in rats before and after exposure to a completely novel maze environment. While both replay and ‘pre-play’ were observed, both were found to be dominated by non-local (different silicon-probe shank) interactions. However, while replay was observed in both pair-wise and higher-order interactions, pre-play was surprisingly specifically restricted to higher-order sequential interactions. This analysis also included the assessment of several traditional as well as novel methods for measuring replay, leading to significant methodological insights into their sensitivity to the established non-stationary nature of excitability in sleep, and to several suggestions for future work.

ACKNOWLEDGEMENTS

I would like to thank the past current members of the laboratory for making it a truly exciting and collaborative place in which to work, investigate and refine ideas - and that truly includes everyone that works in this lab. This research would not be possible without them. I am particularly and immensely grateful to Dr. Eva Pastalkova, Dr. Kenji Mizuseki, and Dr. Kamran Diba for being so generous with their advice, with their support, and with their data. I would like to thank Dr. Adrien Payrache, Dr. John Long, Dr. Mariano Belluscio, and Dr. Shigeyoshi Fujisawa for their sapient and timely advice. I am also deeply grateful to Dr. Marie Vandecasteele, Dr. John Long II, Dr. Erik Schomburg, Dr. Lisa Roux, Dr. Gabrielle Girardeu, and Dr. Josh Callahan each of who's insight is only matched by the pleasure of their company. I also have to thank my parents for everything. I am gratefully indebted to my bright and charismatic girlfriend, Sebnem Tuncdemir, for sticking by me through all these year, and incredibly, through the last month. I would like to thank the Rutgers University Neuroscience Department for their generosity with me, and for educating me as a neuroscientist. I am very grateful to my advisor Dr. György Buzsáki for being an inspiration to me through the depth of his insight, the buoyancy of his humor, and most of all his utterly sincere and tireless enthusiasm for the project of neuroscience. I would like to thank Dr. Bart Krekelberg, for serving as the chair of my committee and also for hosting me at his lab at when I first arrived at Newark. Finally, I am deeply grateful to the members of my committee, Dr. Bart Krekelberg, Dr. Denis Paré, Dr. Francesco Battaglia, for their feedback and support, and for their extreme generosity with the timing of this dissertation.

Table of Contents

Page 1.....	Chapter 1: Introduction
Page 2.....	non-REM and REM Sleep Physiology
Page 3.....	non-REM Sleep Physiology
Page 5.....	REM Sleep Physiology
Page 9.....	Functions of sleep
Page 9.....	Sleep and Synaptic Homeostasis
Page 15.....	Memory Consolidation
Page 19.....	Replay
Page 21.....	Types of replay
Page 22.....	Pair-Wise Rate Replay
Page 23.....	Event-Based Rate Replay
Page 24.....	Pair-Wise Sequence Replay
Page 25.....	Event-Based Sequence Replay
Page 29.....	Bayesian Decoding of Replay
Page 32.....	Pre-Play
Page 34.....	Goal of the Thesis
Chapter 2: Methods	
Page 35.....	TheStateEditor: Behavioral State Scoring and Data Visualization and Annotation Software
Page 36.....	REM Sleep Reorganizes Hippocampal Excitability
Page 36.....	Animals, Surgery and Data Collection
Page 37.....	Identification of Non-REM and REM Episodes
Page 38.....	Unit Clustering and Cell Classification
Page 40.....	Firing Pattern and LFP Activity Changes Across Sleep, Within non-REM and Within REM
Page 42.....	REM Sleep LFPs Correlate with Firing Rate Decrease
Page 42.....	Spike-Weighted Spectra (Sp.W.S)
Page 44.....	Methods for Pre-Play and Replay Analysis
Page 44.....	Animals, Surgery and Data Collection
Page 45.....	Behavioral Procedures and Sleep Scoring
Page 47.....	Place Field Analysis
Page 49.....	Population Activity Events
Page 48.....	Place Cell Pair Activity and Timing Co-Modulation
Page 51.....	Pre-Play and Replay of Sequential Spiking Activity: Rank-order correlations
Page 53.....	Pre-Play and Replay of Sequential Spiking Activity: Paired-Latency Method
Chapter 3: Results	
Page 59.....	TheStateEditor: Behavioral State Scoring and Data Visualization and Annotation Software
Page 59.....	Loading and Pre-processing files

Page 61.....	TheStateEditor: Data Visualization and Navigation
Page 63.....	TheStateEditor: Behavioral State Scoring and Event Selection
Page 66.....	REM Sleep Reorganizes Hippocampal Excitability
Page 66.....	Excitability changes across sleep
Page 68.....	Excitability changes within non-REM and REM episodes
Page 69.....	LFP spectral changes across sleep and within non-REM and REM episodes
Page 69.....	Relationship between non-REM and REM sleep
Page 71.....	Pre-Play and Replay
Page 71.....	Pre-Play and Replay: Population Activity Events
Page 72.....	Place Cell Pair Activity and Timing Co-Modulation
Page 74.....	Pre-Play and Replay of Sequential Spiking Activity: Rank-order correlations
Page 77.....	Pre-Play and Replay of Sequential Spiking Activity: Paired-Latency Method
Page 82.....	Pre-Play and Replay: Bayesian Decoding

Chapter 4: Discussion

Page 87.....	TheStateEditor: Behavioral State Scoring and Data Visualization and Annotation Software
Page 89.....	REM Sleep Reorganizes Hippocampal Excitability
Page 89.....	Models of Sleep Function
Page 90.....	Potential Confounds to a REM Dependent Effect
Page 92.....	REM Represents a Unique Neuromodulatory Regime
Page 92.....	Hippocampal Excitability and Synchrony Are Decoupled During Sleep
Page 93.....	Future Directions
Page 94.....	Hippocampal Pre-Play and Replay of Novel Experiences
Page 96.....	Excitability Changes In Population Activity Events
Page 97.....	Replay Is Robustly Present in Place Cell Pair Activity and Timing Co-Modulation
Page 98.....	Absence of Pre-Play and Replay Using the Rank-order Correlation Method
Page 101.....	Replay and Pre-Play Are Observed in the Non-Local Higher Order Structure
Page 102.....	Changes in Participation Rate Lead to an Increase in the Signal to Noise Ratio During the Post Epoch
Page 104.....	Much of the Variability of the Observed Pre-Play and Replay Signal is Attributable to the Noise Distribution
Page 105.....	Physiological Sources of ‘Noise’
Page 106.....	Experimental Sources of Noise
Page 108.....	Sources of Inter-Animal Variability

Page 110..... The Noisy Nature of Replay and Pre-Play May Explain
the Lack of Within Pre and Post and Within Non-
REM Effects

Page 110..... Partial Dissociation Between Pre to Post Changes in
Bayesian Decoding Quality and Bayesian
Decoded Sequence

Page 115..... Bayesian Decoding Quality Correlates With Bayesian
Decoding Strength in Both the Pre and Post
Epochs and Both are Correlated With Firing Rate

Page 116..... Non-Local Firing Rate, Sequence and Place Coding
Structure

Page 120..... **Conclusion**

List of Figures

Page 122..... **Figure 1.1**

Page 123..... **Figure 1.2**

Page 124..... **Figure 1.4**

Page 125..... **Figure 2.1**

Page 126..... **Figure 2.2**

Page 127..... **Figure 2.3**

Page 128..... **Figure 2.4**

Page 129..... **Figure 3.1**

Page 130..... **Figure 3.2**

Page 131..... **Figure 3.3**

Page 132..... **Figure 3.4**

Page 133..... **Figure 3.5**

Page 134..... **Figure 3.6**

Page 135..... **Figure 3.7**

Page 136..... **Figure 3.8**

Page 137..... **Figure 4.1**

Page 138..... **Figure 4.2**

Page 139..... **Figure 4.3**

Page 140..... **Figure 4.4**

Page 141..... **Figure 4.5**

Page 142..... **Figure 4.6**

Page 143..... **Figure 4.7**

Page 144..... **Figure 4.8**

Page 145..... **Figure 4.9**

Page 146..... **Figure 4.10**

Page 147..... **Figure 4.11**

Page 148..... **Figure 4.12**

Page 149..... **Figure 4.13**

Page 150..... **Figure 4.14**

Page 151..... **Figure 4.15**

Page 152..... **Figure 4.16**

Page 153.....**Figure 4.17**
Page 154.....**Figure 4.18**

List of Tables

Page 155.....**Table 1**
Page 156.....**Table 2**

Page 157.....**Glossary of Terms**

Page 160.....**References**

Page 190.....**Curriculum Vitae**

Introduction

The behavioral patterns of nearly all animals are characterized by alternations between the awake and sleeping states. While most of the overt behaviors canonically associated with survival and reproduction, such as feeding, mating and exploration, occur in the awake state, many animals spend more than half their lives asleep (Siegel, 2005). Moreover, while sleep patterns are often adjusted in response to external pressures, the lack of sleep produces cognitive deficits in the short term (Killgore, 2010) and can prove fatal if prolonged (Everson et al., 1989). The necessity of sleep for survival is also highlighted by the fact that sleep is a period of decreased motor activity and vigilance, thus increasing the risk of predation. Furthermore, animals whose evolutionary niche precludes canonical sleep have evolved mechanisms for preserving this behavioral state. Thus for instance cetaceans such as whales and dolphins who must swim continuously in order not to drown have evolved unihemispheric sleep - a state that allows each hemisphere of the brain to sleep while the other hemisphere remains online, ensuring that swimming is not interrupted (Mukhametov, 1987; Rattenborg et al., 2000). In placental mammals and many birds sleep is further divided into alternating REM and non-REM sleep states. Notably, the presence of REM and non-REM states in both mammals and birds is thought to have evolved independently in each of these lineages, suggesting that the differentiation of sleep into these two states is strongly beneficial for survival (Lee Kavanau, 2002; Siegel, 1995). However, despite the overwhelming amount of evidence suggesting a fundamental role for sleep in general, and REM and non-REM sleep in particular, in survival, the functions of sleep remains little understood. Here I examine some of the basic characteristics of sleep as well as theories regarding its

functional outcomes, with particular emphasis on the hippocampus, a structure extensively implicated in memory and memory consolidation. This chapter will motivate the specific aims of my research in advancing our understanding of the characteristics and functional outcomes of mammalian sleep.

non-REM and REM Sleep Physiology

Mammalian arousal states, and sleep in particular are regulated through the ascending influence of subcortical neuromodulatory systems. The effect of these diverse nuclei are mediated by signals including glutamate, GABA, monoamines, acetylcholine, as well as various neuropeptides (see Jones, 2003 for review). However, while the mechanisms for initiation and maintenance of sleep states, and in particular REM and non-REM sleep, are subcortical, these sleep states produce profoundly unique patterns of activity in forebrain structures. From the dynamic point of view, global sleep can be conceived of as a dampening oscillator in which the opposing forces are represented by non-REM and REM states. In humans sleep begins with a large amplitude cycle, with deep sleep (slow wave sleep) terminated by the first short REM episode. In the course of sleep, non-REM episodes become shorter and more shallow as they lose to the increasing duration of REM episodes (Buzsáki, 2006). In short, sleep is a dynamically organized event with two opposing components.

non-REM Sleep Physiology

While the investigation on the mechanisms for the transition from waking to non-REM sleep is ongoing, a canonical 'wake-sleep' neural switch has been extensively

studied. In this switch GABAergic and glycinergic neurons in the medial and particular ventrolateral preoptic nuclei of the hypothalamus (MnPO and VLPO, respectively) reciprocally inhibit most major subcortical neuromodulatory systems. These include the serotonergic dorsal raphe (DR), the histaminergic tuberomammillary nucleus (TMN), the noradrenergic locus coeruleus (LC), the dopaminergic ventral periaqueductal grey matter (vPAG), the acetylcholinergic lateral dorsal tagmental area (LDT) and pedunculopontine tagmental (PPT) nucleus, as well as several ascending glutamatergic nuclei (Saper et al., 2010). In turn, the suppression of these ascending modulatory systems removes the bulk of the excitatory drive from the basal forebrain, the main acetylcholinergic input to the thalamus and cortical mantle (Zaborszky and Duque, 2003). The decrease in neuromodulatory concentrations, and particularly the near absence of acetylcholine, increases the activity of several mostly inhibitory membrane currents such as the muscarine sensitive current (I_M), the potassium 'leak' current ($I_{K,L}$), as well as the after-hyperpolarization current (I_{AHP} , McCormick, 1993). In turn, this increase in intrinsic inhibitory drive, coupled with the related decrease in excitatory drive from neighboring neurons, results in network dynamics dominated by intrinsic pacemaking currents such as the hyperpolarization current (I_H) and the low-threshold Ca^{2+} current (I_T , McCormick, 1992; Steriade, 2004; Steriade et al., 1993). This un-masked intrinsic oscillatory drive, coupled with complex synaptic interactions (many of which are also oscillatory), results in the synchronization of much of the thalamocortical network into discrete periods of activity and inactivity. This synchrony is reflected in the high-amplitude, low-frequency activity, known as slow oscillations, observed in cortical EEG during non-REM sleep (Riedner et al., 2011). Importantly, this unmasked oscillatory drive causes

thalamocortical neurons to switch from their regular spiking mode in which spiking events reflect sensory and cortical input, to their intrinsic ‘bursting’ mode in which spiking dynamics are dominated by intrinsic oscillations (Steriade et al., 1993). This has the effect of isolating cortical activity from sensory input during sleep, resulting in the familiar loss of consciousness.

In the cortex non-REM periodic activity is organized into two, perhaps three main oscillatory types: the spindle (12-20Hz) oscillation of thalamic origin (Steriade, 1995), as well as the delta (1-4 Hz) and slow (0.1-1Hz) oscillations which are believed to be the result of intrinsic and synaptic mechanisms in both the cortex and the thalamus (Crunelli and Hughes, 2010; Steriade and Amzica, 1998; Steriade et al., 1993, 2001). In the cortex the slow and delta oscillations observed in the EEG reflect the presence of alternating periods of synchronized activity (‘UP’ periods) or hyperpolarization (‘DOWN’ periods) each lasting several hundred milliseconds. Spindles on the other hand, are discrete events which can occur during UP periods and last approximately 1 to 3 seconds in which much of cortex receives strong synchronizing excitatory drive from the thalamus (Steriade et al., 1993, 2001).

In the hippocampus, the suppression of ascending modulatory input likewise results in synchronized, periodic network dynamics. However, consistent with the relative sparseness of thalamic input to the hippocampus, spindles are not generally thought to be a major driver of hippocampal activity during non-REM sleep (though see Sullivan et al., 2014). Instead, hippocampal activity during non-REM is largely organized into brief (~120ms) ripple events with a characteristic LFP frequency of ~150 to 250Hz and which are generated by the recurrent CA3 layer network (Buzsáki, 1986; Buzsáki et

al., 1983, 1992; Csicsvari et al., 1999a; Ylinen et al., 1995). Notably, the low frequency oscillations observed in non-REM are thought to be coherent across much of the cortex, though the extent, dynamics and mechanism of this synchronization remains a subject of debate (Destexhe et al., 1999; Huber et al., 2004; Steriade, 2003; Vyazovskiy et al., 2011a). Consistent with the partly global nature of low frequency oscillations, it has been reported that ripples events are partially predicted by cortical spindles and UP states (Battaglia et al., 2004; Benchenane et al., 2010; Siapas and Wilson, 1998; Siapas et al., 2005; Sirota and Buzsáki, 2005; Wierzynski et al., 2009). However, this coordination, though clearly present, is rather weak. The relatively weak coordination between oscillations in the hippocampus and cerebral cortex during non-REM sleep is consistent with separate, though interacting, mechanisms for the generation of these oscillations.

REM Sleep Physiology

Complementing the diversity of neural dynamics observed in non-REM sleep, REM sleep represents an additional and fundamentally different type of sleep. Notably, it is REM sleep that is thought to be the state most conducive to those periods of bizarre ideation commonly referred to as dreams (Hobson, 1990; Hobson et al., 2000; Siegel, 2011). While varying between species, the amount of sleep spent in REM is typically shorter than that spent in non-REM, making up approximately 24% of human sleep and 19% of rodent (specifically rat) sleep (Lesku et al., 2008). However, in both birds and mammals REM episodes invariably occur only immediately subsequent to a non-REM episode (Siegel, 1995). Paralleling the state transition from the awake to the non-REM

state, REM sleep is established and supported by subcortical circuitry which reciprocally inhibits non-REM sleep promoting cells (Saper et al., 2010). The core REM 'flip-flop' switch consists of the REM promoting sublaterodorsal nucleus (SLD) of the pontine tegmentum and a population of REM inhibiting GABAergic cells in the ventrolateral periaqueductal grey (vlPAG). As non-REM sleep progresses the discharge of noradrenergic neurons of the LC, serotonergic neurons of the DR, as well as the discharge of VLPO neurons containing the neuropeptide orexin gradually decrease (Hassani et al., 2009; Luppí et al., 2012). This decrease has the effect of inhibiting the REM inhibiting vlPAG neurons, thus dis-inhibiting the SLD. When the SLD becomes sufficiently dis-inhibited a rapid state transition occurs in which GABAergic SLD cells inhibit the vlPAG, as well as the LC and DR, and resulting in the near complete elimination of both noradrenergic and serotonergic tone, in turn further disinhibiting the SLD (Gervasoni et al., 1998; Saper et al., 2010). Concurrently, disinhibited REM promoting glutamatergic neurons in the precoeruleus region (PC) and parabrachial nucleus (PB) of the pons strongly activate cholinergic neurons of the basal forebrain as well as other forebrain targets (Saito et al., 1977; Vazquez and Baghdoyan, 2001). Excitation from these nuclei as well as from the cholinergic PPT is thought to be organized into waves first observed in the pons, lateral geniculate nucleus and occipital cortex (PGO waves) of cats (Jouvet, 1967). In rodent models these waves are simply referred to as P-waves, reflecting their origin in the pons. The activation of the basal forebrain results in the strong efflux of acetylcholine, leading to low-amplitude, high-frequency cortical EEG patterns remarkably similar to those observed during active waking (Aserinsky and Kleitman, 1953; Llinás and Paré, 1991). However, unlike waking

or non-REM sleep, locomotion is actively inhibited during REM sleep by the descending projections of glutamatergic cells in the SLD which activate medullar and spinal inhibitory interneurons (Vetrivelan et al., 2009). These interneurons in turn actively inhibit α -motor neurons leading to the loss of muscle tone and paralysis unique to REM (Chase, 2008; Kohyama, 2000). Notably, since the cranial nerves controlling the eye-muscles bypass this circuit, large rapid eye movements (REM) are observed in humans during this state (Lai et al., 2010).

The above REM promoting switch results in a state characterized by desynchronized EEG similar to waking concurrent with the diminution of muscle tone to levels lower than those observed during non-REM sleep. Also unique to this state is the concurrence of high acetylcholinergic and low noradrenergic and serotonergic tone in the forebrain. This unique juxtaposition of characteristics has led this sleep state to also be referred to as *paradoxical sleep* by subsequent researchers (Jouvet and Michel, 1960). The resulting desynchronized state is characterized by tonic neural firing and gamma (30-80Hz) LFP oscillations as well as strong theta (5 – 10Hz) oscillatory activity in the hippocampus, amygdala, and in cortical theta generators such as the medial prefrontal cortex (mPFC) of rats (Buzsáki, 1998; Montgomery et al., 2008; Popa et al., 2010; Siapas et al., 2005). Contrasting non-REM, which is characterized by largely global fluctuations, the coherence of network activity during both waking and REM falls off sharply with distance, suggesting that local interactions dominate during these ‘active’ EEG states (Bullock et al., 1995; Destexhe et al., 1999; Steriade, 2003). In the hippocampus this bias towards local network processing may be reflected in the inhibition of the tri-synaptic (dentate gyrus to CA3 layer to CA1 layer) pathway, in effect

partially isolating each of these layers (Montgomery et al., 2008; Sil'kis, 2009). However, it should be noted that REM sleep is also capable of sustaining longer range interactions as evidenced by the theta frequency coherence observed in neurons and LFPs between the hippocampus and mPFC during REM (Benchenane et al., 2010; Siapas et al., 2005; Wierzynski et al., 2009).

Consistent with their role in the generation of theta band oscillatory activity, hippocampal interneurons show a pronounced increase in firing rate during the REM and active wake states (Buzsáki et al., 1983; Csicsvari et al., 1999a; Mizuseki et al., 2012). Despite the increase in inhibitory activity, CA1 layer pyramidal cell population firing rates are similar across both non-REM and REM and higher during waking (Csicsvari et al., 1999a). However, unlike the transient and sporadic periods of activity observed during non-REM, during REM CA1 layer pyramidal cells can exhibit episodes of sustained elevated firing lasting several hundred milliseconds to seconds and which are reminiscent of the place fields (see below) observed during active waking (Harris et al., 2002; Louie and Wilson, 2001). As I presented in a poster in the 2010 *Society for Neuroscience* conference (Grosmark et al., 2010), hippocampal neural discharge during REM further contrasts that observed during non-REM in that:

- 1) The firing rates of CA1 layer pyramidal cells and interneurons are more variable during individual REM episodes than individual non-REM episodes.
- 2) The firing rates of CA1 layer pyramidal cells and interneurons are more highly preserved across pairs of non-REM episodes than across pairs of REM episodes.

These findings were unaffected by yoking the amount of analyzed non-REM to REM, and were also similar for the pair-wise correlations observed between cells. Consequently, when compared to non-REM, REM sleep represents a period characterized by increases in inhibitory sculpting, decreased synchrony, decreased global connectivity, sustained activation, and variability in hippocampal neural discharge.

Functions of sleep

As alluded to earlier in this chapter, while many questions remain concerning the neurophysiological characteristics of sleep, the most puzzling outstanding question regarding this state may well be its function. Though this question has received many speculative answers from Aristotle to Freud, here we will discuss the two contemporary theories most related to the thesis research: 1) non-REM sleep promotes the homeostatic decrease in excitability, opposing the buildup in excitability from synaptic potentiation during waking, 2) non-REM and/or REM sleep support offline memory consolidation.

Sleep and Synaptic Homeostasis

As we all know the propensity for sleep is governed by both circadian and homeostatic processes. The circadian process, referred to as process C, is simply our heightened propensity to sleep at certain times of the day (Borbély and Achermann, 1999). The homeostatic process, process S, is evident in our heightened propensity for sleep after long bouts of waking, after sleep deprivation and after bouts of taxing physical

or mental activity. An updated version of the homeostatic theory proposed by Giulio Tononi and Chiara Cirelli argues that process S reflects the build-up of cellular excitability from synaptic potentiation during waking (Hanlon et al., 2011; Tononi, 2009; Tononi and Cirelli, 2003, 2006a). During subsequent sleep, this build-up is reversed specifically by non-REM slow-wave activity resulting in synaptic downscaling (depotentialiation) and a reduction of process S. Slow-wave dependent synaptic downscaling thus would have the effect of keeping cellular excitability within a physiologically viable range, as well as limiting the energy expenditure needed for the maintenance of potentiated synapses and increased excitability (Vyazovskiy et al., 2008). This model fits well with the earlier ‘sequential hypothesis’ of sleep-related memory consolidation in which non-REM downscapes all synapses by an approximately equal factor, with the effect of leaving only those synapses most potentiated during awake still potentiated above baseline (Giuditta et al., 1995). Subsequently, REM sleep promotes further potentiation specifically at these ‘surviving’ synapses, effectively reducing the noise (weak synapse) to signal (strong synapse) ratio. However, it should be noted that the sequential hypothesis does not identify process S with the build-up of excitability during awake, and indeed does not require that such net build-up take place.

Converging evidence for the homeostatic hypothesis has been obtained through several different paradigms:

- 1) Cortical expression of several genes associated with synaptic potentiation is highest during waking and lowest during non-REM (Cirelli, 2005; Cirelli and Tononi, 2000a; Pompeiano et al., 1995; Ribeiro et al., 2002).

- 2) Non-circadian sleep propensity (process S) is associated with an increase in cortical delta wave activity both during awake and during ensuing recovery sleep (Vyazovskiy et al., 2011b).
- 3) Extracellularly recorded neurons in the somatosensory cortex of rats increase their firing across waking bouts and decrease their firing rate across sleep (Vyazovskiy et al., 2009).
- 4) Extra-cellular glutamatergic tone, as measured by *in vivo* voltometry, increases during waking and REM sleep and decreases during non-REM sleep (Dash et al., 2009).
- 5) Excitability as measured *in vitro* by the frequency of mEPSCs in neurons of the frontal cortex of rats and mice is highest in sleep deprived animals, intermediate in spontaneously waking animals, and lowest in spontaneously sleeping animals (Liu et al., 2010).
- 6) Cortical brain regions which are specifically recruited (and presumably potentiated) during a behavioral task, show increased delta activity during ensuing sleep (Hanlon et al., 2009; Huber et al., 2004, 2007).
- 7) Low-frequency stimulation mimicking slow-wave activity leads to synaptic depotentiation (Kemp and Bashir, 2001).
- 8) Low neuromodulatory, and particularly noradrenergic tone, as observed during sleep, is conducive to synaptic depotentiation (Cirelli, 2005; Cirelli and Tononi, 2000b).
- 9) In sleep deprived rats, individual cortical neurons can spontaneously display slow-wave like activity patterns during waking, presumably revealing latent

homeostatic influences otherwise hidden by the externally enforced waking state (Vyazovskiy et al., 2011a).

Though supported by an impressive array of evidence, the homeostatic hypothesis faces several largely unresolved challenges. Firstly, it is well documented that homeostatic mechanisms controlling excitability are already present during waking (Abbott and Nelson, 2000; Burrone and Murthy, 2003; Turrigiano, 1999). For instance, lateral inhibition, a nearly universal feature of both cortical and subcortical networks closely links overall excitation and inhibition, actively preventing runaway excitation. Likewise, in a network with stochastic activity, spike-timing dependent plasticity (STDP) is thought to result in approximately equal levels of synaptic potentiation and depotentiation (Lubenov and Siapas, 2008; Turrigiano, 1999). Furthermore, while the evidence above supports the dominance of LTP over LTD during waking in the neocortex and vice-versa during sleep it does not support that this is the case in other brain regions (note that points 1 through 9 above only pertain to the neocortex). In fact, it might be expected that brain regions with distinct network architectures and which support different forms of LTP and LTD, may well display alternate forms of synaptic homeostasis. Consequently, the neocortex may prove to be a special case and other brain regions may not display a dominance of LTP over LTD during waking. This is problematic because in this case it is unclear why the build-up of excitability in the neocortex alone would be the sole determinant of the homeostatic process S. Similarly, it is clear that process S can be built up through seemingly non-mental energy expenditures. For instance, performing a well-known (so as to minimize learning) heavy exercise clearly leads to an increase in process S. Though it is known that exercise stimulates synaptic potentiation (Farmer et al., 2004),

it is not evident that in the marked increase in process S is proportional to (cortical) synaptic potentiation in this case. Furthermore, the homeostatic hypothesis predicts that the blockade of synaptic potentiation, for instance through the administration of NMDA blocker AP-5 or other pharmacological agents inhibiting LTP pathways, should lead to a virtual abolition of process S. Consequently, AP-5 should abolish the need for recovery sleep following sleep deprivation, while no such effect has been reported (Melik et al., 2006; Prospero-García et al., 1994).

Beyond the correspondence between the build-up of excitability during waking and process S, or the generalizability of this build-up to non-neocortical structures, the attribution of synaptic downscaling to the non-REM state is also problematic. While low-frequency stimulation is generally observed to induce LTD (though see also Habib and Dringenberg, 2010), it is not clear that such stimulation carried out *in vivo* or *in vitro* captures the dynamics of observed concurrently with the slow-wave oscillation. Specifically, it is unclear that such stimulation accounts for the presence of spindles or gamma frequency oscillations during non-REM UP periods (Mölle and Born, 2011), while both of these oscillations have been linked to synaptic potentiation (Rosanova and Ulrich, 2005; Whittington et al., 1997). Moreover, low frequency stimulation fails to reflect slow wave activity in that the latter results in periodic, highly synchronous neural firing. Notably, the arrival of coincident inputs as is promoted by the synchronous activity observed during slow-wave sleep, is thought to promote to LTP (Axmacher et al., 2006; Hebb, 1949). Furthermore, the researchers proposing the homeostatic theory of sleep maintain that high noradrenergic tone during waking, and low noradrenergic tone during non-REM sleep are instrumental in build-up and diminution of synaptic

potentiation and thus process S (Cirelli and Tononi, 2000a, 2000b; Cirelli et al., 1996, 2005; Pompeiano et al., 1995; Tononi et al., 1990, 1994). Indeed, noradrenaline's role in promoting neural excitability is well established (Gu, 2002; Mallick et al., 2010; Saar and Barkai, 2009). However, noradrenergic tone is lowest not during non-REM sleep but during REM sleep, when in fact it is virtually absent (Aston-Jones and Bloom, 1981; Hobson et al., 1975). The homeostatic hypothesis does not seem to account for this discrepancy or in fact for REM sleep generally. Finally, only one study (item 4 in the list above) has presented support for a decrease of excitability *within individual* non-REM episodes. However, this study looked at levels of extracellular glutamate, a signal that is known to not only reflect the synaptic activity of neurons but is also known to be extensively regulated by local astrocytes (Chuquet et al., 2010; Hansson and Rönnbäck, 1995; Schousboe and Waagepetersen, 2005). Notably, astrocyte activity is also thought to change with sleep (Frank, 2011; Halassa et al., 2010). Consequently this is at best an indirect measure of excitability changes occurring within individual non-REM episodes. Surprisingly, changes occurring within individual non-REM or REM sleep episodes were either absent or not presented in a recent analysis of extracellular recorded neurons across sleep (item 3 above).

To establish the general validity of the homeostatic model, it is essential to test its predictions in multiple cortical areas. Moreover, since sleep consists of two competing physiological processes, non-REM and REM sleep, it is important to learn how these distinct sleep stages contribute to the hypothesized homeostatic function of sleep. Notably, homeostatic models do not attribute an explicit role to REM sleep, even though alterations of REM sleep are intricately related to cognitive and affective disorders

manifested in the waking brain (Born et al., 2006; Campbell and Gillin, 1987; Gierz et al., 1987; Walker, 2010).

Memory Consolidation

Far from being a unitary event, memory formation and maintenance has been found to be an ongoing process persisting for many days or weeks after initial learning. This ongoing mnemonic processing is termed memory consolidation, reflecting the observation that memories can become strengthened after initial learning without the need for further exposure to the learned stimulus (Stickgold and Walker, 2007). Consolidation is thought to be dependent on the coordinated ‘reactivation’ of those neural activity patterns established during learning (Buzsáki, 1989, 1998; Hebb, 1949). Moreover, for reactivation to accurately reflect learned activity patterns it must be protected from interference from evoked and ongoing neural activity, leading to the hypothesis that ‘off-line’ modes of minimal sensory input, most notably sleep, represent a privileged window for memory consolidation (Buzsáki, 1998; Diekelmann and Born, 2010; Hebb, 1949; Stickgold and Walker, 2007). While reviewing the extensive literature in support of this hypothesis is beyond the scope of this chapter, some pertinent aspects of memory consolidation, and particularly hippocampal-dependent memory consolidation, will be here discussed.

The link between the hippocampus and memory first became evident 1957 through the extraordinary clinical case of patient H.M., whose entire medial temporal lobe, including the hippocampal formation, were surgically removed in an attempt to treat

seemingly intractable epilepsy (Scoville and Milner, 1957). Surprisingly, after the surgery it was found that H.M. displayed severe temporally graded anterograde amnesia, in other words, while H.M. could remember facts learned up to several weeks before the surgery, he failed to remember information learned shortly before the surgery, or indeed to make new memories. Notably, this anterograde amnesia applied to a class of memories broadly described as 'declarative' due to the fact that they can be explicitly stated. However, 'procedural' memory, i.e. implicit motor learning such as that required for typing dexterity, was largely spared in subject H.M. This led to the hypothesis that the hippocampus serves as a temporary storage site for newly formed declarative memories until they are gradually transferred to their final storage site, presumably in the neocortex (Buzsáki, 1998; Eichenbaum, 2000; Quinn et al., 2008; Rauchs et al., 2005; Wiltgen et al., 2004).

Though 'declarative memory' cannot be directly assessed in non-human animals, subsequent lesion work in rats has supported a central role for the hippocampus in 'declarative-like' memory, particularly in memory of temporal and spatial relational information, such as that required for non-cued maze running (Buzsáki, 2005). The importance of the hippocampus in spatial memory as observed by lesion studies is also consistent with the observation that in rats many hippocampal neurons are 'place cells', that is, during maze or open field running these cells become active only when the animal is in a circumscribed spatial location referred to as the cell's 'place field' (Bird and Burgess, 2008; O'Keefe and Dostrovsky, 1971). Notably, similar to the amnesia observed in subject H.M., spatial memories in rats show a temporally graded anterograde dependence on the hippocampus after initial learning, even without further exposure to

the learning context (Squire, 1992). The hippocampo-fugal, cortico-petal transfer of declarative (or ‘declarative-like’) memories, often referred to a ‘system consolidation’, is thought to be a dominant (though not unique) mechanism of memory consolidation. Not surprisingly then, sleep has been implicated in system consolidation. For instance, in an elegant experiment Ribeiro and colleagues showed that unilateral stimulation of the entorhino-hippocampal perforant path during waking, led to waves of up-regulation of synaptic plasticity related immediate early genes (IEGs) in the ipsilateral neocortex during subsequent REM sleep (Ribeiro et al., 2002). Notably, the intimate relationship between the hippocampus and memory consolidation, and memory consolidation and sleep, motivates much of the research, including my own, into the physiology of the hippocampus during sleep.

Though initial behavioral experiments with human subjects failed to show a clear link between declarative memory and sleep (Smith, 2001), subsequent work using more demanding declarative memory tasks has shown that sleep does promote declarative memory consolidation (Barrett and Ekstrand, 1972; Binder et al., 2012; Plihal and Born, 1997). In one such study overnight sleep subsequent to learning a (declarative) word-pairing task led to a modest increase in remembered word-pairs when compared to awake controls. More notably however, when compared to awake controls, subjects allowed to sleep after initial word-pair learning showed markedly improved retention of these word pairs after this memory had been challenged by the learning of a separate, but similar, word-pair list (Ellenbogen et al., 2006a). The protective effect of sleep on newly formed memory from subsequent interference is referred to as memory stabilization. In another study, it was found that sleep promoted the formation transitive inferences from

previously learned item pairings (Ellenbogen et al., 2007). This transitive inference effect as well as the stabilization effect both imply that sleep not only can strengthen, but also modify memory content, thus suggesting an active, rather than passive or permissive function of sleep in memory consolidation (see Ellenbogen et al., 2006b for review).

While much evidence supports a memory consolidating function for sleep, attributing this function to particular sleep states remains a subject of active debate. Reflecting intuitions that memory consolidation would be facilitated by active off-line states, early theories of memory consolidation assumed a dominant role for REM sleep (Horne and McGrath, 1984). However, several preliminary observations that time spent in non-REM correlates with declarative memory (i.e. hippocampal-dependent) memory consolidation, while time spent in REM predicts procedural and emotional memory consolidation were formalized in the ‘dual process’ hypothesis of sleep-dependent memory consolidation (Barrett and Ekstrand, 1972; Borbély, 2009; Gujar et al., 2011; Maquet, 2001; Plihal and Born, 1997; Rauchs et al., 2005; Wagner et al., 2001). Though a majority of experiments agree with this general dichotomy, other experiments showing opposite patterns correlations of declarative and procedural memory with time spent in non-REM and REM complicate the interpretation of the dual process hypothesis (Aeschbach et al., 2008; Gais et al., 2000; Huber et al., 2004). Interpreting these results is further confounded by the fact that time spent in non-REM sleep, particularly at the beginning of the night, is correlated to time spent in REM sleep, particularly at the end of the night (Stickgold et al., 2000). Likewise, though declarative and procedural learning has been linked to the intensity of cortical spindle and slow oscillations (Bódizs et al., 2002; Fogel et al., 2007; Gais et al., 2002), other studies have linked these forms of

learning to REM-related increases in theta-band power (Fogel et al., 2007; Jackson et al., 2008).

Replay

Off-line reactivation of behaviorally relevant neuronal patterns is a central tenet of neural models of system consolidation (Buzsáki, 1996). Given the hippocampus' unique role in temporary storage of memory traces, this structure was one of the first to be examined for evidence of reactivation. The first such evidence was presented in 1989 with the demonstration that single hippocampal cells that display elevated firing rates during running in a maze, also show elevated firing rates during subsequent sleep (Pavlidis and Winson, 1989a). With the advent of multi-electrode recordings techniques, it was shown that pairs of pyramidal cells which showed overlapping place fields during maze running also preferentially co-activated during the subsequent sleep epoch (Dupret et al., 2010; Kudrimoti et al., 1999; O'Neill et al., 2008; Wilson and McNaughton, 1994). Subsequently, it was shown that during sleep the sequential activation of groups of hippocampal pyramidal cells preserve the sequential structure observed during the preceding behavior - a phenomenon termed 'replay' (Davidson et al., 2009; Dragoi and Tonegawa, 2011; Gupta et al., 2010; Ji and Wilson, 2007; Johnson and Redish, 2005; Lansink et al., 2009; Lee and Wilson, 2002; Louie and Wilson, 2001; Nádasdy et al., 1999). Notably, the off-line reinstatement of sequential activity as observed in replay is in line with several models describing the propagation of excitation through the hippocampal network (Battaglia and Pennartz, 2011; Hasselmo, 2008; Káli and Dayan,

2004). Moreover, because of the complexity and specificity of the reactivated sequences involved, replay has been taken as one of the strongest lines of evidence supporting the existence of memory trace reactivation. However, many important questions about replay, and more generally about the mechanisms by which sequential activity is generated in the hippocampus, remain unanswered.

The notable success of replay research is tempered by the presence of several confounds, addressed with varying degree of success in different studies. Firstly, it should be noted that while residual replay signals have been reported up to 24 hours after initial learning (Dupret et al., 2010; Louie and Wilson, 2001), several reports show a rapid decay of replay within the first 15 to 60 minutes of Post sleep (Kudrimoti et al., 1999; Schwindel and McNaughton, 2011; Wilson and McNaughton, 1994). It is unclear how this rapid initial decay can be reconciled with the slow time course of hippocampal-dependent memory consolidation, typically thought to take days to weeks. Consequently, the mechanisms and functional implications of the observed time-course of replay remains an outstanding questions.

Other confounds relate to the specificity of replay to learning related effects. For instance, due to the high degree of non-stationarity in neural discharge during non-REM (i.e. the presence active and inactive states precluding the modeling of non-REM discharge as a simple Poisson process), pairs of pyramidal cells display a baseline positive correlations in the absence of learning specific co-modulation (Battaglia et al., 2005; Peyrache et al., 2010; Tatsuno et al., 2006). Furthermore, this positive correlation is proportional to the cells' firing rates (de la Rocha et al., 2007). Consequently, an increase in pair-wise correlations may reflect an increase in firing rate rather than an

increase in learning-specific co-modulation during Post sleep. Likewise, the use of Pre sleep as a control for Post sleep replay assumes the absence of non-learning specific changes between Pre and Post sleep, an assumption that is not typically explicitly tested. Finally, the observed replay signal, though highly statistically significant, accounts for only a small fraction of the variance observed in Post sleep pair-wise or sequential discharge dynamics. The structure and origin of the remaining, non-task specific dynamics and their relationship to the assumed specific replay of memory remain largely unanswered questions.

Types of Replay

While often considered a unitary phenomena replay has in fact been measured in by variety of methods, which can potentially be dissociated and in turn each of which may be generated by unique physiological processes. In particular there are at least two general types of replay (see also table 1):

1. *Rate replay* - the 'off-line' reactivation of single-cell (Hirase et al., 2001; Pavlides and Winson, 1989), pair-wise (Kudrimoti et al., 1999; Wilson and McNaughton, 1994), or higher-order (Abeles and Gerstein, 1988; Peyrache et al., 2009) firing rate patterns observed during behavior.
2. *Sequence replay* -the 'off-line' reactivation of pairs (Euston et al., 2007; Skaggs and McNaughton, 1996) or larger groups (Diba and Buzsáki, 2007; Dragoi and Tonegawa, 2011; Foster and Wilson, 2006) of neurons in the sequence in which they fired during behavior.

Note that in all cases in which it is specifically *sleep-related* replay that is of interest (rather than wake related replay, Carr et al., 2011) a Pre behavioral epoch is used to control for non-learning specific changes which occur during the Maze epoch. Importantly, note that for both rate replay and sequence replay both pair-wise as well as higher-order (event based) methods have been developed. Moreover, pair-wise and higher-order methods each offer their own unique advantages and disadvantages in the study of replay.

Pair-Wise Rate Replay

Motivated by research predicting and indicating the strengthening of pair-wise connectivity during learning (Bliss and Lømo, 1973; Hebb, 1949), the first instance of (none-single cell) replay was the report by Wilson and McNaughton in 1994 in which place cells that were correlated on a maze (i.e. had overlapping place fields) were selectively correlated (in 100 ms bins) during the subsequent Post sleep. Note that the pair-wise nature of their analysis enabled them to focus on particular *interactions* of interest, particularly, on the interactions between pairs of place cells with overlapping place fields recorded on different tetrodes. In addition, the pair-wise nature of the analysis also increased its statistical power by increasing the n of the analysis (for every i number of cells there are $i*(i - 1)/2$ unique pairs). However, since the correlations were taken *across* many 100 ms bins recorded in the Post sleep, the fine time-scale temporal dynamics of the observed correlations were poorly defined - as is the case for most pair-wise analysis (Perkel et al., 1967a). It should also be noted that the task-dependence of

the pair-wise interactions observed by Wilson and McNaughton were thought to be largely independent (and more robustly task-specific) from the task-dependent single-cell firing rate increases observed by Pavlides and Winson (1989). However, it is now known that pair-wise measures of co-activity are highly sensitive to task-specific and non-specific changes in the firing rates of individual cells (Battaglia et al., 2005; de la Rocha et al., 2007). It is partly due to the ambiguity between single cell and pair-wise (as well as higher-order) effects, that these are here all grouped in the *Rate Replay* category.

Event-Based Rate Replay

While long-term potentiation occurs at the pair-wise level, these pair-wise interactions are thought to aggregate into the recruitment of higher-order (that is, more than two cell) *ensembles* (Hebb, 1949) which function synergistically to affect their downstream targets (Buzsáki, 2010). This has motivated several groups to examine high order structure (Chapin and Nicolelis, 1999; Peyrache et al., 2009; Ribeiro et al., 2004). One such group, Peyrache et al., 2009, used a principal component analysis based technique to reveal the replay of rule-learning specific changes in the medial prefrontal cortex. Crucially, the ability to assign a particular quantifiable 'replay' value to each activity time-bin analyzed during the Post epoch allowed them to determine that the medial prefrontal replay of rule learning was impressively time-locked to ripple/sharp-wave occurrence as detected in the posterior hippocampus. Thus, the *putatively higher order* (i.e. 'ensemble') structure observed, as well as the fine temporal resolution afforded, are two main advantages associated with the event based approach. The term 'putatively

higher-order' is highlighted because putatively higher-order based effects may be caused by a few (or a single) single-cell or second-order (i.e. pair-wise) interactions (Tatsuno et al., 2006). The difficulty associated with determining the particular cells or interactions which are determinate of the overall effect, may thus be considered a primary drawback of event based methods generally¹.

Pair-Wise Sequence Replay

Nearly all the types of memory known to be hippocampus-dependent in the rodent share the commonality of being memories of sequentially experienced events (Wallenstein et al., 1998). Thus memory of the sequence by which a tone predicts a temporally displaced noxious air-puff (Millenson et al., 1977), the sequence of places (i.e. a trajectory) leading to a submerged platform (Morris et al., 1982), and the sequence of odors indicating correct reward (Kesner et al., 2002) all rely on an intact (specifically dorsal) hippocampus. This known hippocampal-dependence of sequential activity motivated Skaggs and McNaughton (1996) to demonstrate that the pair-wise sequence of firing (i.e. A before B, versus B before A) as estimated by using 400 ms - wide crosscorrelograms (both on the maze and during the Pre/Post epochs) was preserved from behavior to the Post epoch. Similarly to Wilson and McNaughton (1994), Skaggs and McNaughton (1996) took advantage of the pair-wise nature of their method to restrict their analysis to pairs of cells which displayed overlapping place-fields and which were

¹To be clear, Peyrache et al. 2009, though not some others (e.g. Ribeiro et al., 2004), went some ways to convincingly show that their effect was not due to a small number of cells, and to specify which pairs of cells contributed the most to the effect. Here we simply note that this analysis was carried out separately to, and was not necessarily implied by, their main result.

recorded on separate tetrodes. Notably, the temporal specificity of the observed effects was taken to be less sensitive to non-specific effects than the earlier pair-wise co-activity measure (though see Moore et al., 1996 and Peyrache et al., 2010) and more reflective of the putatively sequentially structured activity expected in the highly recurrent hippocampal layer CA3 network (Skaggs and McNaughton, 1996).

Event-Based Sequence Replay

Similarly to the logic that motivated research into activity ensembles (see *Event-Based Rate Replay* above), the discovery of pair-wise sequence replay (Qin et al., 1997; Skaggs and McNaughton, 1996) led researchers to theorize the presence of higher-order sequential structure. As noted above the hippocampus is necessary for the coding of memory of sequential events generally, and spatial sequences in particular. In a simple model, hippocampal encoding of a spatial sequence such as A-B-C purely results from B's pair-wise tendency to fire after A, and C's pair-wise tendency to fire after B, with no interaction present between A and C - such as is predicted by 'synfire' chain models (Abeles et al., 1993). Notably, in this case interference between pair-wise interactions could result in corrupted sequences. This is due to the fact that the same cell may be involved in coding different sequences in different mazes. So for instance, hypothetical cell B may also participate in sequence X-B-Z in addition to sequence A-B-C. In this case, cell A's discharge would predict cell B's discharge which in some cases would lead to cell Z's discharge, resulting in the corrupted sequence A-B-Z – a sequence with no behavioral correlate (Nádasy, 2000). Preventing this type of sequence corruption is one

of the reasons why neural models of memory consolidation typically predict network (that is neural-group) level mechanisms for the coordination of sequential reactivation (Hebb, 1949). In addition to this theoretical consideration, work on event-based sequence replay was motivated by the observation of the presence and importance of sharp-wave/ripples in the hippocampus (Buzsáki et al., 1983, 1992). Importantly sharp-wave/ripples dominate neural firing during 'off-line' states and are thought to facilitate the 'binding' of cells that were sequentially but non-contiguously activated during behavior by inducing their co-activity within the small time-constants relevant to learning and memory (Buzsáki, 1989; Girardeau et al., 2009). Importantly, due to the fact that sharp-wave/ripples are typically defined by LFP criteria and may not involve the cells being recorded, most studies of hippocampal event-based sequence replay have instead analyzed population activity events as defined directly from the spiking of the recorded pyramidal cell population using various criteria (these events are also referred to as 'ON events' (Dragoi and Tonegawa, 2011) or 'Frames' (Ji and Wilson, 2007) in the literature).

Notably, most of the studies measuring event-based sequential activity in the hippocampus were carried out in the awake state (Diba and Buzsáki, 2007; Foster and Wilson, 2006; Karlsson and Frank, 2009; Pfeiffer and Foster, 2013) during which hippocampal replay may be stronger (Carr et al., 2011). The first report of event-based sequential structure during non-REM was conducted by Lee and Wilson (2002). They used a 'combinatorial' method by which they examined the distribution of 'words': uninterrupted sequences of place cell activity occurring during population activity events and matching the sequence of their place fields on a linear track. While groundbreaking and thorough the method employed by Lee and Wilson (2002) can be criticized in at least

two regards. First, in stipulating that 'words' have no interruptions (i.e. a place cell spiking out of order) their method is very sensitive to the noise associated with single cells. Indeed, this method is somewhat unique in that, though requiring a minimum number of place cells, its quality would be expected to decrease with increasing cell number, as the 'miss-firing' of any additional cell could cause the termination of the sequence in any given event - leading to temporally shorter and shorter sequences. Secondly, the null distribution Lee and Wilson (2002) used to compute the significance of their effect was based on the permutations of the spike sequences in each event. This is problematic in that this null assumes the lack of any common sequence content across population activity events - an assumption that could statistically bias the results (Diba and Buzsáki, 2007; Pfeiffer and Foster, 2013).

Partly in order to address these concerns regarding the method of Lee and Wilson (also used in Ji and Wilson, 2007), a novel spike-sequence replay method was developed based on the Spearman (rank-order) correlation between sequence of activity of place fields on a maze and their sequence of activity in a given population activity event (Diba and Buzsáki, 2007; Dragoi and Tonegawa, 2011, 2013; Foster and Wilson, 2006). By use of the rank-ordered correlation (rather than longest matching 'word'), this method is more robust against the noise of the single noisiest cell in any given event and is thus expected to become more, rather than less, precise as a function of number of cells in the sample. Secondly, in this method it is the order of place fields on the maze, rather than the order of activity in any given event, that is permuted to generate the null distribution - ensuring that the null distribution of rank-order correlations preserves the effect of any sequence structure intrinsic to the population activity events themselves. The rank-order

correlation method was used to discover the presence of reverse replay, the discharge of place cells during a population activity event in the *opposite* order as the sequence of their place fields on a maze (Diba and Buzsáki, 2007; Foster and Wilson, 2006). Note that since on the linear maze hippocampal neurons tend to display a distinct set of place fields when running in the left and right directions, replay is typically assessed independently in each of these directions (Davidson et al., 2009; Diba and Buzsáki, 2007; Dragoi and Tonegawa, 2011; Foster and Wilson, 2006). In addition, as a consequence of the discovery of reverse replay, researchers studying replay in general also must take the potential for this additional, *reverse*, direction into account for replay of either left or right runs on a linear track (Davidson et al., 2009; Dragoi and Tonegawa, 2011, 2013). As is the case for event-based methods when compared against the pair-wise methods discussed above, rank-order replay has the advantage of providing fine temporal (i.e. event-by-event) resolution in the measurement of the replay signal. However, it is also expected to necessarily be biased towards higher firing rate cells proportionally to the extent to which these are 'over-represented' in the number of events in which they participate.

There are currently only two studies that have used the rank-order method for studying activity patterns during sleep rather than awake: the two pre-play studies (discussed below) performed by Dragoi and Tonegawa (2011, 2013).

Bayesian Decoding of Replay

Another, and more recent, approach to the study of replay is the use of 'decoding' algorithms which, rather than mapping individual neuron's receptive fields, map the activity of populations of neurons onto stimulus space (Chen, 2013; Kloosterman, 2012). While several Bayesian decoding methods for neural data have been proposed (see Chen, 2013 for review), we will here concentrate on a relatively recent uniform-prior probability, memory-less algorithm (Davidson et al., 2009; Kloosterman, 2012) which has been increasingly used in hippocampal memory research (Davidson et al., 2009; Dragoi and Tonegawa, 2011, 2013; Pfeiffer and Foster, 2013). This method, itself a member of a much wider set of dimensionality reduction algorithms, is powerful in several regards. First, the Bayesian method considered here uses a set of previously collected stimulus induced population firing rate patterns whose evoking stimuli are known, as the probability distribution against which the evidence provided by a new set of population firing rate patterns is used to predict the identity of an unknown stimulus. Consequently, the Bayesian algorithm is a transform from population activity data to stimulus space and is thus well suited for the study of memory in which 'unlabeled' activity contains 'evidence' of stimulus-specific memory content. In addition, because its output is in terms of stimulus space, Bayesian decoding can also be usefully validated (and cross-validated) by tracking its performance in decoding neural content of perceived (rather than remembered) stimuli. Finally, by performing this labeling using the evidence provided by all the population activity in a given time bin - it is very robust against the noise associated with individual neurons and is informed by an individual cell's activity as well as it's silence (Kloosterman, 2012).

Yet, the Bayesian technique also has some important draw-backs. In particular it requires the data to first be binned into arbitrary (typically 15 or 20 ms, Davidson et al., 2009; Pfeiffer and Foster, 2013) windows. Perhaps more importantly however, the physiological nature of the transformation involved is often opaque, and so it can be difficult to determine the individual cells contributing to the observed effect, similarly to other methods for dimensionality reduction.

In the Bayesian method variant relevant to hippocampal memory research proposed in Davidson et al., 2009, stimulus induced activity vectors are defined as the average population activity as binned in (2 to 5 cm, Davidson et al., 2009; Dragoi & Tonegawa, 2011) spatial bins along a linear (or linearized, Davidson et al., 2009) track. The population activity within individual population activity events is divided into (15 to 20 ms) temporal bins, the posterior probability of position given each binned activity vector is calculated and normalized to one across all positions (within each bin). For each bin in a population activity event, each element of the posterior probability vector represents the estimate that the animal is directly experiencing or remembering that spatial bin. In other words, for each temporal bin, each element of the posterior probability vector is a measure of the similarity of the population activity vector in that temporal bin to the population activity vector associated with a particular spatial bin.

The sequence replay content of a particular population activity event is calculated across all the temporal bins in that event via a 'line-casting' algorithm. In this algorithm a dense sampling (i.e. exhaustive search) is performed across all linear trajectories (i.e., across 15 to 20 ms temporal bins and 2 to 5 cm spatial bins) to determine the particular linear trajectory associated with the maximal average posterior probability (Davidson et

al., 2009; Kloosterman, 2012). This average posterior probability (i.e. the one associated with the 'best' trajectory) is taken as the event's 'replay score'. A null distribution is similarly derived from permuted versions of the stimulus evoked activity vectors.

The 'line-casting' algorithm is an elegant solution in that the trajectories it derives make use of the entire space of posterior probability of position, rather than just the peak (i.e. best estimate) at each bin. However, while the dense sampling involved is well suited for small number of events such as those that occur during waking maze-running or during short sleep recordings, the computational demands of the dense-sampling algorithm (which must also be run on each shuffled controls) can be prohibitively time consuming when analyzing larger data sets. In addition, and in contrast to the tied-rank correlation method, this algorithm is restricted to linear, rather than simply sequential, trajectories - an assumption made for computationally-relevant (i.e., in order to make full use of the entire distribution of posterior probabilities) rather than physiologically-relevant reasons. Finally, and perhaps most importantly, in this method the posterior probability (a measure of how much each temporal bin's activity resembles a particular spatial bin's activity) is only assessed as a function of trajectory (a measure of spatial sequence). This conflation assumes that rate replay, and sequence replay are physiologically an entirely unitary phenomenon with identical dynamics, though perhaps different noise (see Huxter et al., 2003 for a counter-instance). This assumption, however, is not tested.

As it should be clear from the above summary of the currently available methods, each method has its advantages and shortcomings and the statistical results obtained with the respective methods may identify different, occasionally opposing, relationships and,

thus may lead to unique insights. Without assuming a ‘best’ method(s) (as is traditionally done), I used several of them in my dissertation work and illustrate how the choice of the method can affect the interpretation of complex, spike-based measurements of memory replay and preplay.

Pre-Play

Sequential replay (both pair-wise and event-based) has previously been taken as one of the strongest pieces of available evidence of the maze-specificity of replayed memory content (Lee and Wilson, 2002; Nádasdy et al., 1999; Skaggs and McNaughton, 1996). However, two recent studies by Dragoi and Tonegawa (2011, 2013), show that the sequence of place fields which place cells form on a novel maze can be statistically predicted by their sequential activation during population activity events occurring before first exposure to the maze. The novelty of the maze to the experimental subject is in this case crucial to establishing that the sequential content observed is not in fact replay of a previous exposure. Note that in perhaps most other brain structures, the prediction of stimulus-specific responses during 'off-line' states would be considered rather unsurprising. Indeed, in the sensory neocortex, cells with overlapping receptive fields are known to co-vary even during 'off-line' states (Kenet et al., 2003; Luczak et al., 2009), due to the fact that receptive fields in these arise from largely hard-wired developmental programs (Hubel and Wiesel, 1962; Yoshimura et al., 2005). However, the finding of preplay in the hippocampus was largely unexpected because place coding in this structure is known to be radically modulated by the specific context of the space being coded in that both stimulus-specific single cell responses (Muller and Kubie, 1987) as well as pair-wise

interactions (Hayman et al., 2003; Jeffery, 2011) do not reliably correlate across environments. Dragoi and Tonegawa interpret their results as implying the presence of discrete, hidden attractors within the recursive layer CA3 network (Dragoi and Tonegawa, 2014), though this particular interpretation has yet to be confirmed. Furthermore, in both replay papers Dragoi & Tonegawa (2001, 2013) show sequential higher-order replay using the rank-order and Bayesian decoding methods. Indeed, establishing which other metrics display the pre-play effect may be a useful step in elucidating its origins. It should be stressed, however, that regardless of its mechanisms the discovery of pre-play highlighted our crude understanding of the structure and mechanics governing non-task specific 'off-line' hippocampal firing and how this non-specific component may interact with stimulus coding and memory formation.

Recent work on the statistical properties of firing rate, synaptic weight and the magnitude of population synchrony distributions, however, may shed light to the physiological origin of pre-play. All these, and many other, distributions show long-tailed, typically lognormal distributions. Moreover, the log distributions remain statistically correlated across brain states, environments and situations (Buzsaki and Mizuseki, 2014; Grosmark et al., 2010). An inevitable consequence of such preconfigured networks with a minority of neurons and synapses dominating every brain state include pre-play, preserved activity across multiple sleep sessions and a large part of the replay process. Thus, it is no longer surprising that rates and cell pairs correlate and the task is to demonstrate what aspects of these parameters change during learning *beyond* what is expected from the preconfigured brain dynamic.

Goal of the Thesis

To advance our understanding of both sleep physiology and function I have used both hippocampal layer CA1 silicon-probe based sleep recordings (figure 1.4) previously performed by other members in the lab as well as implanted and recorded my own animals during both sleep and the performance of a novelty task. The goal of these experiments and analyses, and the goal of my thesis, was to further our knowledge regarding the dynamics and functional significance of sleep in the hippocampus. During the course of these studies, it has further been my goal to make a contribution to building a richer analytical vocabulary in which the nuances and complexities of the diverse phenomena observed in sleep can be better understood and more readily related to known hippocampal function.

Methods

TheStateEditor: Behavioral State Scoring and Data Visualization and Annotation Software

In order to perform thorough and efficient behavioral state scoring and large time-scale data visualization and annotation, software was developed that integrates the three primary data types used in behavioral state scoring: 1) whitened, log-transformed time-resolved spectrograms, 2) movement or EMG data, and 3) raw LFP visualization in one portable and end-user friendly application. Furthermore such an application should be stable, self-contained (ideally, not requiring extensive background library of sub-functions), have wide cross-platform support, be adaptable to many potential uses of large time-scale electrophysiological data, be able to interact with custom-made user data and produce easily interpretable outputs. Importantly, such software should present the end-user with an intuitive, interactive and user friendly experience – all features that maximize the efficiency of use. Finally, the software should be able to accept data entry in the widely accessible Matlab vector format, making the utility accessible to other labs which do not regularly use our lab's data formatting conventions (Hazan et al., 2006).

In order to develop such a utility we used Matlab (a widely used, cross-platform, commercially available software (MATLAB and Statistics Toolbox 2012b) for data analysis and visualization) scripts previously developed and tested in our lab. For instance, we employed the excellent, computationally efficient and widely used scripts for the whitening of EEG/LFP signals and for the calculation of time-resolved spectrograms previously developed by Dr. Anton Sirota during his time at our lab. While

an exhaustive elaboration of the intricacies of coding analytical and graphical utilities in Matlab is not here appropriate (see Gilat, 2010), it should be noted that in many instances the native Matlab functionalities were found to be ill-suited for the current purposes. As an example, Matlab's built-in zooming function has the undesired effect of 'grabbing' incoming keystrokes, making the overall utility insensitive to further user input. Consequently, *TheStateEditor's*, zooming utility, as well as its scrolling, resizing, and color-limit selection utilities, as well as others functions such as the different effects of single, double and sustained clicks, were all costume built within the Matlab framework.

REM Sleep Reorganizes Hippocampal Excitability

Animals, Surgery and Data Collection

Five male Long-Evans rats (250-400 g) were implanted with a 4- or 8-shank silicon probe in the right dorsal hippocampus under isoflurane anesthesia (1-1.5%) and recorded from CA1 pyramidal layers, as described earlier (Diba and Buzsáki, 2007; Mizuseki et al., 2009; Pastalkova et al., 2008). Each shank of the silicon probe had eight recording sites (160 μm^2 each site, 1–3-M Ω impedance) and the inter-shank distance was 200 μm . Recording sites were staggered to provide a two-dimensional arrangement (20- μm vertical separation). The silicon probes were attached to micromanipulators and moved slowly to the target over several days/weeks. Two stainless steel screws inserted above the cerebellum were used as indifferent and ground electrodes. The position of the electrodes was confirmed histologically and reported previously in detail (Mizuseki et al., 2009). Sleep recordings were performed in the animal's home cage during the day (i.e., during the sleep cycle of the nocturnal rodent), while the behavior of the rat and LFPs

from several channels were monitored by the experimenter. For the detection of head movements, two small light-emitted diodes (5-cm separation), mounted above the headstage, were recorded by a digital video camera and sampled (at 40 Hz). All protocols were approved by the Institutional Animal Care and Use Committee of Rutgers University.

Identification of Non-REM and REM Episodes

Rapid eye movement sleep (REM) and non-REM episodes were detected offline using the ratio of the power in theta band (5–11 Hz) to delta band (1–4 Hz) of LFP, followed by manual adjustment with the aid of visual inspection of whitened power spectra (using a low-order autoregressive model) and the raw traces (Mizuseki et al., 2009, 2011; Sirota et al., 2008). The manual adjustment was necessary to remove falsely detected short segments of data. REM episodes shorter than 50 sec were discarded. REM epochs were cross-validated with experimenter notes taken during recording to confirm that the rat was immobile and sleeping (Mizuseki et al., 2009). During sleep recording, the rat typically curled up in one of the corners of the home cage. Occasionally, behavioral signs of phasic REM sleep were present, including limb movements and twitching of whiskers. Although electromyogram was not recorded in the present experiments, several previous studies have compared hippocampal and neocortical activities during waking and sleep stages (Ribeiro et al., 2004; Robinson et al., 1977) and demonstrated that the presence of hippocampal theta oscillation in an animal with sleeping posture is sufficient to identify REM sleep (Robinson et al., 1977). Across all rats and sleep sessions, 82 non-REM episodes (mean length 498.3 sec \pm 30.98 SEM) and 45 REM

episodes (mean length 167.8 sec \pm 12.5 SEM) were detected.

Unit Clustering and Cell Classification

Detailed information about the recording system and spike sorting has been described (Diba and Buzsáki, 2007; Mizuseki et al., 2009; Pastalkova et al., 2008). Briefly, signals were amplified (1,000 \times), bandpass-filtered (1 Hz to 5 kHz) and acquired continuously at 20 kHz (RC Electronics) or 32 kHz (NeuraLynx) at 16-bit resolution. After recording, the signals were down-sampled to 1,250 Hz for the local field potential (LFP) analysis. For offline spike sorting, the wideband signals were digitally high-pass filtered (0.8-5 kHz) and the waveforms were resampled. Neurophysiological and behavioral data were explored using NeuroScope (<http://neuroscope.sourceforge.net>). Spike sorting was performed automatically, using KlustaKwik (<http://klustawik.sourceforge.net>), followed by manual adjustment of the clusters (using “Klusters” software package; <http://klusters.sourceforge.net>). Within the remaining data, only units with clear refractory periods and well-defined cluster boundaries were included in the analyses (Harris et al., 2000). While refractory periods of autocorrelograms of single units and cross-correlograms of unit pairs do not guarantee perfect neuron isolation, they indicate spike ‘contamination’ from other neurons (Harris et al., 2000). After spike sorting, we plotted the spike features of units (principal components) as a function of time to assess recording stability, and the units with signs of significant drift over the period of recording were discarded. The amplitudes of the units included in this study, from the first to the last non-REM sleep episodes, were highly

preserved ($r=0.94$, $r=0.98$ for putative pyramidal cells and interneurons, respectively; figure 3.5).

Hippocampal principal cells and interneurons were separated on the basis of their auto-correlograms, combination of trough to peak latency and the asymmetry index of the filtered (0.8 kHz – 5 kHz) spike waveform, bursting properties and mean firing rates (Barthó et al., 2004; Csicsvari et al., 1998; Harris et al., 2000; Mizuseki et al., 2009). Bursting was quantified as the fraction of interspike intervals shorter than or equal to 6 milliseconds (Mizuseki et al., 2009). It should be emphasized that there are no generally accepted methods for the segregation of principal cells and interneurons. However, several previous studies, using either simultaneous intracellular and extracellular recordings (Marshall et al., 2002; Quilichini et al., 2010) or optogenetic identification of interneurons (Royer et al., 2012), indicate that the above single cell features can reliably separate at least the fast spiking interneuron population from the bursting pyramidal cells. In addition, we applied a physiologically identification of the recorded units, using their short-latency temporal interactions with other neurons (Barthó et al., 2004; Mizuseki et al., 2009) to a subgroup of neurons. Monosynaptic connections between pairs of units were detected by a non-parametric significance test based on jittering of spike trains as described previously in detail (Fujisawa et al., 2008). Briefly, for each cell pair, each spike in each neuron in the original data set was randomly and independently perturbed (or "jittered") on a uniform interval of $[-5,+5]$ ms, to form a surrogate data set. The process was repeated independently 1000 times to form 1000 such surrogate data sets. Then, the cross-correlograms were constructed for surrogate data sets as a function of latency across the interval $[-20, +20]$ msec. Global bands at acceptance level 99% were

constructed for the cross-correlogram from the maximum and minimum of each jitter surrogate cross-correlogram across the interval [-20, +20] msec. The short latency peak in the original cross-correlogram was determined to be statistically significant (at $p < 0.01$) when the counts in the cross-correlogram were atypical with respect to the upper global band anywhere at the latency [1,5] msec (Fujisawa et al., 2008). Similarly, short latency significant troughs were considered to be due to inhibition when at least one 1 msec bin was significantly depressed ($p < 0.01$) anywhere at the latency [1,5] msec. Approximately 30% of the units included in the statistical analyses were identified by the short-term cross-correlation method. This physiology-based classification method reliably correlated with unit classification based on single spike features (Mizuseki et al., 2009). Units with mean firing rates of less 0.2 Hz were excluded to avoid ‘noise’ in the statistical tests. A total of 618 CA1 putative principal neurons (mean rate 0.575Hz \pm 0.016 SEM) and 111 putative interneurons (mean rate 15.69 \pm 0.90 SEM) were identified and used for analyses.

Firing Pattern and LFP Activity Changes Across Sleep, Within non-REM and Within REM

Firing pattern changes ‘across sleep’ were defined as changes occurring from the first to the last non-REM episode of each sleep session. To examine rate changes within non-REM or REM episodes, the first or last thirds of each episode were concatenated independently for REM and non-REM within each sleep session.

To assess changes in the alternating periods of enhanced and diminished network activity observed during non-REM sleep, pyramidal layer LFP traces were scored for high-frequency ‘activity’ and ‘inactivity’ epochs. For each session, pyramidal layer LFPs

were filtered in the gamma and epsilon frequency bands (30 to 300 Hz), z-scored, rectified and re-filtered between 0.1 and 5 Hz. 'Activity' events were detected as periods in which filtered high-frequency activity exceeded 0.5 standard deviations from the session mean. Similarly, 'inactivity' events were detected as periods in which filtered high-frequency activity was at least 0.5 SD's below the mean. Events occurring within 50 ms of each other were merged independently for inactivity and activity events. Only events lasting >50 ms were included for further analysis. A complementary assessment of the activity and inactivity periods observed during non-REM was performed on the neural spiking data (figure, 2.6, panel b). For this ON and OFF period detection spikes were pooled across all well-isolated cells (Vyazovskiy et al., 2009). Briefly, OFF periods were defined as all epochs at least 50 ms long in which no spikes (from either pyramidal cells or interneurons) were detected. ON periods were periods immediately subsequent to a detected OFF period. Only ON periods between 50 ms and up to 4000 ms long in which at least 10 spikes from any recorded neuron or a combination of neurons were detected were included in subsequent analysis (Vyazovskiy et al., 2009).

Within-ripple firing rates were assessed from concatenated ripple epochs of a given episode (e.g., first non-REM episode of a session). Synchrony was defined as the mean pair-wise correlations between all pairs of pyramidal cells' firing rates binned in non-overlapping 100 ms bins. To detect the firing rate changes within the same state, non-REM and REM episode lengths were normalized, considering each episode length as 100%. The normalized episodes then were divided into three thirds and the values were averaged across episodes.

REM Sleep LFPs Correlate with Firing Rate Decrease

For further analysis spectrograms were normalized independently for each frequency as the z-score from mean non-REM power for that frequency. The correlation between spectral power and firing rate change across REM was assessed independently for each frequency and taken as the correlation between that frequency's power during REM and the mean population firing rate change from the non-REM episode immediately preceding (non-REM_n) to the non-REM episode immediately following REM across all 45 such $\text{non-REM}_n - \text{REM} - \text{non-REM}_{n+1}$ triplets. These correlations were assessed independently for the pyramidal cell and interneuron populations.

Spike-Weighted Spectra (Sp.W.S)

To quantify a neuron's preference to a particular LFP band, we introduced a 'spike-weighted spectrum' method. The steps of the Sp.W.S computation are illustrated in figure 3.8. The firing rates of pyramidal cells (R) during REM were binned in 1-second bins with 0.5 second overlap and concatenated within session across REM episodes. Likewise, CA1 pyramidal layer LFP spectra were taken using the same time bins and concatenated across REM episodes. The concatenated REM time-resolved spectra were z-scored independently for each frequency across bins (i). Thus, we generated two time series, one for the neuron's firing rate and another for the z-scored power for each frequency band. For each frequency band (f), the binned firing rate (R) and the z-scored spectral power (S_z) were multiplied bin by bin to produce S_{wz} . For each frequency band (f), the Sp.W.S. power ($S(f)$) was defined as the sum of S_{wz} ($f, 1:N$) divided by the sum of the binned firing rate (R). Only pyramidal cells with a REM firing rate (R_e) greater than 0.4

Hz (336/618 cells) were included in the Sp.W.S analyses.

Two methods were used to assess the relationship between Sp.W.S. and ‘across-sleep’ (first to last non-REM episode) changes. First, within each sleep session, a partial correlation (ρ) was computed between the neuron’s Sp.W.S. ($S(f)$) and its rate change across sleep (ΔH) normalizing (i.e. ‘partializing’) by each neuron’s mean within-REM firing rate (Re , figure 3.4, panel e). For each Sp.W.S. frequency $S(f)$, partial correlations were assessed within each session ($n = 22$) across pyramidal cells as:

$$\rho_{S(f),\Delta H|Re} = \frac{\rho_{S(f),\Delta H} - (\rho_{S(f),Re})(\rho_{Re,\Delta H})}{\sqrt{1 - \rho_{S(f),Re}^2}\sqrt{1 - \rho_{Re,\Delta H}^2}}$$

Partial correlations were computed independently for either within-ripple or between-ripple change of firing rate across sleep (figure 3.4, panel e). For each frequency, the 95% confidence intervals of the partial correlations were computed across sleep sessions ($n = 22$) via bootstrap analysis. The relationship was determined to be significant when the confidence intervals excluded the baseline (0 correlation).

In addition, for each session we independently pooled the Sp.W.S. of those cells that belonged to either the top (i.e. most positive) or bottom (i.e. most negative) 20% of the distributions of either between-ripple (figure 3.4, panel b) or within-ripple (figure 3.4, panel c) firing rate change. Each of these 4 groups contained 68 pyramidal cells and the bootstrapped 95% confidence intervals were computed. Lack of overlap between the 95% confidence intervals of the top and bottom 20% of rate changers in a given frequency band was considered a significant effect between frequency preference of neurons during REM and their rate change across sleep. The mean ‘across sleep’ rate changes for the top 20% and bottom 20% subgroups are given in the panels of figure 3.7, panel b and c.

Methods for Pre-Play and Replay Analysis:

Animals, Surgery and Data Collection

Three male Long-Evans rats (250-350g) were bilaterally implanted in the dorsal hippocampus with either two 8 shank ($n = 2$) or two 6 shank ($n = 1$) silicon probes. Each shank of the 8 shank silicon probes had 8 sites while each shank of the 6 shank silicon probes had 10 sites. All sites were vertically staggered along the shank with $20\ \mu\text{m}$ spacing between sites. Each site had an area of $160\ \mu\text{m}^2$ and an impedance of 1–3 M Ω . In each rat $50\ \mu\text{m}$ wires were placed gently abutting the left, right mastoid masseter as well as the back neck muscle for electromyographic (E.M.G.) recordings used in sleep classification. All silicon probes were implanted parallel to the septo-temporal axis of the hippocampus. A bundle of six staggered $50\ \mu\text{m}$ wires with a total vertical extent of 1 mm were placed in the left hippocampus to obtain a constant LFP reference used for sleep scoring. Finally, each rat was fitted with a small 3-dimensional accelerometer (ADXL-330, Analog Devices, Mansfield, Texas) which was used to record the animals' movement, or its absence, during sleep. Two stainless steel screws implanted above the cerebellum were used for referencing and grounding. Implantation was performed under isoflurane (1-1.5%) anesthesia as described in detail in Vandecasteele et al., 2012. Each silicon probe was attached to a micromanipulator and lowered over the course of several days until hippocampal layer CA1 was reached as determined by the appearance of hippocampal CA1 sharp-wave/ripples and pyramidal cell activity. Probe placement was histologically confirmed *post hoc*.

Animals were extensively handled both before and after surgery. Water restriction was initiated one week after the surgery; animals were restricted to 90% of their *ad lib* weight and given one day a week of *ad lib* water access. Animals were well acclimatized and recorded in the ‘familiar’ room (where all sleep recordings were performed) for at least one week prior to novelty maze sessions, this time was used to gradually lower the silicon probes into position. All hippocampal, E.M.G. and accelerometer signals were recorded continuously at 20 kHz using three identical 256-channel Ampliplex Systems (Ampliplex, Szeged, Hungary; 16-bit resolution; analog multiplexing; one in the familiar room and one in each of the two novelty rooms). Cell isolation and classification was performed as previously described above (*Unit Clustering and Cell Classification*). In total 322 well isolated pyramidal cells were included in this study (28, 39, 128, 35 and 92 putative pyramidal cells in each of the five novelty sessions respectively). All protocols were approved by the Institutional Animal Care and Use Committees of Rutgers University and New York University.

Behavioral Procedures and Sleep Scoring

In order to acclimatize the rats to running for reward, rats were pretrained to search for water on a geometrically unrelated open-field 'cheese-board' maze (O'Neill et al., 2008) for several days before novelty maze sessions. Once electrodes reached the CA1 pyramidal layer and the animals were well acclimatized running for water reward as well as to the ‘familiar’ room as determined by the observation that the animals engaged in uninterrupted sleep in this room, a ‘novelty’ session was recorded. A novelty session

consisted of a 'Pre' epoch in the familiar room, a novelty run (lasting between 40 minutes to 1 hour) in one of the two novel rooms and 'Post' epoch back in the familiar room. Only one novelty room was used per novelty session. Note also that in recent studies (Dragoi and Tonegawa, 2011, 2013) novelty to the context of the experimental room was established by occluding the experimental room by placing walls around the home cage during 'Pre' and 'Post' epochs. In contrast, in the present study the animals had never been inside of the novel rooms, ensuring that the animals had no experience of the maze context, even fleeting ones during the plugging of the electrophysiological headstages, prior to novelty exposure.

On novelty days the animals were recorded in their home cage in the familiar room and allowed to sleep for two and a half to four and a half hours, constituting the Pre novelty epoch. At the end of the Pre epoch, the animals were transferred to one of two novelty rooms, one housing a linear maze (n = 3 sessions) and the other a circular maze (n = 2 sessions). The linear maze was 1.9 meters long with 15 cm 'reward areas' on either end where water reward was delivered via an automatic infrared-beam triggered system, and a 1.6 cm 'stem'. Rats only received water reward (~0.2 ml) for trials in which they travelled from one reward site to the other. On circular maze (diameter 1m, circumference 3.14 m) rats were made to run in a clockwise direction by manually preventing the counter-clockwise movement until the rats behavior became stereotyped (~10 minutes). Water reward (~0.2 ml) was delivered in a predetermined 30 cm reward area only when the animals had performed a full clockwise run. The novelty sessions were terminated once the animals were satiated and no longer ran for reward. In both the circular and the linear maze the animal's position was monitored by the continuous

tracking at 39.69 Hz of two LEDs (red and blue) positioned 5 cm apart and clipped onto the animal's headmounted Faraday cage. Subsequent to novelty exposure, the animal was transferred back to its homecage in the familiar room and allowed to sleep for three to four hours, constituting the 'Post' epoch. All sleep and novelty recordings were performed during the animal's day-cycle when rats are prone to sleep.

Sleep scoring was performed using the StateEditor and employing hippocampal LFP, accelerometer (movement), and E.M.G. data as previously described (see *TheStateEditor: Behavioral State Scoring and Event Selection* in the *Results* chapter). For each session all Pre and Post epoch analysis was yoked to the duration of the shortest of the two epochs of each session (mean yoked Pre/Post duration: 3.7 hours, minimum duration: 2.2 hours, maximum duration: 4.4 hours, n = 5 sessions). On any given novelty day the animals were exposed to only one of the two novelty conditions (circular or linear maze). Note however, that due to the corruption of the LED tracking information for one animal's circular maze novelty session only two circular sessions are included in the analysis.

Place Field Analysis

For each well isolated principal cell a spike firing-by-position vector was constructed by binning its spikes in non-overlapping 2 cm bins. This vector was smoothed with a 5 cm Gaussian kernel, and divided by the smoothed (5 cm Gaussian kernel) occupancy-by-position resulting in a smoothed position by-firing rate vector. In the case of the circular maze location was linearized and defined as starting at the edge of

reward area, and increasing clockwise, terminating at the opposing edge of the reward area. The hypothesis of place-selective firing was tested by constructing 5,000 null firing rate vectors in which the cell's spikes were randomly sampled using the un-smoothed occupancy-by-position vector as the probability density function, smoothed with a 5cm Gaussian kernel and divided by the smoothed occupancy-by-position vector. The resulting firing rate vectors were, on average, flat with respect to position, while the between null vector variance at each bin reflected the extent of uncertainty of the rate estimate associated with the overall sampling (occupancy) of each bin. A cell was determined to have a place field (and thus, to be a place cell) if at least 5 consecutive bins were above the 99th percentile of their null distributions. If a cell was determined to have more than one place field, only the place field containing the bin with higher peak firing rate was used for further analysis (Diba and Buzsáki, 2007; Dragoi and Tonegawa, 2011; Foster and Wilson, 2006). For each place field, the location of the within-field peak firing rate bin was used as the place field's location (Diba and Buzsáki, 2007; Dragoi and Tonegawa, 2011; Foster and Wilson, 2006). All place field detection analysis was restricted to epochs during which the animals' velocity was at least 5 cm/s and in which the animal were outside of the reward areas. For linear track maze runs all place field analysis was carried out independently for left and right directions of movement (Diba and Buzsáki, 2007; Dragoi and Tonegawa, 2011; Foster and Wilson, 2006).

Population Activity Events

For each session the combined spiking of all recorded CA1 pyramidal cells were binned in 1ms bins and convolved with a 15 ms Gaussian kernel (Pfeiffer and Foster,

2013). For each session a trigger rate was defined as being 3 standard deviations above the mean of all 1 ms bins within NREM epochs of both Pre and Post epochs combined. Population activity events were detected when the smoothed firing rate vector crossed the trigger rate. The beginning and end of the events were defined as the time points at which the convolved firing rate vector returned to the mean of all within-NREM firing rate bins. Only events lasting between 50 to 500 ms, occurring during drowsy/light, NREM or intermediate sleep and in which at least five distinct pyramidal cells each fired at least one spike were considered for further analysis. For each event the center of mass of all pyramidal cell spiking within the event was taken as its time of occurrence for all analysis requiring a point time estimate. Only cells with firing rates between 0.15 Hz and 2.5 Hz and which participated in between 5% to 50% of the Pre and Post epoch population activity events were considered for further analysis. For each cell active in each event a ‘within-event spike timing’ value was assessed as the center of mass of all the spikes that that cell discharged within that event, relative to the center of mass of all the spikes fired by pyramidal cells within that event.

Place Cell Pair Activity and Timing Co-Modulation

As a first order assay of place specific pre-play and replay three pair-wise measures of cell co-activity were defined. All measures were taken independently from population activity events in the pre and post epochs. For consistency, only those events meeting the criteria for ‘*Pre-Play and Replay of Sequential Spiking Activity: Rank-order correlations*’ (next section below) were included in this analysis. Only pairs that were co-

active in at least 10 events in each epoch, and which were detected on different silicon probe shanks, were included in these analyses. *Pair rate correlation*, was taken as correlation coefficient between place cell pair's within-event rate vectors across events. *Participation correlation* was measured as the correlation of the participation vectors, which were binary vectors indicating whether the cell had fired at least one spike in each event, across events. *Final rate correlation in co-active events* was defined as the correlation of firing rates in events in which both place cells of the pair participated (fired at least one spikes).

Similarly three measures of place cell timing co-modulation were taken for place cell pairs which were co-active in at least 10 events in each epoch and which were detected on different silicon probe shanks. *Spike timing correlations* were assessed as the Pearson's correlation between the within event spike timings (see *Population Activity Events* above) across events in which both cells were active. The mean difference of *within-event spike timing* was taken as the absolute value of the mean difference between the within-event spike timings of place cell pairs in events in which they were co-active. Finally, the *variance of the difference of within event spike timing* was taken across events in which both place cells were co-active. Note that since these last two measures are 'signed' (that is, the value of the difference in timing could be either positive or negative), one member of the pair was taken as the 'reference' and the other as the 'comparison'.

These six measures were compared against the distance between the place field locations on the maze. The null distribution for these comparisons was determined by taking the place field distances of 1,000 shuffled place field vectors in which place fields

were resampled without replacement across cells (Diba and Buzsáki, 2007; Dragoi and Tonegawa, 2011; Foster and Wilson, 2006). For linear maze sessions all measures were taken independently for the right and left run directions.

Pre-Play and Replay of Sequential Spiking Activity: Rank-order correlations

For each event rank-order correlation pre-play and replay (Diba and Buzsáki, 2007; Dragoi and Tonegawa, 2011; Foster and Wilson, 2006) of sequential spiking activity was assessed as the Spearman (rank-order) correlation of the within-event spike timing (center of mass) of place cells active in that event against the location of these cells' place field peaks on the maze, across place cells. In order to account for the possibility of both forward and reverse replay, the absolute value of the correlation coefficient was taken as the rank-order sequential activity score. Since a minimum number of place cells must be active in order to assess the spike-timing to place coding correlation, only events in which at least 5, or 10%, whichever was greater, of all place cells were active were included in this analysis. The hypothesis of place sequence selective firing was tested by comparing the observed distributions of sequential activity scores against the distributions obtained from performing the rank-order correlation analysis on 1,000 shuffled place field vectors. For linear maze sessions, this analysis was carried out independently for the left and right running directions.

In order to measure the contribution of individual place cells to the mean absolute rank-order sequential activity score, we took advantage of the fact that for a given event

with n active place cells the rank-order correlation can be given as the dot product of the z-scored place and timing vectors divided by the number of active cells minus one:

$$1) \quad r = \frac{1}{n-1} \sum_{i=1}^n \left(\frac{P_{ri} - \bar{P}_r}{S(P_r)} \right) \left(\frac{T_{ri} - \bar{T}_r}{S(T_r)} \right)$$

Where P_r and T_r are the rank-ordered place field and spike timing vectors of the n cells active in that event, and S is the standard deviation operator. Note that the first step of computing the correlation coefficient is to compute the product of each z-scored element pair. Consequently, the contribution, c_j , of the j^{th} cell to the rank order correlation of a given event is straightforwardly given by:

$$2) \quad c_j = \frac{1}{n-1} \left(\frac{P_{rj} - \bar{P}_r}{S(P_r)} \right) \left(\frac{T_{rj} - \bar{T}_r}{S(T_r)} \right)$$

In order to account for the fact that in the current case it is the absolute value of the correlation coefficient that is of interest this equation must be normalized by the sign of the correlation coefficient:

$$3) \quad ca_j = \frac{1}{n-1} \left(\frac{P_{rj} - \bar{P}_r}{S(P_r)} \right) \left(\frac{T_{rj} - \bar{T}_r}{S(T_r)} \right) \left(\frac{r}{|r|} \right)$$

for non-zero values of r , and zero otherwise. Finally, the contribution of the j^{th} cell to mean of N absolute correlation coefficients $N^{-1} \sum_1^N |r_N|$ is given by:

$$4) \quad cam_j = \frac{\sum_{e=1}^N ca_{j,e}}{N}$$

Notably the sum all such absolute correlation contribution (*cam*) values for all n place cells equals the mean of the absolute correlation coefficients of the N correlations to which they contribute:

$$5) \quad N^{-1} \sum_{e=1}^N |r_e| = \sum_{j=1}^n cam_j$$

This measure was used to assess the possible biases in the influence of individual cells over the observed distribution of sequential rank-order correlations.

Pre-Play and Replay of Sequential Spiking Activity: Paired-Latency Method

In order to assess the correspondence of the sequential place activity on the maze and within event sequential timing in a non-biased manner (see *Results*) a parallel, pairwise and computationally efficient approach was developed. For a given event with n active and inactive cells, a paired latency vector PL of length $(n)(n - 1)/2$ corresponding to the number of unique non-self pairs was constructed. For each comparison one cell was chosen as the *reference* and the other as the *comparison* – the choice of which cell of the pair is the reference and which the comparison is arbitrary but must be applied consistently across all events and conditions. Any given element of the paired latency vector PL was assigned one of three possible values:

- **1** if the comparison cell's within-event center of mass of spiking (c.o.m.) occurred *after* the reference cell's within-event c.o.m.
- **-1** if the comparison cell's within-event c.o.m. occurred *before* the reference cell's within-event c.o.m.

- **0** if the two members of the pair had the same within-event c.o.m. *or* if either member of the pair did not fire in the event

Paired distance, PD , vectors were similarly constructed from the relationship between linearized place field peaks such that each element of the PD vector had one of three possible values:

- **1** if the comparison cell's place field was to the *right* of the reference cell's place field
- **-1** if the comparison cell's place field was to the *left* of the reference cell's place field
- **0** if place field peaks of both members of the pair occurred on the same bin *or* if either member of the pair was not a place cell

For linear maze sessions separate PD vectors were constructed for left and right directions of movement. Note that starting from these vectors it is possible to fully reconstruct rank-ordered sequence content.

The similarity of within-event sequence to the order of place fields on the maze was calculated as the dot product of PD and PL divided by the dot product of the absolute values of PD and PL :

$$6) \quad Seq_{PL} = \frac{PL \cdot PD}{|PL| \cdot |PD|}$$

In other words, the sum of all the non-zero elements which had the same value in both PL and PD normalized by the sum of the elements that were non-zero in *both* PL and PD . Similarly to the Spearman correlation coefficient this value is a measure of order

similarity and is bounded between -1 (when one sequence is the reverse of the other) to 1 (when the two sequences being compared are the same). However, it is different in at least three salient ways 1) the absence of a z-scoring step reduces the observed bias towards a disproportionate contribution of cells near the edges of the maze to the overall sequence similarity effect, 2) the pair-wise relationship of cells with same c.o.m. or place field peak location values have no contribution on the overall sequence similarity value, and 3) more generally all 0 valued pairs in either the PL or PD vectors do not contribute to the net result, eliminating the necessity for the event by event selection of active place cells and thus increasing the computational efficiency of the analysis.

Paralleling the rank-order analysis, only events in which 5 or 10%, whichever was greater, of the place cells were active were included in the analysis and the absolute value of the paired latency score was used as the final measure of sequence similarity strength.

In order to assess the contribution of individual cells to the overall distribution of absolute PL sequence similarity scores a similar analysis as was carried out in the rank-order case was performed. The contribution k of the j^{th} cell to a particular paired latency comparison is given by:

$$7) \quad k_j = \frac{PL_j \cdot PD_j}{2 \cdot |PL| \cdot |PD|}$$

Where PL_j and PD_j are the sub-vectors consisting of the elements of PL and PD respectively in which the j^{th} cell was used as either the reference or the comparison. The division by two is necessary to avoid the effects of counting each pair twice (once by the reference cell, the other by the comparison cell). The j^{th} cell's contribution to the absolute value of Seq_{PL} is then given by

$$8) \quad ka_j = \left(\frac{PL_j \cdot PD_j}{2|PL| \cdot |PD|} \right) \text{sign}(PL \cdot PD)$$

And the j^{th} cell contribution to the mean of the absolute value of N paired latency scores is given by:

$$9) \quad kam_j = \frac{\sum_{e=1}^N ka_{j,e}}{N}$$

Again, the addition of the by-cell contribution kam_j across all n place cells equals the mean of the absolute value of the paired latency similarity scores across N events:

$$10) \quad N^{-1} \sum_{e=1}^N |Seq_{PL}| = \sum_{j=1}^n kam_j$$

a result which was confirmed empirically.

Where indicated, within-shank pairs were excluded from the analysis by setting their values in the PD vectors to 0.

Pre-Play and Replay: Bayesian Decoding

Memory-less, smooth-prior probability, Bayesian classifiers of position (Davidson et al., 2009) given a time-binned place-cell population spiking vector $spikes$ were constructed from the smoothed firing rate-by-position vectors as:

$$11) \quad Pr(pos|spikes) = \left(\prod_{i=1}^n f_i(pos)^{sp_i} \right) e^{-\tau \sum_{i=1}^n f_i(pos)}$$

Where $f_i(pos)$ is the value of the firing rate-by-position vector of the i^{th} cell at position pos , sp^i is the number of spikes fired by the i^{th} cell in the time bin being decoded, and τ is the duration of the time bin. Posterior probability were subsequently normalized to one:

$$12) \quad Pr(pos|spikes) = \frac{Pr(pos|spikes)}{\sum_{i=1}^{P_n} Pr(pos_i|spikes)}$$

Where P_n is the total number of positions (2-cm bins). All Bayesian analysis was performed independently for the left and right direction of movement.

Five-hundred null Bayesian classifiers were constructed by independent random circular rotations (Davidson et al., 2009) of each cell's unsmoothed firing-rate by position vector and subsequently smoothing with a 5 cm standard deviation Gaussian kernel. This approach was chosen because it selectively disrupts spatial tuning while preserving each cell's firing rate. Smoothing was performed after circular rotation to ameliorate the rate discontinuities expected at the edges of the rotated firing rate-by-position vectors.

In order to test the performance of the classifier's estimates of position, population spiking was binned in non-overlapping 500 ms bins during maze periods in which the animal's velocity was at least 5 cm/s and was not in a reward area. The method was further cross-validated by omitting every fifth lap from the construction of firing rate-by-position vectors and subsequently using these vectors to decode position on these excluded laps.

In order to assess Bayesian decoded pre-play and replay effects the spiking of place cells in population activity events were further divided into 20 ms non-overlapping bins (Davidson et al., 2009). This size of bin was chosen as a compromise between the temporal coarseness of decoded sequential activity, and the number of spikes (and thus the available information content) expected per bin. Bins in which no spiking was detected were excluded from the analysis. Only those events which met the criteria for inclusion in the '*Sequential Spiking Activity*' analysis (see above) and which had at least

five non-zero spiking bins and which showed a decoded trajectory span (see below) of at least 30 cm (Davidson et al., 2009), were included in this analysis. For each bin the decoded position was established as the peak of that posterior probability density of position. The decoded trajectory span was taken as the maximal pair-wise distance between all pairs of decoded positions within an event. For each qualifying population activity event two measures of pre-play and replay were taken. First, Bayesian decoding quality was assessed as the mean of the posterior probability across within-event bins of the peak posterior probabilities of the bins within the event. Second, Bayesian decoded sequence strength was assessed across the within-event decoded positions by using a modified version of the paired latency vector. Bayesian paired distance (PD_{Bayes}) vectors were constructed as described above and compared against a reference paired distance vector (PD_{Ref}) assuming a constant trajectories (that is, constantly increasing position) as:

$$13) \quad Seq_{Bayes} = \frac{PD_{Bayes} \cdot PD_{Ref}}{|PD_{Bayes}| \cdot |PD_{Ref}|}$$

Note that this measures the overall directionality of the decoded trajectory, and in contrast with previous techniques (Davidson et al., 2009; Kloosterman, 2012) does not test the linearity of the trajectory.

Results

TheStateEditor: Behavioral State Scoring and Data Visualization and Annotation Software

Loading and Pre-processing files

TheStateEditor script is a fully stand-alone Matlab function for behavioral state scoring and large-time scale data visualization. In order to use this program, the user need only obtain the script (available upon request) and add it to his or her Matlab 'path'. For users whose data is in the format used by our lab (binary/ASCII file format, Hazan et al., 2006), the user navigates his or her Matlab command prompt to the directory housing the required files: an '*.xml*' parameter file containing information about channel numbering and order as well sampling rate, an '*.eeg*' or '*.lfp*' file which contains the (typically 1250 Hz) LFP data, as well as an optional '*.whl*' LED tracking file (see Hazan et al., 2006 for a full description of these formatting conventions). *TheStateEditor* automatically detects the name of the relevant session from the first '*.xml*' file in the folder and loads the relevant session information. If the directory contains multiple '*.xml*' files the name of the session of interest can be entered manually. Once the correct files are found, the channel selection window is launched (figure 2.1). Here users are asked to choose up to three EEG/LFP channels for visualization. The chosen channels are then loaded into memory, whitened and used for the computation of time-resolved multi-tapered spectrograms (Sirota et al., 2003) in non-overlapping one second bins. This time resolution was chosen as being coarse enough not to cause graphics memory problems when displaying the full spectrograms over long sessions, while still being fine enough

for thorough behavioral state classification. The spectrograms are then \log_{10} -transformed to yield spectrographic power in decibels, and then truncated between their first and 99th percentiles (across all time bins and frequencies) in order to suppress outliers. The channel selection window also prompts the user to optionally select a motion signal. Three motion signal types are supported: 1) head tracking of head mounted LED's in our lab's native *.whl* format, 2) accelerometer or kinetic motion pad data which detects the animal's movements, or 3) MEG channel data. In each case, the motion signals are transformed into one dimensional motion amplitude signals by rectifying, filtering between 0.1 to 10 Hz, and binning into non-overlapping one second bins. Motion outliers were similarly suppressed by truncating the motion signal between its first and 99th percentile. Conversely, users may input motion data directly from *.mat* files. This is useful for those who wish to perform costume pre-processing of motion data, or who wish to use the motion panel in order to display a different type of data altogether. Note that when two channels of a long (~9.8 hour recording) are loaded into the *TheStateEditor* all the pre-processing steps take about three minutes to complete on a standard windows desktop computer.

Users who use different file format must first load (up to three) LFP or EEG channels of interest into Matlab and pass them, as well as other information such file name and sampling rate, into the *TheStateEditor* as part of a structure whose field names are specified in *TheStateEditor*'s help section.

Once the spectrograms are created and the motion signals are processed, *TheStateEditor* saves an auxiliary file (with the default suffix *.eegstates.mat*). This auxiliary file contains the processed spectrograms and motion signal, as well as other

session information such as channel number and sampling rate. In order to reduce redundant disk usage, by default the raw LFP signals are not saved in this auxiliary file and must be loaded from the `‘.eeg’` or `‘.lfp’` file on each use. However, when calling *TheStateEditor* the user may toggle the `‘MakePortable’` flag (which is set to *off* by default), in which case the raw LFP channels will also be saved into this auxiliary file. Importantly, in this latter case, the auxiliary file is fully independent of the original data files, and can be used for launching *TheStateEditor* from different computers or for data sharing. Note that the next time *TheStateEditor* is launched from this directory, if a `‘.eegstates.mat’` is detected it will load all the processed data directly from this auxiliary file, eliminating the need for repeated pre-processing.

TheStateEditor: Data Visualization and Navigation

Once pre-processing is complete and the auxiliary file is saved, the *TheStateEditor*’s main panel is launched (figure 2.2). This is the main window for visualizing and interacting with the data. The main display utilizes two different time-frames, the spectrograms and the motion display (which each have a resolution of one second) display the entire session when the main panel first opens, and are appropriate for visualizing tens to thousands of seconds at a time. The raw LFP displays have a default length of two seconds, appropriate for the visualization of LFP features – but can be resized by the user from 0.5 to 60 seconds. Note that the epochs displayed in the spectrograms and the motion signal display are locked to each other such that they always show data from the same time frame. The LFP display on the other hand, displays the

LFP traces corresponding to the center of the spectrogram displays and is synced automatically.

When first using the *TheStateEditor*, users can press the ‘*H*’ key to load a pop-up screen displaying the different commands (figure 2.3). Users can navigate through the data in one of several ways. Pressing the right or left arrow keys incrementally progresses the display forward or backward, respectively, by 20%. Conversely, users can click and hold any of the displays to ‘drag’ them forward and backward. Note that the syncing between the various panels is automatically preserved in all cases. Zooming in and out of the temporal viewing pane may be achieved either by scrolling the mouse wheel forward or backward, by double-clicking on the display, or by entering the ‘Zoom’ functionality by pressing ‘*Z*’. Users can also jump to a given point or set the temporal extent of the viewing window by manually entering values into the ‘*Go to second*’ or ‘*Window length*’ input boxes on the right side of the main display. Note that in each of these cases, only the temporal extent (x-axis) is affected by zooming. The extent of the color limits for each spectrogram window is manipulated by first clicking on the spectrogram window and then pressing the up (for ‘warmer’ colors) or down (for ‘cooler’ colors) arrow keys. In order to change the extent of the frequency axis (the y-axis of the spectrogram displays) users enter the ‘*Frequency rescale*’ function by pressing ‘*F*’ and then change the extent of displayed frequencies (from 20 to 200 Hz) by pressing the up and down arrows. The default frequency extent shown is from 0.4 to 40 Hz. Note that the lowest frequency (0.4 Hz) is always used as the lower bound of the spectrogram displays.

Note that, while by default, spectrogram data is smoothed (within frequency) with a Gaussian window with a standard deviation of 10 seconds, users may select 0, 10, 15,

20, 30, 45, or 60 second smoothing windows from the ‘*Smoothing Window*’ drop down menu, causing the spectrograms to be dynamically re-smoothed at the desired coarseness. Finally, users may choose to overlay data onto the spectrogram displays through the ‘*Overlay display*’ drop down menu on the right of the main panel. A default overlay type (thick white lines in figure 2.2), the theta (5 to 10 Hz) to delta (0.4 to 4 Hz) ratio, may be selected. Conversely, users can choose the ‘*Choose from file*’ option from the overlay display in order to load costume data. All overlay displays are fitted to the top half of each spectrogram window as shown in figure 2.2.

TheStateEditor: Behavioral State Scoring and Event Selection

Behavioral state scoring using *TheStateEditor* is performed by labeling using five number and color coded possible behavioral states.

- 1) The awake state, color-coded black, is characterized in the hippocampus by the theta (5 to 10 Hz) oscillation, and behaviorally by active movement and EMG activity (Gervasoni et al., 2004; Gottesmann, 1992).
- 2) The drowsy/light sleep stage, color coded yellow, is characterized in the hippocampus by low spectral power with characteristics intermediate between the synchronized and de-synchronized state (Gottesmann, 1992), the presence of sharp-wave/ripples and is behaviorally characterized by a small, though present motion and sustained EMG activity.
- 3) NREM (non-REM), color coded blue, is characterized by large amplitude low frequency spectral activity and sharp-wave ripples in the hippocampus and

behaviorally characterized by the lack of motion, interrupted by occasional small movements or twitches, and a low EMG signal (Crunelli and Hughes, 2010; Steriade and Amzica, 1998; Steriade et al., 1993, 2001).

- 4) Intermediate sleep, color coded green, is a short (~1 to 3 second) sleep state typically observed before REM episodes in which hippocampal spectral shows both low-amplitude theta activity as well as pronounced spindle-frequency (10 to 20 Hz) spectral power, possibly due to volume-conduction from neocortical spindling (Glin et al., 1991; Gottesmann, 1992; Gottesmann et al., 1998).
- 5) REM sleep, color-coded red, is characterized in the hippocampus by persistent theta activity in the virtual absence of motion or EMG activity (Aserinsky and Kleitman, 1953; Jouvet, 1967).

In order to label a given epoch as a given behavioral state, users press the numeric key associated with that behavioral state. The first click (which can be made on any of the data displays) sets the first bound ('edge') of the labeled state, while the second click sets the second bound, and completes that state's selection. Mistakes in state scoring can be undone by overriding with state '0' (no state) or by using the '*Undo State*' or '*Redo State*' buttons at the bottom right of the main display. Behaviorally scored epochs are automatically color labeled on the '*state ribbon*' near the top of the display. Note that in the *state ribbon* (as well as the '*State vector*' output explained below) - all selections are rounded to the nearest second bin and are mutually exclusive (one bin cannot be labeled as two different states). However, *TheStateEditor* also keeps track of the fine-scale information of the precise time at which each state epoch's edges were selected in a '*transition matrix*'. Thus if, for instance, a user desires to terminate an NREM episode at

precisely the time of the last observed sharp-wave/ripple as observed in the LFP displays, he or she can terminate the NREM epoch selection at this time point (chosen from the LFP displays) and recover it from the precise timing from the *transition matrix*, even if for the purposes of the *state ribbon* this time point will be rounded to the nearest second.

Finally, users may add or load up to ten classes of discrete, fine temporal resolution events. Event classes (labeled with the numbers 1 to 10) are chosen from the '*Event #*' drop down menu on the right side of the main panel (figure 2.2). Once an event class is chosen users can add discrete events by pressing the '*E*' key and clicking on a time point, or delete particular events by pressing the '*D*' key and clicking on a particular event. Events are shown as vertical pink dashed lines, and only those events from the selected event class are shown at a given time. Events are defined with high temporal resolution and may be added directly to the LFP display, allowing for a great deal of flexibility in data annotation.

Saving scored behavioral states and selected events is accomplished by pressing the '*S*' key which brings up *TheStateEditor's* saving console (figure 2.4, panel a). This console allows the user to choose the name of the saved output as well as which features to save. The *StateVector* is a vector of n second bins where each entry corresponds to a labeled state (none-labeled bins are labeled '0'). The users may also optionally choose to save a matrix with all the labeled events, as well as separate structure containing the event selection history for that session. States and events can similarly be loaded (by pressing '*L*') from *TheStateEditor's* loading console.

TheStateEditor has been observed to work stably in Matlab running on Windows, Macintosh, and Linux platforms and is now used in our lab for behavioral sleep scoring

as well as other large-time scale data analysis, such as examining the long-lasting spectral effects of various kinds of optogenetic stimulations, or the tracking of firing rates over session-wide time scales.

REM Sleep Reorganizes Hippocampal Excitability

Excitability changes across sleep

Local field potentials (LFP) and spiking activity of isolated CA1 putative pyramidal cells and putative interneurons were recorded in the home cage while the rat was immobile and assumed a characteristic sleep posture. The ratio of theta (5-11 Hz) and delta (1-4 Hz) power was used to identify non-REM and REM episodes (figure 3.1, *Methods*), as described previously (Montgomery et al., 2008). Twenty two sleep sessions (38.2 min \pm 5.8 S.E.M) with at least one non-REM--REM--non-REM cycle were recorded in 5 rats. Mean firing rates of pyramidal cells (n=618) were similar between non-REM and REM episodes, whereas firing rates of interneurons (n=111) were significantly higher during REM ($p < 0.00018$; sign-rank test; (Csicsvari et al., 1999b). In the majority of our analyses, we focused on the following comparisons. First, changes ‘*across sleep*’ were defined as differences between the first and the last non-REM episodes in a sleep session. Second, changes ‘*within non-REM*’ episodes refer to differences between the first and the last thirds of each non-REM. Third, changes ‘*within-REM*’ episodes refer to differences between the first and the last thirds of each REM. Finally, we examined the relationship between these categories.

Since non-REM sleep is characterized by alternating periods of population activity and inactivity in both the neocortex (Steriade et al., 1993) and hippocampus

(Isomura et al., 2006; Ji and Wilson, 2007), we defined active periods as those in which smoothed gamma and epsilon band (30 to 300 Hz) LFP activity was at least 0.5 S.D's above the mean for at least 50 ms. Conversely, inactive periods were detected as those in which gamma and epsilon band activity was 0.5 S.D's below the mean for at least 50ms (see figure 3.6 for an analogous spike-based analysis). The incidence of active periods decreased, whereas the incidence of inactive periods increased significantly from the first to the last non-REM episodes of each session (i.e., across-sleep; figure 3.1, panel b; table 2). The firing rates of both pyramidal cells and interneurons decreased significantly across sleep (figure 3.1, panel b). These findings are in accord with the two-process model of sleep and indicate similarities between sleep-related activity of neurons between the neocortex and hippocampus (Borbély, 1982; Tononi and Cirelli, 2006b; Vyazovskiy et al., 2009).

During sleep hippocampal neurons fire in population synchrony during sharp-wave ripple events and relatively asynchronously between ripples (Buzsáki et al., 1992). The discharge rate of pyramidal neurons between ripples decreased significantly across sleep (figure 3.1, panel b), similar to the decrease in global firing rate. Conversely, the mean firing rate of pyramidal cells within the short-lived ripple events increased during the course of sleep (figure 3.1, panel b). This increase in ripple-related activity across sleep was the result of an increase in the percentage of ripples within which pyramidal cells participated (i.e., fired at least one spike) rather than an increase of the within-ripple firing rates of individual neurons in individual ripples (figure 3.1, panel b, figure 3.7). Concurrent with the increase of within-ripple participation, the coefficient of variation of within-ripple firing rate across cells decreased (figure 3.1, panel b, figure 3.7), suggesting

that the within-ripple participation was more evenly distributed across the population of pyramidal cells at the end compared to the beginning of sleep. Synchrony, as measured by the correlation strength of pyramidal cell pairs in non-overlapping 100 msec bins (Wilson and McNaughton, 1994) also increased across sleep (figure 3.1, panel b), likely due to the more consistent participation of pyramidal cells in ripples. In short, the decreased firing rate across sleep was associated with a ‘paradoxical’ increase in pyramidal cell synchrony and more consistent recruitment of spikes to ripple events (table 2).

Excitability changes within non-REM and REM episodes

Next, we investigated which sleep state might be responsible for the global ‘across-sleep’ changes of firing patterns. Since the duration of individual non-REM and REM episodes vary, their lengths were normalized (see *Methods*) and the pattern of changes within episodes were quantified. In non-REM episodes, we found that firing rates significantly increased between the first and last thirds of the episodes, both in pyramidal cells ($p < 1.99E-14$, $n = 618$) and in interneurons ($p < 4.6E-5$, $n = 111$) (figure 3.2, panel b, figure 3.6). Other measures, such as incidence of active and inactive epochs, the percentage of ripples in which pyramidal cells participated and population synchrony, as measured by pyramidal cell pair-wise correlations, also showed significant and opposite changes within non-REM compared to those observed across sleep (figure 3.2, panel b). In contrast, firing rates significantly decreased within REM epochs, both in pyramidal cells ($p < 0.012$, $n = 618$) and in interneurons ($p < 1.23E-5$, $n = 111$) (Figure 3.2 panel a, figures 2.6).

LFP spectral changes across sleep and within non-REM and REM episodes

In addition to unit firing, the LFP spectral changes across sleep were also calculated. For each sleep session, the LFP spectra in individual non-REM and REM episodes, recorded from the CA1 pyramidal layer, were normalized independently for each frequency by the power of concatenated non-REM episodes and expressed as a z-score. Spectral power decreased significantly in a broad range of frequencies (4-50 Hz) across sleep (i.e., from the first to last non-REM episode; figure 3.3, panel a; $n = 22$ sleep sessions; change in 0-50 Hz integrated power, $p < 0.0024$; sign-rank test). In contrast, a significant increase in power (0-50 Hz) was present within non-REM episodes (Figure 3.3, panel b; $n=82$ non-REM episodes $p < 2.11E-9$; sign rank test). Within REM episodes, a power decrease was observed in the theta-beta (5-20 Hz) and lower gamma (40-50 Hz) band (Figure 3.3, panel c; $n=45$ REM episodes; 0-50Hz power; $p < 2.85E-4$; sign-rank test). Changes in the delta band (1-4 Hz) may reflect changes in the hippocampus or volume-conducted LFP from the neocortex (Isomura et al., 2006; Wolansky et al., 2006).

Relationship between non-REM and REM sleep

Since the evolution of firing patterns and LFP across sleep was similar to those observed within-REM sleep but dissimilar to the changes observed within non-REM episodes, we examined how REM episodes might contribute to the overall reorganization of firing patterns during the course of sleep. The mean firing rate decrease of both the pyramidal cell and interneuron populations from the non-REM episode preceding a REM (non-REM_n) to the non-REM episode following a REM episode (non-REM_{n+1}) was significantly correlated with the theta power of the interleaving REM episode but not the power of other frequencies (figure 3.4, panels a and b), except for the lower gamma band

for pyramidal cells. Similar calculations were performed to examine the relationship between population synchrony (pairwise correlation) during non-REM and spectral power of REM. The increase in synchrony of both pyramidal cells and interneurons from non-REM_n to non-REM_{n+1} was significantly correlated with the theta and gamma (around 40 Hz) power of the interleaving REM episode but not the power of other frequencies (figure 3.4, panels c and d).

To examine how the rate change of individual neurons across sleep was related to their network pattern-related activity during REM sleep, we introduced the method of spike-weighted spectra (Sp.W.S.) by relating the instantaneous firing rates of single cells to the power distribution of the simultaneously detected LFP. LFP spectra and firing rates of individual pyramidal cells were computed in 1-second bins with 0.5 second overlap during REM (figure 3.8, *Methods*). For normalization purposes, the LFP spectrograms were z-scored independently for each frequency band and the LFP power spectrum was multiplied bin-by-bin by the neuron's within-bin firing rate and divided by its overall REM rate (figure 3.8). Since power in each frequency of Sp.W.S is first z-scored, stochastic firing results in power nearing zero, while positive values for a given Sp.W.S. frequency band reflect a cell's selective firing preference in that band. To quantify the relationship between the neuron's frequency preference during REM sleep and its firing pattern change across sleep, the correlation between the neuron's Sp.W.S. in REM and its rate change between the first and last non-REM episodes of sleep was normalized by the neuron's REM mean firing rate (see *Methods* 'partialization' procedure). These partial correlations were computed separately for changes occurring across sleep in either within-ripple or between-ripple firing rates (n=22 sleep sessions). Pyramidal cells with

firing rates less than 0.4 Hz during REM (n=281 of 618 cells) were excluded from the Sp.W.S. analysis. The Sp.W.S analyses (figure 3.4, panel 8; see also figure 3.8) demonstrated that within the same population of simultaneously recorded pyramidal cells the across-sleep decrease of between-ripple firing rate was correlated with the pyramidal neurons' preference to discharge selectively during high power theta (~5 - 10 Hz) and gamma epochs during REM. Similarly, a neuron's theta and gamma power preference reliably predicted its across-sleep firing rate increase within ripples (figure 3.4, panel e).

Pre-Play and Replay

Pre-Play and Replay: Population Activity Events

In order to assess the specific and non-specific contributions to the establishment and replay of hippocampal spatial coding of novel maze experiences, population activity events (see *Methods: Population Activity Events*) were detected during drowsy/light, NREM and intermediate sleep during the Pre and Post novel maze exploration epochs. An average of 4,408.4 (st.d. $\pm 1,324.43$) Pre events and 4,463.6 (st.d. $\pm 1,494.24$) Post events were detected in an average of the 9,521.2 (st.d. $\pm 3,006.51$) seconds and 10,897 (st.d. $\pm 2,749.84$) seconds the animals spent in drowsy/light, NREM and intermediate sleep during the Pre and Post epochs respectively (n = 5 novelty sessions). Notably, in the Post epoch it took the animals an average of 629.58 seconds to settle into drowsy/light sleep (minimum latency to drowsy/light sleep: 269 seconds, maximum latency to drowsy/light sleep: 1394 seconds, n = 5 Post epochs) and an average of 1620.18 seconds to reach NREM sleep (minimum latency to NREM: 906.3 seconds, maximum latency to NREM: 2145 seconds). This is contrast to other studies which have reported to show

sleep-related effects in the first 600 to 900 (10 to 15 minutes) of Post epoch recordings (Kudrimoti et al., 1999; Wilson and McNaughton, 1994).

While the numbers of detected population activity events were similar between the Pre and Post epochs, the composition of these population activity events was found to change. Notably, within population-activity event firing rates increased significantly from the Pre to Post epochs (figure 4.2, panel a, Pre mean: 2.87Hz, Post mean: 2.94Hz, $p \approx 0$, rank-sum test), a trend which was observed in all sessions and significant in four of them (figure 4.2, panel b). Likewise the percentage of pyramidal cells participating (firing at least one spike) in each event was significantly higher during Post compared to Pre for the population of events (figure 4.2, panel c, Pre mean: 24.42%, Post mean: 25.59%, $p \approx 0$, rank-sum test), a trend observed in 4 sessions and significant in 3 of them, with 1 session significantly showing the opposite effect (figure 4.2, panel d). Finally, while the population average showed a slight, but significant, increase in event duration from the Pre to the Post epochs (figure 4.2, panel e, Pre mean: 175.15 ms, Post mean: 178.97 ms, $p < 0.0135$, rank-sum test), this trend was not consistently observed across sessions (figure 4.2, panel f). Consequently, any experience-specific changes in hippocampal CA1 firing from the Pre to Post epochs, must be understood within this wider context of increases in excitability.

Place Cell Pair Activity and Timing Co-Modulation

As a first order confirmation of the validity of our data set in assessing the replay (and by extension, pre-play) phenomenon we measured *pair-wise co-activity* and *spike*

timing co-modulation as a function of the distance of place fields on the novel maze. This approach was chosen as it is closely related to the earliest measures of replay and is a particularly statistically robust measurements of the replay phenomena (see *Introduction*). This analysis was carried out independently for the linear ($n = 7,738$ place cell pairs) and circular ($n = 3,802$ place cell pairs) maze conditions. It was observed that in accordance with previous measurements of pair-wise reactivation of place cells with overlapping place fields (Kudrimoti et al., 1999; Wilson and McNaughton, 1994), elevated correlations were observed between the within-event firing rate vectors for pairs of place cells with place field peaks that were within 20 cm of each other on the novel maze (i.e. strongly overlapping pairs; figure 4.3, panels a and b). This effect was only observed during Post epoch events and was also true for the correlation between these place cell's activity vectors (figure 4.3, panels c and d). Notably, it was also found that place cell pairs with peaks more than 140 cm apart (i.e. non-overlapping pairs), showed a pronounced deficit in correlation strength in both of these measures, again, only during the Post epochs. While we also measured the modulation strength of rate correlations only within events in which both place cells were active, only the circular maze condition showed this effect (figure 4.3, panels e and f), and this measure is included in order to illustrate some of the observed across condition variability.

Similarly we conducted pair-wise measures of within-event place cell *spike timing co-modulation* by place field peak distance. Consistent with previous findings (Qin et al., 1997; Skaggs and McNaughton, 1996) and the pair-wise co-activity effects described above, we found striking co-modulation of spike-timing as a function of place field peak distance on the maze. The spike timing of pairs of place cells with nearby place fields

tended to be highly correlated across events (figure 4.4, panels a and b). Moreover, these place cells with nearby place fields tended to fire in closer temporal proximity to each other (figure 4.4, panels c and d) and the latency between their firing was subject to reduced variability (figure 4.4, panels e and f). With one exception (figure 4.4 panel d) the inverse of these effects, that is, reduced temporal correlation, increased spike timing latencies, and increased spike timing variability, was seen for place cells with far (~140 cm or greater) place fields peaks on the novel maze. Notably, the appearance of robust and specific pair-wise activity and timing changes in Post epoch events is strongly supportive of the presence of the replay phenomenon in our data set. By contrast, activity and spike timing in Pre epoch events was rarely, and then only very weakly, modulated by place field peak distances on the subsequent exposure to the novel maze. Notably, this pre-play-like activity, when observed, was limited to place cell pairs with place field peaks within 20 cm of each other on the linear maze, and was in all three instances opposite in direction to the Post epoch effects (figure 4.3. panel c and figure 4.4 panels a and e). Thus, through the use of pair-wise measure it was established that classical replay was present in our data set, while pre-play was found to be nearly negligible under this pair-wise approach.

Pre-Play and Replay of Sequential Spiking Activity: Rank-order correlations

Next we sought to examine the event-resolved characteristics of pre-play and replay using an established measure in which pre-play and replay strength is taken as the rank-order correlation between the sequence of place fields of cells on a maze to their

temporal sequence of activation during offline events (Diba and Buzsáki, 2008; Dragoi and Tonegawa, 2011, 2013; Foster and Wilson, 2006). Note that since these correlations are only taken over the place cells active in each event, only events in which at least 10% or 5, whichever was higher, of the isolated place cells fired at least one spike each were included in this analysis (52.0% (\pm 27.7% st.d.) of Pre events and 57.2% (\pm 28.0% st.d.) Post events, $n = 5$ Pre and 5 Post epochs, respectively). The results of this analysis were compared (rank-sum tests, significance threshold: $p < 0.025$) with a null (shuffle) distributions derived by randomly resampling the location of place fields on the maze without replacement (that is, random re-assignment of place field peaks to place cells) 1,000 times and recalculating each event's sequence score for each shuffle. In order to account for both forward and reverse pre-play or replay the absolute value of each of the correlation coefficients was taken as the sequence correlation score. Using this measure, we observed only two cases of significant Pre versus Post sequence score differences, and each of these showed effects in opposite direction (figure 4.5, panel a). In contrast to the recent literature concerning the pre-play phenomena and also utilizing the rank-order correlation methodology, the observed sequence correlation scores were found to be higher than the null in only three of five Pre epochs, and significant in only one of these (figure 4.5, panel b). Finally, and perhaps most surprisingly, given the limited amount known from the literature (see *Introduction*, Dragoi and Tonegawa, 2011, 2013) and our findings concerning pair-wise replay effects, though all five Post epochs showed absolute sequence correlation scores higher than their associated shuffled distributions, this effect was only significant in one Post epoch (figure 4.5, panel c).

In order to examine the origin of the observed rank-order pre-play and replay effects (or their absence), we calculated the per-cell contribution to the observed overall mean of the absolute rank-order correlation values across the events of each Pre and Post epochs (see *Methods*). It was found that those cells with place-fields near the ends of the maze contributed much more heavily to the mean absolute rank-order sequence effect than those cells near the center of the maze (figure 4.6, panel a). While there is physiological evidence to support selective coding for salient maze features such as rewards (around which the ends of the mazes were defined (Diba and Buzsáki, 2008; Gothard et al., 1996; Hollup et al., 2001; Redish et al., 2000)), the same effect was found in the per-cell contribution distributions of 100 shuffled place cell vectors (figure 4.6, panel b). Thus, while there may be relevant saliency effects at play in determining per-cell contributions to the overall sequence correlation, it was found that the rank-order correlation methodology itself, when applied to the estimation of sequence similarity is mostly determined by the activity of those cells with place fields near the end of the maze. This is in line with our intuitions of the working of correlation coefficients in that it is known, for instance, that one outlier can have a larger effect on the position of the line of best (and thus the resulting correlation coefficient) than a large number of points clustered near each other (Gideon and Hollister, 1987). Conversely, this effect can be thought of algebraically in that, since z-scoring is the first step in the taking of a correlation coefficient (see *Methods*), values that tend to fall into the middle of the distribution will have a z-scores near 0, and a similarly small contribution to the overall correlation coefficient.

Pre-Play and Replay of Sequential Spiking Activity: Paired-Latency Method

An alternative method for event-resolved sequential pre-play and replay estimation was developed as an event-based version of the method of Skaggs et al., 1996. Briefly, in the *Paired Latency* (P.L.) method, for each event a template was constructed in which each element indicated the relative order (i.e. before or after) of a given pair of active cells. Similarly, *Paired Distance* vectors were constructed from the paired relative positions of place fields on the novel maze. Note that these vectors constitute an expansion of the original sequence into pair-wise elements and thus contain all the rank-ordered information of the original sequence. Importantly, while each element of these templates conveys information about a particular pair-wise interaction, the templates were assessed on a per-event basis. For a given event, a P.L. sequence similarity measure was taken between its P.L. vector and a P.D. vector as the number of similarly ordered co-active pairs minus the number of differently ordered pairs, normalized by the total number of co-active pairs. Note that only cells that were both active in both event and that had place fields (and that thus contributed to the P.L. and P.D. vectors, respectively) influenced the overall measure. Consequently, the same criteria for event inclusion (5 each or 10% percent of place cells, whichever was greater, active in the event) were used as in the rank-order replay analysis.

The per-cell contribution to the overall absolute mean P.L. score for each Pre and Post epochs was measured for both for the observed and place-field shuffled P.L. sequence scores (figure 4.7, panel a, see *Methods*). Note that, when compared to the rank-order correlations the distribution of contributions to the overall absolute sequence similarity effect is much more uniform in the P.L. case. In addition, note that the P.L.

sequence similarity measure is highly correlated with sequence rank-order correlation in all sessions (figure 4.8, mean correlation coefficient: 0.983, \pm 0.004 st.d.). This is consistent with the fact that these are two closely related measures of sequence similarity.

In line with the similarity between these two measures, it was found that the overall pre-play and replay estimates content were also qualitatively similar between these two measure (compare tied-rank pre-play and replay in figure 4.3 to P.L. sequence pre-play and replay in figure 4.9). However, taking advantage of the paired nature of the P.L. sequence similarity analysis, and following earlier work (Skaggs and McNaughton, 1996; Wilson and McNaughton, 1994) in which pairs of cells isolated on the same tetrodes were excluded from analysis, we defined non-local P.L. sequence pre-play and replay scores by excluding all pairs of place cells detected on the same silicon probe shank (mean percentage of same shank pairs: 18.5% (\pm 8.6% st.d.), $n = 5$ sessions). In contrast to the previous conditions, using this non-local measure we found that all five Pre epoch and all five Post epochs means of absolute non-local P.L. sequence similarity scores tended to be higher than their respective shuffle controls (figure 4.10, panels b and c). Under this non-local measure, two out of five cases of significant pre-play (figure 4.10, panel b) and four out of five cases of significant replay (figure 4.10, panel c) were observed. Notably, this difference suggests that the replay is more consistent and robust than is pre-play. Interestingly, however, the Post distribution of absolute P.L. sequence scores was found to be higher than the Pre distribution in only three sessions, and significant in one of these cases (figure 4.10, panel a), while another session (figure 4.10, panel a, fourth row) showed significantly higher Pre than Post values. Thus though we were able to observe both pre-play and replay (with the latter being found to be stronger

and more consistent than the former) when each Pre and Post distributions were compared to their respective shuffled distributions, the relationship between these, and the comparison of Pre directly to Post, was not straight-forward.

In order to elucidate this relationship we examined the effect of a variable known to increase from the Pre to Post epochs, participation rate (that is, the number of place cells which fired at least one spike in a given event, and thus were included in that event's sequence statistics) on the null (shuffled) P.L. sequence similarity scores. If the entire distribution is considered (i.e. if the absolute value is not taken) only the variance is found to increase with participation (figure 4.11, panel a, compare green to black line). However, after taking the absolute value, both the variance and the mean, are strongly modulated by the number of participants (figure 4.11, panel b). For instance, the mean shuffled absolute sequence score of events with 5 to 9 participating cells was found to be 0.362, approximately 2.7 times higher than the sequence score for shuffled events with 35 to 39 participants (mean: 0.134). Note that this effect does not apply to any particular event, but rather to the average of the distribution of many random events. One consequence of this is that if within-event sequences were random, given that the majority of sessions showed a significant increase in participation rate from the Pre to the Post epochs (figure 4.2, panel c) we would expect that Pre epochs would show significantly higher P.L. sequence scores than did the Post sessions – the fact that this is usually not the case may implies the presence of at least a higher proportion of maze specific structured sequence content in the Post compared to the Pre epochs.

Subsequently, we examined the correlation of event duration, within-event place cell firing rate, within-event number of participants, and within event participating firing

rate (that is, the within-event firing rate across place cells which fired at least one spike in that event) against non-local absolute P.L. sequence similarity scores in the Pre and Post conditions (figure 4.12, panels a, b, c and d, respectively). Event duration, place cell firing rate, and number of participants were all found to correlate negatively with absolute PL. sequence score. Note that these correlations, some of them quite robust (figure 4.12, a and c) were observed in both the Pre and Post epochs and largely are predicted by the interaction between participant number and mean absolute P.L. sequence score previously described in the null case (figure 4.12). The one positive correlation observed, that between participating (non-zero) place cell firing rate and absolute P.L. score is harder to attribute to the null participant and sequence strength negative correlation. However, the possibility that short duration events both tend to have fewer cells (leading to a higher null sequence score mean) and higher firing rates due to the smaller value of the denominator in assessing the rate cannot be discarded. These cases are illustrative how similarly the observed and null distributions are modulated by number of participants implying that a large part of the variance of sequence similarity scores across events is attributable to random, non-maze specific sequence content. These cases are also illustrative of how, without taking into account the behavior of the null distribution, it can be possible to mistake these random effects for a physiological signal (Tatsuno et al., 2006).

Importantly, since only place field identities were shuffled to construct the null (shuffled) sets and subsequently tested against the observed (unperturbed) events the participation rate is fully preserved between the observed and each of its null comparisons. In order to test how changes in the null distribution from Pre to Post affect

the observed pre-play and replay effects, we compared the within session null distributions to each other as well comparing the observed distributions against the null derived from the opposing epoch of the same session (figure 4.13). A majority (three out of five) sessions show a significant decrease from Pre to Post in the mean of their null distributions (figure 4.13, panel d). In fact, while comparing the observed Pre distribution to shuffled Post distribution shows that the same two sessions remain significant, only two out of five observed Post distributions are higher than the observed Pre null distribution. Consequently, it may be that both changes in the strength of sequential content as well as changes in the context of the expected null distribution in which these content-specific sequences are observed, that drive the increased significance of the replay as compared to pre-play.

Next we examined the evolution of the observed absolute P.L sequence correlation values as well as the within-event firing place cell firing rates over the course of the Pre and Post epochs. Across all five Post epochs only one instance of significant replay decay (that is, a decrease in the replay signal) was observed (figure 4.14, panels e and f). No significant changes in non-local sequence strength were observed across any of the five Pre epochs (figure 4.14, panels a and b). However, within-event firing rates of place cells were found to decrease both across the Pre epochs (three out of five Pre epochs with significant correlations, figure 4.14, panels c and d) and across Post epochs (four out of five significant correlations, figure 4.14, panels g and h).

Next, we examined changes in firing rate and non-local absolute P.L sequence scores within NREM episodes lasting at least 100 seconds. Consistent with previous results (Grosmark et al., 2012), normalized pyramidal firing rates were found to increase

within NREM (figure 4.15, panels a, b and c). While within-events changes in firing rates were not as pronounced as those previously presented for within-ripple changes (Grosmark et al., 2012), a consistent decrease of normalized within-event firing rates was observed within NREM session (figure 4.15, panels d, e and f). However, no significant changes in non-local absolute P.L. sequence score were observed during within NREM episodes (figure 4.15, panels h, i, and j). We need to emphasize though, that the possibility that this dissociation between changes in excitability and pre-play/replay strength is attributable to the divergent statistical power of the two measurements cannot be discarded.

Pre-Play and Replay: Bayesian Decoding

Pre-play and replay were further examined using a Bayesian decoding technique (Davidson et al., 2009) which is complementary to the pair-wise and sequence similarity approaches (see *Introduction*). For each session Bayesian classifiers were constructed from the population of smoothed firing-rate by-position vectors (see *Methods*) with a smooth prior distribution for position. For linear track sessions this analysis was carried out independently for each run direction. Five-hundred null Bayesian classifiers were constructed by randomly circularly rotating each place cell's un-smoothed firing rate vector independently, and subsequently smoothing. Five-hundred null (shuffled) Bayesian decoders were constructed by random circular rotations of each place cell's firing-rate by position vectors and re-smoothing.

In order to establish the quality of the classifiers, the animals' trajectories were reconstructed from place cell population activity vectors binned in non-overlapping 500 ms bins during epochs in which the animal's velocity was at least 5 cm/sec. In order to validate the classifiers, trajectories were decoded from place cell population activity in these same bins using either the real or shuffled decoders. Prediction accuracy was found to robustly exceed the shuffled decoding both for linear maze (figure 4.16, panel b, mean decoded position error: 10.88 cm, mean shuffled decoder error: 50.19 cm, rank-sum test $p \approx 0$, $n = 1,145$ bins) and circular maze (figure 4.16, panel c, mean decoded position error: 14.18 cm, mean shuffled decoder error: 94.72 cm, rank-sum test $p \approx 0$, $n = 1,820$ bins). Note that the magnitude of the decoding errors is in line with those previously reported (Davidson et al., 2009; Dragoi and Tonegawa, 2013; Pfeiffer and Foster, 2013). To ascertain that the data was not being over-fit, we cross-validated the quality of the decoder on one fifth of all laps (see *Methods*). Cross validated trajectory reconstructions robustly out-performed shuffled trajectory reconstructions in both the linear maze (figure 4.16, panel e, mean decoder error: 11.33 cm, mean shuffled decoder error: 49.77 cm, $p \approx 0$, $n = 207$ bins) and the circular maze (figure 4.16, panel f, mean decoder error: 13.01 cm, mean shuffled decoder error: 93.71 cm, $p \approx 0$, $n = 334$ bins).

After justifying our method for coding the trajectory of the animal, we used the Bayesian decoder to examine pre-play and replay in population activity events divided into non-overlapping 20 ms bins. In addition to the event inclusion criteria articulated above for the '*Pre-Play and Replay of Sequential Spiking Activity*' analysis, only events which had at least five non-zero firing rate 20 ms bins, and whose decoded trajectory extended a minimum distance of 30 cm were included in this analysis (see *Methods*, $n =$

11,640 Pre and 11,818 Post events in total). For each of these 20 ms bins, the position corresponding to the maximal decoded posterior probability, as well the value of this maximal posterior probability were obtained. Using templates derived from the novel run epochs two measures were obtained for each population event 1) Bayesian decoding quality was assessed as the mean maximal posterior probability across all the bins in the events (figure 4.17, panels a-d) , 2) Bayesian decoded sequence Pre-play and Replay strength was assessed as the sequence score, using a modified version paired-latency method across the sequences of decoded positions in each event (figure 4.17, panels e-f, see *Methods* for details of the modified P.L. method). Note that similarly to the analysis above, the absolute value of P.L. sequence similarity was taken to account for both forward and reverse sequential content. The measures were also taken using each of 500 shuffled decoders, establishing the null distributions. Bayesian decoding quality (mean posterior probability) increased significantly (significance threshold: $p < 0.025$, rank-sum test) from the Pre to Post epochs in four out five novelty sessions (figure 4.17, panels a and b). However, the observed Bayesian decoding quality was only significantly better than the null distribution of Bayesian decoding qualities in two of the Post epochs and in none of the Pre epochs (figure 4.17, panels a, c and d). Conversely, only two sessions showed significant increases Bayesian decoded sequence strength (figure 4.17, panels e and f). However, all five Post and three Pre epochs showed Bayesian decoded sequence strength higher than their shuffled controls (figure 4.17, panels e, g and h). These findings confirm that both replay and pre-play are detectable using the Bayesian framework (Dragoi and Tonegawa, 2011, 2013), and that, similarly to the *Pre-Play and Replay of Sequential Spiking Activity* analysis, replay is the more robust phenomenon. The fact that

it is the Bayesian decoding quality which shows the most robust changes from Pre to Post, while decoded sequence strength shows the most robust differences against the null distributions should be interpreted cautiously, but may be suggestive of the underlying specific and non-specific structure underlying pre-play and replay. Moreover, it should be noted that these two measures (decoding quality and decoded sequence strength) are not typically measured separately and are in fact conflated in the recent literature (Davidson et al., 2009; Kloosterman, 2012).

Finally, we examined the correlations of Bayesian decoding quality and sequence strength against each other as well as against non-local P.L. absolute sequence strength (see, '*Pre-Play and Replay of Sequential Spiking Activity: Paired-Latency Method*' above). Bayesian decoding strength and Bayesian sequence strength were significantly correlated ($p < 0.025$, Fisher Z-Test) in four out of five Pre epochs and in three out of five post epochs (figure 4.18, panel a). Notably, while this effect might be expected in the Post epoch, the strength of this coupling during the Pre epoch may be suggestive of possible mechanisms of pre-play. Bayesian sequence strength and non-local P.L. sequence strength were found to be significantly correlated in all conditions (figure 4.18, panel b), confirming that these two measures of sequence content are in fact related. Notably however, Bayesian decoding quality (which also correlates with Bayesian sequence strength, see above) did not show consistent correlations with P.L. spike sequence strength (figure 4.18, panel c).

In our final comparison we measured the correlations of these three measures (Bayesian decoding quality, Bayesian sequence strength, and non-local P.L. spike sequence strength) against within-event mean firing rate. The strongest of these

relationships, was that between Bayesian decoding quality and within event mean-firing rate in which all five Pre and all five Post epochs show positive and significant correlations (figure 4.18, panel e). Given that spiking is the evidence used by the Bayesian decoder to estimate position this is an expected result. However, it was found that contrary to the generally negative correlation between P.L absolute spike sequence strength and firing rate (figure 4.18, panel f and figure 4.12, panel b), Bayesian sequence strength showed a positive correlation with within-event firing rate (figure 4.18, panel d, five and three significant correlations for the Pre and Post epochs respectively). The strong coupling between Bayesian sequence strength and P.L. strength (figure 4.18, panel b), combined with the fact that they show opposite interactions with within-event firing rate is of particular interest, and may have applications in understanding the physiological relationships between firing rate and sequential activity.

Discussion

TheStateEditor: Behavioral State Scoring, Data Visualization and Annotation Software

Historically the field of neuroscience has relied on a wide and increasing array of analysis and visualization techniques for small-timescale (millisecond to second resolution) phenomena. For electrophysiological data in particular, with its emphasis on single or clusters of spikes, the most common analysis and visualization techniques, including raster plots, auto and cross-correlograms, peri-event time histograms, evoked response potentials amongst many others (Ostojic et al., 2009; Perkel et al., 1967a, 1967b), are most readily interpretable and interpreted in sub-second time scales. This emphasis on small-timescales is often straight-forwardly explained by the many cellular and network phenomena of interest which occur at sub-second time scales. However, the visualization and analysis of electrophysiological and behavioral phenomena that occur at large time-scales (tens to thousands of seconds), and particularly the analysis of behavioral state dynamics, poses a distinct set of challenges. Particularly, precisely because many cellular and network dynamics occur at small time scales, ‘raw’ physiological signals are typically not easily interpretable at large time-scales. For instance, there is typically little information that can be gleaned from a raster or LFP plot displayed over thousands of seconds, except perhaps in the overall amplitude of the activity. Large-time scale analysis is typically performed over second-order data such as time-resolved spectrograms of LFP data, with the drawback of the loss of finer-time scale dynamics. Consequently, researchers are often faced with a choice of interacting with representations of large time-scale data that are either faithful or intelligible.

TheStateEditor presents one attempt to bridge this gap in a practical and intuitive way. The program presents both large (in the form of spectrogram and motion signal display) and small (in the LFP displays) time-scale data in one coherent, readily navigable and interactive display.

While many of the features of *TheStateEditor*, such as the *StateRibbon* and the motion display, are designed to facilitate the scoring of behavioral states it can also be used to visualize and annotate many types of large time-scale data. This is particularly facilitated by the program's ability to load properly formatted Matlab vectors of arbitrary data types and to load and edit up to ten different event classes. Thus it may be used in contexts far removed from the current study. For instance, a researcher performing an EEG study on infant speech processing may find *TheStateEditor* a relatively easy to use method for examining attention related changes that occur during the course of a particular session. He or she would be able to load simultaneously recorded galvanic skin response data from a '.mat' (Matlab format) file as well as load events, perhaps representing the presentation of certain word stimuli, as well as edit these events. Importantly, if this theoretical researcher was versed in the Matlab programming language and wished to change aspects of *TheStateEditor* to better suit their particular needs, they would be free to modify the script to better suit their interests.

In addition, *TheStateEditor* is useful in that it allows the extraction of particular (up to three) EEG channels, as well as the related time-resolved spectrograms, motion, state and event data. These data can all be saved in two files (the 'auxiliary' file and the 'state' file discussed in the *Results*) and readily shared. The sharing of data in this format, which contains both the 'raw' data as well as the processed state and event labels, may be

of particular use for the cross-validation of analyses such as state scoring in which the hand-labeling of data by experts, with or without a prior unsupervised step, remains the current best practice (Schulz, 2008; Silber et al., 2007).

REM Sleep Reorganizes Hippocampal Excitability

In our study of putatively homeostatic changes in excitability we found that firing rates in hippocampal layer CA1 change during sleep and that these changes display a sawtooth pattern, so that the modest increase in discharge activity within non-REM episodes are overcome by the larger rate deceleration within the intervening REM episodes, resulting in an overall rate decrease during the course of sleep. Theta power of REM sleep is coupled with an increase in synchrony and decrease in rate variability of pyramidal cells during the brief ripple events across sleep. REM mechanisms are thus implicated in both the rate and synchrony changes. These findings suggest that different stages of sleep have different contributions to firing pattern changes. Moreover, a simple global discharge rate measure in the hippocampus does not faithfully characterize the firing pattern reorganization that takes place during the course of sleep.

Models of Sleep Function

There are two dominant views on the role of sleep in firing pattern regulation. According to the ‘consolidation’ model, neurons that are activated by recent waking experience remain selectively active during sleep, firing mainly within hippocampal

ripples and neocortical sleep spindles (Born et al., 2006; Buzsáki, 1989; Carr et al., 2011; McClelland et al., 1995; Sejnowski and Destexhe, 2000; Stickgold, 2005). The increased firing of the active neurons is balanced by a commensurate decrease in the remaining neuronal population so that the global firing rates and population excitability remain relatively constant (Dragoi et al., 2003). In contrast, ‘homeostatic’ models suggest that waking experience-related neurons add to the overall excitability of the cortical networks and sleep (i.e., non-REM) serves to equalize and reduce rates (Borbély, 1982; Lubenov and Siapas, 2009; Tononi and Cirelli, 2006b). Thus, both models attribute importance to sleep-related plasticity, as manifested in the rate changes of individual neurons and/or synaptic weight changes. While our findings do not provide direct information on these issues, they show that rate and synchrony effects should be treated separately (Wilson and McNaughton, 1994) and that it is REM sleep that may be instrumental in bringing about both rate effects and increased synchrony in the hippocampus.

Potential Confounds to a REM Dependent Effect

As is the intrinsic limit of all studies dealing with observations rather than manipulations, we cannot unambiguously ascertain the mechanism of the observed rate increase during non-REM or its decrease during REM. One potentially linked factor to the observed firing rate changes during sleep is a parallel change in core and brain temperature. As observed in rabbits, the temperature of the brain decreases during sleep, interrupted by rapid increases of up to 0.4°C during REM episodes (Baker and Hayward, 1967; Kawamura and Sawyer, 1965). However, temperature change is unlikely to be the

sole cause of the saw-tooth discharge pattern of non-REM and REM, since in the waking, exploring rat, elevation of brain temperature during running is associated with increased neuronal discharge rate and higher excitability (Moser et al., 1993).

Perhaps a more serious confound however, is our inability to determine from the current data whether the changes observed are intrinsic to the hippocampus or are due to changes in firing rate in one or more of layer CA1's input structures. In addition, REM and non-REM are treated here as being stationary states imposing a uniform effect on excitability over time, each of these states are known to host a diversity of micro-states and associated network oscillations. Indeed, frontal spindling is known to increase within non-REM epochs (Gottesmann, 1992; Terrier and Gottesmann, 1978) and amygdala activity displays the elevated metabolic activity during REM in humans (Maquet et al., 1996). However, while non-hippocampal sources of excitation certainly affect CA1 firing during sleep (Hahn et al., 2012; Isomura et al., 2006; Sirota and Buzsáki, 2005; Sullivan et al., 2011), it would be un-parsimonious to assume that these non-local sources completely mask and override local changes. Furthermore, even if the sources of the observed changes are indeed found to be non-local and opposite of local changes, then these latter upstream structures, rather than the hippocampus proper, would still be found to be in disagreement with the standard homeostatic theory (Tononi and Cirelli, 2006b). Finally, while the methods used in this study preclude the definitive establishment of a cellular basis for the observed effects, we stress that spiking remains the *sine qua non* of neural information flow (Barlow, 1972). Consequently, the observed dynamics in spiking during both non-REM and REM, while perhaps not uniquely determinant, are at the very least relevant to the dynamics of information processing during these states.

REM Represents a Unique Neuromodulatory Regime

Of the two brain states here considered (non-REM and REM), only REM episodes are associated with decreasing firing rates in the hippocampus (Montgomery et al., 2008). Although both active waking and REM sleep are associated with similar network states, characterized by theta oscillations and sustained neuronal firing, these states are fundamentally different when viewed from the perspective of the brain stem (McCarley, 2007; Vertes, 1984). Exploration is strongly linked to elevated activity of cholinergic, serotonergic, histaminergic and noradrenergic neurons, whereas during REM sleep only the cholinergic tone is high (Steriade, 2004). It is thus possible that serotonin and/or norepinephrine are responsible for producing different directions of rate and excitability changes during waking and REM, especially because these neuromodulators have been shown to strongly affect long-term synaptic plasticity (Bliss et al., 1983) and REM sleep deprivation results in impaired synaptic plasticity (McDermott et al., 2006).

Hippocampal Excitability and Synchrony Are Decoupled During Sleep

Another unexpected observation in our experiments was the parallel changes of decreased global firing rates and increased synchrony during sharp wave ripples across sleep (Diekelmann et al., 2011). Increased firing rates are typically accompanied by spurious increases in synchrony measures (Perkel et al., 1967a). However, in the hippocampus, large, non-linear increases in population synchrony are brought about by

ripples (Buzsáki et al., 1992), and increased synchrony in our experiments occurred almost exclusively during hippocampal ripples. In fact, within non-REM episodes firing rates between ripples decreased in parallel with the increased participation of neurons in ripples. We hypothesize that the two types of changes, i.e., decreasing firing rates and increased synchrony during the course of sleep, are due to the same mechanism(s) since both changes were significantly correlated with the power of theta oscillations during REM episodes.

Future Directions

It remains to be demonstrated whether the described sleep-related firing patterns changes are unique to the hippocampal CA1 region or can be generalized to other cortical regions. According to a current influential model, the most important role of non-REM sleep is to decrease firing rates (Tononi and Cirelli, 2006b). Since this prediction is opposite to the present observations in the hippocampus, one potential outcome is that firing rate regulations in the neocortex and hippocampus follow different rules. Another alternative is that downscaling of neocortical firing rates are also brought about by the intervening REM episodes, as observed in the hippocampus.

In addition to testing the whether the effects observed in the current study are also found outside of the hippocampus the mechanisms which bring it about may also be usefully examined through various manipulations. The ever-progressing field of optogenetics offers a broad array of tools suitable for testing the precise nature of the

observed changes in firing rate and structure (Anikeeva and Deisseroth, 2012; Boyden et al., 2005; Stark et al., 2012). Indeed, in principle, intrinsic excitability can be examined in channel rhodopsin expressing neurons by testing the robustness over time of the local neural response to a set light stimulus intensity (Chen et al., 2012). In addition, the composition of sleep itself can be usefully manipulated through a variety of pharmacological (Datta et al., 2004, 2008; Gais and Born, 2004; Gottesmann et al., 1998) and optogenetic (Adamantidis et al., 2010; Carter et al., 2010; Rolls et al., 2011) techniques, allowing the dissection of the specific functional contribution of each of its components.

Hippocampal Pre-Play and Replay of Novel Experiences

Having developed an understanding of the general framework of hippocampal excitability across and within behavioral states, we next sought to study how hippocampal firing patterns relate to the learning of novel environments. In other words, having gained an understanding of the general context of firing rate changes during sleep, we next sought to understand how this context interacted with the coding of the specific memory content associated with the learning of a novel set of stimuli. This is a critical issue because learning-induced specificity and the role of sleep in memory consolidation were often contrasted to homeostatic models of memory. Our research covering both topics indicates that the two models are complementary rather than exclusionary. Furthermore, our study was methodologically distinct from previous work on the pre-play and replay phenomena in at least three salient features. First, and most important, in all previous studies concerning these phenomena all Pre and Post recordings were carried

out in the same, familiar, room in which maze epoch (whether novel or familiar) was also recorded (Dragoi and Tonegawa, 2011, 2013; Euston et al., 2007; Hirase et al., 2001; Ji and Wilson, 2007; Kudrimoti et al., 1999; Lansink et al., 2009; Lee and Wilson, 2002; Louie and Wilson, 2001; Nádasdy et al., 1999; Nakashiba et al., 2009; O'Neill et al., 2008; Pavlides and Winson, 1989b; Pennartz et al., 2004; Peyrache et al., 2009; Poe et al., 2000; Qin et al., 1997; Ribeiro et al., 2004; Skaggs and McNaughton, 1996; Tatsuno et al., 2006; Wikenheiser and David Redish, 2013; Wilson and McNaughton, 1994). By contrast, our Pre and Post epoch recordings were carried out in a different (and familiar) room than that in which novelty maze exposure was recorded, thus presenting a totally novel context for the animal. Notably, the novelty of 'novel maze' exposure condition is critical in order to establish that the activity observed during the Pre epoch (i.e. pre-play, (Dragoi and Tonegawa, 2011, 2013) is truly non-specific with respect to the novel scenario into which the animal is introduced. Testing an animal in a different room (context) and not only in a novel arm of a familiar maze or a different apparatus in the same room is critical, since research has shown that strong changes in firing rate representations occur only in different rooms (Leutgeb et al., 2005). Secondly, in order to contextualize any novel maze-specific content within more general putatively homeostatic changes in hippocampal activity, our Pre and Post recordings were, on average, each more than 3 hours long, while most previous studies on the subject only analyzed fifteen minutes to an hour of Pre and Post maze activity (Dragoi and Tonegawa, 2011, 2013; Euston et al., 2007; Hirase et al., 2001; Nakashiba et al., 2009; O'Neill et al., 2008; Pennartz et al., 2004; Peyrache et al., 2009; Skaggs and McNaughton, 1996; Wilson and McNaughton, 1994). Finally, though not strictly unique, the current study

was carried out with a particular focus on how background patterns of excitability (and particularly firing and event participation rates) affect measurements of the relationship between the temporal sequences of activity observed in ‘off-line’ states and the spatial distribution of place fields on a novel maze. These features in the design of our experimental and analysis methods, were included in order to obtain a more global perspective of the dynamics of both the non-specific context of 'off-line' hippocampal activity as well as the putative maze-specific content nested within it.

Excitability Changes In Population Activity Events

In accordance with the previous work relating to changes in excitability across behavioral states (Vyazovskiy et al., 2009) we found an increase in firing rate and participation in population activity events from the Pre to the Post epochs. While in our previous study (see ‘*REM Sleep Reorganizes Hippocampal Excitability*’ above; Grosmark et al., 2012) we did not directly test hippocampal firing rate changes during the waking state, our observation that firing rate decreases across sleep strongly implicated the waking state in their commensurate increase, a hypothesis supported by the current findings (figure 4.2). Importantly, these expected, putatively non-task specific, increases in the levels of excitability informed much of our subsequent analysis into the nature of maze-related activity content.

Replay Is Robustly Present in Place Cell Pair Activity and Timing Co-Modulation

Given the novel aspects of our study, we first sought to replicate earlier work concerning (non-local) pair-wise measures of place cell co-activity (Kudrimoti et al., 1999; O'Neill et al., 2008; Pennartz et al., 2004; Wilson and McNaughton, 1994) and timing co-modulation (Euston et al., 2007; O'Neill et al., 2008; Skaggs and McNaughton, 1996) as a function of place field peak distance (figures 3.3 and 3.4, respectively). Notably, while, as expected, place cells with overlapping place-fields (20 cm or less pair-field peak distance) were found to be strongly co-modulated both in terms of activity and spike-timing, we also observed that cells with strongly segregated place fields (120 cm or more place field distance) were selectively segregated in terms of these measures. Interestingly, the spike-timing segregation of pairs with far place fields, appeared more robust on the linear track condition (figure 4.4, left column) than in the circular maze condition (figure 4.4, right column). The linear and circular mazes are distinct in at least two important respects. First, in the linear case, the animals were made to run in both the left and right direction to earn reward, while in the circular maze running was generally limited to one (clockwise) running direction and only runs in this direction were rewarded. All the pre-play and replay analysis presented deals with the bi-directionality of the linear maze condition by taking all measurements separately for the left and right running directions (Diba and Buzsáki, 2007; Dragoi and Tonegawa, 2011, 2013; Foster and Wilson, 2006). Secondly, running on the linear maze was strongly modulated by the maze's two ends at which points the animals were rewarded and had to change running direction. This is in contrast to the 'endless' nature of circular maze runs. Consequently, it may be that as compared to the circular maze condition, the replay of linear maze

experience is biased such that the two ends of the maze are selectively temporally segregated. Notably, this hypothesis is supported by the strong bimodality we observed in the place field distance effects on spike-timing on the linear maze (figure 4.4, left column, note that relatively little modulation is observed at intermediate place-field distances). Importantly, a unique role for the ends of the linear track in replay (or pre-play) is also suggested by findings showing that hippocampal place cells tend to aggregate around behaviorally salient landmarks (Gothard et al., 1996; Hollup et al., 2001; Redish et al., 2000), and specifically around the ends of linear mazes (Diba and Buzsáki, 2008). The relationship between the putative over-representation of the ends of the linear maze and the segregation of the timing of spikes corresponding to opposite ends of the maze, and in turn the relationship of each to pre-play and replay signals, presents an interesting subject for future detailed analysis.

Another interesting result regarding the pair-wise analysis was the virtual absence of a pre-play signal which will be discussed in more detail below.

Absence of Pre-Play and Replay Using the Rank-order Correlation Method

Having established the presence of pair-wise replay (but not pre-play) in our data set, we next sought to examine this phenomena using the event-based rank-order correlation method, where the template of the neuronal sequences is obtained from the place field sequences on the maze (Diba & Buzsáki, 2007; Dragoi & Tonegawa, 2011, 2013; Foster & Wilson, 2006). To our surprise, we found that contrary to both previous work (Dragoi and Tonegawa, 2011, 2013) and our observation concerning pair-wise

effects, neither replay nor pre-play was observed using this method. This discrepancy may be attributable to several substantive differences between the previously employed methods and our own. While the rank-order method has been employed by several groups (Diba and Buzsáki, 2007; Dragoi and Tonegawa, 2011, 2013; Foster and Wilson, 2006) all but one of these groups (Dragoi and Tonegawa, 2011, 2013) restricted their analysis to the awake state, while in turn one of these (Diba and Buzsáki, 2007) was carried out mostly in the CA3 rather than CA1 hippocampal layer. Importantly, in the aforementioned studies carried out Dragoi and Tonegawa, and in contrast to our own study, both the Pre and Post sleep sessions were recorded in the same room as the novelty run. In those studies the animal's view of the room was occluded during Pre and Post epochs by placement of walls around the animal's home cage during recordings. However, the possibility that the animal had, even a fleeting, view of the experimental room during transportation into the room or during the plugging in of its head-stage cannot be excluded. Moreover, the animal would be expected to have had a clear view of the experimental room's ceiling and any cues, such as the placement of lights, shelving, etc., found on it. In addition, during the transfer onto the maze the animal would presumably have had ample visual and vestibular input to indicate that it was moving entirely within the context of the experimental room. While these stimuli may each seem insignificant, hippocampal place coding is known to be radically modulated by the context in which a place is experienced (this, indeed, is part of the motivation for the replay experiments, see *Introduction*, Leutgeb et al., 2005). In turn, the establishment of the context of hippocampal place maps, as well as the 'anchoring' of the head direction system is known to rely on vestibular cues (Taube, 2007) and more importantly on visual

cues and specifically, distal visual cues (Cressant et al., 1997; Zugaro et al., 2001, 2004). Consequently, it may be that in the pre-play studies, as well as in all the studies of replay of which we are aware, at least part of the similarity found in the 'off-line' and maze conditions may be explained by the common context in which they were both recorded. Importantly, the claim here is not that these studies are thus invalid, rather that the specific contribution of contextual cues to pre-play and replay remains largely unknown, and cannot be excluded based on the methods of these studies.

Two further differences between the present study and previous work on pre-play (Dragoi and Tonegawa, 2011, 2013) should also be considered. Firstly, in order to exclude potential behavioral variability attributable to the reward areas, and in contrast to the pre-play studies, in our study the reward areas at the ends of the maze were not included in the place cell analysis (Diba and Buzsáki, 2007). This difference is particularly notable because, as previously shown (*Results*, figure 4.6), the rank-order method is heavily biased by the activity of cells with fields near the edges of the maze. Finally, another difference comes from the differing recording apparatuses employed in these studies. Specifically, in the present studies hippocampal neurons were recorded using either 6 or 8 shank silicon probes, which offer geometrically precise, spatially dense sampling of the hippocampal CA1 layer (Bragin et al., 2000). In contrast, in the pre-play studies hippocampal neurons were recorded using multi-tetrode arrays, which offer low geometrical precision and more diffuse spatial sampling. As a consequence, the current study contained a higher percentage of neural pairs recorded at proximate (local) anatomical locations. The possible local versus non-local pair recordings are discussed below.

Replay and Pre-Play Are Observed in the Non-Local Higher Order Structure

Having developed and verified (*Methods, Results*, figures 3.7 and 3.8) the novel *Paired Latency* (P.L.) method for assessing the higher-order similarity of rank-ordered sequence content we proceeded to examine pre-play and replay using this new method. As expected from the results of the P.L. method's similarity to the tied-rank correlation approach (figure 4.8), a similar absence of pre-play and replay effects was observed (figure 4.9). However, the P.L. method also enabled us to follow previous work (Skaggs and McNaughton, 1996; Wilson and McNaughton, 1994) in restricting our analysis to pairs recorded on different (non-local) silicon probe shanks. Interestingly, when only non-local pairs were considered we found four Post epochs with significant replay and two Pre epochs with significant pre-play (figure 4.10). Furthermore, the observed replay and pre-play strength was higher than that of the shuffle controls in all five Post epochs and all five Pre epochs, respectively. The fact that non-local pairs contributed to the wake-sleep sequence similarities is in line with a previous study which used a method based on principal component analysis to detect replay of rule learning (Peyrache et al., 2009). Notably, the authors demonstrate that nearly all of the observed replay effect is accounted for by pairs of cells recorded on different tetrodes, even though interactions between pairs of cells recorded on the same tetrode dominate the correlation structure. However, it should be noted that this study was carried out in the prefrontal cortex, whose anatomical organization is quite different for that of the hippocampus. While the putative roles of anatomically distributed coding in the non-columnar hippocampus will be discussed further below, it is here worth noting that in both Peyrache et al., 2009 as well as in the current study differences between distal and local interactions should be

interpreted cautiously in light of the fact that the latter are subject to spike cross-contamination. However, the use of silicon probe based recordings in the current study, together with the increased sampling density and spike waveform resolution afforded by this technique may enable us to fruitfully test the null hypothesis of spike cross-contamination in future analysis.

Changes in Participation Rate Lead to an Increase in the Signal to Noise Ratio During the Post Epoch

In order to assess how the observed non-local sequential pre-play and replay signals co-varied with the previously observed changing context of excitability we first sought to characterize the way that the null (shuffled order) sequence similarity structure co-varied with within-event participation (that is, number of pyramidal cells that fired at least one spikes). Participation was chosen both because in the P.L. sequence similarity method (as well as in the rank-order correlation method) each cell that fires at least one spike contributes only one value (i.e. the timing of the temporal center of mass of all the spikes it emitted within that event) and because participation was a measure of excitability observed to change from the Pre to the Post epochs (figure 4.2). Notably, we observed a strong inverse relationship between the number of participants in an event, and the expected absolute value of the null P.L. sequence similarity strength of that event (figure 4.11). Importantly, given that within-event participation rates tended to increase from the Pre to the Post epoch, a commensurate decrease in the null distribution was expected and observed (figure 4.13). Moreover, without this observed decrease in the null

distribution from the Pre to the Post epoch, only two out of the previously four epochs displaying significant replay, retained this effect (figure 4.13). Consequently, our analysis implies that both changes in the signal content of sequential activity as well as changes in the noisy context of sequential activity from the Pre to the Post may play part in the emergence of replay.

While the particular signal and noise components described here are specific to the method employed and its relationship to participation rate, we emphasize that increases in neuronal excitability as expressed by firing rate were reliably observed from the Pre to the Post epochs (figure 4.2). Whether the increase in excitability (intrinsic or synaptic) is entirely restricted to task-specific network components or affect both task-specific and non-specific components of network activity needs to be investigated further. We consider the task-specific excitability hypothesis less plausible due to several considerations. Firstly, wake-related increases (Vyazovskiy et al., 2009) as well as commensurate sleep-related decreases in excitability have been observed to occur in the absence of any particular learning task (Grosmark et al., 2012). This observation was confirmed in the current study by our result showing that similar sleep-related decreases in hippocampal firing rates occurred both within the Post sleep, which occurred subsequent to a novel and salient spatial task as well as within the Pre sleep, which occurred following the absence of such novel stimuli (figure 4.14 panels g and c, respectively). Furthermore, if the observed increases in excitability were entirely restricted to task-specific cells, the increased firing of these cells, together with their enhanced excitation of local interneurons (Buzsáki and Chrobak, 1995; Pelletier and Lacaille, 2008), would be expected to lead to an unobserved drastic increase in the signal

to noise characteristics of task-specific network activity. Finally, given that individual hippocampal cells as well as individual hippocampal cell pairs may contribute to the encoding of more than one set of stimuli (Hayman et al., 2003; Jeffery, 2011), it is unclear that wholly task-specific changes in excitability are theoretically possible. Given our current knowledge, we hypothesize that the consistently observed increases in excitability from the Pre to the Post epochs are due to both task-related as well as unrelated, putatively homeostatic, changes in excitability. While the none-specific component of these changes has been traditionally treated as a confound in much of the replay literature (Tatsuno et al., 2006; Wilson and McNaughton, 1994), we emphasize that given the observation that following learning both task-related and none-related changes in excitability are found to be physiologically reliable phenomena, a physiological role for both in the processes governing hippocampal-dependent memory consolidation cannot be excluded (Buzsáki, 1989; Ermentrout et al., 2008). While the precise nature and contribution to memory of each these components is beyond the scope of a single study, we note that a careful cataloging of each of these putative components is a necessary first step in such an exercise.

Much of the Variability of the Observed Pre-Play and Replay Signal is Attributable to the Noise Distribution

While the null distribution absolute P.L. sequence similarity strength was found to be strongly modulated by the participation rate, this effect was also robustly present in the observed (non-shuffled) distributions of absolute P.L. sequence similarity strengths

(figure 4.12). Consequently, much of the variance in the observed pre-play and replay signals may be straight-forwardly explained as noise with regard to the sequence of place cells on the novel maze. The sources of this noise, both physiological and experimental are worth taking into consideration.

Physiological Sources of 'Noise'

Firstly, it should be remembered that, though reproducible (including in the present study), the presence of sequential replay (or pre-play) of the sequence of place fields on a maze in discrete 'off-line' events, remains hypothetical and is not without its confounds or competing hypotheses (Lubenov and Siapas, 2008; Moore et al., 1996; Tononi and Cirelli, 2006b). Moreover, even if the sequential replay hypothesis is accepted, it should not be assumed that this effect would be observed in all events and by all analyses equally. Particularly, there may be different types of population activity events (Carr et al., 2012; Patel et al., 2013) which preferentially support sequence reinstatement. More importantly however, even if it is assumed that all population activity events carry memory-related sequence content, the experimental paradigm inevitably constitutes only small fraction of the totality of the animal's memories, and thus perhaps only a small fraction of the sequence variance within population activity events (Dupret et al., 2010).

Experimental Sources of Noise

In addition to physiological sources of noise in detecting pre-play and replay, there are many other sources which may contribute to our imperfect ability to sample, control and decode on-going hippocampal activity. Firstly, replay is the re-instantiation of stimulus-specific neural activity patterns during ‘off-line’ states. Consequently, replay (and by extension, pre-play) is constrained by the accuracy with which neural patterns can be attributed as being specific to a particular stimuli – in our case, the attribution of a CA1 cell’s place field to a particular place. Furthermore, the methods we employed impose certain assumptions about the behavior of this place coding. For instance, we assume that within the cell’s place-field the spatial bin with the highest firing rate is a suitable discrete representation of the place-cell’s place coding (Kim et al., 2012). Perhaps more importantly, pyramidal cells can have more than one place field, however, only the place field with the highest firing rate is usually considered (see *Methods*, (Diba and Buzsáki, 2007; Dragoi and Tonegawa, 2011, 2013). Moreover, since the analyses are carried out on the rank-ordered sequences, the actual spatial or temporal distances between consecutive pairs are not taken into account. However, it is not clear that, for instance, two place cells which have place field peaks on adjacent 2 cm bins on the maze truly code for different consecutive places in a physiologically meaningful sense. Note that this last critique also applies to the rank-ordering of the centers of mass of spike timing of the participating cells within population activity events: it is unclear whether the fact that one cell discharged a spike one sample (our sampling rate for spike detection was 20,000 Hz) before or after the spiking of another cell can have physiological relevance in terms of hippocampal neural function. Notably, the paired nature of the P.L.

methodology may be well suited for answering some of these questions, which involve the particular subset of pairs which contribute to the replay signal.

While some of these issues, particularly the discrete nature of place preference and spike timing estimation for the rank-order correlation and P.L. methods, do not affect replay or pre-play estimates obtained through Bayesian decoding methods (Kloosterman, 2012), the Bayesian method as here employed, suffers from other draw-backs. The most obvious of these, is that it requires the division of population activity events into physiologically arbitrary bins for place decoding. In the present work a bin size of 20 ms was chosen based on the tradeoff between decoding precision (temporal resolution) and decoding accuracy (which is related to the amount of spiking activity within each bin). While it would be interesting to delineate bins in a more physiologically relevant way, if the resulting bins were not of a uniform length, such binning would be expected to introduce noise due to Bayesian decoders' sensitivity to rate, since in this latter case rate would be expected to co-vary with the length of each bin (Kloosterman, 2012).

Other experimental contributors to noise affect both spike-sequence and Bayesian decoding approaches. One is the assumption that all cells with place fields on the novel maze are part of the replayed (or pre-played) 'ensemble' during the 'off-line' states. Finally, the possibility of artifactual neuronal drift or spike cross-contamination is a pervasive concern when dealing with long time epoch electrophysiological data. While in the present study we only considered stable, well isolated, principal cells that fired at at least 0.15 Hz in both the Pre and Post epochs, issues relating to isolation quality and stability can never be entirely discarded in the context of extracellular electrophysiology.

Sources of Inter-Animal Variability

Another source of noise evident when examining the data arises from putative differences between animals and sessions. Much of this, in turn, is directly attributable to number of well isolated cells obtained in each case. Particularly, since both the rank-order correlation and P.L. similarity methods rely a minimum number of cells being active in each event (for instance, no correlation can be obtained between two points) differing number of recorded cells will result in a different number of analyzable events. Furthermore, the number of participants in each event was also shown to affect the null distribution of P.L. similarity strength. More generally, given the numerous sources of noise listed above, for instance the potential for misattribution of a cell's firing selectivity to a particular place, which would be expected *a priori* to be random with respect to the replay or pre-play signal being analyzed, a significant number of 'noise' neurons must be expected to reside within the sample. Consequently, an increase in the number of neurons recorded would, in general, be expected to increase the signal to noise characteristics of the replay and pre-play signals. Obviously, addressing the issue of variability as a function of sample size quantitatively would require much larger numbers of simultaneously recorded neurons. With current methods, this is not realistic though. We need to emphasize that the number of neurons recorded in the current study are similar to those reported in previous studies of replay and pre-play (Dragoi and Tonegawa, 2011, 2013; Lee and Wilson, 2002; Wilson and McNaughton, 1994).

Other contributors to inter-animal and inter-session variability may not be clearly attributable to the numbers of cells recorded and many remain unknown. For instance, the reasons explaining why some animals perform or learn tasks faster than others remains an

active area of research (Ray and Hansen, 2004; Ray et al., 2006). Likewise, the depth and composition of sleep may be expected to vary between animals. Indeed, in our own experience, we have observed that certain animals fall asleep much more readily than others. The causes of these differences, but more importantly for the current purposes, the impact they have on memory replay and pre-play presents an interesting topic for future research.

Notably, given the observed differences, both experimental (for instance cell number) as well as physiological ('temperament' and depth of sleep amongst others), between animals, and the relative weakness of the replay and pre-play signals, perhaps the preferred approach in replay research should be to perform statistical tests across animals rather than within them. The fact that this is not the preferred approach can be largely explained by the prohibitively expensive and time-consuming nature of implanting silicon probes, and testing a sufficient number of animals, each of with enough well isolated units and satisfactory behavioral performance, to pass standard statistical tests. For instance, current neuron clustering methods involving 100 potential neurons with manual adjustment may take up to a week. However, the increasing prevalence of data sharing may, in the future, enable individual laboratories to pool their resources to such an end, and in turn, allow individual researchers access to data from many more animals.

The Noisy Nature of Replay and Pre-Play May Explain the Lack of Within Pre and Post and Within Non-REM Effects

In the previous section, we described the excitability changes that occur within non-REM episodes and across sleep (see above, Grosmark et al., 2012). However, we did not find similar changes in either P.L. sequence similarity strength (figure 4.14) or Bayesian decoded sequence strength (figure 4.17) in either the Pre or the Post epochs. This may be attributable to the different statistical properties of the two measures. In the excitability case, comparisons were performed across simultaneously recorded cells. Within each session the replay and pre-play signals provide a single measurement for each event, thus greatly reducing the statistical power of measurements designed to test whether the signal significantly changes or is stationary. This statistical consideration, combined with the above-discussed numerous sources of noise associated with the replay and pre-play signals may thus explain the negative results. However, with the detailed knowledge concerning the mechanics governing the pre-play and replay signals, it may be possible to develop a more reliable detector of the replay and pre-play phenomena and re-examine this issue in the future.

Partial Dissociation Between Pre to Post Changes in Bayesian Decoding Quality and Bayesian Decoded Sequence Strength

The bulk of our analysis dealt with the rank-order and P.L. sequence similarity score measurements of pre-play and replay since these measurements relate to the

underlying neural data in a (more) direct way. On the other hand, the results obtained from the Bayesian decoding analysis offered a convergent and complimentary second examination of the pre-play and replay phenomena. One important aspect in which the spike-sequence and the Bayesian decoding methods diverge is that the latter method is based on the decoded estimate of position, the estimated quality of which is given by the Bayesian posterior probability of position. Furthermore, this estimate of position, as well as its estimated quality, can be measured independently of the presence or absence of trajectory (i.e., sequential position) content across the time bins being decoded. Importantly, in the Bayesian-based method which is increasingly used for decoding replay and pre-play, these two measures are typically conflated into one overall replay score based on the linear trajectory which captures the maximal amount of posterior probability of position across the decoded time bins (see *Introduction*, (Davidson et al., 2009; Kloosterman, 2012; Pfeiffer and Foster, 2013)). By contrast, the current study examined these two measures separately. Note that the Bayesian decoder calculates the posterior probability of position by comparing the population activity vectors (a vector of within-bin firing rates across place cells) of the 20 ms bin being decoded against the template of population activity vectors as binned by the animal's position on the novel maze (see *Methods*, equation 11). The posterior probability of the animal 'being' at a given spatial bin (i.e. decoded position probability) is subsequently normalized to one across all spatial bins (see *Methods*, equation 12). Consequently, Bayesian decoding quality represents the specificity of the similarity of an observed population activity vector in a given time bin to the population activity vector in the given spatial (2 cm) bin to which it is most similar (in other words, the best estimate of position), normalized by

its 'non-specific' similarity to all spatial bins. Bayesian decoded sequence strength was subsequently measured across all (non-zero firing rate) bins as the constancy of the pairwise directionality of the decoded position across all pairs of bins within an event (see *Methods*). To put it slightly more intuitively, this method measured how consistently all the pairs of decoded position of 20 ms bins in a given event indicated that the animal was 'moving' (in decoded position space) forwards *or* backwards across the bins of that event. Note that in contrast with previous methods for Bayesian decoding of sequence content (Davidson et al., 2009; Kloosterman, 2012; Pfeiffer and Foster, 2013), our method does not assume that the 'movement' through decoded position space is linear with respect to position and time.

Using these methods, we found that both Bayesian decoding quality as well as Bayesian decoded sequence strength tended to increase from the Pre to the Post epochs (figure 4.17). However, the nature of this increase was notably different for the two measurements. While Bayesian decoding quality increased significantly in four of the five sessions, a significant increase in Bayesian sequence strength was only observed in two out of the five sessions. Conversely, Bayesian decoding quality was significantly greater than the null distribution (generated by random circular shifting of each place cell's firing-rate by position vectors, see *Methods*) in two Post epochs and in none of the Pre epochs. The observed Bayesian decoded sequence strength was greater than the null distribution in all five Post epochs and in three Pre epochs (figure 4.17). Note that this last result - the observation of significant maze-related sequence content during Pre epochs - presents our second confirmation of the previously discussed reports of hippocampal pre-play (Dragoi and Tonegawa, 2011, 2013). However, the observation of

a divergence between Bayesian decoding quality and sequence content in their Pre to Post differences versus their relationships to their respective null distributions was facilitated by the novel Bayesian method employed in the current study and may bear further examination. Indeed, on first consideration, the observed divergence may seem to be in the opposite direction as would be expected. Consider that novel maze-specific sequential activity was, until the recent discovery of pre-play, thought to be almost entirely due to the *de novo* arrangement of place fields on the novel maze (Lee and Wilson, 2002; Louie and Wilson, 2001; Skaggs and McNaughton, 1996). And on the other hand, both firing rates (Hirase et al., 2001; Pavlides and Winson, 1989) and pair-wise interactions (Kudrimoti et al., 1999; Wilson and McNaughton, 1994) (which are in turn intimately related; Luczak et al., 2007; Peyrache et al., 2010) are known to be highly preserved across brain states (Grosmark et al., 2010; Mizuseki and Buzsáki, 2013). However, note that in this latter context firing rates and pair-wise interactions are assessed across time while in the current context of Bayesian decoding the population activity vectors are assessed with respect to stimulus (i.e. maze position) space (Kloosterman, 2012). In other words, while it is known that, for instance, cells with high firing rates tend to be *temporally* highly correlated at small (e.g., 100 ms) time-scales across brain states (de la Rocha et al., 2007), this does not imply that they are correlated in terms of their *place selectivity*.

One interpretation of the current result, which will be examined more closely below, is that hippocampal sequential activity is largely hard-wired and similar during place coding and 'off-line' states (Dragoi and Tonegawa, 2011, 2013). On the other hand, population activity (i.e. firing rate) across the hippocampal CA1 network is able to

change dynamically, facilitating novel and maze-specific firing patterns during the 'off-line' state specifically in the Post epoch (Hirase et al., 2001; Wilson and McNaughton, 1994). While the confirmation pre-play in the current study lends support to the presence of a common hard-wired component in the sequence structure of place coding, the robust increases in Bayesian decoding quality may be thought to be wholly attributable to putatively non-specific increases in firing rate (figures 4.2 and 4.17), particularly as these two measures were shown to be tightly linked (figure 4.18, panel e). However, the observation that two Post epochs showed a significantly greater degree of Bayesian decoding quality than their null controls (which preserve the firing rates of place cells on the maze) as well as previous work (e.g. Hirase et al., 2001; Wilson and McNaughton, 1994) suggests that at least part of the firing patterns observed during the Post epoch arise *de novo* during exploration of the novel maze. Nonetheless, the observation that the *de novo* maze-specific component of activity during the Post epoch is dominated by changes in the firing rate and co-activity of individual place cells rather than in their sequence of activity, should be interpreted cautiously. In particular this finding must be confirmed with other, perhaps more transparent, methods than Bayesian decoding and with a particular emphasis on the contribution of specific and non-specific changes in population activity (including single-cell, pair-wise and higher order activity) to this effect.

Bayesian Decoding Quality Correlates With Bayesian Decoding Strength in Both the Pre and Post Epochs and Both are Correlated With Firing Rate

While we have elaborated on the subtle divergence in the behavior of Bayesian decoding quality and Bayesian decoding, the overall concurrence of these two measures is also of interest. Bayesian decoding quality and sequence strength were found to correlate significantly in three Post epochs (figure 4.18, panel a). This should be expected given that the estimation of the sequential position content signal is necessarily dependent on the quality of the estimate of position. Perhaps more surprisingly however was that these two measures were also significantly correlated in four out of the five Pre epochs (figure 4.18, panel a). Moreover, this coupling increased significantly from the Pre to the Post epochs in only one session. Notably, these findings suggest the expression of sequential pre-play is facilitated by population activity (i.e. population firing rate vectors) resembling those which will subsequently be expressed on the maze (Grosmark, 2013, *Society for Neuroscience* poster presentation), a theme that will be returned to below.

Other strong correlations were observed between both Bayesian decoding quality and sequence strength and within-event firing rate (figure 4.18, panels d and e). This relationship can be straight-forwardly explained by the dependence of Bayesian decoding quality on firing rate in position estimation. Yet, this relationship is of particular interest because the P.L. sequence similarity score was observed to have a tendency toward the opposite correlation with within-event firing (figure 4.18, panel f) rate, yet is strongly coupled to Bayesian decoded sequence strength (figure 4.18, panel b). Since these two are distinct measures of the same phenomena (i.e. replay and pre-play) the fact that they are found to co-vary in opposite directions with firing rate may be usefully leveraged to

construct a method for the detection of replay and pre-play that is on the whole less biased in either direction by firing rate variations and thus potentially more sensitive to the underlying replay and pre-play signals.

Non-Local Firing Rate, Sequence and Place Coding Structure

When considered as a whole, the findings elaborated here may suggest some significant refinements for current models of the specific and non-specific contributions to hippocampal memory formation (Battaglia and Pennartz, 2011; Blum and Abbott, 1996; Hasselmo, 2008; Káli and Dayan, 2004). We found that activity and co-activity of hippocampal layer CA1 neurons replay their activity on a novel maze in the neurons' pair-wise co-firing patterns (*rate replay*, figure 4.3) as well as in the similarity between 20 ms population activity vectors during the Post epoch and the activity vectors at particular 2 cm-long positions on the novel maze (figure 4.17). We also found sequential replay both in the pair-wise co-modulation of spike-timing by place field distance, as well as in the higher order sequence content as assessed by the non-local P.L. similarity method as well as the Bayesian decoding method. Replay of sequential activity patterns was observed both in the non-local P.L. sequence similarity measure as well as in the Bayesian decoded sequence strength measure. Interestingly however, pre-play was not observed in pair-wise measures of rate or spike-timing co-modulation or in Bayesian decoding quality (a measure of the stimulus specificity of firing rate patterns across the population of place cells). Taking these results into account, our data favors a model in which sharp-wave/ripples (the dominant source of population activity events in the

hippocampus, (Buzsáki et al., 1992)) spread along the septo-temporal extent (parallel to our placement of the silicon probes) of the hippocampal formation (Patel et al., 2013). In turn, in a given sharp-wave/ripple the particular path of the anatomically constrained spread of excitation is guided by the coordinated activity of interneurons (Basu et al., 2013), as well as high-firing rate principal cells. These high firing cells are known to be anatomically segregated, have strong efferent connections, and are thought to be crucial for the anatomically distributed flow of information in cortical structures (Buzsáki and Mizuseki, 2014; Mizuseki and Buzsáki, 2013; Yassin et al., 2010). Notably, this spread, like sharp wave bursts themselves, would originate in the uniquely recurrent layer CA3 and then propagate to layer CA1 (Buzsáki, 1986; Buzsáki et al., 1983, 1992; Csicsvari et al., 1999a; Ylinen et al., 1995). In the 'naive' condition (Pre epoch) the spread of activity, though anatomically constrained, would recruit both future place-selective and non-selective cells. However, the spread of the activity would be limited to either septal to temporal or temporal to septal direction (Patel et al., 2013). Under this hypothesis, any given cell pair along a given possible path of excitation would experience only a very weak interaction, proportional to the stochasticity of both the anatomical flow of activity and their recruitment on any given pre-play event. The average pair-wise interaction in this highly stochastic (though still constrained) context might be expected to lie below the detectable threshold. However, this structure (though weak) would be expected to be more apparent in the higher-order interactions between many, specifically anatomically distributed, cells - as we observed in non-local P.L. sequence similarity pre-play.

Furthermore, since high-firing rate cells tend to discharge more spikes per event, these cells would be expected to be more reliable marker of any underlying spread of

excitation (Buzsáki and Mizuseki, 2014; Peyrache et al., 2010). Notably, as discussed in the *Introduction*, event-based measures are inevitably biased towards higher firing rate cells since highly active neurons tend to participate in more events. This latter bias may also account for the presence of pre-play in event-based metrics but not in the un-biased pair-wise metrics. Recall that pre-play was also observed using Bayesian decoding of sequence, a method that does not differentiate between local and distal interactions. However, two factors may help explain this discrepancy. Firstly, since Bayesian decoding methods simultaneously 'weigh' all the 'evidence' (i.e., spiking) within each 20 ms bins, the estimates generated are expected to be more robust, and specifically less sensitive to the noise generated by individual cells or spikes to which the spike-sequence metrics are quite sensitive. Secondly, the binning involved in Bayesian decoding may generally discount short-time scale interactions (specifically below 20 ms) and is thus biased towards longer time scale interactions which are more likely to be anatomically distributed (Patel et al., 2013).

During novel maze running excitability spreads from the septal to temporal poles of the hippocampus (Lubenov and Siapas, 2009; Patel et al., 2012) - recruiting an ensemble of neurons, each in a place specific manner. This recruitment in turn imposes a lasting *de novo* structure of potentiation, as reflected by both the intrinsic biophysical properties of the neurons (Daoudal and Debanne, 2003; Disterhoft et al., 1986; Zhang and Linden, 2003) and synaptic plasticity (Kentros et al., 1998; Nakazawa et al., 2004). As a result, cells recruited to place coding ensembles are more readily excitable (Hirase et al., 2001; Pavlides and Winson, 1989) and discharge more synchronously (Kudrimoti et al., 1999; Wilson and McNaughton, 1994) during the subsequent Post epoch. In this *de novo*

regime, a putatively non-specific increase in global excitability during the Post sleep, in turn leads to the preferential activation of the pre-potentiated ensemble (O'Neill et al., 2008; Wilson and McNaughton, 1994).

The sequential firing content of the events may be explained by two complementary models. In the first, the anatomical spread of sharp-wave ripples is largely unresponsive to changes in excitability within the network, and remains the same between the Pre and Post sleep. However, the maze-specific sequence structure arises from the fact that cells participating in the ensemble are preferentially co-active, thus increasing the signal to noise characteristics of sequential activity (Dragoi and Tonegawa, 2014). In the second model, changes in excitability, and particularly the strengthening of recurrent connections between neurons in CA3 (and particularly, of high-firing rate neurons) leads to a bias in the spread of sharp-wave/ripples towards paths which preferentially recruit the neurons of the ensemble in the order in which they became active on the maze (Abeles et al., 1993; Hasselmo, 2008). Notably, this second model may better account for the coupling between the population of activity and sequence content observed both in the Post and more importantly, in the Pre epoch (Bayesian decoding quality versus Bayesian sequence strength correlation, figure 4.18, panel b).

My thinking regarding the origin of organized events in both Pre and Post sleep is admittedly speculative at this stage. While several aspects of the current study support the plasticity model, I need to emphasize that a crucial tenet of the model, i.e., the presence and dynamics of the anatomical spread of excitability, has been analyzed only indirectly in the current work. The analysis of excitability in the spatial dimension and its relationship to sequence structure, thus represents an interesting direction for future work.

Conclusion

I have shown in my dissertation work that sleep is an active and functionally important aspect of hippocampal physiology, far from being a stationary 'off-line' state. First, I developed appropriate tools for visualizing and annotating long time-scale data relevant to brain state dynamics. In addition to being well suited for brain state scoring, the intuitive and open source aspects of this software make it useful for a wide range of long time-scale electrophysiological data applications.

In my analysis of putatively homeostatic changes in excitability during sleep, I found that the overall firing rates of both pyramidal cells and interneurons of hippocampal layer CA1 decrease over the course of sleep. Surprisingly however, within ripple events pyramidal cell firing rates increased over the course of sleep, due to the fact that, on average, more of them participated per ripple, leading to an overall increase in synchrony. Contrary to models of homeostatic regulation of excitability during sleep (Tononi and Cirelli, 2006b), the observed firing rate decrease occurred during REM, and not non-REM, sleep, and both the overall firing rate decrease and the increased recruitment of pyramidal cells to ripples were predicted by theta (5-10 Hz) LFP power during REM sleep. These findings place the role of REM sleep into a novel perspective.

Lastly, in the study concerning specific and non-specific contributions to the memory content observed during sleep, our experiments and extensive analyses demonstrated that sleep related firing patterns in the CA1 hippocampal region can both predict and be modified by hippocampal spatial coding in a novel environment. However, the findings also show that various measures introduced for the quantification of 'replay'

and ‘pre-play’ phenomena have numerous drawbacks and should be interpreted carefully. Our detailed analyses and the comparison of traditional or more novel methods spell out some of the specific issues related to each of these methods and make several suggestions for future work.

While many questions, specifically pertaining to cellular mechanisms, and how the current results relate to other brain structures remain outstanding, the experiments performed in this thesis provide a framework into future research related to the structure of memory content during sleep. My results also bear importance regarding the relationship between, specific (memory) and relatively non-specific (homeostatic) roles of sleep in the hippocampus.

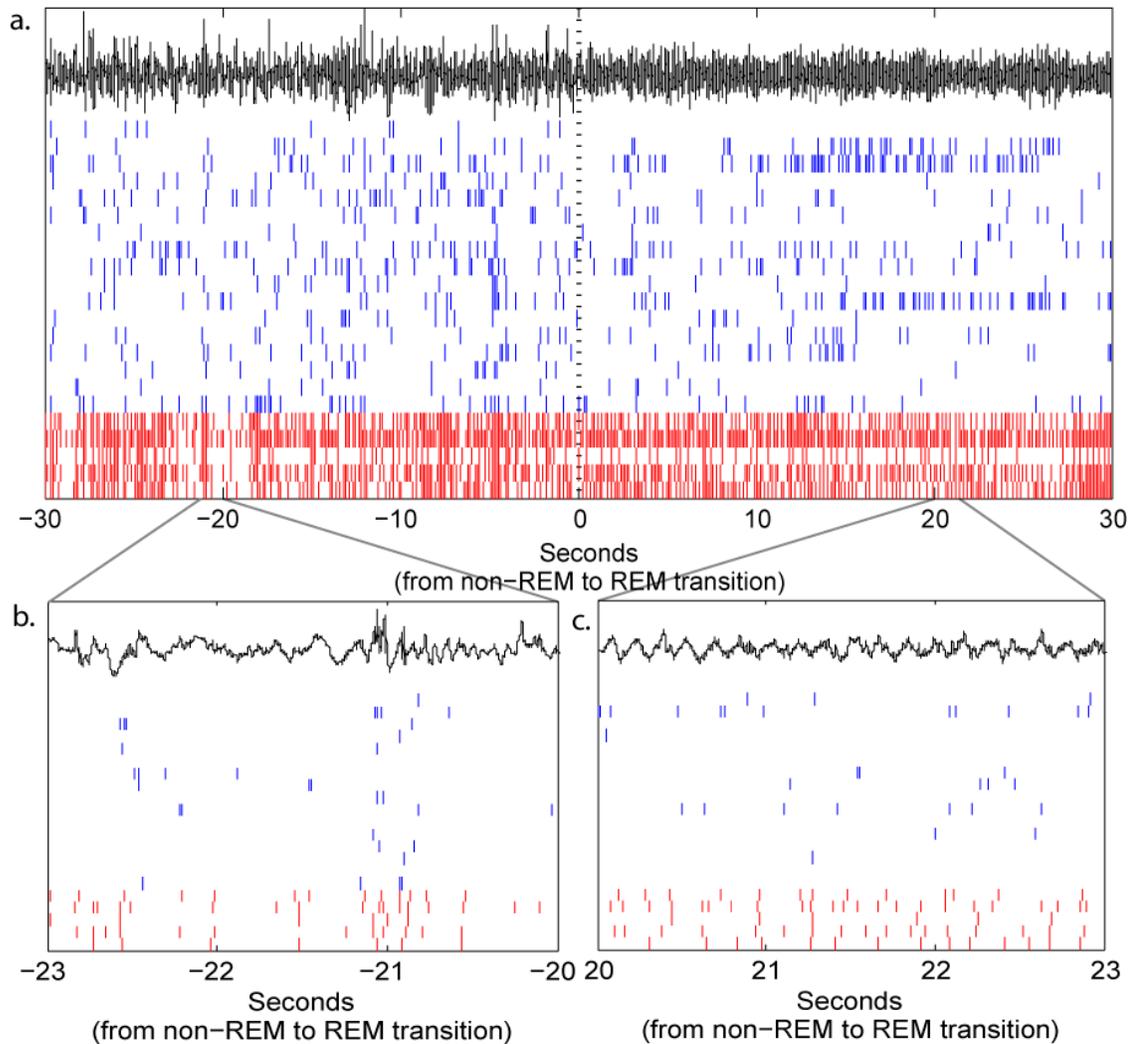


Figure 1.1: The non-REM to REM transition. **A.** Pyramidal layer unfiltered LFP (black), and the rasters for 17 putative pyramidal cells (blue) and 5 interneurons (red) are shown for the sixty seconds around a non-REM to REM transition. **B** and **C** show 3 seconds in non-REM and REM respectively. Note that the LFP during non-REM is irregular and high amplitude while the REM LFP is lower amplitude (here only slightly) and displays highly regular theta (5-10 Hz) oscillations. Also evident in REM is the appearance of ‘episode’ fields of sustained firing in some pyramidal cells (two such fields can be seen at the top right of **A**). Conversely, while pyramidal cell discharge looks highly unstructured at the resolution of tens of seconds, when viewed a few seconds at a time it is apparent that neural firing is structured into brief sporadic synchronous population firing events which are coupled to high-amplitude and frequency ripple events in the LFP (**B**, second -21). Finally, though not immediately evident in this example, it should be noted that interneuron firing rates increase during REM.

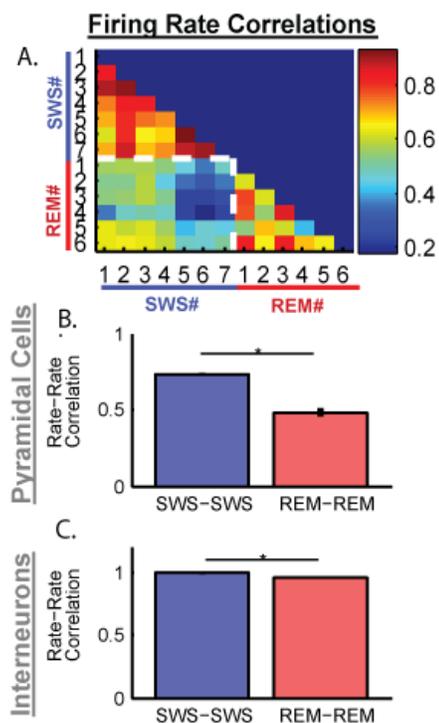


Figure 1.2: Firing rates are persevered across non-REM but not REM episodes (from Grosmark et al., 2010). **A.** Correlation matrix for pairs of firing rate vectors from 7 non-REM (SWS) and 6 REM episodes from 1 uninterrupted sleep session. Firing rates are less preserved between pairs of REM than non-REM episodes and are least similar between these states. **B.** Group data for pyramidal cell firing rate vectors between all pairs of non-REM (blue) and REM (red) episodes. **C.** Interneuron firing rate vectors were likewise more highly preserved between non-REM than REM episodes, despite the fact that the firing rate of these interneurons was significantly greater during REM (data not shown).

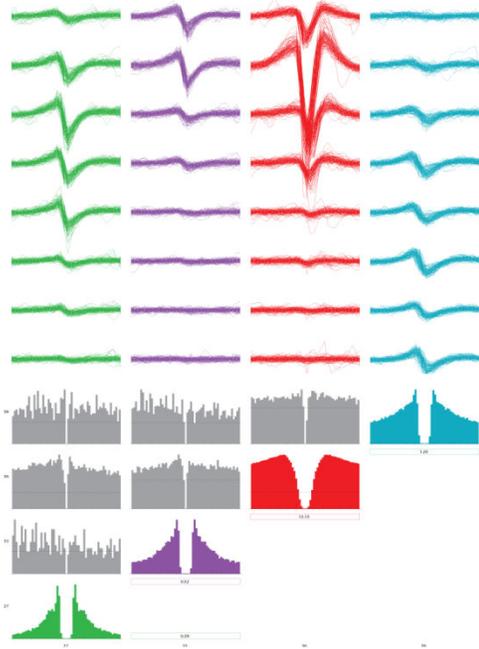


Figure 1.3: Extra-cellular silicon-probe recordings. Waveforms (top) across eight of a single silicon probe shanks, and auto- and cross-correlograms (bottom), shown for one interneuron (left, purple) and 3 putative pyramidal cells recorded from hippocampal CA1 pyramidal layer. Note that while waveforms between cells may look similar on one site, they display distinct waveforms on other sites.

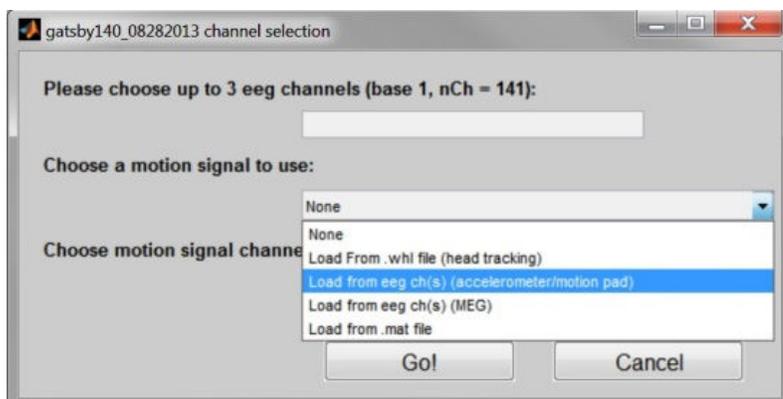


Figure 2.1: *TheStateEditor* Channel Selection Window. This window is used for selecting up to three channels for visualization, annotation and state classification as well as an optional motion signal. Note that this window is skipped if the *TheStateEditor* auxiliary file, created automatically on the first loading of a session in *TheStateEditor*, is detected.

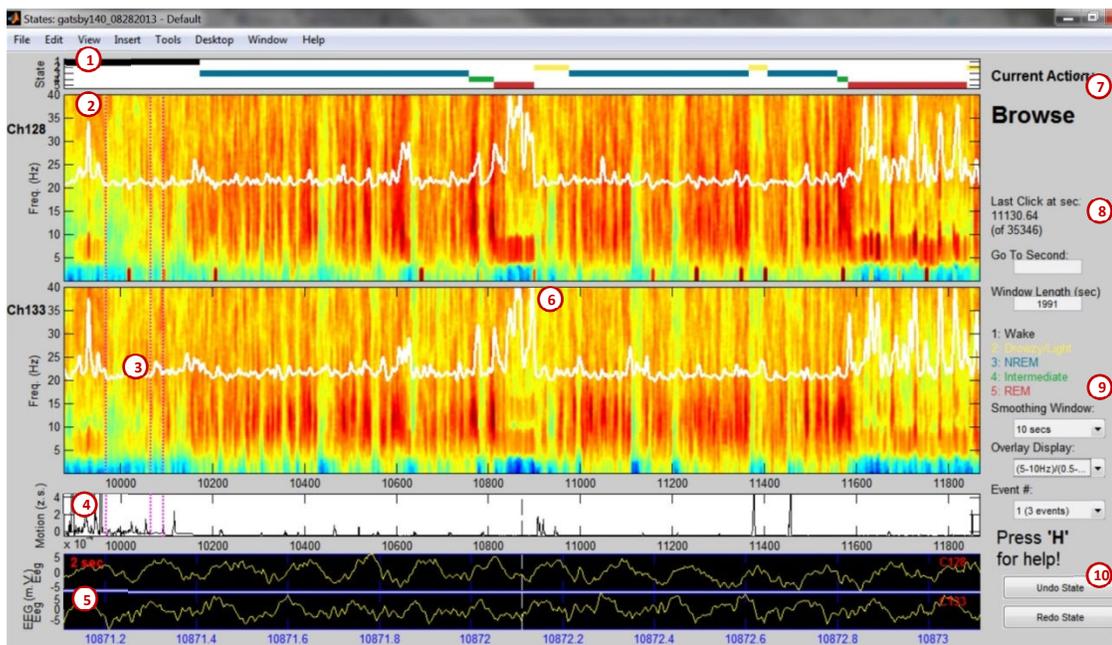


Figure 2.2: TheStateEditor Main Console: 1) The *StateRibbon* showing five labeled states. 2) Time resolved spectrogram for one of the two selected LFP channels (up to three channels may be selected). 3) Three events (vertical pink dashed lines) selected during waking period. 4) Motion display - this example shows data acquired from a head-mounted accelerometer. 5) Raw LFP display of the two selected channels. 6) Theta/Delta ratio overlay display (selected from the '*Overlay Display*' drop down menu on the right). 7) Current selected action. 8) Time point of last click as well as '*Go To Second*' and '*Window Lengths*' boxes for manual specification of these parameters. 9) '*Smoothing Window*', '*Overlay Display*' and '*Event #*' selection (note that Event class 1 is currently selected). 10) '*Undo*' and '*Redo*' buttons which can be useful for fixing mistakes during state labeling.

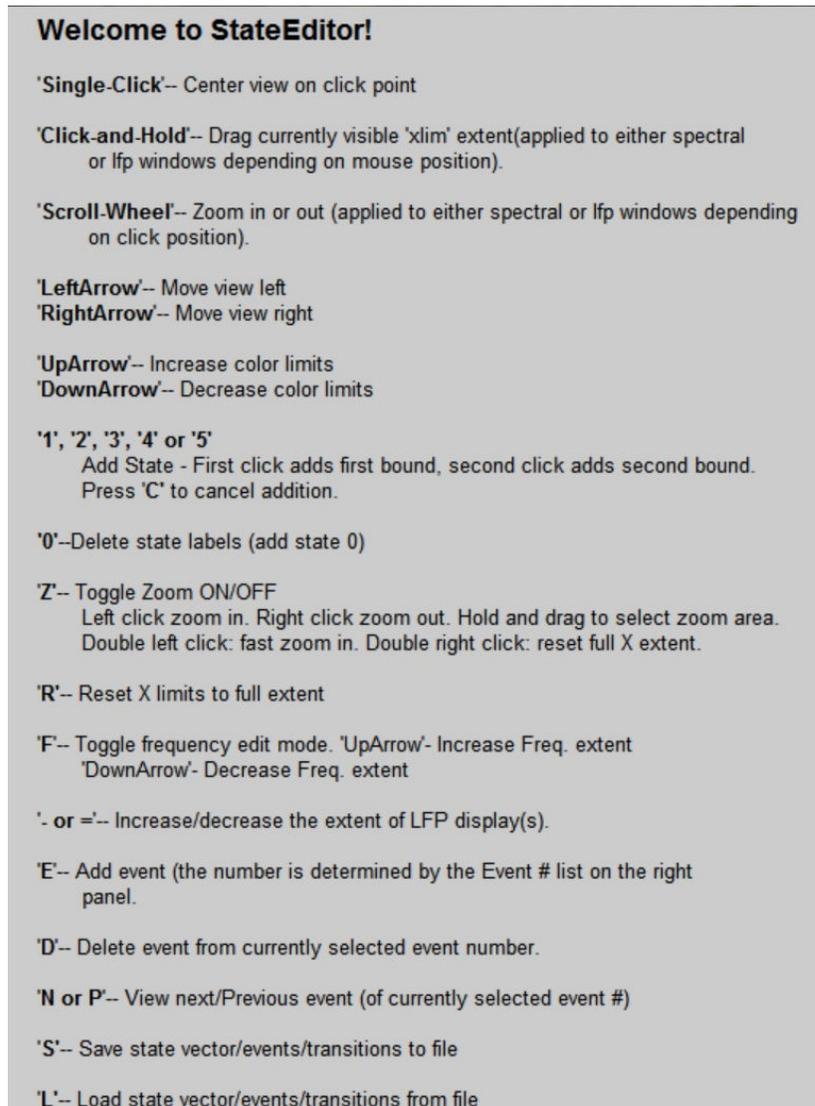


Figure 2.3: *TheStateEditor* Help Panel: This panel summarizes *TheStateEditor's* main controls and is accessed from the main console by pressing 'H'.

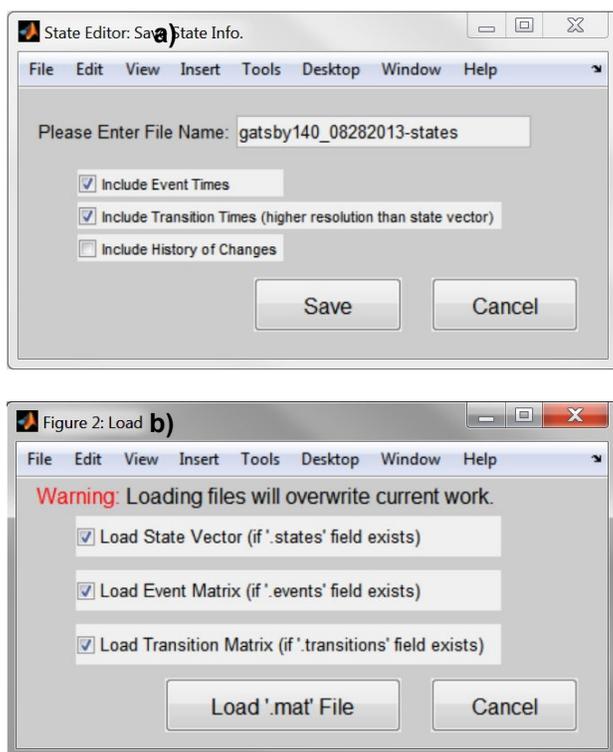


Figure 2.4: *TheStateEditor* Saving and Loading Consoles: *TheStateEditor* outputs up to four different data object (see *Results*) which are saved into one *.mat* file (panel a) and used for further analysis or loaded back into the session for further editing (panel b).

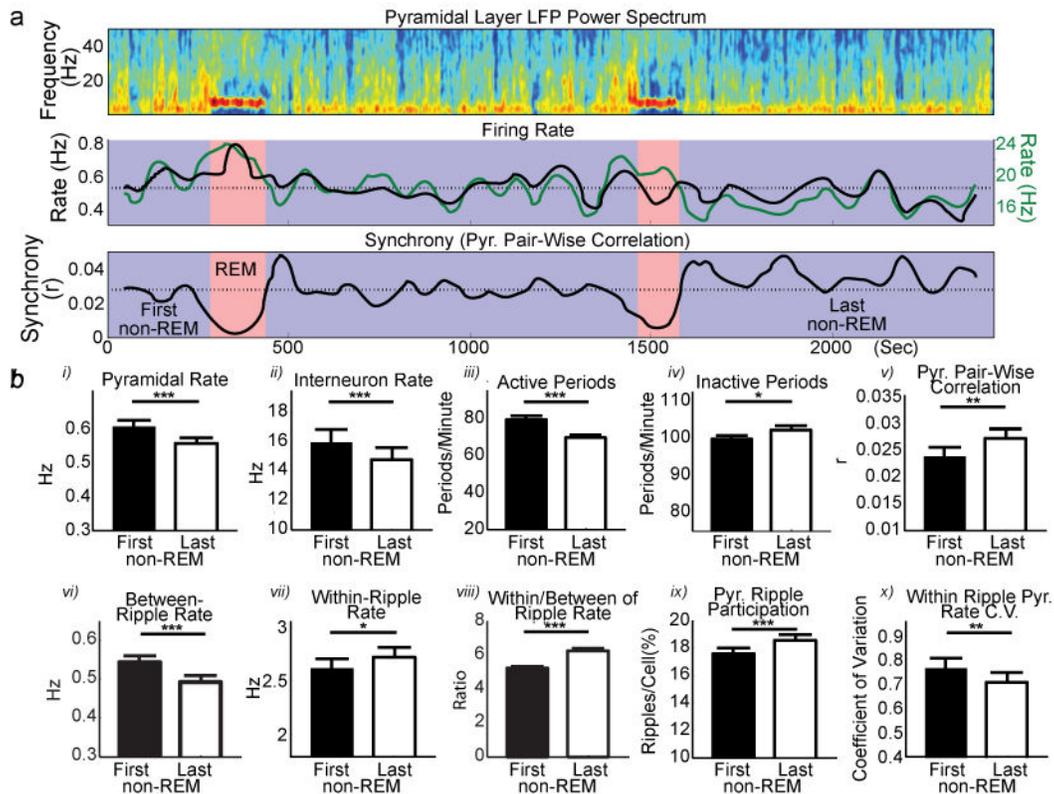


Figure 3.1. Excitability changes across sleep. a. Time-resolved spectrum (top), and smoothed (60-sec) mean firing rate changes of pyramidal cells (black) and interneurons (green; middle panel) and mean pyramidal pair-wise correlation ('synchrony', bottom) across representative non-REM (blue) and REM (pink) episodes from one uninterrupted sleep session (dotted lines show session means). b. Significant changes across sleep, calculated between the first (black bars) and last (white bars) non-REM episodes of sleep, firing rates (panels *i* and *ii*), incidence of periods of high-frequency LFP activity (*iii*) and inactivity (*iv*), synchrony (*v*), firing rate of pyramidal cells between ripples (*vi*) and within ripples (*vii*), ripple-induced firing rate modulation (*viii*), percentage of ripples in which individual pyramidal cells participated (i.e., fired at least one spike, *ix*), and coefficient of variation of within-ripple firing rate across cells (*x*). Note that the decrease in firing rates across sleep is concomitant with increasing synchrony. All comparisons were carried out as sign-rank (paired) tests. * $p < 0.05$; ** $p < 0.005$; *** $p < 0.0005$.

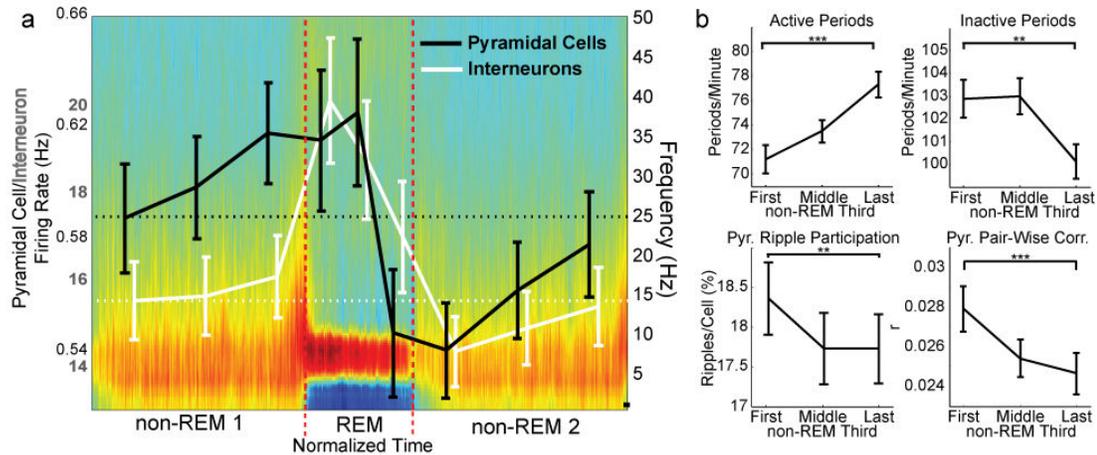


Figure 3.2: Excitability changes within non-REM episodes are opposite to those across sleep. a. Time normalized power spectra of adjacent non-REM_n--REM--non-REM_{n+1} episodes (mean of n=45 non-REM_n--REM--non-REM_{n+1} cycles) and corresponding firing rates (\pm SEM) of pyramidal cells and interneurons shown within thirds of non-REM and REM episodes. Note that firing rates increase within non-REM episodes and decrease within REM episodes (Figure S2). Horizontal lines, mean rates at the beginning of the non-REM_n--REM--non-REM_{n+1} cycle. b. Incidence of LFP high-frequency activity and inactivity epochs, percentage of ripples in which pyramidal cells participate, and pair-wise correlation of pyramidal cells across thirds of non-REM. Note opposite changes as those observed across sleep (Fig. 1). * $p < 0.05$; ** $p < 0.005$; *** $p < 0.0005$.

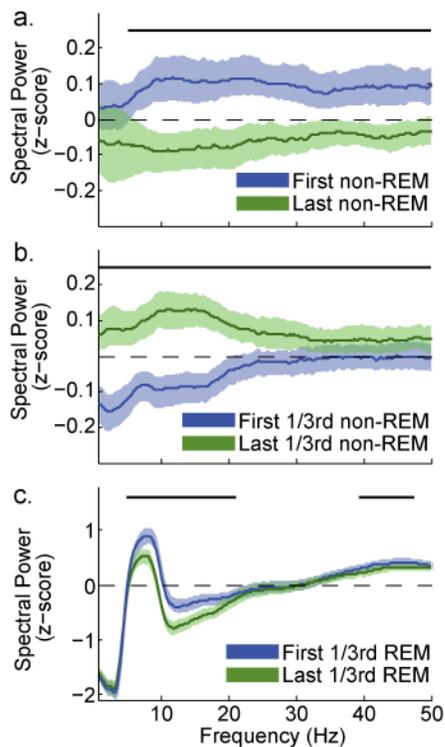


Figure 3.3: LFP spectral power decreases across sleep and within REM episodes and increases within non-REM episodes. For each session pyramidal layer LFP spectra were normalized for each frequency bin as the z-score of non-REM power. a. Spectral power decreases from the first to last non-REM episode of a sleep session across a wide range of frequencies ($n = 22$ sleep sessions). Shaded regions show 95% confidence intervals. Top black bars show frequencies for which sign-rank test $p < 0.05$. b. Spectral power changes within non-REM episodes ($n=82$). c. Spectral power changes within REM episodes ($n = 45$).

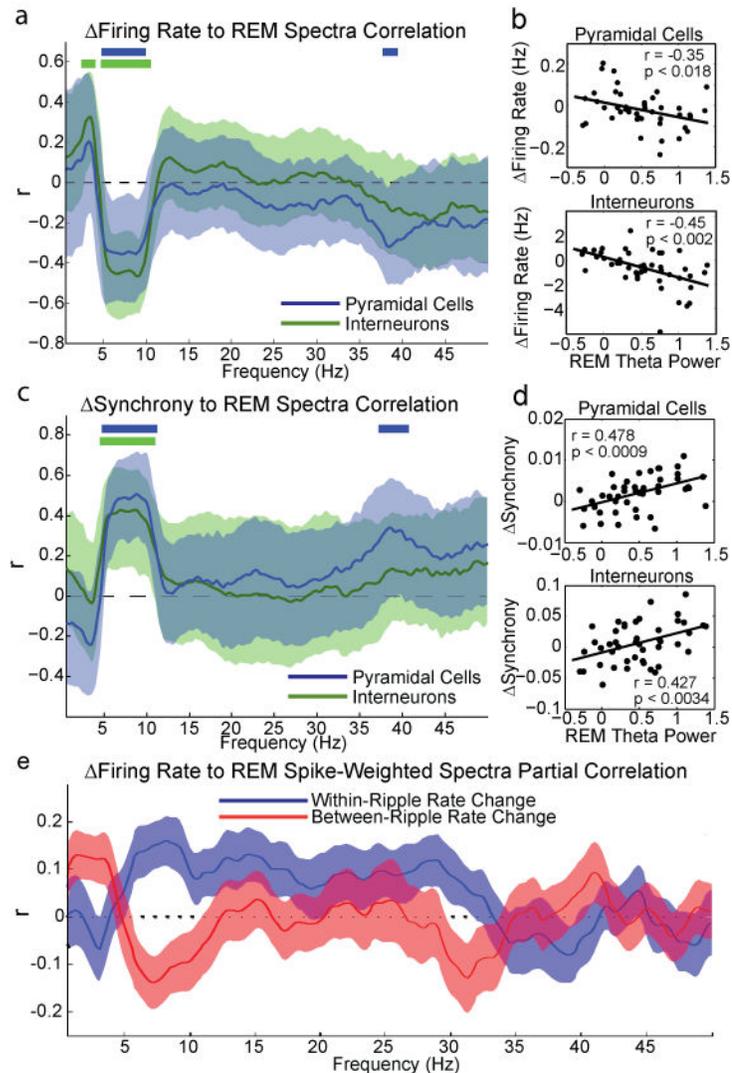


Fig. 3.4. REM sleep affects firing patterns in non-REM. a. Correlation values between firing rate changes between non-REM_n and non-REM_{n+1} episodes versus LFP power of the intervening REM (mean and 95% confidence intervals) in the 0-50 Hz frequency range. Top solid bars indicate frequency bands in 0.25Hz steps of significant correlation. Note significant effect of REM theta power on rate changes between successive non-REM episodes for both pyramidal cells and interneurons. b. Mean firing rate changes (Hz) between non-REM_n and non-REM_{n+1} episodes (y axis) as a function of the theta power during the interleaving REM episode. Power during REM was normalized by the power of concatenated non-REM episodes (z score). c and d, Same as a and b but for synchrony change (pairwise correlation) between successive non-REM episodes. e. Correlation values between firing rate changes across sleep and spike-weighted spectra (Sp. W. S.; see Online Methods) during REM for pyramidal cells in the 0-50 Hz frequency range. Results are shown separately for spikes that occurred between ripples (red) and within-ripple (blue) events.

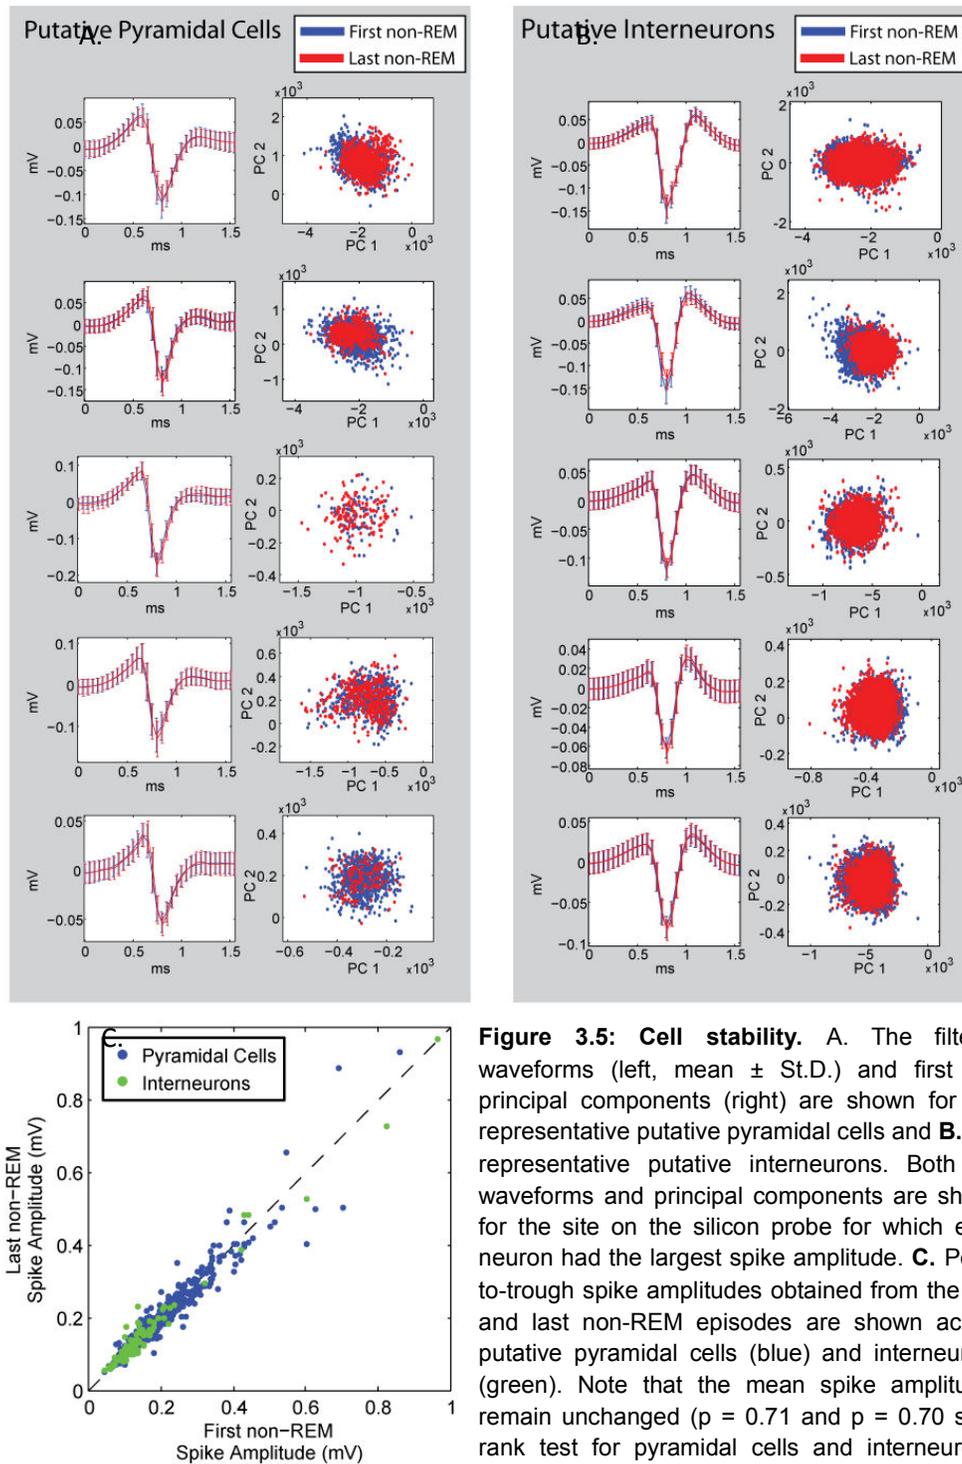


Figure 3.5: Cell stability. A. The filtered waveforms (left, mean \pm St.D.) and first two principal components (right) are shown for five representative putative pyramidal cells and B. five representative putative interneurons. Both the waveforms and principal components are shown for the site on the silicon probe for which each neuron had the largest spike amplitude. C. Peak-to-trough spike amplitudes obtained from the first and last non-REM episodes are shown across putative pyramidal cells (blue) and interneurons (green). Note that the mean spike amplitudes remain unchanged ($p = 0.71$ and $p = 0.70$ sign-rank test for pyramidal cells and interneurons, respectively) and amplitudes are highly preserved across cells from the first to the last non-REM of each session ($r = 0.94$, $r = 0.98$ for pyramidal cells and interneurons, respectively).

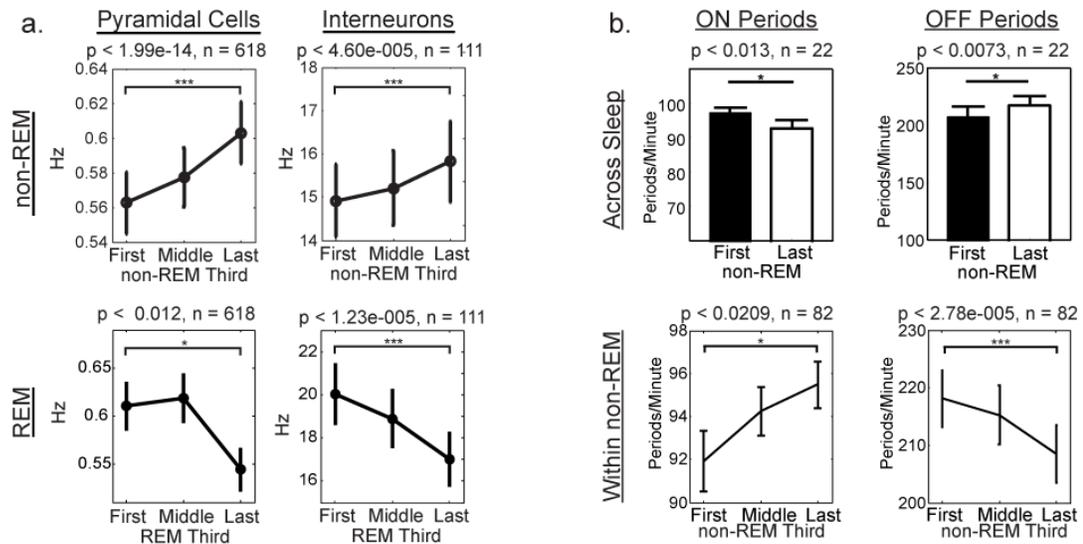
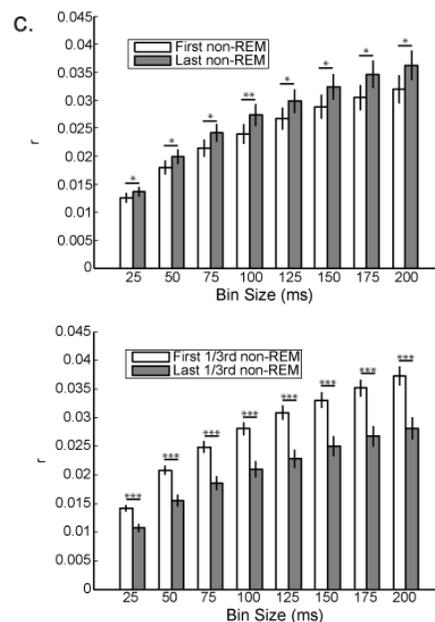


Figure 3.6. Firing rate and synchrony changes across sleep and within non-REM or REM episodes. **a.** Firing rate increases of pyramidal cells and interneurons from the first to last third of non-REM episodes (top row) and rate decreases from the first to last third of REM episodes (bottom row). The non-REM data shown here reflect or combination of all the pre-REM and post-REM non-REM episodes shown in Fig. 2a. **b.** 'OFF' periods were defined as periods lasting at least 50ms during non-REM in which no spikes were detected from any neuron or combination of neurons (see Supplementary Methods, Vyazovskiy et al. 2009). 'ON' periods were defined as periods between 50 and 4000ms long in which at least 10 spikes were detected from any neuron or combination of neurons. Only ON periods occurring immediately subsequent to OFF periods were considered. Note that the incidence of ON periods decreased across sleep (**b**, top left panel) but increased within non-REM episodes (**b**, bottom left panel), while the inverse pattern of changes was observed for OFF periods (**b**, right column). **c.** Changes of mean pair-wise correlation between pyramidal cell pairs ('synchrony', as in Fig. 2b) across sleep (i.e. first to last non-REM episode of each session, $n = 22$, **c**, top-panel) and within non-REM episodes (first to last third of each non-REM episode, $n = 82$, **c**, bottom panel). Note that synchrony changes are present across a wide range of bin sizes. (* $p < 0.05$, ** $p < 0.005$, *** $p < 0.0005$, sign-rank test).



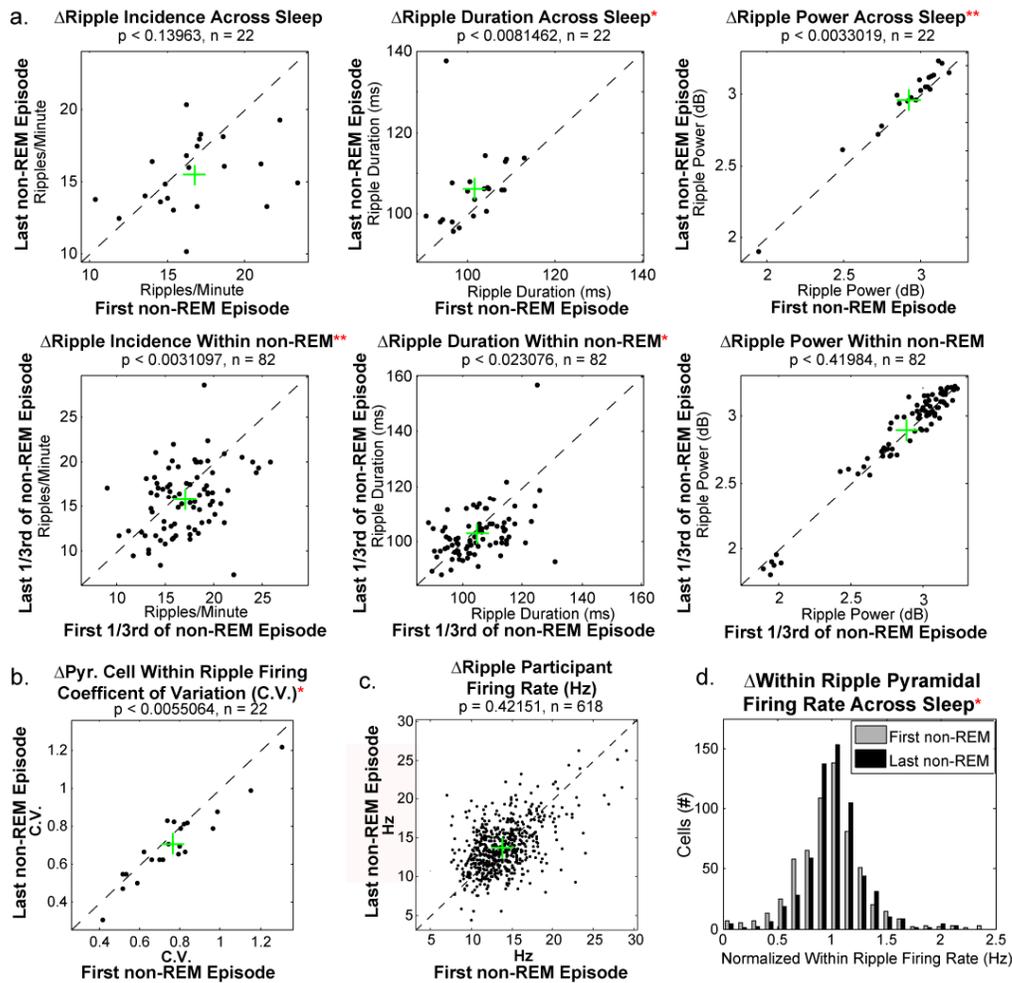


Figure 3.7: Ripple changes across sleep and within non-REM episodes. LFP ripples show distinct changes across sleep sessions (a, top row, green crosses show population mean) and within non-REM episodes (bottom row). Each dot represents either a sleep session (top row) or a non-REM episode (bottom). **b.** Changes in the coefficient of variation (c.v.) of firing rates within ripples across sleep. **c.** Firing rates of each of 618 pyramidal cells within those subsets of ripples in which they participated (i.e. fired at least one spike) do not change significantly between the first and last non-REM of each sleep session. Consequently, increases in within-ripple pyramidal firing rates across sleep (Fig. 1b) were due to individual cells participating in more ripples rather than discharging at a higher rate in those subset of ripples in which they participated. **d.** Distribution of firing rate changes within ripples across sleep (c, $n = 618$, $p < 0.029$, sign-rank test). For panel d, within-ripple firing rates were normalized for each cell as the ratio of the within-ripple rate to the mean within-ripple firing rate of the neuron in the entire sleep session. The increase of within-ripple firing rate and the concurrent decrease in C.V. imply a preferential within-ripple rate increase for those pyramidal cells which show low within-ripple firing rates in early sleep (* $p < 0.05$, ** $p < 0.005$, sign-rank tests).

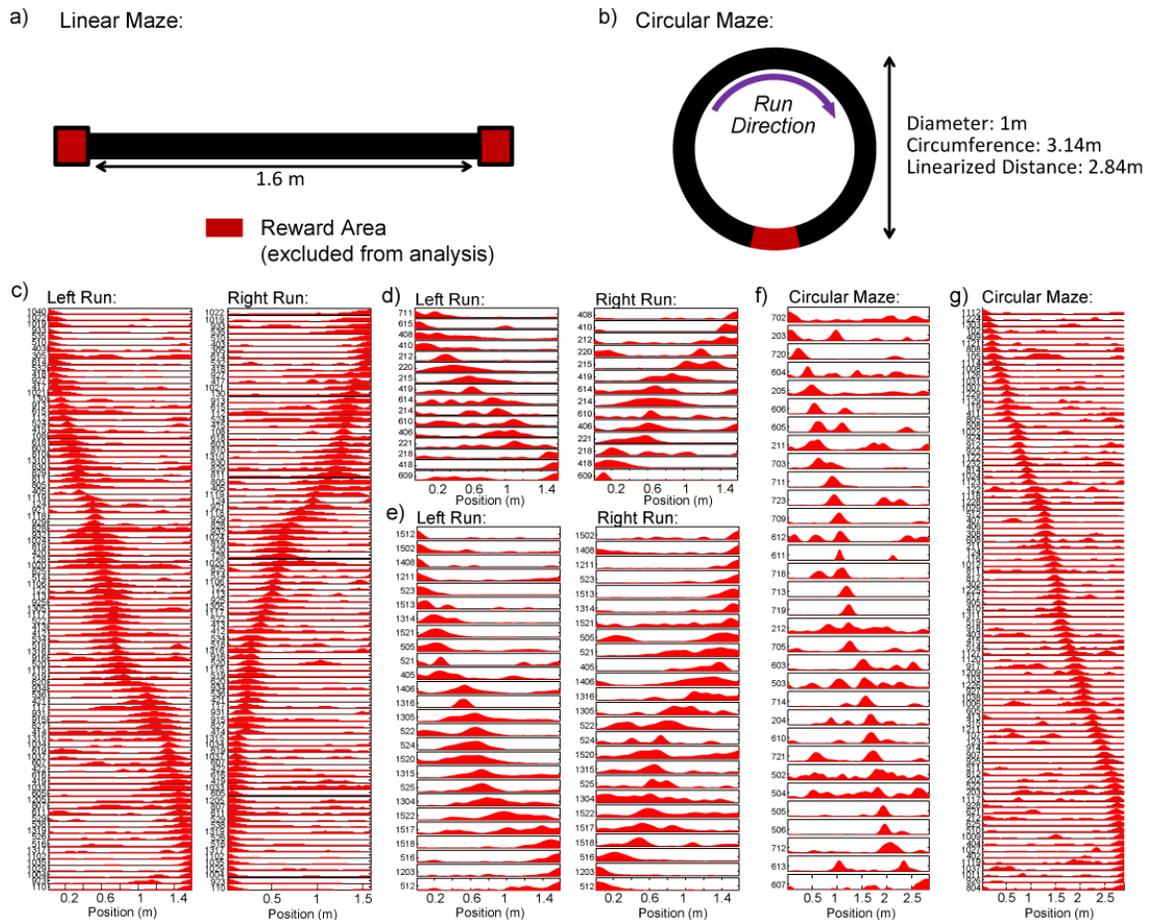


Figure 4.1: Novelty run place fields: water-deprived CA1 implanted rats were made to run either a linear (a) or a circular (b) maze for water reward. In each case the rat had no previous experience with either the experimental room or the maze. The reward was delivered at both ends of the linear arm, or at a constant position along the circular track. Panels c, d, e, show the smooth firing rate-by-position vectors for the three linear maze novelty run sessions. Place cells are shown as aligned by the peak place preference. For the linear maze sessions (c, d, e) 94, 16, and 26 left direction place cells and 93, 14 and 25 right direction place cells were found from a total of 128, 28 and 39 well isolated pyramidal cells for the session shown in c, d, and e, respectively. For the circular maze sessions (f and g), a total of 32 and 92 place cells were found out of a total of 35 and 92 well isolated pyramidal cells for f and g, respectively. Note that while many place cells show more than one place field, only the place field with the highest peak firing rate was used for the Pre-play/Replay analysis.

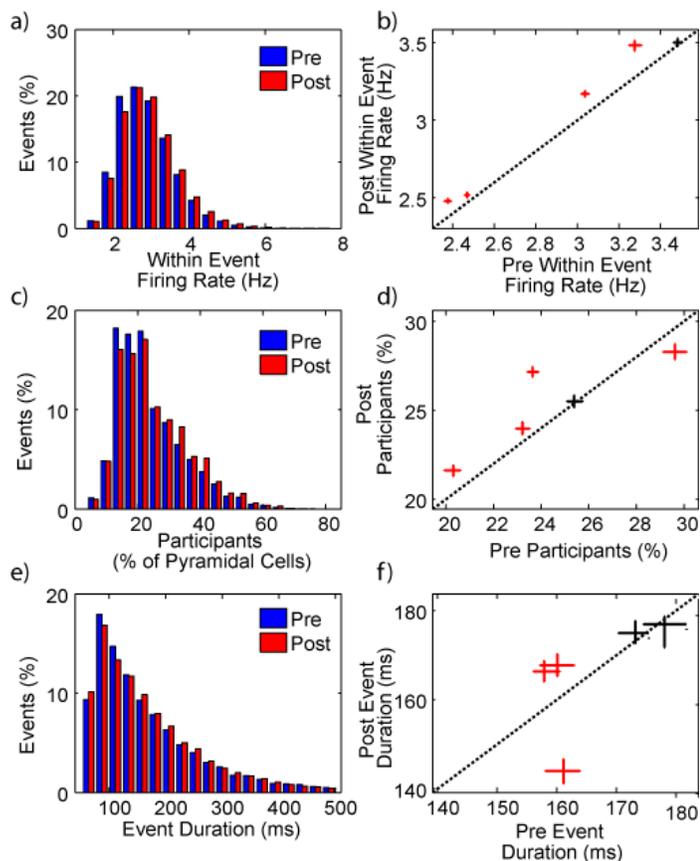


Figure 4.2: Population Events During Pre and Post Sleep. Panels **a**, **c**, and **e** show the Pre and Post distributions across all Pre and Post events ($n = 22,042$ Pre and $22,318$ Post events). In panels **b**, **d**, and **f** the comparisons are shown by session, horizontal and vertical bars show the 95% bootstrapped confidence intervals for the Pre and Post respectively. Significant Pre, Post differences ($p < 0.025$ rank-sum test) are colored in red. Notably, an increase in within-event firing rates was consistently observed from the Pre to Post epochs (**a** and **b**). Likewise the percentage of pyramidal cells participating (firing at least one spike) in each event increased from the Pre to the Post epoch (**c** and **d**). Finally, while the population average showed a slight increase in event duration from the Pre to the Post epochs this trend was not as consistent as that observed in the above changes in excitability observed across sessions (**f**).

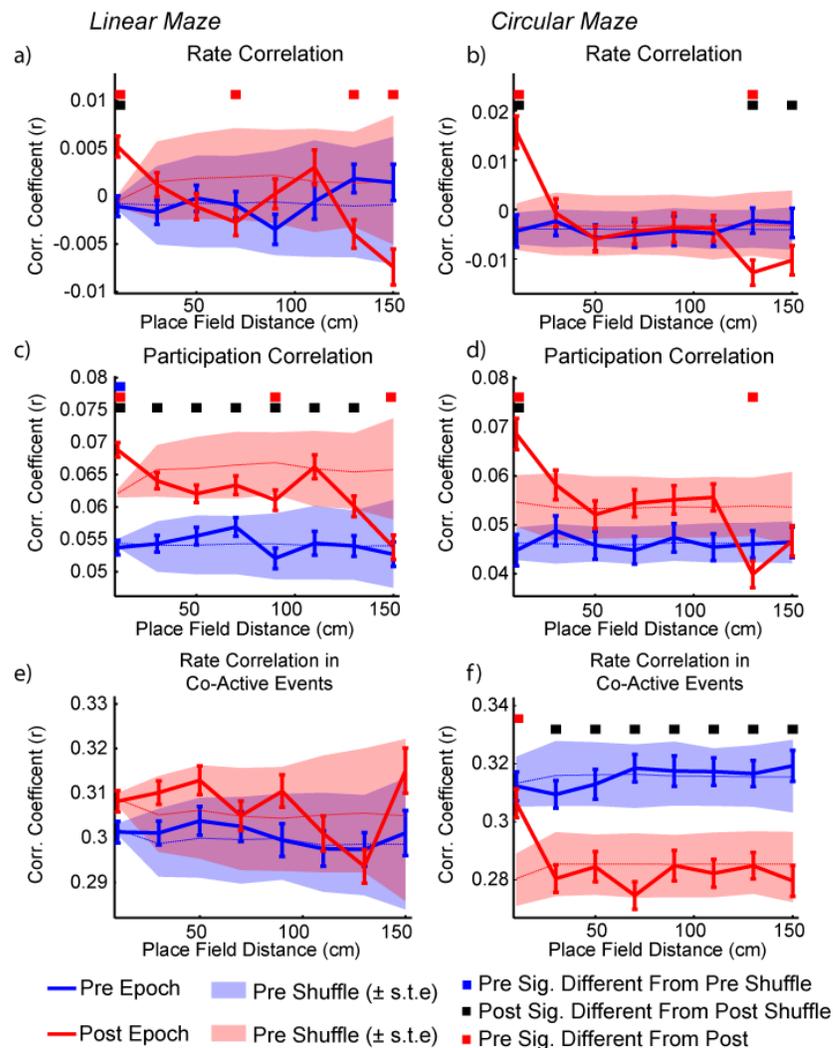


Figure 4.3: Cell Activity Co-Modulation By Place Cell Distance. Three measures of activity co-modulations were derived for all non-local shank place cell pairs: 1) correlations of each cell's within-event firing rate vector (**a** and **b**), 2) correlation of each cell's binary participation vector (an estimate of how often the cells were co-active, **c** and **d**), and 3) the correlation of firing rates restricted to events in which both cells were active (**e** and **f**). Only events that met the criteria for inclusion in the sequential replay analysis were included in this analysis. Linear track ($n = 7,738$ pairs) and circular track ($n = 3,802$ pairs) sessions were analyzed separately. These pairs were subsequently binned by the distance between their place fields (20 cm bins). For circular maze sessions this distance was assessed as the shortest linearized path (either forward or backward) between the two place fields. For each 20 cm bin, blue and red lines show the mean of the Pre and Post distribution respectively. Error bars show \pm standard error for each bin. Dashed blue and red lines show the mean of the null distribution of 1000 shuffles of the place field distance for the Pre and Post distributions respectively. Shaded areas show the 99% confidence interval of shuffled distributions. Three measures of significance were used 1) for each bin the Pre was determined to be significant if its mean value was outside of the 99% confidence interval of its shuffled distribution (blue square), 2) the same comparison was made between the Post and Post shuffled distribution (red squares), and 3) the Pre and Post distributions were compared using rank-sum tests Bonferroni corrected for bin number ($p < 0.05$, after Bonferroni correction: $p < 0.00625$). Note the significant and opposite effects of close and far place fields on the co-modulation of firing rate (**a**, **b**) and participation (**c**, **d**). Also note the small, but statistically significant suppression in co-activity during Pre for cells which will subsequently show nearby place fields on the maze.

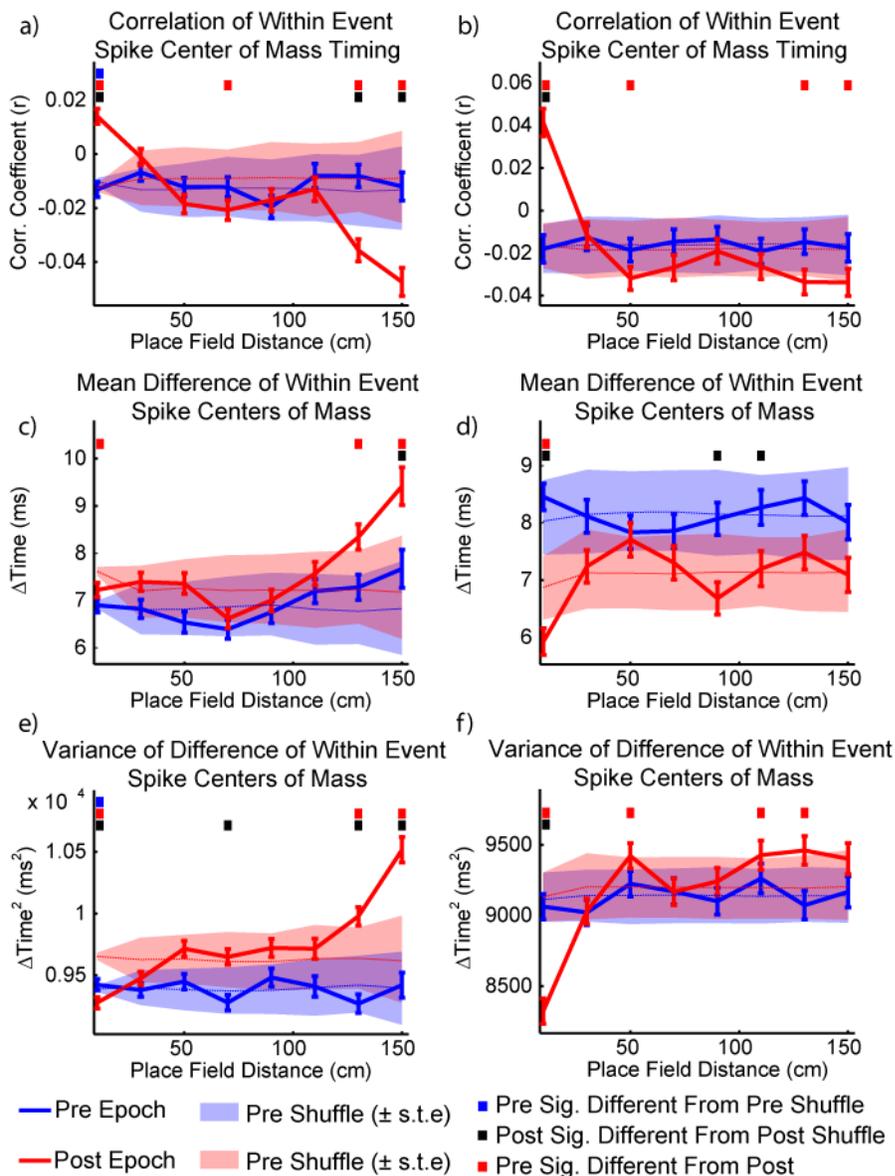


Figure 4.4: Cell Timing Co-Modulation By Place Cell Distance. To assess the co-modulation of place cell spike timing by cell distance all three measures were taken 1) the correlation of the timing each cell's center of mass of spiking across all events in which they were coactive (**a** and **b**), 2) the absolute value of the mean within-event temporal difference between each cell's center of mass of spiking (**c** and **d**), and 3) the variance of this difference. For first of these measures, for each event the timing of each cell's center of mass was referenced to the center of mass of all detected pyramidal cell spikes in that event (**e** and **f**). All graphical conventions and statistical tests are the same as in the last figure. Note that Post epoch cells with nearby place fields tend to fire in close temporal proximity to each other (**a** and **b**), with decreased temporal variability (**c** and **d**), and this temporal proximity is co-modulated (**a** and **b**), while the opposite is true for pairs of place cells with far place fields and absent during Pre epoch. Also potentially intriguing is the small but significant suppression of co-modulation of timing and elevated variability of pairs of cells during Pre which will have nearby place fields during the subsequent maze epoch, though this effect only observed in the linear track condition (first bins of **a** and **e**).

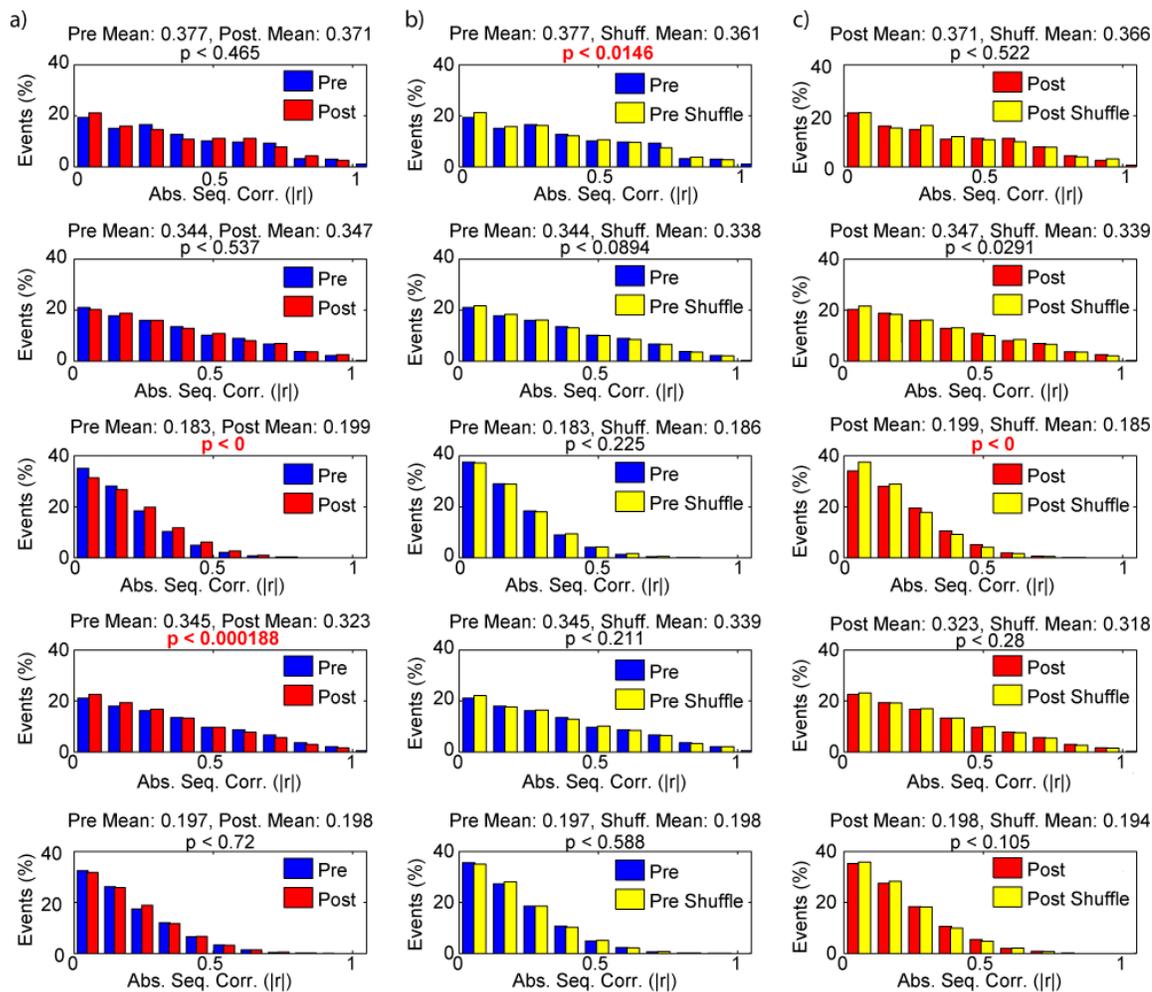


Figure 4.5: Rank-Order Sequence Correlation Pre-Play and Replay: The correlation derived pre-play or replay score for each event in each session was established as the Spearman (rank-order) correlation between the order of place fields on the maze and the time of the center of mass of each cell's spiking activity in that event across all place cells which fired at least one spike in that event. For linear maze sessions, all analyses were carried out independently for the left and right run directions and subsequently re-combined to yield session-wide pre-play and replay scores. Column **a** shows the sequence correlation score comparison across all events in the Pre (blue) and Post (red) epochs, column **b** shows the comparison between the Pre (blue) and its corresponding shuffled distribution (yellow), and finally column **c** shows the Post (red) versus null comparison. Rows correspond to sessions, with the last two rows in each column corresponding to the circular maze sessions. For each panel, the y-axis indicates the percentage of events falling within that bin. All significance tests were performed via rank-sum tests, significant ($p < 0.025$) differences are highlighted in red type font. Note that only two sessions show significant Pre to Post changes in sequence correlation score strength (with these changes occurring in opposite directions), while only one instance of significant pre-play (**b**, first row) and only one instance of significant replay (**c**, third row) were observed using the sequence correlation methodology.

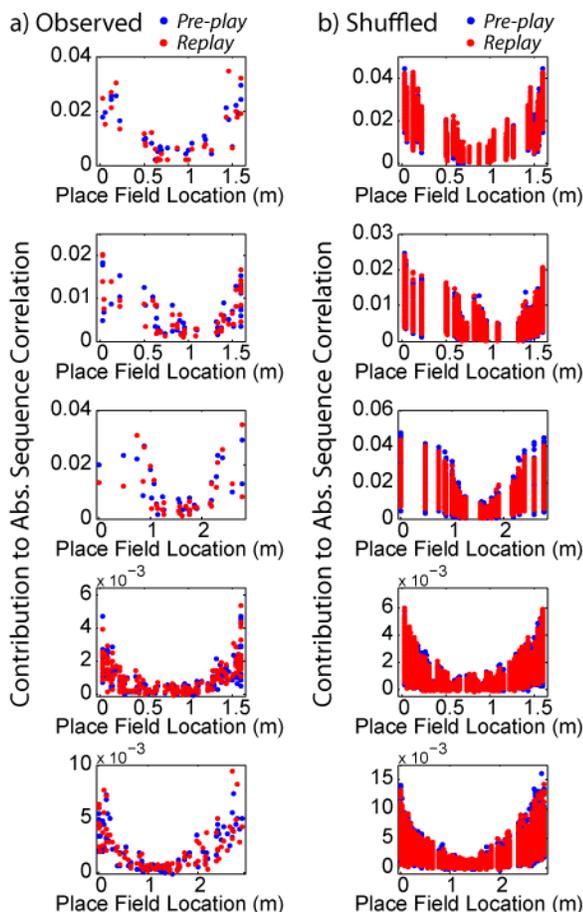


Figure 4.6: Per Cell Contribution To Overall Abs. Sequence Pre-Play and Replay. The contribution of each cell to the correlation coefficient can be determined by taking the mean of the $(n - 1)$ normalized z-product of each comparison across all comparisons for each cell (see Methods). Note that for a given panel in column **a**, summing all the values in either the pre-play (blue) or replay (red) conditions gives the value of the mean absolute sequence correlation coefficient for that condition. Note that in all cases, cells with place fields near the boundaries of the maze contribute much more heavily than do cells near the middle of the maze. Furthermore, this effect is also present if the place vectors are shuffled 100 times (**b**) and so cannot be physiological in origin.

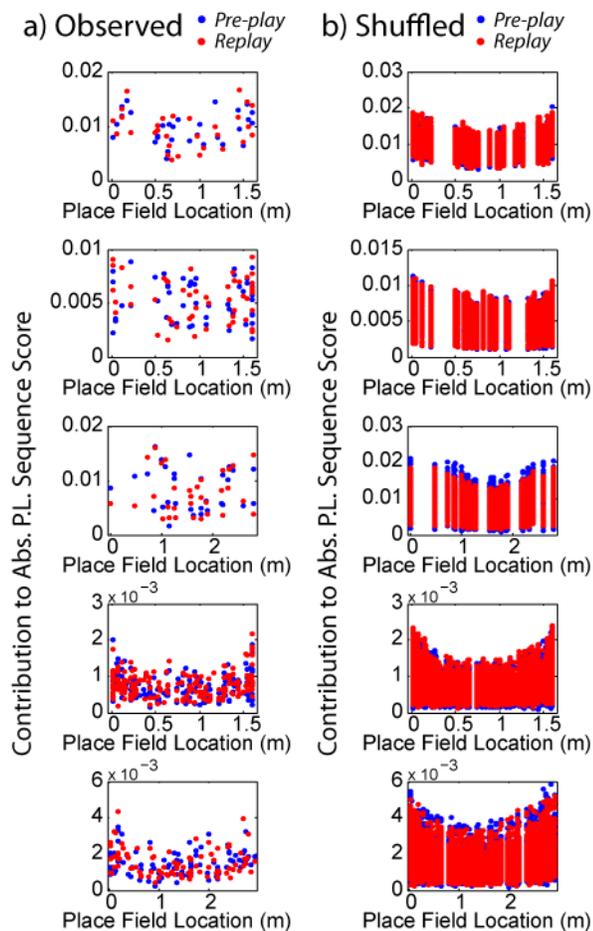


Figure 4.7: Per-Cell Contribution to Overall P.L. Abs. Sequence Pre-Play and Replay. The observed (a) per-cell contribution to absolute P.L. sequence score is shown for both the Pre (blue) and Post (epochs). Note that the summation of, for instance, all the blue values for a panel in a equals the mean of the absolute P.L. sequence similarity scores for Pre epoch events in that session. Per-cell contributions to mean absolute P.L. sequence similarity scores for 100 shuffles are shown in b.

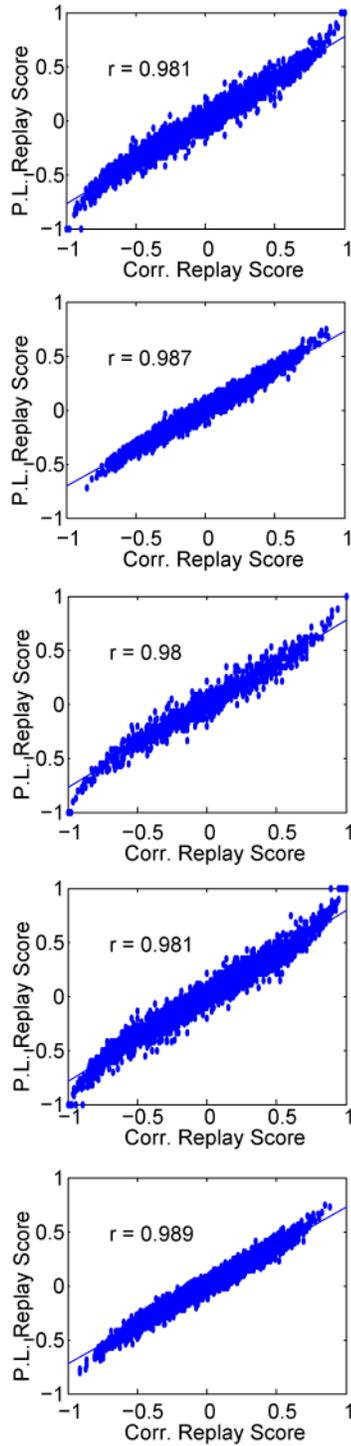


Figure 4.8: Correlation of Seq. Correlation with P.L. Sequence Score. Values are shown across both Pre and Post epochs. Note that while not identical, sequence score obtained using rank-order and P.L. sequence correlation methods highly correlated.

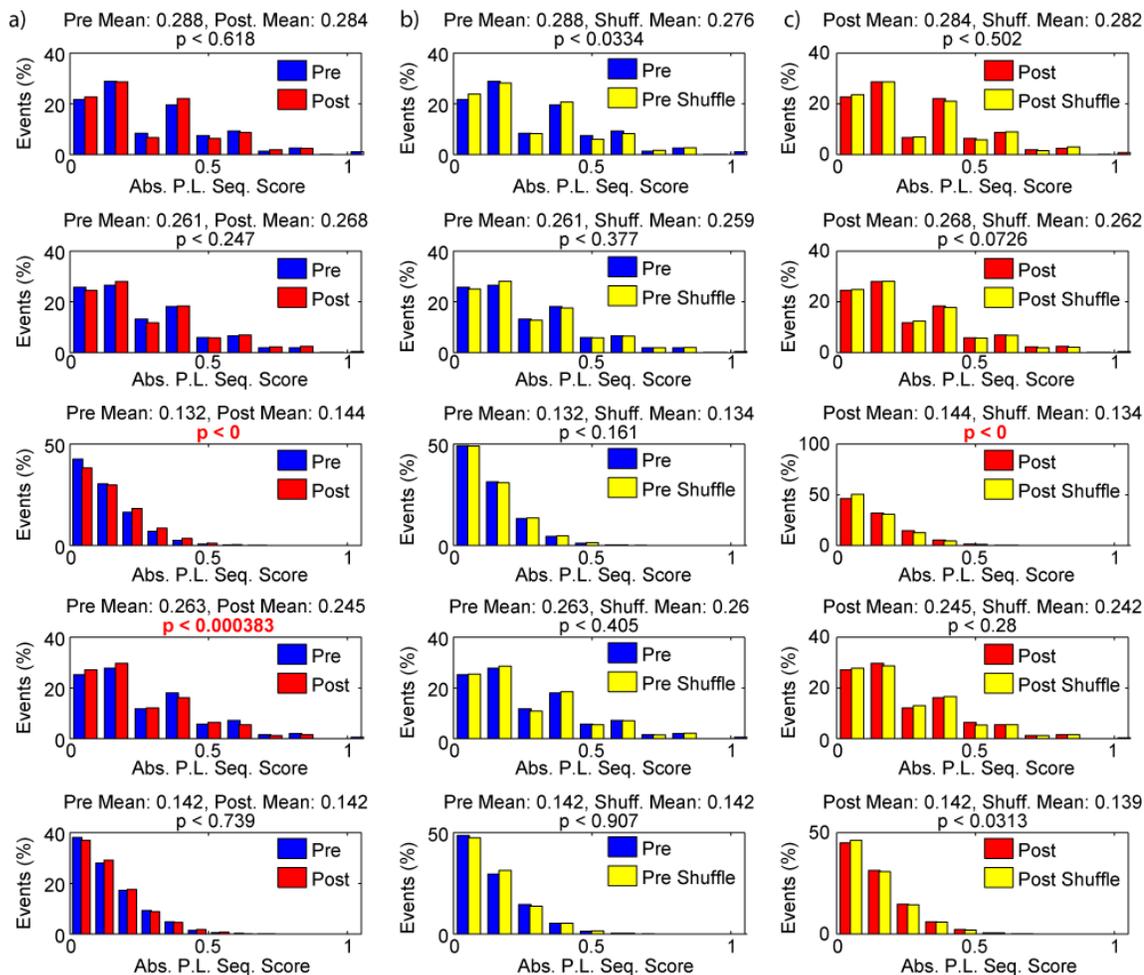


Figure 4.9: Paired Latency Sequence Pre-Play and Replay: Using the paired latency (P.L.) method, pre-play and replay were assessed on the same events that met the criteria for inclusion in the rank-order correlation sequence analysis. Similarly, the null distribution was estimated by randomly resampling the location of place fields on the maze without replacement 1,000 times. For linear maze sessions all analyses were carried out independently for each run direction and then re-combined. The Pre versus Post, Pre versus Shuffle, and Post versus Shuffle comparisons are shown in columns **a**, **b**, and **c** respectively. All significance tests were performed using rank-sum tests with the threshold for significance set at $p < 0.025$. Significant effects are highlighted in red font. Note that the only major qualitative difference in the effects observed between the rank-order correlation and paired latency method is that no case is significant pre-play observed.

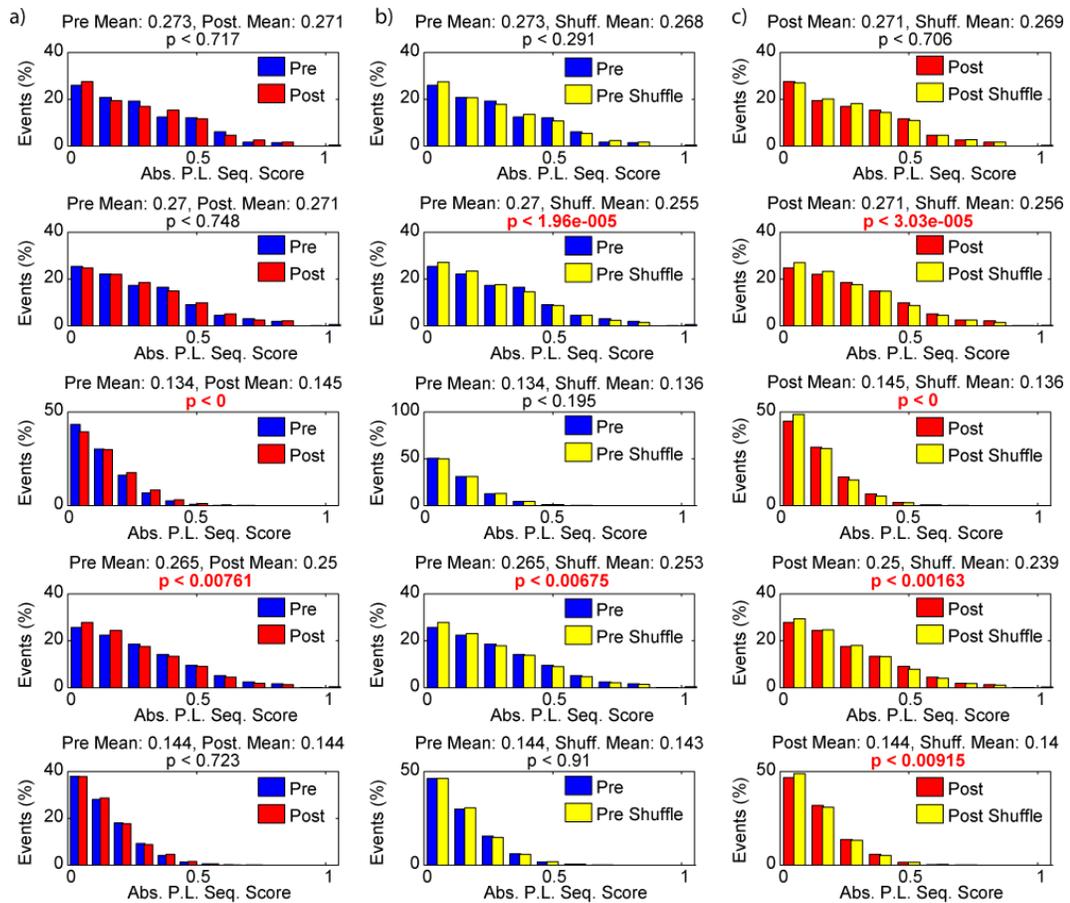


Figure 4.10: Paired Distance Sequence Pre-Play and Replay On Non-Local Pairs:

Taking advantage of the paired nature of the P.L. analysis, all the comparisons in which both place cells were isolated on the same shank of the silicon probe were censored by setting their paired latency values to zero. Note that by restricting our analysis to non-local pairs, two significant instances of pre-play (b, second and fourth rows) and four significant instances of replay (c, bottom four rows) were observed.

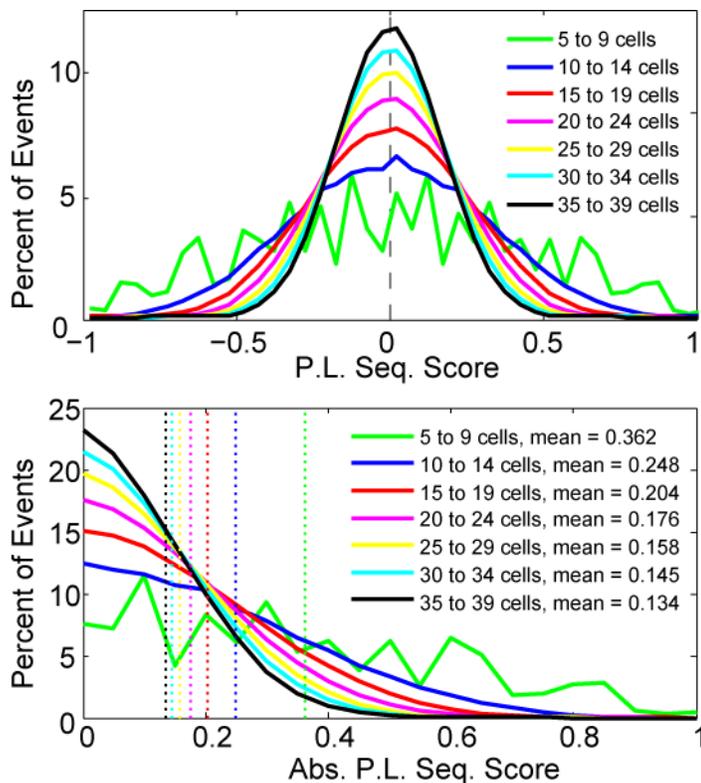


Figure 4.11: Null Distribution By Number of Participating Cells. The distribution of shuffled P.L. sequence score values are shown plotted by the number of participants (the number of cells firing at least one spike during the event). Note that if the absolute value is not taken (**a**), the decreases with number of participants while the mean stays constant at zero. However, both the mean and the variance of the absolute value of non-local P.L. score increase with number of participants (**b**). Vertical dashed lines show the mean of each absolute distribution. Note that this same effect is also present if the rank-order correlation method is used instead (data not shown). Notably, this effect, together with the observed elevation in within-event firing rate and participation in the Pre to the Post epochs, would be expected to bias Post events towards lower sequence correlation values, and may thus explain the relative absence of significant increases in mean sequence scores from the Pre to the Post epochs.

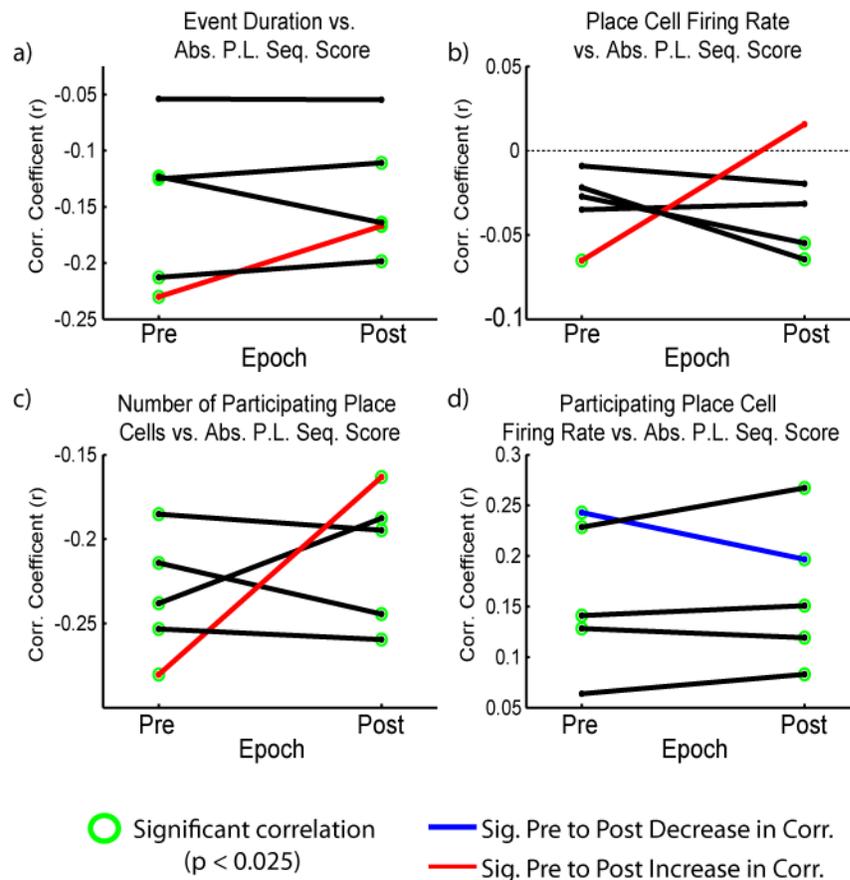


Figure 4.12: Pre-Play and Replay Correlations With Event Duration and Rate. Non-local shank absolute P.L. sequence pre-play or replay negatively correlates across population activity events with event duration (a) significantly in four Pre and four Post epochs. Likewise, within-event place cell firing rate and number of participating place cells both also negatively correlate with absolute P.L. sequence score either weakly (b) or strongly (c). Conversely, the firing rates of participating place cells (that is, excluding cases of 0Hz rate), strongly positively correlates with absolute P.L. sequence scores. Note that the correlations observed in (a, b, and c) are directly predicted by the null distribution of sequence score by number of participants. The significance of individual correlation coefficients or changes in correlation coefficient were assessed using the Fisher Z-test.

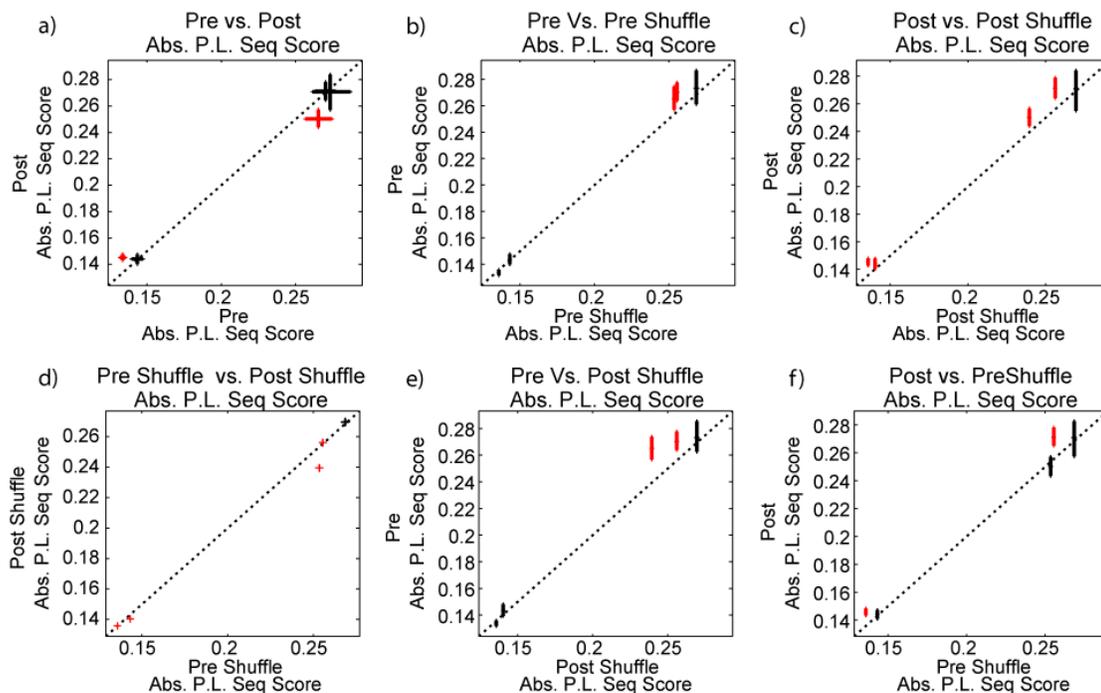


Figure 4.13: Null Distribution Comparison: Each panel shows the comparison between two conditions, error bars show the bootstrapped 95% confidence interval for the corresponding axis. Significant differences ($p < 0.025$, rank-sum test) are shown in red. Panels **a**, **b** and **c** show the comparison of Pre versus Post, Pre versus Pre Shuffled, and Post versus Post Shuffled distributions of non-local shank absolute P.L. sequence scores, recapitulating the three columns of figure N and are shown for reference. Panel **d** shows comparison between Pre Shuffled and Post Shuffled distributions, note that four out of 5 of these are significantly different, with Pre distributions tending to have a higher mean. When the observed Pre sequence score distributions are compared against the Shuffled Post distributions (**e**), the same two Pre epochs continue showing significant ‘pre-play’. However, when the Post distribution is compared against the Pre Shuffled distribution, only two Post epochs retain a significant ‘replay’ effect (**f**), which may suggest that changes from Pre to Post in the background activity determinate of the null distribution, may contribute to the enhancement of the signal to noise characteristics of sequential replay, during the post novelty epoch.

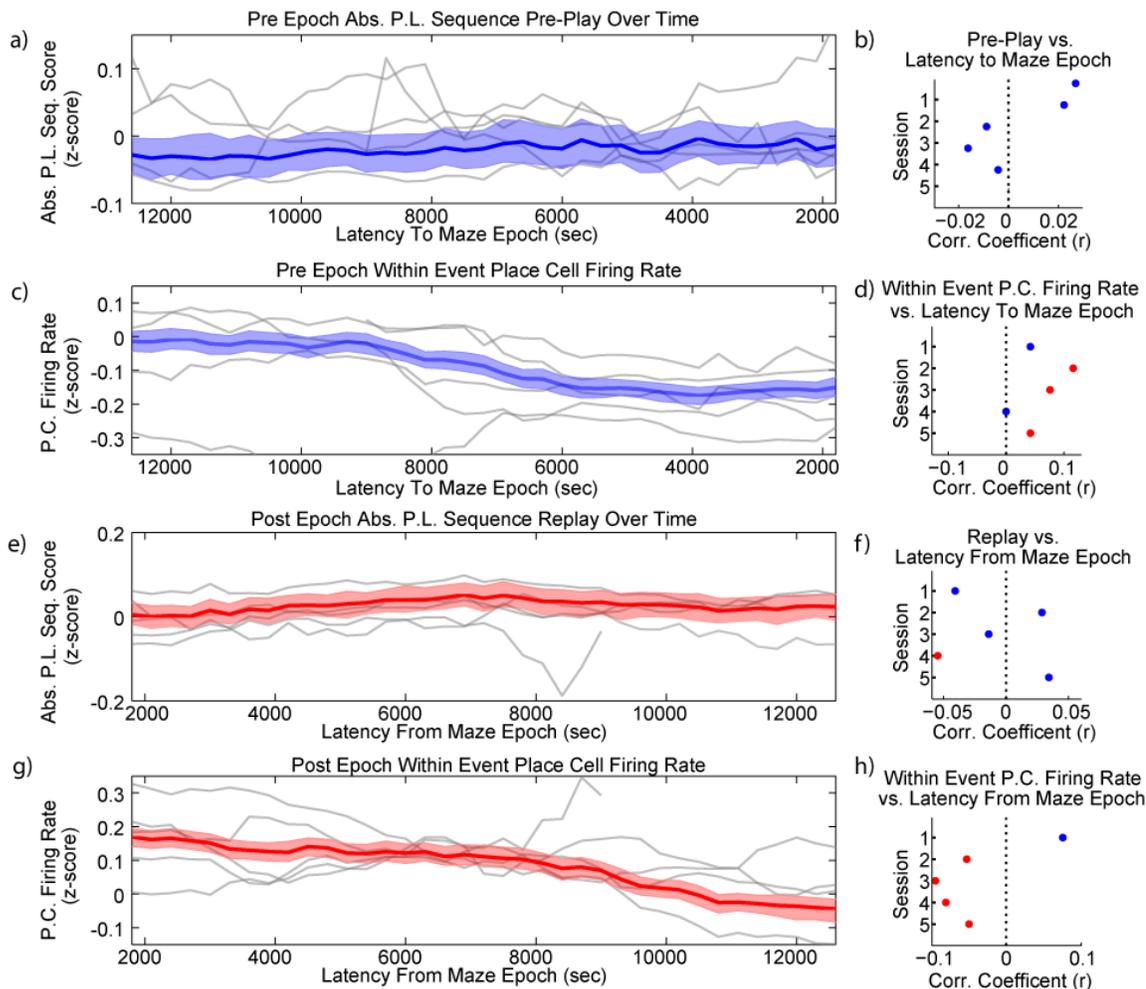


Figure 4.14: P.L. Pre-Play and Replay and Place Cell Firing Rate Over Time: non-local P.L. sequence score (panels **a** and **b** for Pre and Post respectively) and within event place cell firing rates (panels **c** and **g**) were first z-scored across both Pre and Post epochs in order to reduce intra-session variability, and binned in sliding one hour bins with a step size of five minutes. For each of the panels on the left (**a**, **c**, **e**, and **g**) the mean and bootstrapped 95% confidence interval are shown in blue (Pre) or red (Post). Note that the x-axis of panels **a** and **b** are inverted to reflect that time is measured as the latency from the start (Pre) or end (Post) of the maze epoch. For each session within-bin means are shown in light grey. Right panels (**b**, **d**, **f**, and **h**) display the per session correlation coefficients for each of the comparisons derived from the none-binned distributions. Note that while neither pre-play (**a** and **b**) or replay (**c** and **g**) consistently change over time, the within-event firing rates show a bias towards decreased firing over the course of either the Pre or Post epochs.

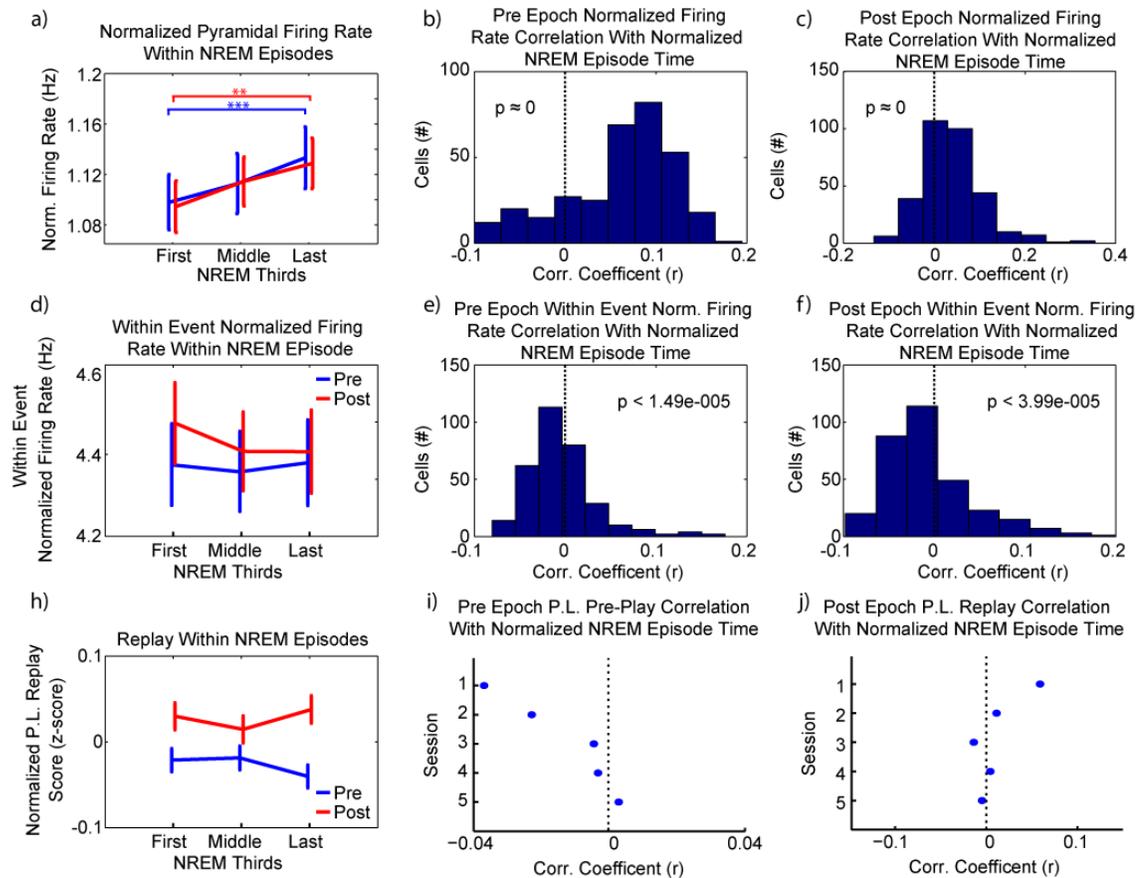


Figure 4.15: Within NREM Firing Rate and Replay/Pre-play Effects: For each cell firing rates in both the Pre and Post epochs were binned in 1 second bins. In order to decrease inter-cell and inter session variability, all firing rates were normalized by cell by dividing each bin by that cell's overall firing rate for the entire session. For each bin occurring during an NREM episode with a durations of at least 100 seconds a 'normalized NREM episode time' value was defined as the corresponding percentage of that NREM episode's duration such that the beginning of that NREM episode corresponded to 0 percent and the end of that episode corresponded to 100%. The mean normalized firing rate was taken for the first, middle and last NREM thirds (corresponding to values 0-33%, 33%-66%, and 66%-100% of normalized NREM episode time, respectively) during the Pre and Post sleep epochs (a). In accordance with previous finding, pyramidal cell firing was found to increase within NREM episodes (first NREM third to last NREM third comparison, Pre: $p < 4.17 \times 10^{-5}$, Post $p < 0.0012$, sign-rank (paired) test). Consistent with this effect individual pyramidal cells were found to correlate positively with normalized NREM episode time in both the Pre (b) and Post (c) epochs ($n = 322$ cells, sign-rank test). While normalized within-event firing rates were not found to change significantly from the first to last third of NREM episodes (d), within-event firing rates correlated negatively with normalized episode time in both the Pre (e) and Post (f) epochs ($n = 322$ cells, sign-rank test) which is consistent with the previously observed within ripple firing rate changes. However, no significant first to last third NREM effects were observed for absolute P.L. sequence pre-play or replay (h), and the correlation between normalized NREM episode time and absolute sequence score was not significant in any Pre (i) or Post (j) epoch.

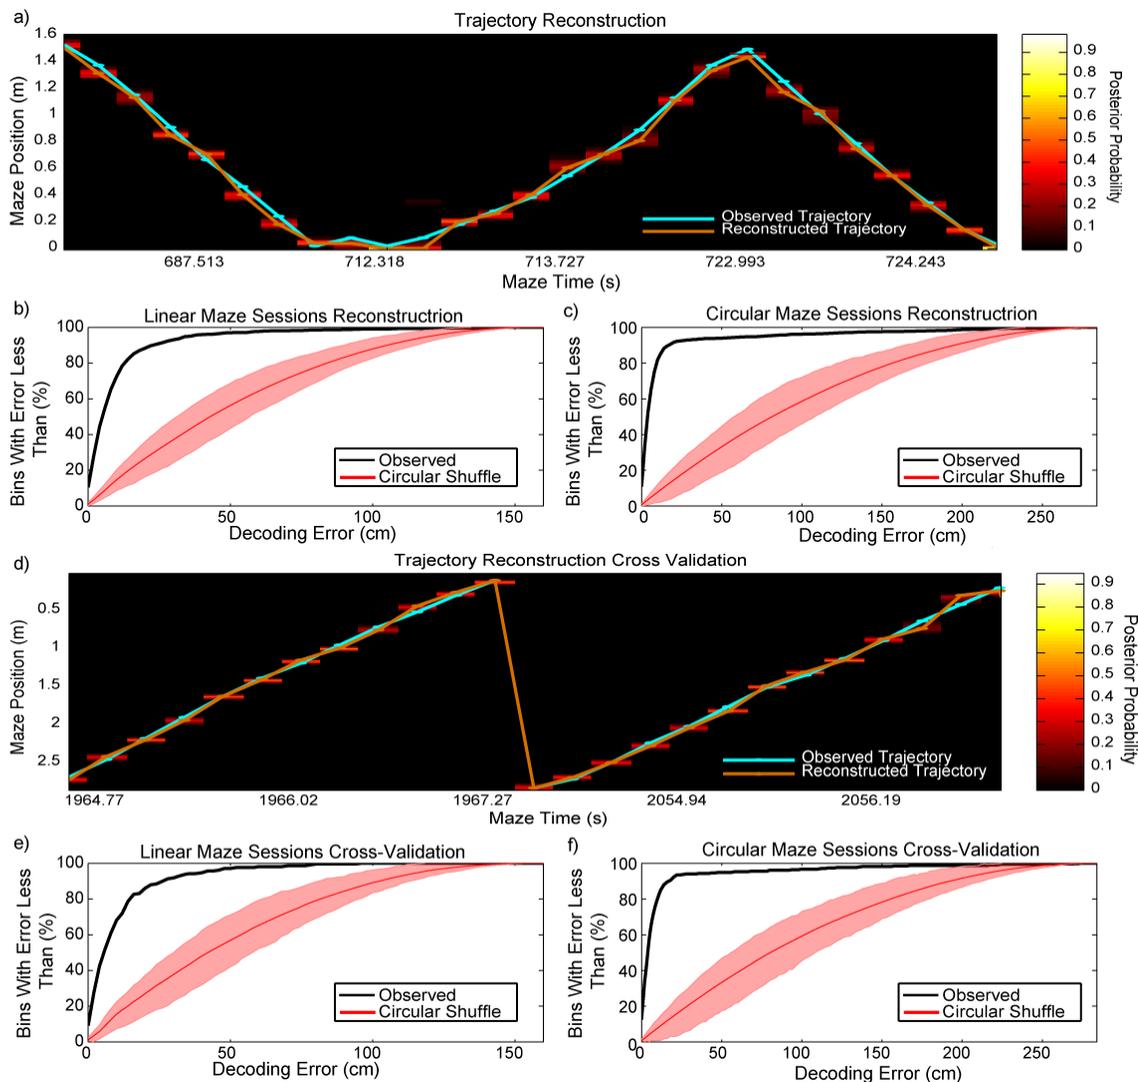


Figure 4.16: Bayesian Decoding and Reconstruction: Panel **a** shows the Bayesian reconstruction of four linear maze trials, note that for making this panel as well as panel **d** 250 ms bins were used and time spent at the reward (non-analyzed) locations were omitted for illustrative purposes. The cumulative error distribution is shown across all 500 ms bins on the linear (**b**, $n = 1145$ bins, mean error: $10.88 \text{ cm} \pm 0.505 \text{ cm}$ standard error), and circular (**c**, $n = 1820$ bins, mean error: $14.18 \text{ cm} \pm 0.401 \text{ cm}$ standard error) mazes. Red shaded regions show the the 95 percent confidence interval of the shuffled cumulative error probability distributions. In order to cross-validate these measures, for each session, Bayesian classifiers were constructed after excluding every fifth lap on the maze and subsequently tested on the excluded laps. Panel **d** shows the reconstruction of two such circular maze laps. Bayesian decoding performance was found to be significantly better than chance in both the linear (**e**, $n = 207$ bins, mean error: $11.33 \text{ cm} \pm 1.02 \text{ cm}$ standard error) and circular (**f**, $n = 334$ bins, mean error: $13.08 \pm 1.86 \text{ cm}$ standard error) maze conditions.

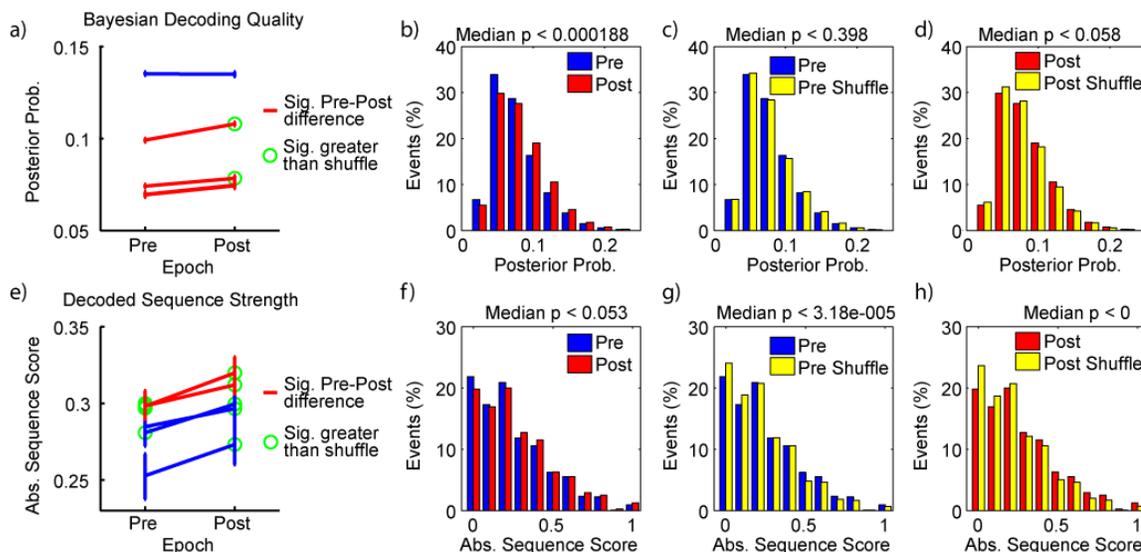


Figure 4.17: Bayesian Decoding of Pre-Play and Replay Summary: Bayesian decoding quality (panels **a-d**) and Bayesian decoded sequence Pre-play and Replay strength (panels **e-h**) were calculated within events divided into 20 ms bins. For each comparison in this figure significance ($p < 0.025$) was determined using rank-sum tests. For panels **a** and **e** red lines indicate significant differences between Pre and Post, while blue lines indicate non-significant effects, error bars show the boot-strapped 95% confidence interval for each condition. Green circles indicate conditions in which the experimental distribution was significantly greater than the null (shuffle) distribution. In order to visualize the effects shown in **a** and **e**, for each of the two measures for each of the three comparisons (pre vs. post, pre vs. shuffle, and post vs. shuffle) mean distributions were obtained by averaging the normalized distributions across all 5 sessions. For each measure and comparison the median of the p-values across all 5 comparisons is displayed in order to indicate effect robustness. While Bayesian decoding quality increased significantly in 4 of the 5 session (panels **a** and **b**), Pre epoch Bayesian decoding quality was not significantly greater than the null (0 out of 5 significant epochs, panels **a** and **c**), and was greater than the null distribution in 2 out of 5 Post epochs (panels **a** and **d**). Conversely, decoded sequence matching strength increased significantly from the Pre to Post epochs in 2 out of 5 sessions (panels **e** and **f**), while 3 out of 5 Pre epochs (panels **e** and **g**) showed significant decoded Preplay and all 5 Post epochs showed decoded sequence Replay highly significantly greater than the null distribution (panels **e** and **h**).

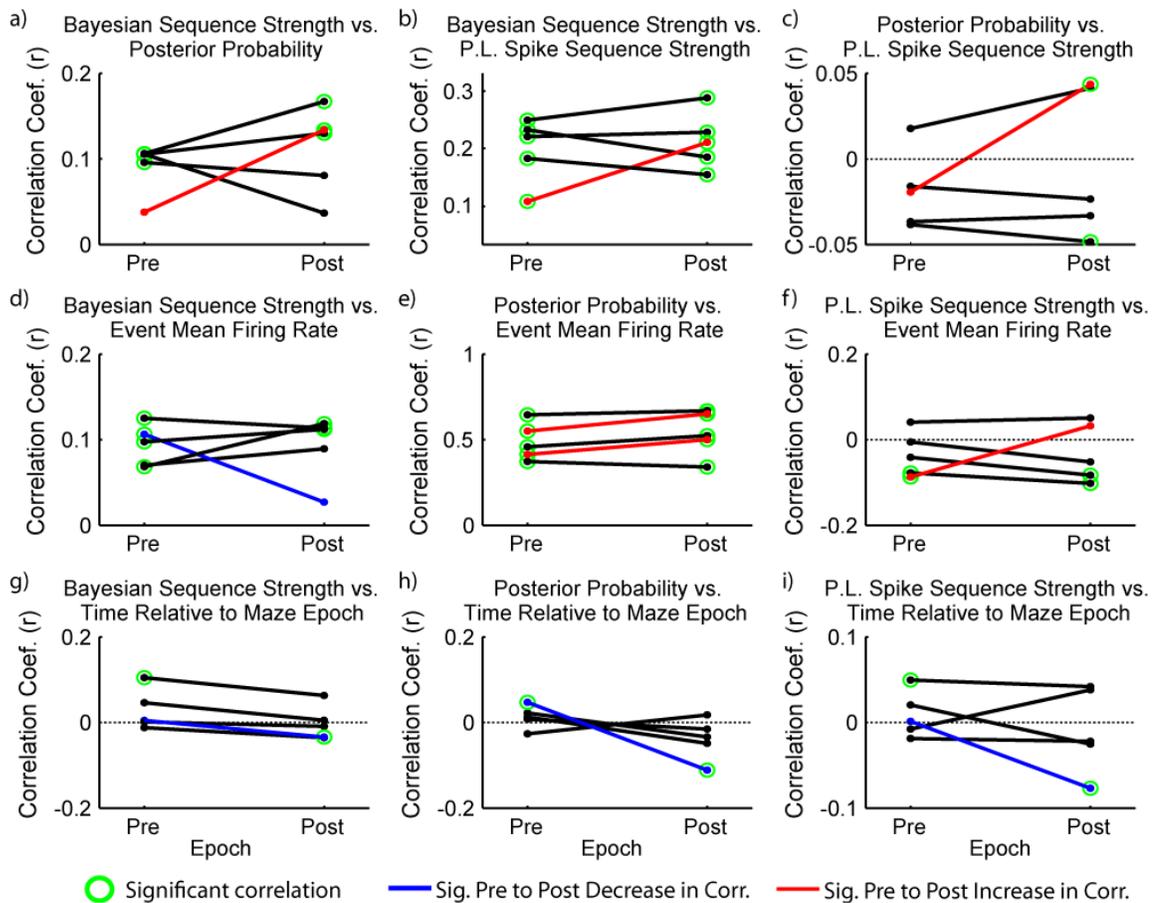


Figure 4.18: Bayesian Decoding Relative to P.L. Sequence Score, Firing Rate and Time: All significance testing of correlation coefficients. Note that for this analysis, only those events that met the criteria for inclusion in the Bayesian analysis were included for all conditions. Mean posterior probability is significantly correlated with Bayesian sequence strength (a), which in turn is highly significantly correlated with P.L. spike sequence strength (b), however posterior probability and P.L. sequence strength do not show consistent correlations. Notably, while both Bayesian sequence strength, and especially mean posterior probability are strongly correlated with event rate (d and e, respectively), P.L. sequence strength showed the opposite trend (f, see also figure N), suggesting that while both of Bayesian and spike sequence analyses are highly influenced by rate, this influence is highly divergent, and indeed opposite. Notably, no consistent relationship (for instance, decay) was found with time for any of the measurements (g, h, and i).

	Pair-Wise Methods	Event-Based Methods
Rate replay**	<ul style="list-style-type: none"> Correlation of firing rates in 100 ms bins (Wilson and McNaughton, 1994) 	<ul style="list-style-type: none"> PCA based projection (Peyrache et al., 2009) The <i>posterior probability</i> metric of Bayesian population activity decoding (Davidson et al., 2009; Foster and Wilson, 2006; Kloosterman, 2012; Pfeiffer and Foster, 2013)***
Sequence replay	<ul style="list-style-type: none"> Bias in pair-wise cross correlograms (Skaggs and McNaughton, 1996) 	<ul style="list-style-type: none"> Template matching (Louie and Wilson, 2001) Combinatorial 'word' analysis of place cell spike-sequences (Lee and Wilson, 2002) Rank-order correlations between place field sequence on a maze and spiking sequence during an 'off-line' event (Diba and Buzsáki, 2007; Dragoi and Tonegawa, 2011; Foster and Wilson, 2006) Bayesian decoded 'trajectories' (Davidson et al., 2009; Pfeiffer and Foster, 2013)***

Table 1: Types of replay*

* Note that this table is meant to provide a heuristic for the methods most relevant to this thesis and is not an exhaustive list of replay methodologies

** Note that one type of relevant *rate replay*, the modulation of firing rates of one or more cells (Hirase et al., 2001; Pavlides and Winson, 1989b) is not included in this table.

*** Most methods for assessing Bayesian decoded replay confound the posterior probability measure (*activity replay*) and the trajectory measure (*sequence replay*) into one over-all replay score, and the individual contributions of each component are not reported.

Excitability changes across sleep (supplementary to Fig. 1b)

	First non- REM	Last non- REM	Change	Units	P <	n
Pyramidal Firing Rate	0.6	0.55	↓***	Hz	3.58E-016	618
Interneuron Firing Rate	15.92	14.74	↓***	Hz	7.1282E-021	111
Active Period Incidence	79.51	69.61	↓***	Periods/Minute	0.00026	22
Inactive Period Incidence	99.69	102.23	↑*	Periods/Minute	0.0203	22
Within Ripple Pyr. Rate	2.62	2.73	↑*	Hz	0.029	618
Between Ripple Pyr. Rate	0.543	0.493	↓***	Hz	3.589E-16	618
Pyr. Pair-Wise Corr.	0.024	0.027	↑**	r (100ms bins)	0.0033	22
Pyr. Ripple Participation	17.56	18.56	↑***	%Ripples/Pyr Cell	0.00022	618
Within Ripple Pyr. Rate	0.76	0.71	↓**	C.V. (across cells)	0.0045	22

Excitability changes within non-REM (supplementary to Fig. 2 a,b)

	First 1/3rd of non- REM	Last 1/3rd of non- REM	Change	Units	P <	n
Pyramidal Firing Rate	0.55785	0.596	↑***	Hz	5.21E-14	618
Interneuron Firing Rate	14.782	15.647	↑***	Hz	5.03E-05	111
Active Period Incidence	71.213	73.503	↑***	Periods/Minute	3.87E-06	82
Inactive Period Incidence	102.87	208.27	↓**	Periods/Minute	0.0014	82
Within Ripple Pyr. Rate	2.7248	2.633	↓	Hz	0.146	618
Between Ripple Pyr. Rate	0.491	0.541	↑***	Hz	9.83E-23	618
Pyr. Pair-Wise Corr.	0.0279	0.0246	↓***	r (100ms bins)	5.08E-05	82
Pyr. Ripple Participation	18.36	17.73	↓**	%Ripples/Pyr Cell	0.003	618
Within Ripple Pyr. Rate	0.739	0.7213	↓	C.V. (across cells)	0.3548	22

Excitability changes within REM (supplementary to Fig. 2a)

	First 1/3rd of REM	Last 1/3rd of REM	Change	Units	P <	n
Pyramidal Firing Rate	0.61314	0.54577	↓*	Hz	0.012836	618
Interneuron Firing Rate	20.052	16.959	↓***	Hz	7.22E-06	111
Pyr. Pair-Wise Corr.	0.0074	0.015	↑	r (100ms bins)	0.225	45

Table 2: Excitability changes across sleep, within non-REM episodes, or within REM episodes (n = 618 pyramidal cells, 111 interneurons, 22 sleep sessions, 82 non-REM episodes or 45 REM episodes, all significance values from (paired) sign-rank tests, *p<0.05; **p<0.005; ***p<0.0005.).

Glossary of Terms

Across Sleep Changes – in our study these are assessed by comparing the first and last non-REM episodes of each session.

Bayesian Decoding Quality - assessed as the mean of the posterior probability across within-event bins of the peak posterior probabilities of the bins within the event. This is a measure of the specificity of the similarity of an observed population activity vector in a given time bin to the population activity vector in the given spatial (2 cm) bin to which it is most similar (in other words, the best estimate of position), normalized by its 'non-specific' similarity to all spatial bins – and is thus a method for *Rate Replay*.

Bayesian Decoded Sequence Strength – This measure was assessed across the within-event decoded positions by using a modified version of the paired latency vector. To put it simply, this measure reflects how consistently all the pairs of decoded position of 20 ms bins in a given event indicated that the animal was 'moving' (in decoded position space) forwards *or* backwards across the bins of that event. It is thus a measure of *Sequence Replay*.

Excitability - used broadly to refer to several statistical aspects of neural activity, including firing rate and synaptic strength.

Novel Maze – a maze environment which the experimental subject has never previously experienced. Uniquely, in our experiment both the experimental maze as well as the experimental room were novel to the animal.

'Off-Line' State – generally this term reflects to behavioral states in which the animal is not actively engaging with its environment. In our study, we restrict the 'off-line' state analysis to drowsy/light-sleep, non-REM and intermediate sleep epochs occurring during either the Pre or the Post sleep recordings.

Paired-Latency Method – a novel event-based method for assessing *Sequence Replay*. It is based on the similarity of two vectors each of whose elements reflect the directionality (i.e. before or after *or* to the left or to the right of) of a particular pair-wise interaction. It may thus be thought of as an event-based adaptation of the method employed by Skaggs and McNaughton, 1996.

Participant – a cell which discharged at least one spike in a given event.

Place Cell – any hippocampal layer CA1 cell which showed at least one *place field*.

Place Field – a series of at least five consecutive bins in which a principal cell's firing rate was above the 99th percentile of the shuffled firing rate. Note that while many neurons displayed more than one place field, only the place field which contained the highest *place field peak* firing rate was considered for further analysis.

Place Field Peak – the spatial (2 cm) bin within a place field which displayed the highest firing rate. Note that this bin was taken as the place cell's *location* for all the (non-Bayesian) replay and pre-play analysis.

Population Activity Event – a short (50 to 500 ms) epoch in which population activity synchronously rises above a baseline. In the current study these were defined as epochs occurring in the 'off-line' state, in which the smoothed (Gaussian window,

15 ms st.d.) activity vector assessed across all pyramidal cell spikes rose 3 standard deviation of its mean during non-REM (see *Methods*).

Posterior Probability of Position – the Bayesian decoded estimate that the animal is ‘perceiving’ (or ‘remembering’) a particular position (i.e. 2 cm bin) on the maze. This estimate is based on the similarity of the population activity being decoded to the population activity at each spatial bin on the maze. Note that posterior probability is normalized so that it sums to one across all spatial bins.

Rate replay - the 'off-line' reactivation of single-cell, pair-wise, or higher-order firing rate patterns observed during behavior.

Sequence Replay -the 'off-line' reactivation of pairs or larger groups of neurons in the sequence in which they fired during behavior.

Synchrony – in general this term refers to the propensity of a population of neurons to discharge together and to be silent together in time. In our particular study it was defined as the mean pair-wise correlations between all pairs of pyramidal cells’ firing rates binned in non-overlapping 100 ms bins.

Within non-REM Changes - in our study these are assessed by comparing the first and last thirds of each non-REM epoch.

References

- Abbott, L.F., and Nelson, S.B. (2000). Synaptic plasticity: taming the beast. *Nat. Neurosci.* *3 Suppl*, 1178–1183.
- Abeles, M., and Gerstein, G.L. (1988). Detecting spatiotemporal firing patterns among simultaneously recorded single neurons. *J. Neurophysiol.* *60*, 909–924.
- Abeles, M., Bergman, H., Margalit, E., and Vaadia, E. (1993). Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *J. Neurophysiol.* *70*, 1629–1638.
- Adamantidis, A., Carter, M.C., and de Lecea, L. (2010). Optogenetic deconstruction of sleep-wake circuitry in the brain. *Front. Mol. Neurosci.* *2*, 31.
- Aeschbach, D., Cutler, A.J., and Ronda, J.M. (2008). A role for non-rapid-eye-movement sleep homeostasis in perceptual learning. *J. Neurosci. Off. J. Soc. Neurosci.* *28*, 2766–2772.
- Anikeeva, P., and Deisseroth, K. (2012). Photothermal genetic engineering. *ACS Nano* *6*, 7548–7552.
- Aserinsky, E., and Kleitman, N. (1953). Regularly occurring periods of eye motility, and concomitant phenomena, during sleep. *Science* *118*, 273–274.
- Aston-Jones, G., and Bloom, F.E. (1981). Activity of norepinephrine-containing locus coeruleus neurons in behaving rats anticipates fluctuations in the sleep-waking cycle. *J. Neurosci. Off. J. Soc. Neurosci.* *1*, 876–886.
- Axmacher, N., Mormann, F., Fernández, G., Elger, C.E., and Fell, J. (2006). Memory formation by neuronal synchronization. *Brain Res. Rev.* *52*, 170–182.
- Baker, M.A., and Hayward, J.N. (1967). Autonomic basis for the rise in brain temperature during paradoxical sleep. *Science* *157*, 1586–1588.
- Barlow, H.B. (1972). Single units and sensation: a neuron doctrine for perceptual psychology? *Perception* *1*, 371–394.
- Barrett, T.R., and Ekstrand, B.R. (1972). Effect of sleep on memory. 3. Controlling for time-of-day effects. *J. Exp. Psychol.* *96*, 321–327.
- Barthó, P., Hirase, H., Monconduit, L., Zugaro, M., Harris, K.D., and Buzsáki, G. (2004). Characterization of neocortical principal cells and interneurons by network interactions and extracellular features. *J. Neurophysiol.* *92*, 600–608.
- Basu, J., Srinivas, K.V., Cheung, S.K., Taniguchi, H., Huang, Z.J., and Siegelbaum, S.A. (2013). A cortico-hippocampal learning rule shapes inhibitory microcircuit activity to enhance hippocampal information flow. *Neuron* *79*, 1208–1221.

- Battaglia, F.P., and Pennartz, C.M.A. (2011). The construction of semantic memory: grammar-based representations learned from relational episodic information. *Front. Comput. Neurosci.* 5, 36.
- Battaglia, F.P., Sutherland, G.R., and McNaughton, B.L. (2004). Hippocampal sharp wave bursts coincide with neocortical “up-state” transitions. *Learn. Mem. Cold Spring Harb. N 11*, 697–704.
- Battaglia, F.P., Sutherland, G.R., Cowen, S.L., Mc Naughton, B.L., and Harris, K.D. (2005). Firing rate modulation: a simple statistical view of memory trace reactivation. *Neural Netw. Off. J. Int. Neural Netw. Soc.* 18, 1280–1291.
- Benchenane, K., Peyrache, A., Khamassi, M., Tierney, P.L., Gioanni, Y., Battaglia, F.P., and Wiener, S.I. (2010). Coherent theta oscillations and reorganization of spike timing in the hippocampal- prefrontal network upon learning. *Neuron* 66, 921–936.
- Binder, S., Baier, P.C., Mölle, M., Inostroza, M., Born, J., and Marshall, L. (2012). Sleep enhances memory consolidation in the hippocampus-dependent object-place recognition task in rats. *Neurobiol. Learn. Mem.* 97, 213–219.
- Bird, C.M., and Burgess, N. (2008). The hippocampus and memory: insights from spatial processing. *Nat. Rev. Neurosci.* 9, 182–194.
- Bliss, T.V.P., and Lømo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J. Physiol.* 232, 331–356.
- Bliss, T.V., Goddard, G.V., and Riives, M. (1983). Reduction of long-term potentiation in the dentate gyrus of the rat following selective depletion of monoamines. *J. Physiol.* 334, 475–491.
- Blum, K.I., and Abbott, L.F. (1996). A Model of Spatial Map Formation in the Hippocampus of the Rat. *Neural Comput.* 8, 85–93.
- Bódizs, R., Békésy, M., Szucs, A., Barsi, P., and Halász, P. (2002). Sleep-dependent hippocampal slow activity correlates with waking memory performance in humans. *Neurobiol. Learn. Mem.* 78, 441–457.
- Borbély, A.A. (1982). A two process model of sleep regulation. *Hum. Neurobiol.* 1, 195–204.
- Borbély, A.A. (2009). Refining sleep homeostasis in the two-process model. *J. Sleep Res.* 18, 1–2.
- Borbély, A.A., and Achermann, P. (1999). Sleep homeostasis and models of sleep regulation. *J. Biol. Rhythms* 14, 557–568.

- Born, J., Rasch, B., and Gais, S. (2006). Sleep to remember. *Neurosci. Rev. J. Bringing Neurobiol. Neurol. Psychiatry* 12, 410–424.
- Boyden, E.S., Zhang, F., Bamberg, E., Nagel, G., and Deisseroth, K. (2005). Millisecond-timescale, genetically targeted optical control of neural activity. *Nat. Neurosci.* 8, 1263–1268.
- Bragin, A., Hetke, J., Wilson, C.L., Anderson, D.J., Engel, J., Jr, and Buzsáki, G. (2000). Multiple site silicon-based probes for chronic recordings in freely moving rats: implantation, recording and histological verification. *J. Neurosci. Methods* 98, 77–82.
- Bullock, T.H., McClune, M.C., Achimowicz, J.Z., Iragui-Madoz, V.J., Duckrow, R.B., and Spencer, S.S. (1995). Temporal fluctuations in coherence of brain waves. *Proc. Natl. Acad. Sci. U. S. A.* 92, 11568–11572.
- Burrone, J., and Murthy, V.N. (2003). Synaptic gain control and homeostasis. *Curr. Opin. Neurobiol.* 13, 560–567.
- Buzsáki, G. (1986). Hippocampal sharp waves: their origin and significance. *Brain Res.* 398, 242–252.
- Buzsáki, G. (1989). Two-stage model of memory trace formation: a role for “noisy” brain states. *Neuroscience* 31, 551–570.
- Buzsáki, G. (1996). The hippocampo-neocortical dialogue. *Cereb. Cortex N. Y. N* 1991 6, 81–92.
- Buzsáki, G. (1998). Memory consolidation during sleep: a neurophysiological perspective. *J. Sleep Res.* 7 *Suppl* 1, 17–23.
- Buzsáki, G. (2005). Theta rhythm of navigation: link between path integration and landmark navigation, episodic and semantic memory. *Hippocampus* 15, 827–840.
- Buzsáki, G. (2006). *Rhythms of the Brain* (Oxford University Press).
- Buzsáki, G. (2010). Neural syntax: cell assemblies, synapsembles, and readers. *Neuron* 68, 362–385.
- Buzsáki, G., and Chrobak, J.J. (1995). Temporal structure in spatially organized neuronal ensembles: a role for interneuronal networks. *Curr. Opin. Neurobiol.* 5, 504–510.
- Buzsáki, G., and Mizuseki, K. (2014). The log-dynamic brain: how skewed distributions affect network operations. *Nat. Rev. Neurosci.* 15, 264–278.
- Buzsáki, G., Leung, L.W., and Vanderwolf, C.H. (1983). Cellular bases of hippocampal EEG in the behaving rat. *Brain Res.* 287, 139–171.

- Buzsáki, G., Horváth, Z., Urioste, R., Hetke, J., and Wise, K. (1992). High-frequency network oscillation in the hippocampus. *Science* 256, 1025–1027.
- Campbell, S.S., and Gillin, J.C. (1987). Sleep measures in depression: How sensitive? How specific? *Psychiatr. Ann.* 17, 647–653.
- Carr, M.F., Jadhav, S.P., and Frank, L.M. (2011). Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nat. Neurosci.* 14, 147–153.
- Carr, M.F., Karlsson, M.P., and Frank, L.M. (2012). Transient slow gamma synchrony underlies hippocampal memory replay. *Neuron* 75, 700–713.
- Carter, M.E., Yizhar, O., Chikahisa, S., Nguyen, H., Adamantidis, A., Nishino, S., Deisseroth, K., and de Lecea, L. (2010). Tuning arousal with optogenetic modulation of locus coeruleus neurons. *Nat. Neurosci.* 13, 1526–1533.
- Chapin, J.K., and Nicolelis, M.A.L. (1999). Principal component analysis of neuronal ensemble activity reveals multidimensional somatosensory representations. *J. Neurosci. Methods* 94, 121–140.
- Chase, M.H. (2008). Confirmation of the consensus that glycinergic postsynaptic inhibition is responsible for the atonia of REM sleep. *Sleep* 31, 1487–1491.
- Chen, Z. (2013). An overview of Bayesian methods for neural spike train analysis. *Comput. Intell. Neurosci.* 2013, 251905.
- Chen, S., Mohajerani, M.H., Xie, Y., and Murphy, T.H. (2012). Optogenetic Analysis of Neuronal Excitability during Global Ischemia Reveals Selective Deficits in Sensory Processing following Reperfusion in Mouse Cortex. *J. Neurosci.* 32, 13510–13519.
- Chuquet, J., Quilichini, P., Nimchinsky, E.A., and Buzsáki, G. (2010). Predominant enhancement of glucose uptake in astrocytes versus neurons during activation of the somatosensory cortex. *J. Neurosci. Off. J. Soc. Neurosci.* 30, 15298–15303.
- Cirelli, C. (2005). A molecular window on sleep: changes in gene expression between sleep and wakefulness. *Neurosci. Rev. J. Bringing Neurobiol. Neurol. Psychiatry* 11, 63–74.
- Cirelli, C., and Tononi, G. (2000a). Differential expression of plasticity-related genes in waking and sleep and their regulation by the noradrenergic system. *J. Neurosci. Off. J. Soc. Neurosci.* 20, 9187–9194.
- Cirelli, C., and Tononi, G. (2000b). On the functional significance of c-fos induction during the sleep-waking cycle. *Sleep* 23, 453–469.

- Cirelli, C., Pompeiano, M., and Tononi, G. (1996). Neuronal gene expression in the waking state: a role for the locus coeruleus. *Science* 274, 1211–1215.
- Cirelli, C., Huber, R., Gopalakrishnan, A., Southard, T.L., and Tononi, G. (2005). Locus ceruleus control of slow-wave homeostasis. *J. Neurosci. Off. J. Soc. Neurosci.* 25, 4503–4511.
- Cressant, A., Muller, R.U., and Poucet, B. (1997). Failure of centrally placed objects to control the firing fields of hippocampal place cells. *J. Neurosci. Off. J. Soc. Neurosci.* 17, 2531–2542.
- Crunelli, V., and Hughes, S.W. (2010). The slow (<1 Hz) rhythm of non-REM sleep: a dialogue between three cardinal oscillators. *Nat. Neurosci.* 13, 9–17.
- Csicsvari, J., Hirase, H., Czurko, A., and Buzsáki, G. (1998). Reliability and state dependence of pyramidal cell-interneuron synapses in the hippocampus: an ensemble approach in the behaving rat. *Neuron* 21, 179–189.
- Csicsvari, J., Hirase, H., Czurkó, A., Mamiya, A., and Buzsáki, G. (1999a). Fast network oscillations in the hippocampal CA1 region of the behaving rat. *J. Neurosci. Off. J. Soc. Neurosci.* 19, RC20.
- Csicsvari, J., Hirase, H., Czurkó, A., Mamiya, A., and Buzsáki, G. (1999b). Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving Rat. *J. Neurosci. Off. J. Soc. Neurosci.* 19, 274–287.
- Daoudal, G., and Debanne, D. (2003). Long-Term Plasticity of Intrinsic Excitability: Learning Rules and Mechanisms. *Learn. Mem.* 10, 456–465.
- Dash, M.B., Douglas, C.L., Vyazovskiy, V.V., Cirelli, C., and Tononi, G. (2009). Long-term homeostasis of extracellular glutamate in the rat cerebral cortex across sleep and waking states. *J. Neurosci. Off. J. Soc. Neurosci.* 29, 620–629.
- Datta, S., Mavanji, V., Ulloor, J., and Patterson, E.H. (2004). Activation of phasic pontine-wave generator prevents rapid eye movement sleep deprivation-induced learning impairment in the rat: a mechanism for sleep-dependent plasticity. *J. Neurosci. Off. J. Soc. Neurosci.* 24, 1416–1427.
- Datta, S., Li, G., and Auerbach, S. (2008). Activation of phasic pontine-wave generator in the rat: a mechanism for expression of plasticity-related genes and proteins in the dorsal hippocampus and amygdala. *Eur. J. Neurosci.* 27, 1876–1892.
- Davidson, T.J., Kloosterman, F., and Wilson, M.A. (2009). Hippocampal replay of extended experience. *Neuron* 63, 497–507.
- Destexhe, A., Contreras, D., and Steriade, M. (1999). Spatiotemporal analysis of local field potentials and unit discharges in cat cerebral cortex during natural wake and sleep states. *J. Neurosci. Off. J. Soc. Neurosci.* 19, 4595–4608.

- Diba, K., and Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nat. Neurosci.* *10*, 1241–1242.
- Diba, K., and Buzsáki, G. (2008). Hippocampal network dynamics constrain the time lag between pyramidal cells across modified environments. *J. Neurosci. Off. J. Soc. Neurosci.* *28*, 13448–13456.
- Diekelmann, S., and Born, J. (2010). The memory function of sleep. *Nat. Rev. Neurosci.* *11*, 114–126.
- Diekelmann, S., Büchel, C., Born, J., and Rasch, B. (2011). Labile or stable: opposing consequences for memory when reactivated during waking and sleep. *Nat. Neurosci.* *14*, 381–386.
- Disterhoft, J.F., Coulter, D.A., and Alkon, D.L. (1986). Conditioning-specific membrane changes of rabbit hippocampal neurons measured in vitro. *Proc. Natl. Acad. Sci.* *83*, 2733–2737.
- Dragoi, G., and Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature* *469*, 397–401.
- Dragoi, G., and Tonegawa, S. (2013). Distinct preplay of multiple novel spatial experiences in the rat. *Proc. Natl. Acad. Sci. U. S. A.* *110*, 9100–9105.
- Dragoi, G., and Tonegawa, S. (2014). Selection of preconfigured cell assemblies for representation of novel spatial experiences. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* *369*, 20120522.
- Dragoi, G., Harris, K.D., and Buzsáki, G. (2003). Place representation within hippocampal networks is modified by long-term potentiation. *Neuron* *39*, 843–853.
- Dupret, D., O’Neill, J., Pleydell-Bouverie, B., and Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nat. Neurosci.* *13*, 995–1002.
- Eichenbaum, H. (2000). A cortical-hippocampal system for declarative memory. *Nat. Rev. Neurosci.* *1*, 41–50.
- Ellenbogen, J.M., Hulbert, J.C., Stickgold, R., Dinges, D.F., and Thompson-Schill, S.L. (2006a). Interfering with theories of sleep and memory: sleep, declarative memory, and associative interference. *Curr. Biol. CB* *16*, 1290–1294.
- Ellenbogen, J.M., Payne, J.D., and Stickgold, R. (2006b). The role of sleep in declarative memory consolidation: passive, permissive, active or none? *Curr. Opin. Neurobiol.* *16*, 716–722.

- Ellenbogen, J.M., Hu, P.T., Payne, J.D., Titone, D., and Walker, M.P. (2007). Human relational memory requires time and sleep. *Proc. Natl. Acad. Sci. U. S. A.* *104*, 7723–7728.
- Ermentrout, G.B., Galán, R.F., and Urban, N.N. (2008). Reliability, synchrony and noise. *Trends Neurosci.* *31*, 428–434.
- Euston, D.R., Tatsuno, M., and McNaughton, B.L. (2007). Fast-forward playback of recent memory sequences in prefrontal cortex during sleep. *Science* *318*, 1147–1150.
- Everson, C.A., Bergmann, B.M., and Rechtschaffen, A. (1989). Sleep deprivation in the rat: III. Total sleep deprivation. *Sleep* *12*, 13–21.
- Farmer, J., Zhao, X., van Praag, H., Wodtke, K., Gage, F.H., and Christie, B.R. (2004). Effects of voluntary exercise on synaptic plasticity and gene expression in the dentate gyrus of adult male Sprague-Dawley rats in vivo. *Neuroscience* *124*, 71–79.
- Fogel, S.M., Smith, C.T., and Cote, K.A. (2007). Dissociable learning-dependent changes in REM and non-REM sleep in declarative and procedural memory systems. *Behav. Brain Res.* *180*, 48–61.
- Foster, D.J., and Wilson, M.A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* *440*, 680–683.
- Frank, M.G. (2011). Beyond the neuron: astroglial regulation of mammalian sleep. *Curr. Top. Med. Chem.* *11*, 2452–2456.
- Fujisawa, S., Amarasingham, A., Harrison, M.T., and Buzsáki, G. (2008). Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nat. Neurosci.* *11*, 823–833.
- Gais, S., and Born, J. (2004). Low acetylcholine during slow-wave sleep is critical for declarative memory consolidation. *Proc. Natl. Acad. Sci. U. S. A.* *101*, 2140–2144.
- Gais, S., Plihal, W., Wagner, U., and Born, J. (2000). Early sleep triggers memory for early visual discrimination skills. *Nat. Neurosci.* *3*, 1335–1339.
- Gais, S., Mölle, M., Helms, K., and Born, J. (2002). Learning-dependent increases in sleep spindle density. *J. Neurosci. Off. J. Soc. Neurosci.* *22*, 6830–6834.
- Gervasoni, D., Darracq, L., Fort, P., Soulière, F., Chouvet, G., and Luppi, P.H. (1998). Electrophysiological evidence that noradrenergic neurons of the rat locus coeruleus are tonically inhibited by GABA during sleep. *Eur. J. Neurosci.* *10*, 964–970.

- Gervasoni, D., Lin, S.-C., Ribeiro, S., Soares, E.S., Pantoja, J., and Nicolelis, M.A.L. (2004). Global forebrain dynamics predict rat behavioral states and their transitions. *J. Neurosci. Off. J. Soc. Neurosci.* *24*, 11137–11147.
- Gideon, R.A., and Hollister, R.A. (1987). A Rank Correlation Coefficient Resistant to Outliers. *J. Am. Stat. Assoc.* *82*, 656–666.
- Gierz, M., Campbell, S.S., and Gillin, J.C. (1987). Sleep disturbances in various nonaffective psychiatric disorders. *Psychiatr. Clin. North Am.* *10*, 565–581.
- Gilat, A. (2010). *MATLAB: An Introduction with Applications*, 4th Edition (John Wiley & Sons, Inc.).
- Girardeau, G., Benchenane, K., Wiener, S.I., Buzsáki, G., and Zugaro, M.B. (2009). Selective suppression of hippocampal ripples impairs spatial memory. *Nat. Neurosci.* *12*, 1222–1223.
- Giuditta, A., Ambrosini, M.V., Montagnese, P., Mandile, P., Cotugno, M., Grassi Zucconi, G., and Vescia, S. (1995). The sequential hypothesis of the function of sleep. *Behav. Brain Res.* *69*, 157–166.
- Glin, L., Arnaud, C., Berracochea, D., Galey, D., Jaffard, R., and Gottesmann, C. (1991). The intermediate stage of sleep in mice. *Physiol. Behav.* *50*, 951–953.
- Gothard, K.M., Skaggs, W.E., Moore, K.M., and McNaughton, B.L. (1996). Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. *J. Neurosci. Off. J. Soc. Neurosci.* *16*, 823–835.
- Gottesmann, C. (1992). Detection of seven sleep-waking stages in the rat. *Neurosci. Biobehav. Rev.* *16*, 31–38.
- Gottesmann, C., Gandolfo, G., Arnaud, C., and Gauthier, P. (1998). The intermediate stage and paradoxical sleep in the rat: influence of three generations of hypnotics. *Eur. J. Neurosci.* *10*, 409–414.
- Grosmark, A.D. (2013). A “gestalt” rule for spontaneous hippocampal sequences during sleep (San Diego, CA.).
- Grosmark, A.D., Pastalkova, E., Amarasingham, A., Diba, K., and Buzsáki, G. (2010). Dynamics of pair-wise correlation patterns across sleep states (San Diego, CA.).
- Grosmark, A.D., Mizuseki, K., Pastalkova, E., Diba, K., and Buzsáki, G. (2012). REM sleep reorganizes hippocampal excitability. *Neuron* *75*, 1001–1007.
- Gu, Q. (2002). Neuromodulatory transmitter systems in the cortex and their role in cortical plasticity. *Neuroscience* *111*, 815–835.

- Gujar, N., McDonald, S.A., Nishida, M., and Walker, M.P. (2011). A role for REM sleep in recalibrating the sensitivity of the human brain to specific emotions. *Cereb. Cortex N. Y. N 1991* *21*, 115–123.
- 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*, 695–705.
- Habib, D., and Dringenberg, H.C. (2010). Low-frequency-induced synaptic potentiation: a paradigm shift in the field of memory-related plasticity mechanisms? *Hippocampus* *20*, 29–35.
- Hahn, T.T.G., McFarland, J.M., Berberich, S., Sakmann, B., and Mehta, M.R. (2012). Spontaneous persistent activity in entorhinal cortex modulates cortico-hippocampal interaction in vivo. *Nat. Neurosci.* *15*, 1531–1538.
- Halassa, M.M., Dal Maschio, M., Beltramo, R., Haydon, P.G., Benfenati, F., and Fellin, T. (2010). Integrated brain circuits: neuron-astrocyte interaction in sleep-related rhythmogenesis. *ScientificWorldJournal* *10*, 1634–1645.
- Hanlon, E.C., Faraguna, U., Vyazovskiy, V.V., Tononi, G., and Cirelli, C. (2009). Effects of skilled training on sleep slow wave activity and cortical gene expression in the rat. *Sleep* *32*, 719–729.
- Hanlon, E.C., Vyazovskiy, V.V., Faraguna, U., Tononi, G., and Cirelli, C. (2011). Synaptic potentiation and sleep need: clues from molecular and electrophysiological studies. *Curr. Top. Med. Chem.* *11*, 2472–2482.
- Hansson, E., and Rönnebeck, L. (1995). Astrocytes in glutamate neurotransmission. *FASEB J. Off. Publ. Fed. Am. Soc. Exp. Biol.* *9*, 343–350.
- Harris, K.D., Henze, D.A., Csicsvari, J., Hirase, H., and Buzsáki, G. (2000). Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *J. Neurophysiol.* *84*, 401–414.
- Harris, K.D., Henze, D.A., Hirase, H., Leinekugel, X., Dragoi, G., Czurkó, A., and Buzsáki, G. (2002). Spike train dynamics predicts theta-related phase precession in hippocampal pyramidal cells. *Nature* *417*, 738–741.
- Hassani, O.K., Lee, M.G., and Jones, B.E. (2009). Melanin-concentrating hormone neurons discharge in a reciprocal manner to orexin neurons across the sleep-wake cycle. *Proc. Natl. Acad. Sci. U. S. A.* *106*, 2418–2422.
- Hasselmo, M.E. (2008). Temporally structured replay of neural activity in a model of entorhinal cortex, hippocampus and postsubiculum. *Eur. J. Neurosci.* *28*, 1301–1315.

- Hayman, R.M.A., Chakraborty, S., Anderson, M.I., and Jeffery, K.J. (2003). Context-specific acquisition of location discrimination by hippocampal place cells. *Eur. J. Neurosci.* *18*, 2825–2834.
- Hazan, L., Zugaro, M., and Buzsáki, G. (2006). Klusters, NeuroScope, NDManager: a free software suite for neurophysiological data processing and visualization. *J. Neurosci. Methods* *155*, 207–216.
- Hebb, D.O. (1949). *The Organization of Behavior: A Neuropsychological Theory* (Psychology Press).
- Hirase, H., Leinekugel, X., Czurkó, A., Csicsvari, J., and Buzsáki, G. (2001). Firing rates of hippocampal neurons are preserved during subsequent sleep episodes and modified by novel awake experience. *Proc. Natl. Acad. Sci. U. S. A.* *98*, 9386–9390.
- Hobson, J.A. (1990). Sleep and dreaming. *J. Neurosci. Off. J. Soc. Neurosci.* *10*, 371–382.
- Hobson, J.A., McCarley, R.W., and Wyzinski, P.W. (1975). Sleep cycle oscillation: reciprocal discharge by two brainstem neuronal groups. *Science* *189*, 55–58.
- Hobson, J.A., Pace-Schott, E.F., and Stickgold, R. (2000). Dreaming and the brain: toward a cognitive neuroscience of conscious states. *Behav. Brain Sci.* *23*, 793–842; discussion 904–1121.
- Hollup, S.A., Molden, S., Donnett, J.G., Moser, M.B., and Moser, E.I. (2001). Accumulation of hippocampal place fields at the goal location in an annular watermaze task. *J. Neurosci. Off. J. Soc. Neurosci.* *21*, 1635–1644.
- Horne, J.A., and McGrath, M.J. (1984). The consolidation hypothesis for REM sleep function: stress and other confounding factors--a review. *Biol. Psychol.* *18*, 165–184.
- Hubel, D.H., and Wiesel, T.N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.* *160*, 106–154.2.
- Huber, R., Ghilardi, M.F., Massimini, M., and Tononi, G. (2004). Local sleep and learning. *Nature* *430*, 78–81.
- Huber, R., Esser, S.K., Ferrarelli, F., Massimini, M., Peterson, M.J., and Tononi, G. (2007). TMS-induced cortical potentiation during wakefulness locally increases slow wave activity during sleep. *PloS One* *2*, e276.
- Huxter, J., Burgess, N., and O'Keefe, J. (2003). Independent rate and temporal coding in hippocampal pyramidal cells. *Nature* *425*, 828–832.

- Isomura, Y., Sirota, A., Ozen, S., Montgomery, S., Mizuseki, K., Henze, D.A., and Buzsáki, G. (2006). Integration and segregation of activity in entorhinal-hippocampal subregions by neocortical slow oscillations. *Neuron* 52, 871–882.
- Jackson, C., McCabe, B.J., Nicol, A.U., Grout, A.S., Brown, M.W., and Horn, G. (2008). Dynamics of a memory trace: effects of sleep on consolidation. *Curr. Biol. CB* 18, 393–400.
- Jeffery, K.J. (2011). Place Cells, Grid Cells, Attractors, and Remapping. *Neural Plast.* 2011, 1–11.
- Ji, D., and Wilson, M.A. (2007). Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nat. Neurosci.* 10, 100–107.
- Johnson, A., and Redish, A.D. (2005). Hippocampal replay contributes to within session learning in a temporal difference reinforcement learning model. *Neural Netw. Off. J. Int. Neural Netw. Soc.* 18, 1163–1171.
- Jones, B.E. (2003). Arousal systems. *Front. Biosci. J. Virtual Libr.* 8, s438–451.
- Jouvet, M. (1967). Neurophysiology of the states of sleep. *Physiol. Rev.* 47, 117–177.
- Jouvet, M., and Michel, F. (1960). [New research on the structures responsible for the “paradoxical phase” of sleep]. *J. Physiol. (Paris)* 52, 130–131.
- Káli, S., and Dayan, P. (2004). Off-line replay maintains declarative memories in a model of hippocampal-neocortical interactions. *Nat. Neurosci.* 7, 286–294.
- Karlsson, M.P., and Frank, L.M. (2009). Awake replay of remote experiences in the hippocampus. *Nat. Neurosci.* 12, 913–918.
- Lee Kavanau, J. (2002). REM and NREM sleep as natural accompaniments of the evolution of warm-bloodedness. *Neurosci. Biobehav. Rev.* 26, 889–906.
- Kawamura, H., and Sawyer, C.H. (1965). Elevation in brain temperature during paradoxical sleep. *Science* 150, 912–913.
- Kemp, N., and Bashir, Z.I. (2001). Long-term depression: a cascade of induction and expression mechanisms. *Prog. Neurobiol.* 65, 339–365.
- Kenet, T., Bibitchkov, D., Tsodyks, M., Grinvald, A., and Arieli, A. (2003). Spontaneously emerging cortical representations of visual attributes. *Nature* 425, 954–956.
- Kentros, C., Hargreaves, E., Hawkins, R.D., Kandel, E.R., Shapiro, M., and Muller, R.V. (1998). Abolition of Long-Term Stability of New Hippocampal Place Cell Maps by NMDA Receptor Blockade. *Science* 280, 2121–2126.

- Kesner, R.P., Gilbert, P.E., and Barua, L.A. (2002). The role of the hippocampus in memory for the temporal order of a sequence of odors. *Behav. Neurosci.* *116*, 286–290.
- Killgore, W.D.S. (2010). Effects of sleep deprivation on cognition. *Prog. Brain Res.* *185*, 105–129.
- Kim, S.M., Ganguli, S., and Frank, L.M. (2012). Spatial information outflow from the hippocampal circuit: distributed spatial coding and phase precession in the subiculum. *J. Neurosci. Off. J. Soc. Neurosci.* *32*, 11539–11558.
- Kloosterman, F. (2012). Analysis of Hippocampal Memory Replay Using Neural Population Decoding. In *Neuronal Network Analysis*, T. Fellin, and M. Halassa, eds. (Humana Press), pp. 259–282.
- Kohyama, J. (2000). REM sleep atonia: responsible brain regions, quantification, and clinical implication. *Brain Dev.* *22 Suppl 1*, S136–142.
- Kudrimoti, H.S., Barnes, C.A., and McNaughton, B.L. (1999). Reactivation of hippocampal cell assemblies: effects of behavioral state, experience, and EEG dynamics. *J. Neurosci. Off. J. Soc. Neurosci.* *19*, 4090–4101.
- Lai, Y.-Y., Kodama, T., Schenkel, E., and Siegel, J.M. (2010). Behavioral response and transmitter release during atonia elicited by medial medullary stimulation. *J. Neurophysiol.* *104*, 2024–2033.
- Lansink, C.S., Goltstein, P.M., Lankelma, J.V., McNaughton, B.L., and Pennartz, C.M.A. (2009). Hippocampus leads ventral striatum in replay of place-reward information. *PLoS Biol.* *7*, e1000173.
- Lee, A.K., and Wilson, M.A. (2002). Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron* *36*, 1183–1194.
- Lesku, J.A., Roth, T.C., Rattenborg, N.C., Amlaner, C.J., and Lima, S.L. (2008). Phylogenetics and the correlates of mammalian sleep: a reappraisal. *Sleep Med. Rev.* *12*, 229–244.
- Leutgeb, S., Leutgeb, J.K., Barnes, C.A., Moser, E.I., McNaughton, B.L., and Moser, M.-B. (2005). Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science* *309*, 619–623.
- Liu, Z.-W., Faraguna, U., Cirelli, C., Tononi, G., and Gao, X.-B. (2010). Direct evidence for wake-related increases and sleep-related decreases in synaptic strength in rodent cortex. *J. Neurosci. Off. J. Soc. Neurosci.* *30*, 8671–8675.
- Llinás, R.R., and Paré, D. (1991). Of dreaming and wakefulness. *Neuroscience* *44*, 521–535.

- Louie, K., and Wilson, M.A. (2001). Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron* 29, 145–156.
- Lubenov, E.V., and Siapas, A.G. (2008). Decoupling through synchrony in neuronal circuits with propagation delays. *Neuron* 58, 118–131.
- Lubenov, E.V., and Siapas, A.G. (2009). Hippocampal theta oscillations are travelling waves. *Nature* 459, 534–539.
- Luczak, A., Barthó, P., Marguet, S.L., Buzsáki, G., and Harris, K.D. (2007). Sequential structure of neocortical spontaneous activity in vivo. *Proc. Natl. Acad. Sci.* 104, 347–352.
- Luczak, A., Barthó, P., and Harris, K.D. (2009). Spontaneous Events Outline the Realm of Possible Sensory Responses in Neocortical Populations. *Neuron* 62, 413–425.
- Luppi, P.-H., Clement, O., Sapin, E., Peyron, C., Gervasoni, D., Léger, L., and Fort, P. (2012). Brainstem mechanisms of paradoxical (REM) sleep generation. *Pflug. Arch. Eur. J. Physiol.* 463, 43–52.
- Mallick, B.N., Singh, S., and Singh, A. (2010). Mechanism of noradrenaline-induced stimulation of Na-K ATPase activity in the rat brain: implications on REM sleep deprivation-induced increase in brain excitability. *Mol. Cell. Biochem.* 336, 3–16.
- Maquet, P. (2001). The role of sleep in learning and memory. *Science* 294, 1048–1052.
- Maquet, P., Péters, J., Aerts, J., Delfiore, G., Degueldre, C., Luxen, A., and Franck, G. (1996). Functional neuroanatomy of human rapid-eye-movement sleep and dreaming. *Nature* 383, 163–166.
- Marshall, L., Henze, D.A., Hirase, H., Leinekugel, X., Dragoi, G., and Buzsáki, G. (2002). Hippocampal pyramidal cell-interneuron spike transmission is frequency dependent and responsible for place modulation of interneuron discharge. *J. Neurosci. Off. J. Soc. Neurosci.* 22, RC197.
- McCarley, R.W. (2007). Neurobiology of REM and NREM sleep. *Sleep Med.* 8, 302–330.
- 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. *Psychol. Rev.* 102, 419–457.
- McCormick, D.A. (1992). Neurotransmitter actions in the thalamus and cerebral cortex and their role in neuromodulation of thalamocortical activity. *Prog. Neurobiol.* 39, 337–388.

- McCormick, D.A. (1993). Actions of acetylcholine in the cerebral cortex and thalamus and implications for function. *Prog. Brain Res.* 98, 303–308.
- McDermott, C.M., Hardy, M.N., Bazan, N.G., and Magee, J.C. (2006). Sleep deprivation-induced alterations in excitatory synaptic transmission in the CA1 region of the rat hippocampus. *J. Physiol.* 570, 553–565.
- Melik, E., Babar, E., and Guven, M. (2006). Effects of AP5 infusion into the lateral ventricle on the activities and hippocampal electrical patterns of sleep episodes in rats. *Physiol. Behav.* 87, 377–382.
- Millenson, J.R., Kehoe, E.J., and Gormezano, I. (1977). Classical conditioning of the rabbit's nictitating membrane response under fixed and mixed CS-US intervals. *Learn. Motiv.* 8, 351–366.
- Mizuseki, K., and Buzsáki, G. (2013). Preconfigured, skewed distribution of firing rates in the hippocampus and entorhinal cortex. *Cell Rep.* 4, 1010–1021.
- Mizuseki, K., Sirota, A., Pastalkova, E., and Buzsáki, G. (2009). Theta oscillations provide temporal windows for local circuit computation in the entorhinal-hippocampal loop. *Neuron* 64, 267–280.
- Mizuseki, K., Diba, K., Pastalkova, E., and Buzsáki, G. (2011). Hippocampal CA1 pyramidal cells form functionally distinct sublayers. *Nat. Neurosci.* 14, 1174–1181.
- Mizuseki, K., Royer, S., Diba, K., and Buzsáki, G. (2012). Activity dynamics and behavioral correlates of CA3 and CA1 hippocampal pyramidal neurons. *Hippocampus*.
- Mölle, M., and Born, J. (2011). Slow oscillations orchestrating fast oscillations and memory consolidation. *Prog. Brain Res.* 193, 93–110.
- Montgomery, S.M., Sirota, A., and Buzsáki, G. (2008). Theta and gamma coordination of hippocampal networks during waking and rapid eye movement sleep. *J. Neurosci. Off. J. Soc. Neurosci.* 28, 6731–6741.
- Moore, G.P., Rosenberg, J.R., Hary, D., and Breeze, P. (1996). Replay“ of hippocampal ”memories. *Science* 274, 1216–1217.
- 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.
- Moser, E., Mathiesen, I., and Andersen, P. (1993). Association between brain temperature and dentate field potentials in exploring and swimming rats. *Science* 259, 1324–1326.

- Mukhametov, L.M. (1987). Unihemispheric slow-wave sleep in the Amazonian dolphin, *Inia geoffrensis*. *Neurosci. Lett.* *79*, 128–132.
- Muller, R.U., and Kubie, J.L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J. Neurosci. Off. J. Soc. Neurosci.* *7*, 1951–1968.
- Nádasdy, Z. (2000). Spike sequences and their consequences. *J. Physiol. Paris* *94*, 505–524.
- Nádasdy, Z., Hirase, H., Czurkó, A., Csicsvari, J., and Buzsáki, G. (1999). Replay and time compression of recurring spike sequences in the hippocampus. *J. Neurosci. Off. J. Soc. Neurosci.* *19*, 9497–9507.
- 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*, 781–787.
- Nakazawa, K., McHugh, T.J., Wilson, M.A., and Tonegawa, S. (2004). NMDA receptors, place cells and hippocampal spatial memory. *Nat. Rev. Neurosci.* *5*, 361–372.
- O’Keefe, J., and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* *34*, 171–175.
- O’Neill, J., Senior, T.J., Allen, K., Huxter, J.R., and Csicsvari, J. (2008). Reactivation of experience-dependent cell assembly patterns in the hippocampus. *Nat. Neurosci.* *11*, 209–215.
- Ostojic, S., Brunel, N., and Hakim, V. (2009). How Connectivity, Background Activity, and Synaptic Properties Shape the Cross-Correlation between Spike Trains. *J. Neurosci.* *29*, 10234–10253.
- Pastalkova, E., Itskov, V., Amarasingham, A., and Buzsáki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. *Science* *321*, 1322–1327.
- Patel, J., Fujisawa, S., Berényi, A., Royer, S., and Buzsáki, G. (2012). Traveling theta waves along the entire septotemporal axis of the hippocampus. *Neuron* *75*, 410–417.
- Patel, J., Schomburg, E.W., Berényi, A., Fujisawa, S., and Buzsáki, G. (2013). Local generation and propagation of ripples along the septotemporal axis of the hippocampus. *J. Neurosci. Off. J. Soc. Neurosci.* *33*, 17029–17041.
- Pavlidis, C., and Winson, J. (1989a). Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. *J. Neurosci. Off. J. Soc. Neurosci.* *9*, 2907–2918.

- Pavlidis, C., and Winson, J. (1989b). Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. *J. Neurosci. Off. J. Soc. Neurosci.* *9*, 2907–2918.
- Pelletier, J.G., and Lacaille, J.-C. (2008). Long-term synaptic plasticity in hippocampal feedback inhibitory networks. *Prog. Brain Res.* *169*, 241–250.
- Pennartz, C.M.A., Lee, E., Verheul, J., Lipa, P., Barnes, C.A., and McNaughton, B.L. (2004). The ventral striatum in off-line processing: ensemble reactivation during sleep and modulation by hippocampal ripples. *J. Neurosci. Off. J. Soc. Neurosci.* *24*, 6446–6456.
- Perkel, D.H., Gerstein, G.L., and Moore, G.P. (1967a). Neuronal spike trains and stochastic point processes. II. Simultaneous spike trains. *Biophys. J.* *7*, 419–440.
- Perkel, D.H., Gerstein, G.L., and Moore, G.P. (1967b). Neuronal spike trains and stochastic point processes. I. The single spike train. *Biophys. J.* *7*, 391–418.
- Peyrache, A., Khamassi, M., Benchenane, K., Wiener, S.I., and Battaglia, F.P. (2009). Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nat. Neurosci.* *12*, 919–926.
- Peyrache, A., Benchenane, K., Khamassi, M., Wiener, S.I., and Battaglia, F.P. (2010). Sequential Reinstatement of Neocortical Activity during Slow Oscillations Depends on Cells' Global Activity. *Front. Syst. Neurosci.* *3*, 18.
- Pfeiffer, B.E., and Foster, D.J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* *497*, 74–79.
- Plihal, W., and Born, J. (1997). Effects of Early and Late Nocturnal Sleep on Declarative and Procedural Memory. *J. Cogn. Neurosci.* *9*, 534–547.
- Poe, G.R., Nitz, D.A., McNaughton, B.L., and Barnes, C.A. (2000). Experience-dependent phase-reversal of hippocampal neuron firing during REM sleep. *Brain Res.* *855*, 176–180.
- Pompeiano, M., Cirelli, C., Arrighi, P., and Tononi, G. (1995). c-Fos expression during wakefulness and sleep. *Neurophysiol. Clin. Clin. Neurophysiol.* *25*, 329–341.
- Popa, D., Duvarci, S., Popescu, A.T., Léna, C., and Paré, D. (2010). Coherent amygdalocortical theta promotes fear memory consolidation during paradoxical sleep. *Proc. Natl. Acad. Sci. U. S. A.* *107*, 6516–6519.
- Prospero-García, O., Criado, J.R., and Henriksen, S.J. (1994). Pharmacology of ethanol and glutamate antagonists on rodent sleep: a comparative study. *Pharmacol. Biochem. Behav.* *49*, 413–416.

- 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*, 1525–1533.
- Quilichini, P., Sirota, A., and Buzsáki, G. (2010). Intrinsic circuit organization and theta-gamma oscillation dynamics in the entorhinal cortex of the rat. *J. Neurosci. Off. J. Soc. Neurosci.* *30*, 11128–11142.
- Quinn, J.J., Ma, Q.D., Tinsley, M.R., Koch, C., and Fanselow, M.S. (2008). Inverse temporal contributions of the dorsal hippocampus and medial prefrontal cortex to the expression of long-term fear memories. *Learn. Mem. Cold Spring Harb. N* *15*, 368–372.
- Rattenborg, N.C., Amlaner, C.J., and Lima, S.L. (2000). Behavioral, neurophysiological and evolutionary perspectives on unihemispheric sleep. *Neurosci. Biobehav. Rev.* *24*, 817–842.
- Rauchs, G., Desgranges, B., Foret, J., and Eustache, F. (2005). The relationships between memory systems and sleep stages. *J. Sleep Res.* *14*, 123–140.
- Ray, J., and Hansen, S. (2004). Temperament in the Rat: Sex Differences and Hormonal Influences on Harm Avoidance and Novelty Seeking. *Behav. Neurosci.* *118*, 488–497.
- Ray, J., Hansen, S., and Waters, N. (2006). Links between temperamental dimensions and brain monoamines in the rat. *Behav. Neurosci.* *120*, 85–92.
- Redish, A.D., Rosenzweig, E.S., Bohanick, J.D., McNaughton, B.L., and Barnes, C.A. (2000). Dynamics of hippocampal ensemble activity realignment: time versus space. *J. Neurosci. Off. J. Soc. Neurosci.* *20*, 9298–9309.
- Ribeiro, S., Mello, C.V., Velho, T., Gardner, T.J., Jarvis, E.D., and Pavlides, C. (2002). Induction of hippocampal long-term potentiation during waking leads to increased extrahippocampal zif-268 expression during ensuing rapid-eye-movement sleep. *J. Neurosci. Off. J. Soc. Neurosci.* *22*, 10914–10923.
- Ribeiro, S., Gervasoni, D., Soares, E.S., Zhou, Y., Lin, S.-C., Pantoja, J., Lavine, M., and Nicolelis, M.A.L. (2004). Long-lasting novelty-induced neuronal reverberation during slow-wave sleep in multiple forebrain areas. *PLoS Biol.* *2*, E24.
- Riedner, B.A., Hulse, B.K., Murphy, M.J., Ferrarelli, F., and Tononi, G. (2011). Temporal dynamics of cortical sources underlying spontaneous and peripherally evoked slow waves. *Prog. Brain Res.* *193*, 201–218.
- Robinson, T.E., Kramis, R.C., and Vanderwolf, C.H. (1977). Two types of cerebral activation during active sleep: relations to behavior. *Brain Res.* *124*, 544–549.

- De la Rocha, J., Doiron, B., Shea-Brown, E., Josić, K., and Reyes, A. (2007). Correlation between neural spike trains increases with firing rate. *Nature* *448*, 802–806.
- Rolls, A., Colas, D., Adamantidis, A., Carter, M., Lanre-Amos, T., Heller, H.C., and de Lecea, L. (2011). Optogenetic disruption of sleep continuity impairs memory consolidation. *Proc. Natl. Acad. Sci. U. S. A.* *108*, 13305–13310.
- Rosanova, M., and Ulrich, D. (2005). Pattern-specific associative long-term potentiation induced by a sleep spindle-related spike train. *J. Neurosci. Off. J. Soc. Neurosci.* *25*, 9398–9405.
- Royer, S., Zemelman, B.V., Losonczy, A., Kim, J., Chance, F., Magee, J.C., and Buzsáki, G. (2012). Control of timing, rate and bursts of hippocampal place cells by dendritic and somatic inhibition. *Nat. Neurosci.* *15*, 769–775.
- Saar, D., and Barkai, E. (2009). Long-lasting maintenance of learning-induced enhanced neuronal excitability: mechanisms and functional significance. *Mol. Neurobiol.* *39*, 171–177.
- Saito, H., Sakai, K., and Jouvet, M. (1977). Discharge patterns of the nucleus parabrachialis lateralis neurons of the cat during sleep and waking. *Brain Res.* *134*, 59–72.
- Saper, C.B., Fuller, P.M., Pedersen, N.P., Lu, J., and Scammell, T.E. (2010). Sleep state switching. *Neuron* *68*, 1023–1042.
- Schousboe, A., and Waagepetersen, H.S. (2005). Role of astrocytes in glutamate homeostasis: implications for excitotoxicity. *Neurotox. Res.* *8*, 221–225.
- Schulz, H. (2008). Rethinking sleep analysis. *J. Clin. Sleep Med. JCSM Off. Publ. Am. Acad. Sleep Med.* *4*, 99–103.
- Schwindel, C.D., and McNaughton, B.L. (2011). Hippocampal-cortical interactions and the dynamics of memory trace reactivation. *Prog. Brain Res.* *193*, 163–177.
- Scoville, W.B., and Milner, B. (1957). Loss of Recent Memory After Bilateral Hippocampal Lesions. *J. Neurol. Neurosurg. Psychiatry* *20*, 11–21.
- Sejnowski, T.J., and Destexhe, A. (2000). Why do we sleep? *Brain Res.* *886*, 208–223.
- Siapas, A.G., and Wilson, M.A. (1998). Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep. *Neuron* *21*, 1123–1128.
- Siapas, A.G., Lubenov, E.V., and Wilson, M.A. (2005). Prefrontal phase locking to hippocampal theta oscillations. *Neuron* *46*, 141–151.
- Siegel, J.M. (1995). Phylogeny and the function of REM sleep. *Behav. Brain Res.* *69*, 29–34.

- Siegel, J.M. (2005). Clues to the functions of mammalian sleep. *Nature* 437, 1264–1271.
- Siegel, J.M. (2011). REM sleep: a biological and psychological paradox. *Sleep Med. Rev.* 15, 139–142.
- Sil'kis, I.G. (2009). Characteristics of the functioning of the hippocampal formation in waking and paradoxical sleep. *Neurosci. Behav. Physiol.* 39, 523–534.
- Silber, M.H., Ancoli-Israel, S., Bonnet, M.H., Chokroverty, S., Grigg-Damberger, M.M., Hirshkowitz, M., Kapen, S., Keenan, S.A., Kryger, M.H., Penzel, T., et al. (2007). The visual scoring of sleep in adults. *J. Clin. Sleep Med. JCSM Off. Publ. Am. Acad. Sleep Med.* 3, 121–131.
- Sirota, A., and Buzsáki, G. (2005). Interaction between neocortical and hippocampal networks via slow oscillations. *Thalamus Relat. Syst.* 3, 245–259.
- Sirota, A., Csicsvari, J., Buhl, D., and Buzsáki, G. (2003). Communication between neocortex and hippocampus during sleep in rodents. *Proc. Natl. Acad. Sci. U. S. A.* 100, 2065–2069.
- Sirota, A., Montgomery, S., Fujisawa, S., Isomura, Y., Zugaro, M., and Buzsáki, G. (2008). Entrainment of neocortical neurons and gamma oscillations by the hippocampal theta rhythm. *Neuron* 60, 683–697.
- 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.
- Smith, C. (2001). Sleep states and memory processes in humans: procedural versus declarative memory systems. *Sleep Med. Rev.* 5, 491–506.
- Squire, L.R. (1992). Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychol. Rev.* 99, 195–231.
- Stark, E., Koos, T., and Buzsáki, G. (2012). Diode probes for spatiotemporal optical control of multiple neurons in freely moving animals. *J. Neurophysiol.* 108, 349–363.
- Steriade, M. (1995). Thalamic origin of sleep spindles: Morison and Bassett (1945). *J. Neurophysiol.* 73, 921–922.
- Steriade, M. (2003). *Neuronal Substrates of Sleep and Epilepsy* (Cambridge University Press).
- Steriade, M. (2004). Acetylcholine systems and rhythmic activities during the waking--sleep cycle. *Prog. Brain Res.* 145, 179–196.
- Steriade, M., and Amzica, F. (1998). Coalescence of sleep rhythms and their chronology in corticothalamic networks. *Sleep Res. Online SRO* 1, 1–10.

- Steriade, M., McCormick, D.A., and Sejnowski, T.J. (1993). Thalamocortical oscillations in the sleeping and aroused brain. *Science* 262, 679–685.
- Steriade, M., Timofeev, I., and Grenier, F. (2001). Natural waking and sleep states: a view from inside neocortical neurons. *J. Neurophysiol.* 85, 1969–1985.
- Stickgold, R. (2005). Sleep-dependent memory consolidation. *Nature* 437, 1272–1278.
- Stickgold, R., and Walker, M.P. (2007). Sleep-dependent memory consolidation and reconsolidation. *Sleep Med.* 8, 331–343.
- Stickgold, R., Whidbee, D., Schirmer, B., Patel, V., and Hobson, J.A. (2000). Visual discrimination task improvement: A multi-step process occurring during sleep. *J. Cogn. Neurosci.* 12, 246–254.
- Sullivan, D., Csicsvari, J., Mizuseki, K., Montgomery, S., Diba, K., and Buzsáki, G. (2011). Relationships between hippocampal sharp waves, ripples, and fast gamma oscillation: influence of dentate and entorhinal cortical activity. *J. Neurosci. Off. J. Soc. Neurosci.* 31, 8605–8616.
- Sullivan, D., Mizuseki, K., Sorgi, A., and Buzsáki, G. (2014). Comparison of sleep spindles and theta oscillations in the hippocampus. *J. Neurosci. Off. J. Soc. Neurosci.* 34, 662–674.
- Tatsuno, M., Lipa, P., and McNaughton, B.L. (2006). Methodological considerations on the use of template matching to study long-lasting memory trace replay. *J. Neurosci. Off. J. Soc. Neurosci.* 26, 10727–10742.
- Taube, J.S. (2007). The head direction signal: origins and sensory-motor integration. *Annu. Rev. Neurosci.* 30, 181–207.
- Terrier, G., and Gottesmann, C. (1978). Study of cortical spindles during sleep in the rat. *Brain Res. Bull.* 3, 701–706.
- Tononi, G. (2009). Slow wave homeostasis and synaptic plasticity. *J. Clin. Sleep Med. JCSM Off. Publ. Am. Acad. Sleep Med.* 5, S16–19.
- Tononi, G., and Cirelli, C. (2003). Sleep and synaptic homeostasis: a hypothesis. *Brain Res. Bull.* 62, 143–150.
- Tononi, G., and Cirelli, C. (2006a). Sleep function and synaptic homeostasis. *Sleep Med. Rev.* 10, 49–62.
- Tononi, G., and Cirelli, C. (2006b). Sleep function and synaptic homeostasis. *Sleep Med. Rev.* 10, 49–62.
- Tononi, G., Pompeiano, M., and Ronca-Testoni, S. (1990). Noradrenergic receptor binding during sleep-waking states in the rat. *Arch. Ital. Biol.* 128, 67–76.

- Tononi, G., Pompeiano, M., and Cirelli, C. (1994). The locus coeruleus and immediate-early genes in spontaneous and forced wakefulness. *Brain Res. Bull.* 35, 589–596.
- Turrigiano, G.G. (1999). Homeostatic plasticity in neuronal networks: the more things change, the more they stay the same. *Trends Neurosci.* 22, 221–227.
- Vandecasteele, M., M, S., Royer, S., Belluscio, M., Berényi, A., Diba, K., Fujisawa, S., Grosmark, A., Mao, D., Mizuseki, K., et al. (2012). Large-scale recording of neurons by movable silicon probes in behaving rodents. *J. Vis. Exp. JoVE* e3568.
- Vazquez, J., and Baghdoyan, H.A. (2001). Basal forebrain acetylcholine release during REM sleep is significantly greater than during waking. *Am. J. Physiol. Regul. Integr. Comp. Physiol.* 280, R598–601.
- Vertes, R.P. (1984). Brainstem control of the events of REM sleep. *Prog. Neurobiol.* 22, 241–288.
- Vetrivelan, R., Fuller, P.M., Tong, Q., and Lu, J. (2009). Medullary circuitry regulating rapid eye movement sleep and motor atonia. *J. Neurosci. Off. J. Soc. Neurosci.* 29, 9361–9369.
- Vyazovskiy, V.V., Cirelli, C., Tononi, G., and Tobler, I. (2008). Cortical metabolic rates as measured by 2-deoxyglucose-uptake are increased after waking and decreased after sleep in mice. *Brain Res. Bull.* 75, 591–597.
- Vyazovskiy, V.V., Olcese, U., Lazimy, Y.M., Faraguna, U., Esser, S.K., Williams, J.C., Cirelli, C., and Tononi, G. (2009). Cortical firing and sleep homeostasis. *Neuron* 63, 865–878.
- Vyazovskiy, V.V., Olcese, U., Hanlon, E.C., Nir, Y., Cirelli, C., and Tononi, G. (2011a). Local sleep in awake rats. *Nature* 472, 443–447.
- Vyazovskiy, V.V., Cirelli, C., and Tononi, G. (2011b). Electrophysiological correlates of sleep homeostasis in freely behaving rats. *Prog. Brain Res.* 193, 17–38.
- Wagner, U., Gais, S., and Born, J. (2001). Emotional memory formation is enhanced across sleep intervals with high amounts of rapid eye movement sleep. *Learn. Mem. Cold Spring Harb. N* 8, 112–119.
- Walker, M.P. (2010). Sleep, memory and emotion. *Prog. Brain Res.* 185, 49–68.
- Wallenstein, G.V., Eichenbaum, H., and Hasselmo, M.E. (1998). The hippocampus as an associator of discontinuous events. *Trends Neurosci.* 21, 317–323.
- Whittington, M.A., Traub, R.D., Faulkner, H.J., Stanford, I.M., and Jefferys, J.G. (1997). Recurrent excitatory postsynaptic potentials induced by synchronized fast cortical oscillations. *Proc. Natl. Acad. Sci. U. S. A.* 94, 12198–12203.

- Wierzynski, C.M., Lubenov, E.V., Gu, M., and Siapas, A.G. (2009). State-dependent spike-timing relationships between hippocampal and prefrontal circuits during sleep. *Neuron* *61*, 587–596.
- Wikenheiser, A.M., and David Redish, A. (2013). The balance of forward and backward hippocampal sequences shifts across behavioral states. *Hippocampus* *23*, 22–29.
- Wilson, M.A., and McNaughton, B.L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science* *265*, 676–679.
- Wiltgen, B.J., Brown, R.A.M., Talton, L.E., and Silva, A.J. (2004). New circuits for old memories: the role of the neocortex in consolidation. *Neuron* *44*, 101–108.
- Wolansky, T., Clement, E.A., Peters, S.R., Palczak, M.A., and Dickson, C.T. (2006). Hippocampal slow oscillation: a novel EEG state and its coordination with ongoing neocortical activity. *J. Neurosci. Off. J. Soc. Neurosci.* *26*, 6213–6229.
- Yassin, L., Benedetti, B.L., Jouhanneau, J.-S., Wen, J.A., Poulet, J.F.A., and Barth, A.L. (2010). An embedded subnetwork of highly active neurons in the neocortex. *Neuron* *68*, 1043–1050.
- Ylinen, A., Bragin, A., Nádasdy, Z., Jandó, G., Szabó, I., Sik, A., and Buzsáki, G. (1995). Sharp wave-associated high-frequency oscillation (200 Hz) in the intact hippocampus: network and intracellular mechanisms. *J. Neurosci. Off. J. Soc. Neurosci.* *15*, 30–46.
- Yoshimura, Y., Dantzker, J.L.M., and Callaway, E.M. (2005). Excitatory cortical neurons form fine-scale functional networks. *Nature* *433*, 868–873.
- Zaborszky, L., and Duque, A. (2003). Sleep-wake mechanisms and basal forebrain circuitry. *Front. Biosci. J. Virtual Libr.* *8*, d1146–1169.
- Zhang, W., and Linden, D.J. (2003). The other side of the engram: experience-driven changes in neuronal intrinsic excitability. *Nat. Rev. Neurosci.* *4*, 885–900.
- Zugaro, M.B., Berthoz, A., and Wiener, S.I. (2001). Background, but not foreground, spatial cues are taken as references for head direction responses by rat anterodorsal thalamus neurons. *J. Neurosci. Off. J. Soc. Neurosci.* *21*, RC154.
- Zugaro, M.B., Arleo, A., Déjean, C., Burguière, E., Khamassi, M., and Wiener, S.I. (2004). Rat anterodorsal thalamic head direction neurons depend upon dynamic visual signals to select anchoring landmark cues. *Eur. J. Neurosci.* *20*, 530–536.
- (2012). MATLAB and Statistics Toolbox Release 2012b (Natick, Massachusetts: The MathWorks Inc.).

Curriculum Vitae

Andres Grosmark

Born: January 18th, 1984 in Buenos Aires, Argentina

Education

Secondary School: Long Reach High School, Columbia, Maryland, USA 1998-2002

Degree: High School Diploma

College: Duke University

Degrees: -Bachelor of Science with Honors in Biology, with a concentration in Neuroscience

Cum Laude

-Bachelor of Arts in Philosophy

Graduate School: Rutgers University, Behavioral and Neural Sciences Program, The Center for Molecular and Behavioral Neuroscience, Newark, NJ, USA.

Ph.D. Advisor: Dr. György Buzsáki

Previous Work Experience:

Aug. 2001 - Dec. 2001: Intern with Dr. Gloria Su, John Hopkins University

May 2002 - Aug. 2002: Intern with Dr. Jane Trepel, NCI, NIH, Bethesda

May 2003 - Aug. 2003: Intern with Dr. Spyro Mouses, CDDL

Jan. 2005 - June 2006: Student and Research Assistant with Dr. Miguel Nicolelis
Duke University, Neurobiology Department

June 2006- Sept., 2007: Research Assistant with Dr. Sidney Simon, Duke
University, Neurobiology Department

Publications:

Sahin, F., Qiu, W., Wilentz, R. E., Iacobuzio-Donahue, C. A., Grosmark, A., & Su, G. H. (2005). RPL38, FOSL1, and UPP1 are predominantly expressed in the pancreatic ductal epithelium. *Pancreas*, 30(2), 158–167.

Dzirasa, K., Ribeiro, S., Costa, R., Santos, L. M., Lin, S.-C., Grosmark, A., ... Nicolelis, M. A. L. (2006). Dopaminergic control of sleep-wake states. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 26(41), 10577–10589. doi:10.1523/JNEUROSCI.1767-06.2006

Grosmark, A. D., Mizuseki, K., Pastalkova, E., Diba, K., & Buzsáki, G. (2012). REM sleep reorganizes hippocampal excitability. *Neuron*, 75(6), 1001–1007. doi:10.1016/j.neuron.2012.08.015

Vandecasteele, M., M, S., Royer, S., Belluscio, M., Berényi, A., Diba, K., ... Buzsáki, G. (2012). Large-scale recording of neurons by movable silicon probes in behaving rodents. *Journal of Visualized Experiments: JoVE*, (61), e3568. doi:10.3791/3568