Prefrontal-Hippocampal Interactions

Thesis by

Casimir M. Wierzynski

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

2009

(Defended May 6, 2009)

For Jean.

Acknowledgments

First I would like to thank my advisor, Thanos Siapas, for his exceedingly generous and unswerving support for the last six years, and for creating a fantastic lab environment that fosters openness, teamwork, and scientific rigor. Next I would like to thank the rest of my thesis committee, Mark Konishi, Gilles Laurent, and Erin Schuman, for all of their thoughtful and penetrating suggestions, always delivered with great warmth and encouragement. I am grateful to my colleagues in the Siapas Lab, Ming Gu and Eugene Lubenov, with whom I have been privileged to collaborate on some of the work in this thesis, for their friendship and for sharing their insights on neuroscience and life. In addition, I have benefited greatly from interactions with our neighbors in the Laurent Lab, past and present. Last, and by no means least, I would like to thank Mike Walsh for his selfless dedication and technical wizardry that has touched every experimental aspect of this work.

Chapter 3 is reprinted from Wierzynski, C.M., Lubenov, E.V., Gu, M. and Siapas, A.G., State-dependent spike-timing relationships between hippocampal and prefrontal circuits during sleep, *Neuron* **61**, 587-596, Copyright 2009, with permission from Elsevier.

This work was supported by a Department of Defense National Defense Science and Engineering Graduate Fellowship (C.W.), the Caltech Information Science and Technology Center for Biological Circuits Design, the James S. McDonnell Foundation, the Bren Foundation, the McKnight Foundation, the Whitehall Foundation, and the National Institutes of Health.

Abstract

Long-term memories are established in the neocortex under the influence of hippocampal activity. The precise circuit mechanisms underlying this process, however, remain poorly understood. According to the dominant paradigm, memories are formed in two stages: first, neocortical activity during awake behavior embeds traces in hippocampal circuits; second, spontaneous hippocampal activity during offline periods, such as sleep, drives synaptic changes across cortical circuits so as to produce a stable, long-term memory trace. Evidence for this two-stage model at the level of neural activity, however, is incomplete. In this thesis we study interactions between the hippocampus and medial prefrontal cortex (mPFC) to elucidate the basic principles of how these brain circuits work in concert in support of long-term memory. Using recordings of single-unit actvity from multi-tetrode arrays in the hippocampus and mPFC of freely behaving rats, we performed two sets of experiments, each addressing one stage of the two-stage model. First, during awake behavior, we find a class of mPFC cells whose firing reflects the strength of a learned association and show that these tend to be strongly modulated by the hippocampus. Second, during sleep, we identify precise spike timing relationships between single mPFC and hippocampal cells that are consistent with information flow from the hippocampus to the prefrontal cortex, and show that these timing relationships are highly dependent on sleep stage. Taken together, these results provide key constraints on the circuit mechanisms of long-term memory formation.

Contents

Acknowledgments **iv**

Abstract **v**

1 Introduction **1**

1.1 Memory Consolidation 1

1.2 The Two-Stage Model of Memory Formation 2

1.3 The Link between Sleep and Memory 3

1.4 The Prefrontal-Hippocampal Pathway 4

1.5 Preview 5

2 Prefrontal-Hippocampal Interactions during Learning **6**

2.1 Introduction 6

2.2 Results 7

2.2.1 Task-dependent firing in mPFC neurons 7

2.2.2 Evolution of prefrontal responses during learning 8

2.2.3 Hippocampal modulation of CS-excited cells 13

2.3 Discussion 13

2.4 Methods 18

2.4.1 Electrophysiological recordings. 18

2.4.2 Spike and local field analysis 20

2.4.3 Behavioral training 20

2.4.4 EMG analysis 21

2.4.5 Analysis of unit responses 21
2.4.6 Analysis of phase-locking 22

3 Prefrontal-Hippocampal Interactions during Sleep 23

3.1 Introduction 23
3.2 Results 24
3.2.1 Directionality in prefrontal-hippocampal spike timing 24
3.2.2 State-dependence of prefrontal-hippocampal spike timing . . . 30
3.2.3 Biphasic structure of prefrontal responses 34
3.3 Discussion 37
3.4 Methods 42
3.4.1 Electrophysiological recordings 42
3.4.2 Sleep sessions 43
3.4.3 Spike and local field analysis 43
3.4.4 Sleep stage identification 43
3.4.5 Cross-covariance analysis 44
3.4.6 Multiple Comparison Corrections 45
3.4.7 Population tests of interactions across sleep stages 45
3.4.8 Sharp-wave/Ripple (SWR) event identification 46
3.4.9 Hippocampal burst analysis 46

4 Discussion 48

4.1 The two-stage model revisited 48
4.1.1 Awake behavior 48
4.1.2 Sleep 50
4.2 Directions for Future Work 50
4.2.1 Experience dependence of mPFC-CA1 correlations 50
4.2.2 Theta phase-locking of recent versus remote memories 52

List of Figures

1.1 Opposite temporal gradients for dependence on hippocampal and cortical circuits 2
1.2 The two-stage model of memory formation 3

2.1 Examples of eyelid EMG responses during conditioning 8
2.2 Trace eyeblink conditioning learning curve 9
2.3 Examples of CS-excited cells in the mPFC 10
2.4 Distribution of peak firing times for CS-excited cells 11
2.5 Evolution of firing rates in response to the CS as a function of training 12
2.6 Verification that evolution of CSe firing is not due to non-specific excitability 14
2.7 Evolution of CS-triggered neural responses with learning 15
2.8 CS-excited cells in the mPFC are predominantly phase-locked to hippocampal theta oscillations 16
2.9 Classification of prefrontal cells engaged in learning a hippocampus-dependent task. 19

3.1 Hippocampal and prefrontal spiking activity during sleep 25
3.2 State-dependent cross-covariances between single cells in the hippocampus and prefrontal cortex 26
3.3 Population analysis of state-dependent prefrontal-hippocampal interactions 28

3.4 Absence of relationship between cortico-hippocampal cross-covariances and respective auto-covariances 29
3.5 Illustrations of discrete interactions between mPFC and CA1 across sleep stages 31
3.6 Mean firing rates across SWS and REM sleep 32
3.7 Randomization procedure for testing the effect of unequal sample sizes in SWS versus REM 33
3.8 Incidence of prefrontal-hippocampal interactions during SWS 35
3.9 Multi-phase prefrontal response to hippocampal spiking 36
3.10 Relationship between burst strength and ripple band power 38
3.11 Prefrontal and hippocampal responses to hippocampal bursts as a function of burst strength 39
3.12 Inhibitory prefrontal responses to CA1 bursts 40

4.1 A model of consolidation during eyeblink conditioning 49
4.2 A model of prefrontal responses to hippocampal bursts 51

Chapter 1

Introduction

1.1 Memory Consolidation

Like many brain faculties, memory provides important clues to its normal function when it fails. Anterograde amnesia has been particularly instructive since it stems from a failure of memory formation. This pathology is often due to hippocampal insult, as in the case of the patient Henry Molaison (H.M.), famously studied by Brenda Milner[1]. Since this revolutionary work, a large body of lesion experiments in animals has confirmed that the hippocampus is required to form new memories but is not the long-term site of memory storage[2], thought instead to be widely distributed across the neocortex.

A number of studies have measured the time course of hippocampal dependence for a given task by performing hippocampal lesions at varying points in the animal's training. These typically reveal complete learning deficits after early lesions and nearly intact learning relative to control animals after longer time intervals[2]. This timeline suggests that memories have a life cycle: they are initially dependent on the hippocampal circuits, but after a critical period—weeks to months in rodents, years in humans—they can be recalled by the neocortex even after the hippocampus is removed (Figure 1.1). The neural mechanisms of this transformation, known as *consolidation*, are poorly understood and will be the general subject of this thesis.

The consolidation literature is vast[3] and inconsistent in its nomenclature. To

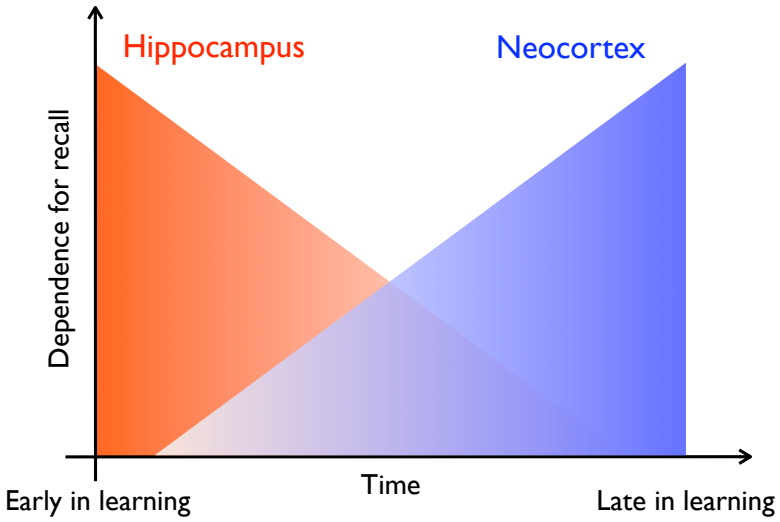

Figure 1.1: **Opposite temporal gradients for dependence on hippocampal and cortical circuits.** Adapted from[6].

avoid confusion, we note that what we call consolidation is sometimes referred to as *systems-* or *network consolidation,* to distinguish it from *cellular consolidation* and *reconsolidation*, which are other processes that modulate the strength of recently acquired or reactivated memories over much shorter time scales (minutes to hours). Cellular consolidation refers to how short-term changes in synaptic efficacy induced by conditioning or tetanic stimulation are prolonged through transcriptional regulation, protein synthesis, and other molecular cascades[4]. Reconsolidation refers to the process by which previously consolidated memories become labile when they are retrieved and require *de novo* protein synthesis to remain stable, long-term memories[5].

1.2 The Two-Stage Model of Memory Formation

The essential but time-limited role of the hippocampus has motivated a widely held model whereby memories are formed in two stages[7]. In the first stage, sensory stimuli reach the neocortex during awake behavior, and the resulting cortical activity embeds traces in hippocampal circuits (Figure 1.2, left). In the second stage, spontaneous hippocampal activity in the absence of sensory stimuli drives synaptic changes across

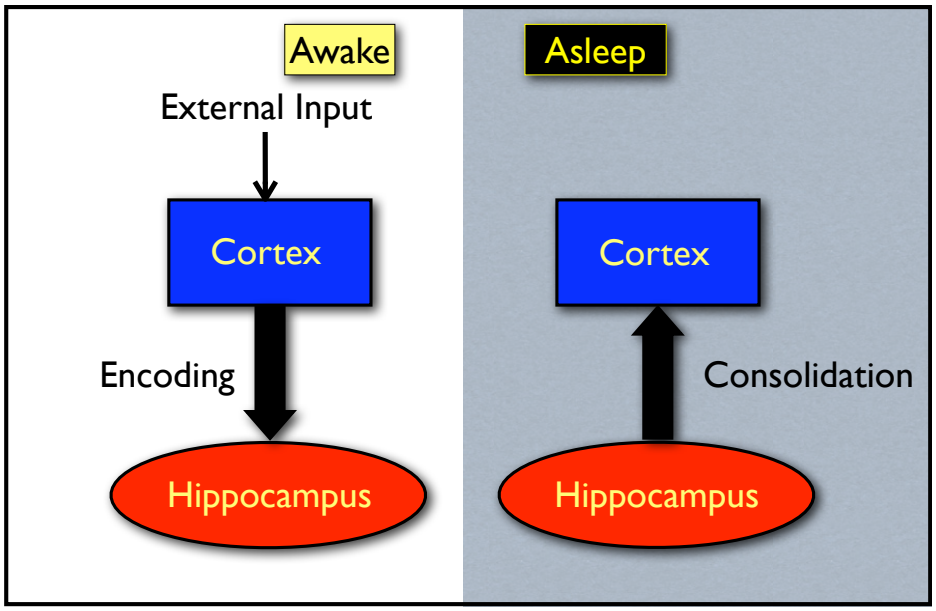

Figure 1.2: **The two-stage model of memory formation.**

cortical circuits so as to produce a stable, long-term memory trace (Figure 1.2, right). This latter stage is postulated to occur during offline periods such as sleep.

Besides providing a fit to lesion data, and perhaps an explanation of why we sleep, this two-stage model also solves a computational problem endemic to auto-associative storage: *catastrophic interference*. Neural networks that learn quickly, after just one or a few training rounds, tend to overfit the data, so that adding another item to the store can destroy earlier memories. Mating a fast-learning hippocampus to the slower-learning neocortex solves this problem by having the hippocampus, in effect, train the neocortex using repeated presentations. With this slower training schedule, the cortical network can discover higher-level regularities in the data that enable more robust storage patterns.[8].

1.3 The Link between Sleep and Memory

A key prediction of the two-stage model is that brain activity during sleep is critical for memory consolidation. Testing this prediction is complicated by the fact that experi-

mentally regulating sleep, either up or down, can modulate many other physiological variables, such as vigilance, that are not specific to memory but can still affect the performance of memory tasks. Nevertheless, a large number of behavioral studies have found relationships between sleep and memory performance[9], although the link remains controversial[10].

At the level of neural activity, however, it is clear that the sleeping brain is highly active and plastic[11]. Moreover, mammalian sleep is comprised of several distinct stages, collectively referred to as rapid eye movement (REM) sleep and slow-wave sleep (SWS), featuring dramatically different neural activity patterns across the neocortex and hippocampus. Coupled to spike timing dependent plasticity mechanisms[12], this activity during sleep would modify synaptic strengths throughout the brain, perhaps in different ways depending on sleep stage. Thus, even without definitive behavioral evidence that sleep supports memory, it is highly likely that brain circuits are reorganized during sleep—a key requirement of the two-stage model—and therefore important to understand the precise nature of cortico-hippocampal activity patterns across all sleep stages.

1.4 The Prefrontal-Hippocampal Pathway

While the two-stage model is an attractive theory for memory formation, experimental evidence and possible mechanisms at the level of neural activity have been difficult to obtain. Focusing on interactions between the prefrontal cortex and the hippocampus provides a useful starting point for this investigation for several reasons:

1. The prefrontal cortex receives direct projections from CA1 pyramidal cells, principally to the prelimbic and infralimbic cortices[13], that are excitatory[14,15], glutamatergic, and plastic[16,17].

2. The prefrontal cortex in rats has been implicated in a number of hippocampus-dependent memory tasks, including associative learning[18], and fear condition-

ing[19]. Moreover, for these tasks, the involvement of the hippocampus and prefrontal cortex follow the opposite temporal gradients illustrated in Figure 1.1.

3. Imaging studies have found differential expression of immediate-early genes in the prefrontal cortex and hippocampus during the retrieval of recent versus remote memories consistent with these gradients[20].

4. There is electrophysiological evidence for coordinated activity between area CA1 and the prefrontal cortex during both awake behavior[21] and slow-wave sleep[22,23].

This combination of direct connectivity, differential engagement in learning tasks and electrophysiological coordination strongly favor the chances of detecting neural activity patterns across the hippocampus and the neocortex that could provide mechanistic understanding of the consolidation process. In this work, therefore, we have focused on the prefrontal-hippocampal pathway. To be sure, it is possible that these patterns, should they exist, may apply only to the prefrontal cortex and not to the rest of the neocortex. Even in this case, however, the importance of the prefrontal cortex in learning and memory would make these regularities of general interest.

1.5 Preview

The rest of this thesis is divided into three parts. First we will focus on awake behavior and show that prefrontal cells that interact strongly with the hippocampus are preferentially recruited to form a neural correlate of associative learning, providing a possible electrophysiological signature of consolidation. Second, we will present a series of results concerning the existence of precise spike timing relationships between the hippocampus and prefrontal cortex during sleep. Finally, we will synthesize these findings and suggest directions for future work.

Chapter 2

Prefrontal-Hippocampal Interactions during Learning

2.1 Introduction

According to the predominant model of memory formation, long-term memories are gradually consolidated in the neocortex under the influence of the hippocampus[2]. Evidence for this model has come from a range of experimental approaches, from gene expression patterns[24,6,20] to lesion studies[25,26,18,19]. Precisely how hippocampal activity reorganizes cortical circuits to form stable memories, however, remains unknown. Understanding this process requires directly observing the neural activity corresponding to a specific neocortical memory formed under the influence of the hippocampus in a freely behaving animal. Two crucial steps towards this goal are first, to find cells in the cortex whose responses change systematically during training in a well-defined learning paradigm; and second, to show that these cortical cells are influenced by hippocampal activity.

Eyeblink conditioning is a form of associative learning that engages the hippocampus across a wide range of species and parameters[27]. When the conditioned and unconditioned stimuli do not overlap in time (trace eyeblink conditioning), learning the task requires an intact hippocampus in rabbits[26], rats[28], mice[29], and humans[30]. Furthermore, lesions to the medial prefrontal cortex in rats disrupt the recall of the conditioned eyeblink response if they are performed one month, but not one day,

post-learning[18]. Hence, trace eyeblink conditioning provides what may be the simplest model system for hippocampus-prefrontal mediated memory consolidation.

The prefrontal cortex has received particular attention in the study of cortico-hippocampal interactions because the hippocampus projects to this area via a plastic, monosynaptic pathway[17,13]. In addition, electrophysiological activity patterns in the mPFC are tightly coupled to the hippocampus during sleep[22,23,31] and awake behavior[21,32,33]. One form of coupling is through the theta rhythm, a prominent 4-10 Hz local field oscillation in the hippocampus[34]. Most principal cells in the hippocampus fire around a preferred phase of the theta oscillation[35], a phenomenon known as phase-locking. Modulation with respect to hippocampal theta has also been identified in mPFC[21], which suggests theta phase-locking as a natural measure of hippocampal influence over a particular cortical cell. Indeed, theta phase-locking has been shown to be a necessary condition for a prefrontal cell to have significant cross-correlations with cells in the CA1 subfield of the hippocampus[21].

The reliance of trace eyeblink conditioning on the prefrontal cortex and hippocampus, combined with the existence of hippocampus-modulated cells in the prefrontal cortex, raises two key questions. First, do prefrontal cells alter their firing to a conditioned stimulus (CS) in a manner that is consistent with a long-term cortical memory? Second, what role do phase-locked cells play in this process?

2.2 Results

2.2.1 Task-dependent firing in mPFC neurons

In this study, we recorded single-unit activity from the mPFC of three freely behaving rats over the entire course of trace eyeblink conditioning (89 sessions; $N = 851$ cells total). Figure 2.1 shows examples of the behavioral response to paired presentations of a tone (CS) and mild eyelid shock (US) early and late in learning. All three animals developed robust conditioned eyelid responses within 1,200–1,800 trials (Fig. 2.2).

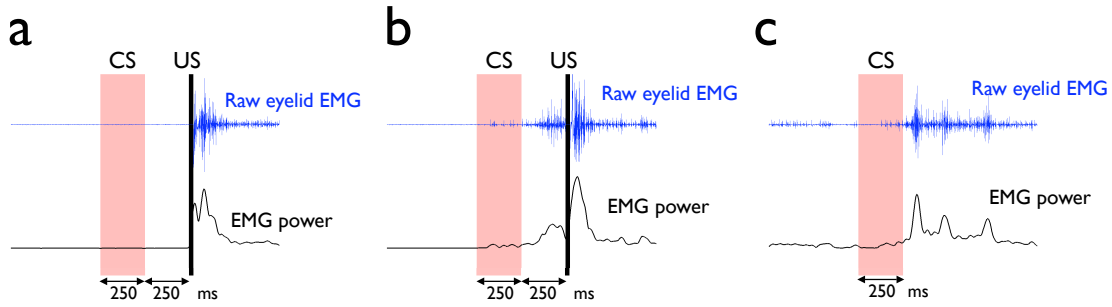

Figure 2.1: **Examples of eyelid during conditioning** . For each subpanel, the lower trace shows the electromyogram (EMG) activity from the *orbicularis oculi* muscle of the eyelid during the presentation of the conditioned stimulus (CS), a 250 ms tone, followed after 250 ms by the unconditioned stimulus (US), a 10 ms bipolar shock. The stimulus artifact is blacked out. The upper trace is the RMS power of the lower trace using a 50 ms Hann window. **(a)** Early in training, there is an unconditioned response (UR) to the US but no response to the CS; **(b)** Late in training, a robust conditioned response (CR). **(c)** A CS-only "probe" trial late in learning reveals a CR without the stimulus artifact.

Of the 851 mPFC cells recorded, 63 (7.4%) exhibited excitatory responses during the period between the onset of the CS and the US. These cells were designated CS-excited (CSe). Figure 2.3a,b shows examples of event-triggered rasters and peri-event spike histograms for two CSe cells. These CSe cells also illustrate the phenomenon of phase-locking to hippocampal theta oscillations (Fig. 2.3c,d).

Unlike the conditioned EMG response, which increased just before the onset of the US (Fig. 2.2a,b), the majority of CSe responses peaked around 125 ms after the CS onset, with a second mode around 125 ms after the CS offset (Fig. 2.4). Thus, the prefrontal CSe responses do not form a model of the eyelid EMG profile, as in the case of hippocampal responses during delay eyeblink conditioning[36], nor do they uniformly bridge the temporal gap between CS and US.

2.2.2 Evolution of prefrontal responses during learning

Given prefrontal cells whose firing was modulated by the learning task, we first asked whether these task-dependent responses varied over the course of learning. To this end, we assigned to each training session (and to each cell recorded in that session)

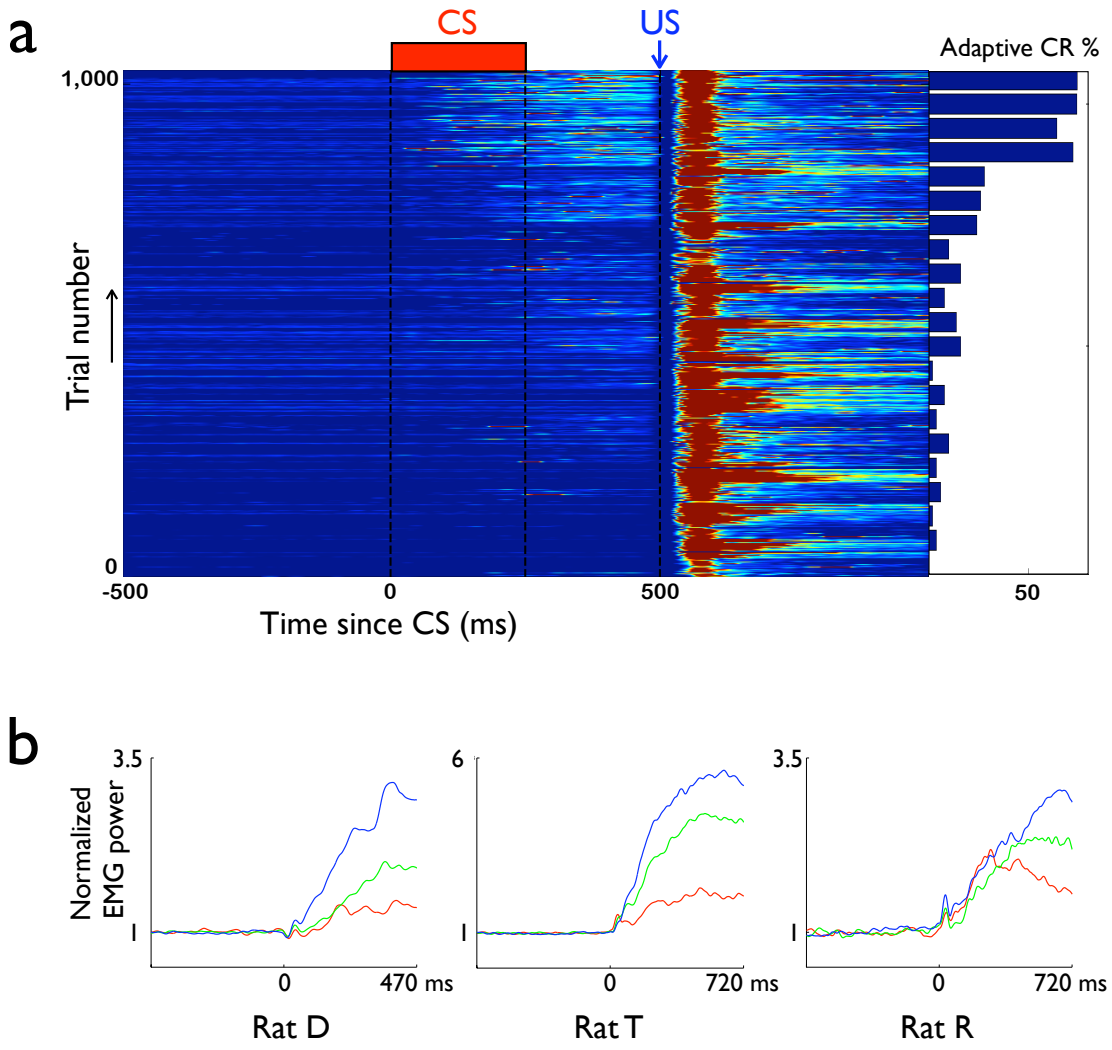

Figure 2.2: **Trace eyeblink conditioning learning curve** . **(a)** An example of the learning curve for one of the animals. Each row is one trial. The color intensity represents EMG power as calculated above. EMG power has been set to zero during the stimulus artifact. Sidebar: Fraction of trials in 50-trial blocks where EMG power exceeds the CR threshold (see Methods). **(b)** Evolution of eyelid responses with training. All training trials were ordered and divided into three equal subsets. For each rat, the red, green, and blue curves show the mean EMG power in the first, second, and third subsets, corresponding to early, middle, and late trials. The CS onset is at time zero and the US onset is at 500 ms for rat D and 750 ms for rats T,R.

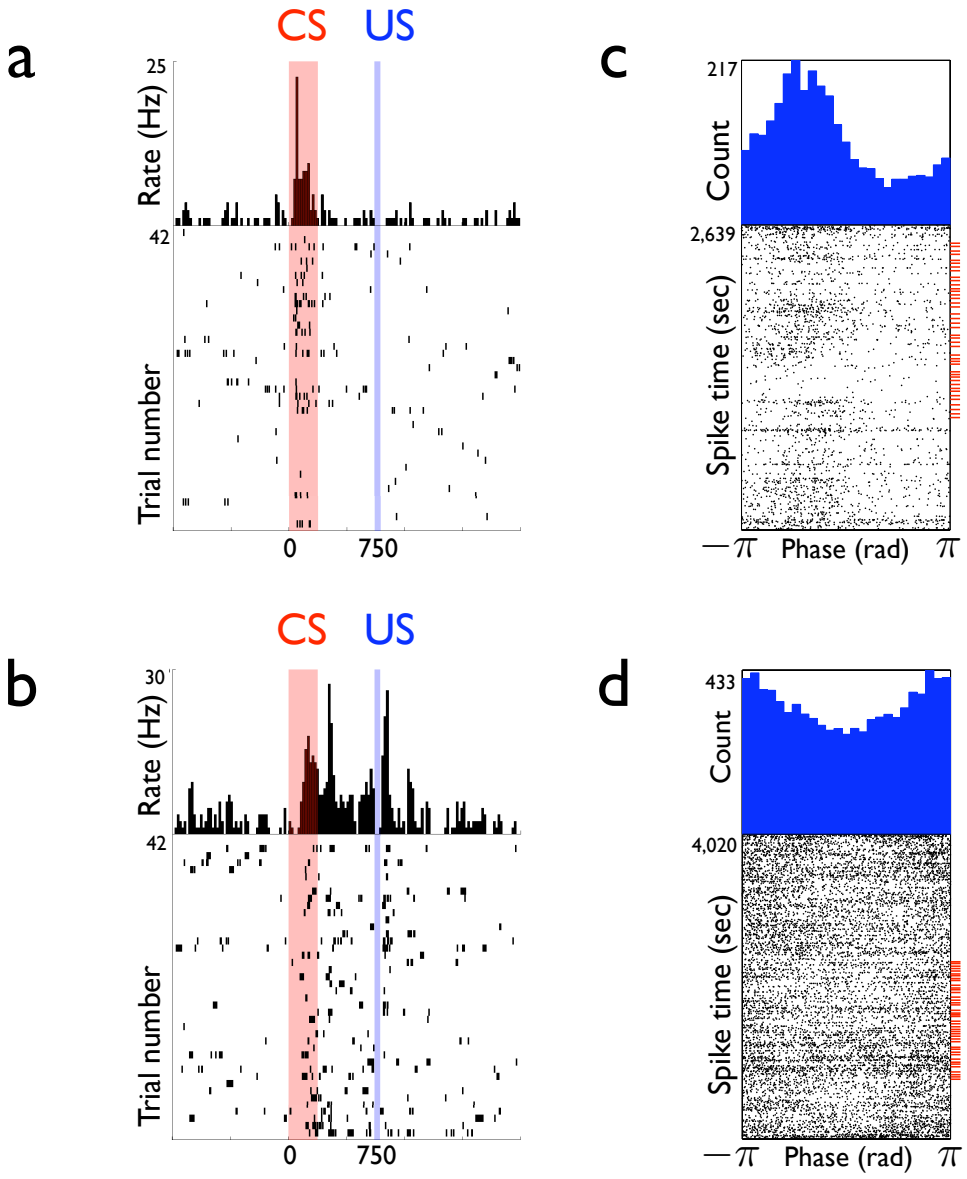

Figure 2.3: **Examples of CS-excited cells in the mPFC. (a),(b)** The lower panel show rasters of each trial during a training session, aligned to the onset of the CS. The upper panel shows the peri-event spike histogram with 20 ms bins. **(c),(d)** Theta phase histograms for the cells shown in **a,b** respectively. The bottom rasters show all spikes during the training session. Spike times (y-axis) are plotted against the phase of the hippocampal theta rhythm (x-axis). Red hashmarks on the y-axis indicate CS-onsets.

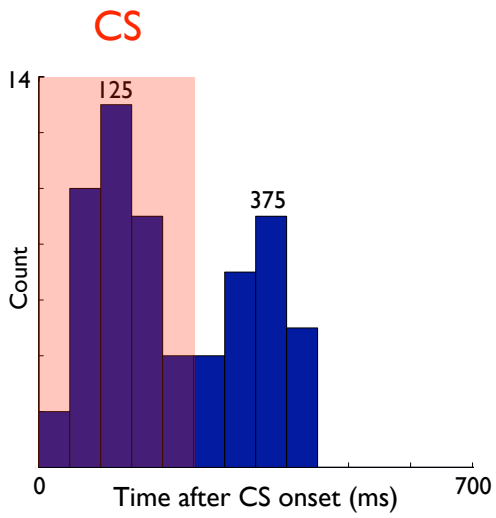

Figure 2.4: **Distribution of peak firing times for CS-excited cells.** Peak firing times were computed by convolving the peri-event spike histogram of each cell with a Gaussian kernel with $\sigma = 6$ ms and locating the maximum. Note that the distribution is bimodal, with one mode around 125 ms and a second around 375 ms.

a rank, defined as the fraction of training, measured in trials, completed up to and including that session. We then measured the mean firing rates of each CSe cell during the period between the CS and US onsets, and plotted these as a function of rank. As shown in Figure 2.5a, the firing rate of CSe cells tended to increase with training, and this relationship was highly significant ($R^2 = 0.26, p < 10^{-4}$; linear regression). In order to show the evolution of CSe responses, we divided training into three equal intervals and computed the average CSe response during those intervals (Fig. 2.5b). The evolution of CSe responses is also revealed by plotting all of the CSe rasters in series, ordered by rank (Fig. 2.7).

To verify that growing CSe responses were not merely due to a general increase in neural or behavioral excitability, we performed two sets of control analyses. First, as shown in Figure 2.6, we plotted the mean firing rates of CSe cells outside CS/US presentations throughout the training session as a function of rank and found no significant increase ($R^2 = .06$, $p > .05$; linear regression). Similarly, we found no relationship between firing rate and rank for non-CSe cells ($R^2 = .002$, $p > 0.2$; linear

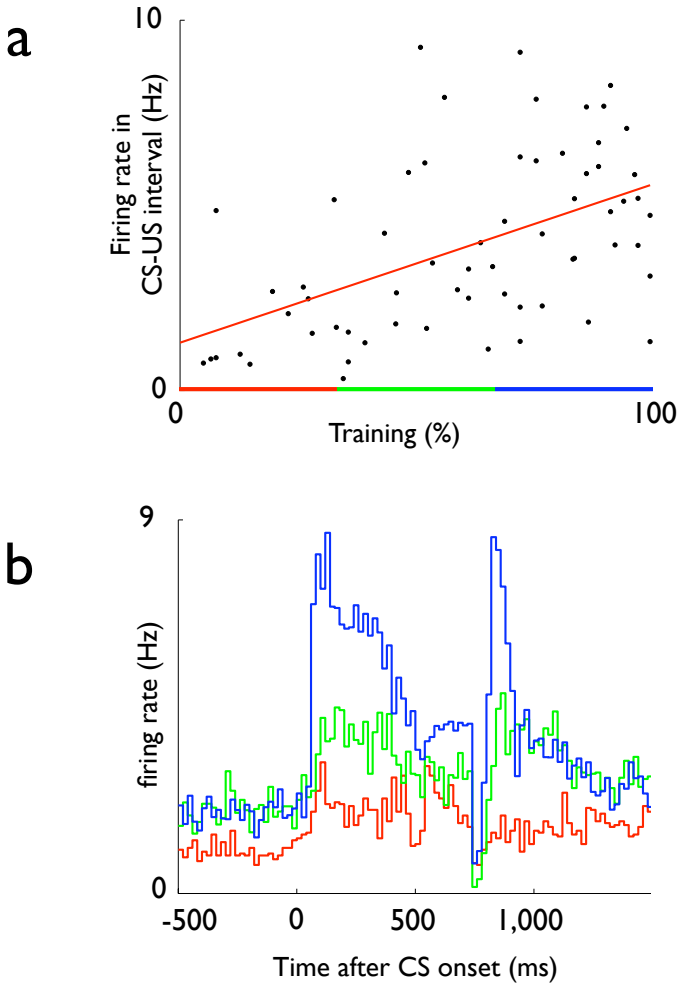

Figure 2.5: **Evolution of firing rates in response to the CS as a function of training. (a)** Mean firing rates between CS and US onsets versus training. Each point represents one CSe cell in one training session (63 points total). Firing rates increase significantly as a function of learning (linear regression, $p < 10^{-4}$). **(b)** Mean rasters during the early (red), middle (green), and late (blue) thirds of training, indicated on the x-axis of **(a)**. Note the increase of CS responses with learning. Because of possible stimulus artifacts, spikes between $t_{US} - 10$ and $t_{US}+40$ were removed, where t_{US} is the US time in miliseconds, leading to dips at $t = 500$ and $t = 750$.

regression). Second, we computed the average acceleration of the animal around the CS onset as a function of rank and found no relationship between the two ($R^2 = .001$, $p > 0.8$; linear regression).

2.2.3 Hippocampal modulation of CS-excited cells

Having identified a set of prefrontal cells whose CS-response increased with learning, we turned to our second question, namely, whether cortical cells that encode learned associations have a special relationship to the hippocampus. As a metric of hippocampal modulation, we used the degree of phase-locking to the hippocampal theta rhythm[21]. In general, we found that CSe cells tended to be significantly more phase-locked than non-CSe cells. In particular, the distribution of the Rayleigh Z-statistic, a measure of circular unimodality, for CSe cells was significantly higher (more phase-locked) than non-CSe cells ($p < 2 \times 10^{-6}$, one-sided Kolmogorov-Smirnov test). Figure 2.8a–d illustrates the difference between these distributions. Specifically, 39% of non-CSe cells were significantly phase-locked, compared to 68% of CSe cells.

Finally, as shown in Figure 2.8e, we found that the fractions of CS-excited or phase-locked cells did not vary significantly with training. The fraction of CSe cells that were phase locked did decrease with training, although this decrease was not statistically significant.

2.3 Discussion

We have identified in freely behaving animals a class of prefrontal cells that encode a fundamental building block of learning: the association between a CS and US. Moreover, we have found that the strength of the encoding, as measured by firing rates, systematically increases with the strength of the association, as measured by the behavioral output. Finally, we have shown that the prefrontal cells that encode the CS are predominantly phase-locked to the hippocampal theta rhythm.

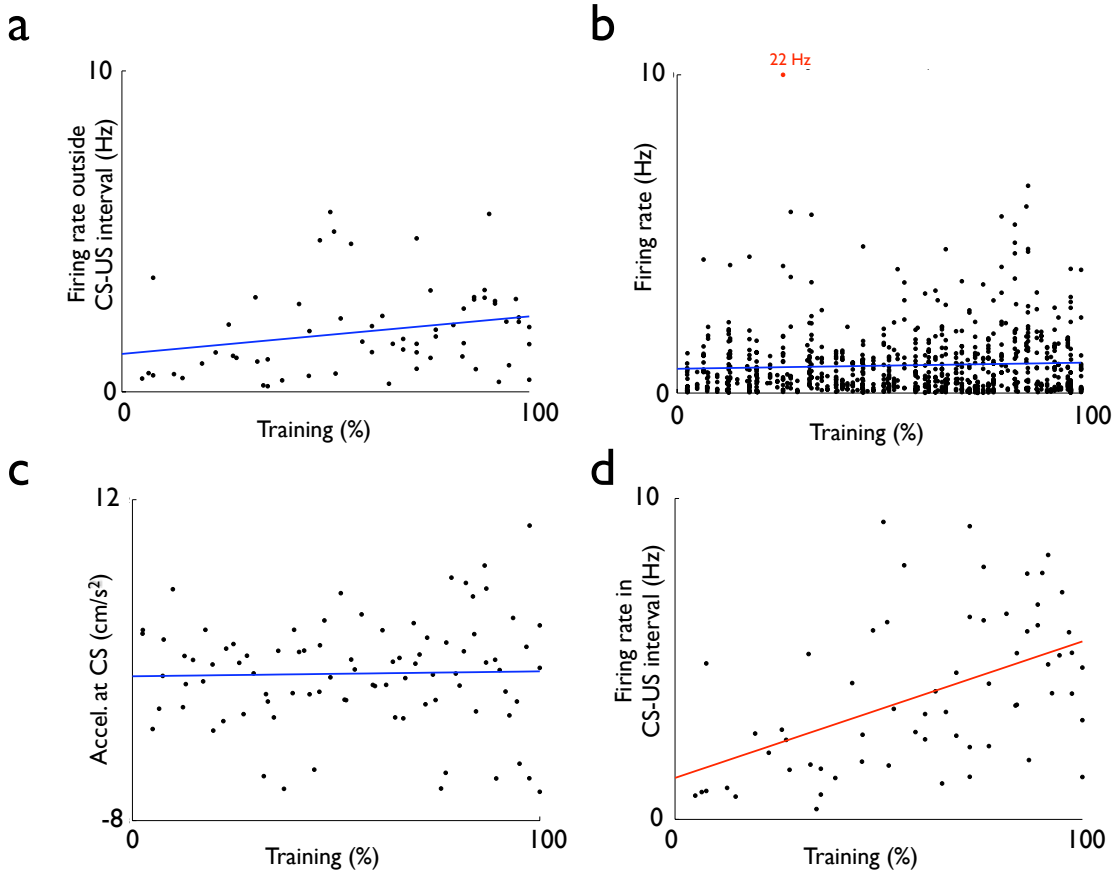

Figure 2.6: **Verification that evolution of CSe firing is not due to non-specific excitability. (a)** Firing rates of CSe cells outside CS-US intervals as a function of training. The small positive trend is not statistically significant ($p > 0.05$). **(b)** Firing rates of non-CSe cells during the entire session as a function of training. There is no significant relationship between general excitability and training. Note that one data point (red) has been displaced in order to keep the same y-axis as A. **(c)** Average acceleration in a 400 ms window centered at the CS-onset time for each dataset as a function of training. Note that there is an overall tendency for the animal to increase its speed after the CS, but no trend. **(d)** Mean firing rate of CSe cells as a function of training. Same as Figure 2.7a, included here for comparison with **a** and **b**.

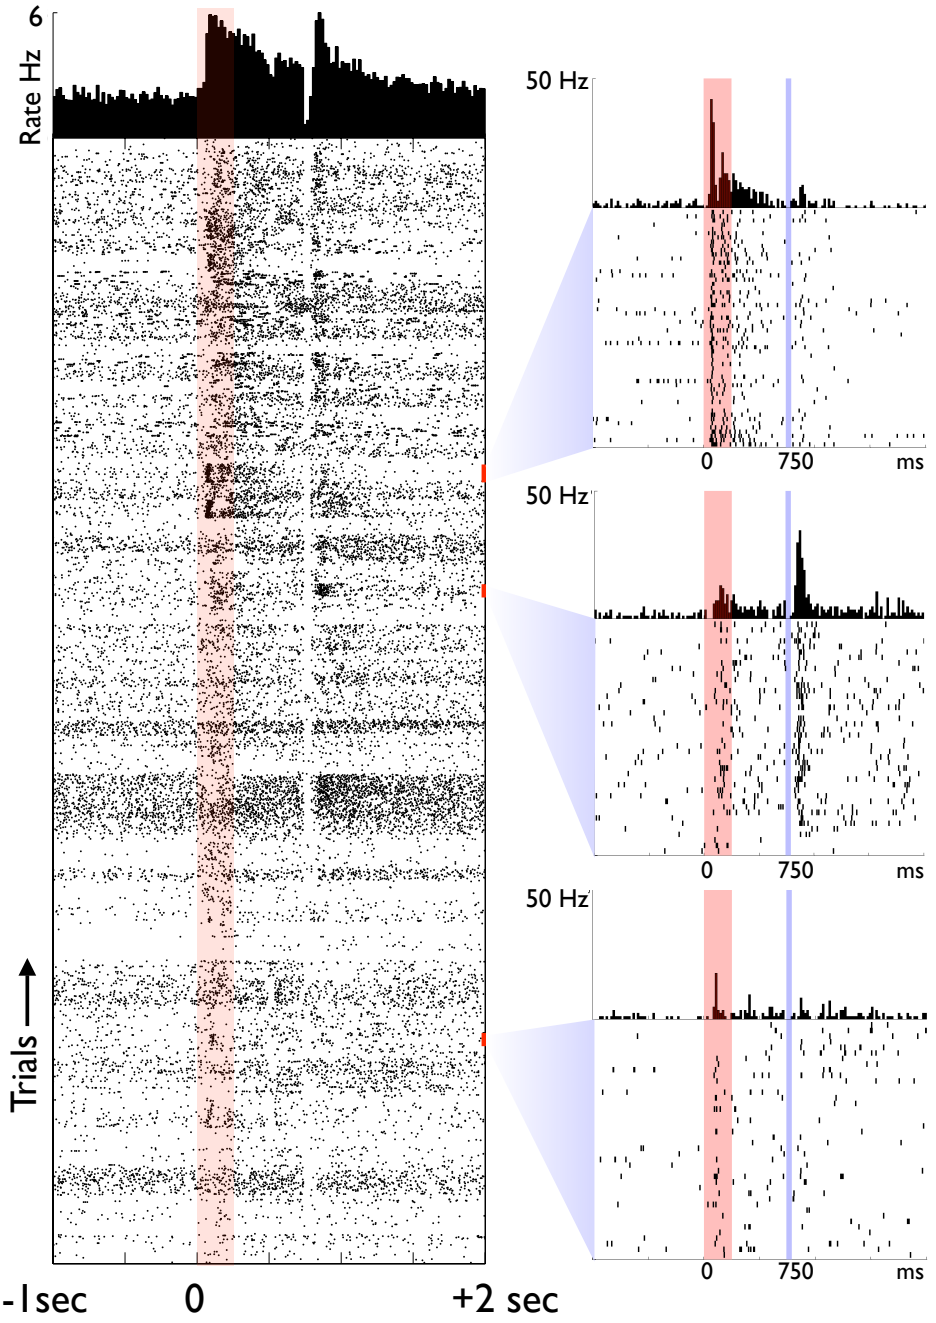

Figure 2.7: **Evolution of CS-triggered neural responses with learning.** Event-triggered rasters from all CS-excited cells from all 3 animals are ordered by their position in the training sequence. Rasters and peri-event spike histograms of three selected cells from the learning sequence are shown at right. As in Figure 2.5, spikes between $t_{US} - 10$ and $t_{US} + 40$ were removed.

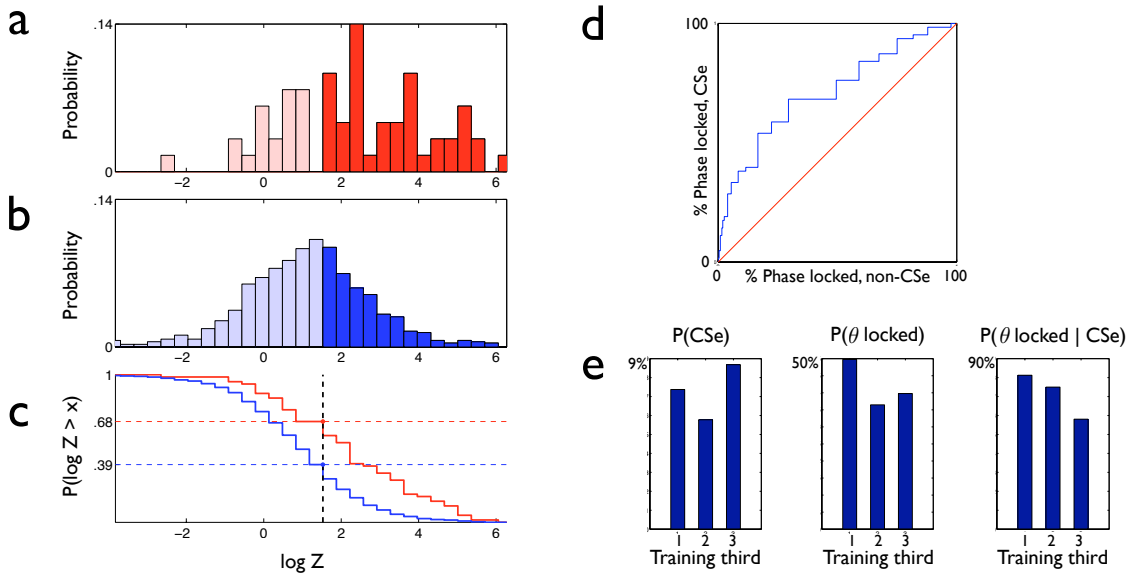

Figure 2.8: **CS-excited cells in the mPFC are predominantly phase-locked to hippocampal theta oscillations. (a),(b)** Distribution of phase-locking of CSe and non-CSe cells, as measured by $\log Z$, where Z is the Rayleigh Z-statistic. For 50 or more spikes, $p = e^{-Z}$, where p is the significance level for rejecting the hypothesis that spikes are uniformly distributed with respect to theta phase[37]. Vertical breaks at $\log Z = 1.52$ correspond to $p = 0.01$. The distribution of phase-locking for CSe cells, shown in **a**, lies to the right (more phase locked) than that for non-CSe cells, shown in **b**. **(c)** Cumulative densities for the distributions in **a,b** showing the fraction of cells phase-locked at a significance level of $p < e^{-Z}$. Note that 68% of CSe cells are phase-locked at the $p < 0.01$ level, while for non-CSe cells the fraction is only 39%. **(d)** The fraction of phase-locked CSe cells versus non-CSe cells for a range of p values. Note that a higher fraction of CSe cells are phase-locked than non-CSe cells for any p value chosen to define phase-locking. **(e)** From left to right: for each third of training, fractions of all prefrontal cells that are CSe; phase-locked; and the fraction of CSe cells that are phase-locked.

Previous work with trace eyeblink conditioning has found CS-excited cells in the anterior cingulate cortex of restrained rabbits[38]. In contrast to the present study, these excitatory responses peaked in the first training session and decayed to baseline levels as the animal became fully trained. Consistent with this response profile, pre-training lesions of the same brain area in rabbits have been found to impair learning of trace eyeblink conditioning[39]. The CS-excited responses reported here were recorded in the prelimbic and infralimbic regions of the mPFC, which receive monosynaptic projections from the hippocampus[13], as opposed to the anterior cingulate, which does not[40]. Taken together, these results suggest a functional specialization within the mPFC, where anterior cingulate neurons signal the novelty and salience of the conditioned stimulus while prelimbic and infralimbic circuits encode the association itself. How this specialization might depend on differential innervation from the hippocampus remains an open question.

Two previous studies have analyzed the behavioral modulation of theta phase-locking in prefrontal cortex during linear track traversals[32] and a spatial working memory task[33]. Neither study recorded from prefrontal cells during the learning process, and as a result the relationship between phase-locking and learning was not previously explored.

In the awake state, theta oscillations are most prominent during exploratory behaviors such as locomotion and generally absent during immobility[41]. The phase-locking statistics in our study were computed during periods of elevated theta power as the animal moved freely within its environment during each training session. As a result, they are likely to reflect an intrinsic property of a prefrontal cell, perhaps due to actual or effective connectivity between that cell and the hippocampus, rather than a purely task-driven property. This observation suggests that phase-locked cells in prefrontal cortex, tuned to theta-modulated hippocampal input, are recruited to form the initial representation of the CS (Fig. 2.9). This idea is supported by the observation that presenting CS-US pairs during periods of ongoing theta oscillations accelerates the learning of trace eyeblink conditioning[42], since, by the definition of phase-locking,

theta waves would increase temporal precision specifically across phase-locked circuits.

By showing a link between hippocampal phase-locking and CS representations, we have identified a hippocampus-dependent task (trace eyeblink conditioning) as well as an electrophysiological tag (theta phase-locking) that are well suited for probing the dynamics of the consolidation process, both during awake behavior as well as preceding and subsequent sleep. Future studies could include comparing the hippocampal dependence, via phase-locking, of recent versus remote cortical memories over time spans relevant to consolidation as identified by lesion studies, typically one or more months. Conversely, tracing the co-evolution of CS-responsiveness and phase-locking after asymptotic learning could provide animal-specific milestones during the consolidation process by which to guide lesions or other manipulations.

2.4 Methods

2.4.1 Electrophysiological recordings.

Electrophysiological signals were acquired using tetrode recordings[43]. Three male Long-Evans rats from 3-5 months old (weight = 350-450g) were implanted with a custom-built microdrive array allowing the independent adjustment of 24 individual tetrodes. Twelve tetrodes targeted the prelimbic and infralimbic regions of the mPFC (AP: 1.5-3.5mm from bregma; ML: 1-1.75mm, angled at 15 degrees from the saggital plane) and twelve tetrodes targeted the dorsal CA1 subfield of the hippocampus (AP: -3.75 to -4.75mm from bregma; ML: 1.5-3.5mm). Individual tetrodes were gradually lowered to their targets over several days and further micro-adjusted to optimize yield and stability. Each tetrode signal was buffered by a unity-gain headstage preamplifier and further differentially amplified with a gain of 2000. The broadband amplified signals were digitally acquired at 25kHz as 24-bit samples (National Instruments PXI-4472) and stored to disk using custom acquisition software that we have developed. In addition, four 0.005" stainless steel wires were implanted in the animal's contralateral

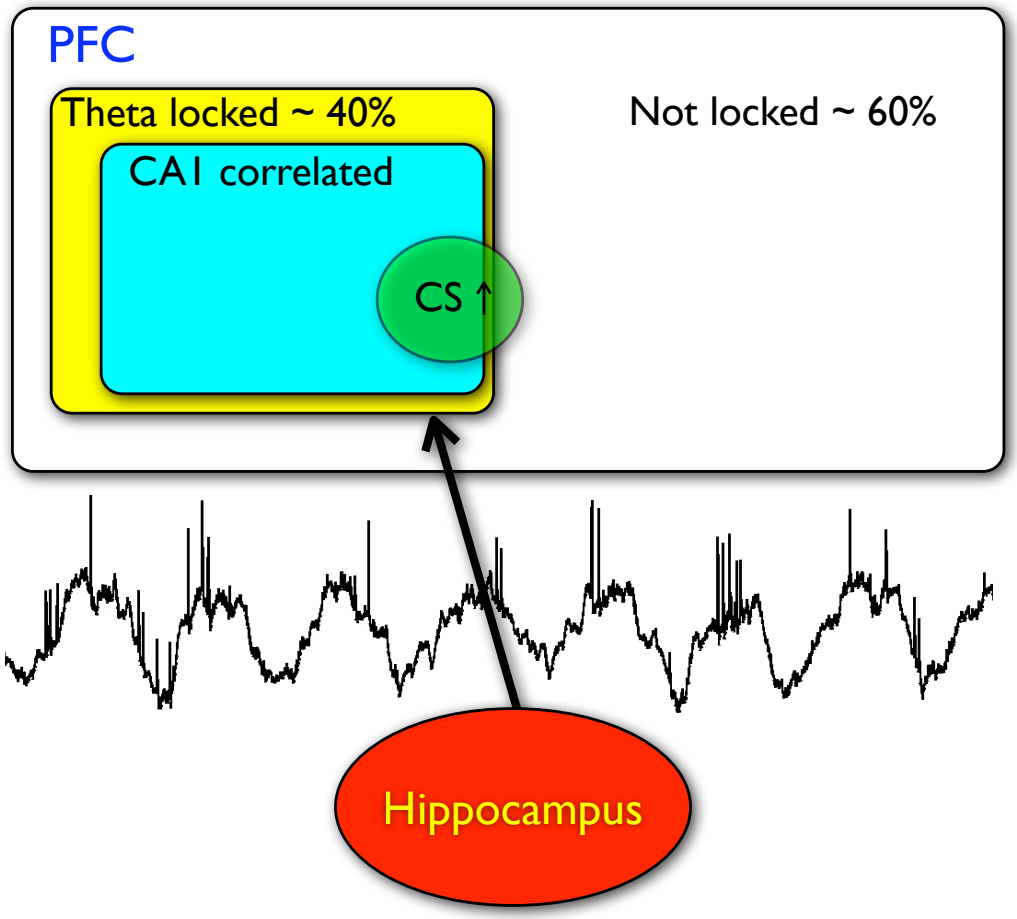

Figure 2.9: **Classification of prefrontal cells engaged in learning a hippocampus-dependent task.** Approximately 40% of prefrontal cells are phase-locked to hippocampal theta oscillations (yellow), and the subset of CA1-correlated cells in prefrontal cortex is entirely phase-locked (blue)[21]. It is reasonable to believe that the converse is also true, i.e., that phase-locked prefrontal cells are CA1-correlated, but difficult to verify because of finite sampling of hippocampal cells. In this work we show that cells that form the initial encoding of the CS-US association are recruited from the theta phase-locked subset (green), which have a special relationship with the hippocampus.

upper eyelid. The caudal pair were used for bipolar stimulation (see Behavioral Training below); the rostral pair, used for eyelid EMG measurements, were buffered by a unity-gain headstage preamplifier, highpass filtered at 100Hz, differentially amplified with a gain of 2000, and fed to the same acquisition system as the neural signals. Three light-emitting diodes were fixed to the top of the microdrive array to allow tracking of the animal's position from video recordings. Each frame of video, all tone and stimulation events, and the acquisition system sample clock were timestamped by a 10MHz clock to synchronize position, behavioral, and neuronal data. All animal procedures were done in accordance with NIH guidelines and with approval of the Caltech Institutional Animal Care and Use Committee.

2.4.2 Spike and local field analysis

Spikes and LFP traces were obtained by digitally filtering the broadband signal. For spikes, a bandpass filter was designed using the Parks-McClellan algorithm with transition bands of 500-600Hz and 6000-6100Hz and maximal ripple of 10^{-5} in the stopband and 10^{-3} in the passband. LFPs were computed by downsampling the broadband signal by a factor of 12 in three stages (2,2,3); each stage used a 500-tap FIR linear-phase lowpass filter designed using the window method. Spikes were clustered into single units on the basis of their amplitudes recorded on each of the four tetrode channels.

2.4.3 Behavioral training

Animals were trained in either a 50 cm by 70 cm box or on a 170 cm by 10 cm track, both of which were in the same room as the animal's sleep box. The walls and ceiling were covered with anechoic foam, and the room was electromagnetically shielded and acoustically sealed. The conditioned stimulus was a 250ms, 5kHz tone delivered through a Fostex FX-120 speaker at 80 dB SPL above the environment. The unconditioned stimulus was a 10 ms bipolar current pulse (5 ms for each cycle) delivered though an isolated current source (WPI A360). This relatively short stim-

ulation time was used to minimize the temporal extent of the stimulation artifact[44]. The magnitude of the stimulation current was set during a habituation session before training at a level that produced a reliable unconditioned response, between 2 and 5 mA, and fixed throughout conditioning. For two animals (Rats T,R), the US onset was 500 ms after the CS offset; for one animal (Rat D), the gap was 250 ms. Both of these intervals have been shown to be hippocampus-dependent in rats[28,45]. Intertrial intervals were uniformly random between 20 and 40 seconds. Every fifth trial, the animal received a "probe" trial consisting of a tone alone. Tones for probe trials were either 5kHz (training frequency) or 11kHz, 7.42kHz, 3.37kHz, or 2.27kHz (testing frequencies), counterbalanced. Animals performed two sessions per day of 50 or 100 trials, separated by at least four hours of sleep. Each animal's performance reached a behavioral plateau where CRs were performed at an average rate of at least 60% over 250 consecutive trials, which was considered the stopping criterion for this analysis. Number of trials to reach criterion were 1642, 1693, and 1281 for rats T, R, and D, respectively. CRs were defined as trials where the eyelid EMG power (see EMG analysis below) exceeded threshold for an interval longer than 10 ms between T-360 and T-10 ms, where T is the onset time for the unconditioned stimulus. The threshold was defined as the mean plus one standard deviation of the EMG power calculated over the 2 seconds immediately preceding the CS onset.

2.4.4 EMG analysis

Differential EMG signals were downsampled by a factor of 12. We then computed EMG power as the RMS value of the downsampled EMG over a sliding 50 ms Hann window.

2.4.5 Analysis of unit responses

All trials in a given session at the training frequency were aligned at the time of the tone presentation. To eliminate the possibility of stimulus artifacts, spikes between

US-onset - 10 ms and US-onset + 40 ms were dropped. Remaining spikes were counted in 20 ms bins, and the counts summed across trials to yield a peri-event spike histogram. To assess responsiveness to the CS, the bins between the CS-onset and US-onset - 10 ms were compared to the same number of bins immediately preceding the CS-onset using an unpaired t-test. To assess responsiveness to the US, the bins between US-onset + 40 ms and US-onset + 290 ms were compared to the same number of bins immediately preceding the CS-onset using an unpaired t-test. In both cases, the one-sided p-value of the t-test was used a measure of CS excitedness. Unless otherwise stated, we used $p < 0.01$ as a definition of excited or inhibited.

2.4.6 Analysis of phase-locking

We analyzed phase preferences of prefrontal units using methods described here[21]. We selected a single hippocampal tetrode per animal on the basis of the most robust theta oscillations; all theta phases for were calculated relative to that tetrode throughout all sessions. To measure phase-locking properties independently of CS/US responses, spikes that fell within 4 seconds of a CS-onset were excluded from phase-locking calculations. Including all spikes did not materially change any of the results and is therefore not shown.

Chapter 3

Prefrontal-Hippocampal Interactions during Sleep

3.1 Introduction

Many lines of evidence have shown that the hippocampus is critical for the formation of long-term memories and that this hippocampal involvement is time-limited [2,46,26]. The predominant conjecture is that memories are gradually established across neocortical circuits under the influence of the hippocampus[47,48]. This circuit reorganization is believed to result from coordinated activity between and within the hippocampus and the neocortex not only during awake behavior, but also during sleep[49,22,50,51,52]. Consistent with this conjecture, cortical and hippocampal networks remain highly active and plastic during sleep. One of the most striking features of mammalian sleep is the existence of discrete stages—slow-wave sleep (SWS) and rapid eye movement (REM) sleep—with different electrical and biochemical profiles. In particular, SWS and REM sleep differ drastically in the level of synchronous firing in the hippocampus (Fig. 3.1). Given the importance of synchrony and spike timing in synaptic plasticity, and given the putative role of sleep in learning and memory, a key open question is whether there exist consistent spike timing relationships across cortico-hippocampal circuits during sleep, and whether these differ in SWS versus REM sleep.

The hippocampal-prefrontal circuit is of particular interest given its importance in spatial and associative learning[53,18], as well as the known interactions between the

two areas during awake behavior[21,32,33] and sleep[22,23]. Previous work has shown that hippocampal and prefrontal multi-unit activity are significantly correlated during SWS, with the hippocampus leading the prefrontal cortex[22,23]. However, key questions remain open concerning the interactions between these areas: How common are directional interactions across prefrontal-hippocampal cell pairs? Is there diversity in their directionality, time lag, and strength? How are these interactions structured relative to prominent electrophysiological events in the sleeping brain, such as hippocampal ripples and neocortical spindles? Do they differ during SWS and REM sleep? The present study addresses these questions.

3.2 Results

3.2.1 Directionality in prefrontal-hippocampal spike timing

In order to assess hippocampal-prefrontal timing relationships at the single cell-pair level during all stages of sleep, we used chronic multi-tetrode recordings to monitor the simultaneous activity of CA1 and medial prefrontal (mPFC) cells of freely behaving rats during long intervals of natural sleep (Fig. 3.1). We computed cross-covariances between all pairs of simultaneously recorded prefrontal and hippocampal single units (219 CA1, 76 mPFC cells). We restricted our analysis to putative pyramidal cells in the hippocampus (183 CA1 cells), using a mean firing rate criterion of less than 1 Hz, and we considered only pairs where the firing rates of both cells exceeded 0.05 Hz in SWS and REM sleep (2779 total mPFC-CA1 pairs).

Figure 3.2 shows an example of the cross-covariances computed between one mPFC cell and all of the simultaneously recorded CA1 cells in one dataset during SWS and REM sleep. Each row in Figure 3.2(ii) represents the cross-covariance between the given mPFC cell and one CA1 cell as a function of lags ranging from -500 to 500 ms, where positive lags signify that prefrontal activity follows hippocampal activity.

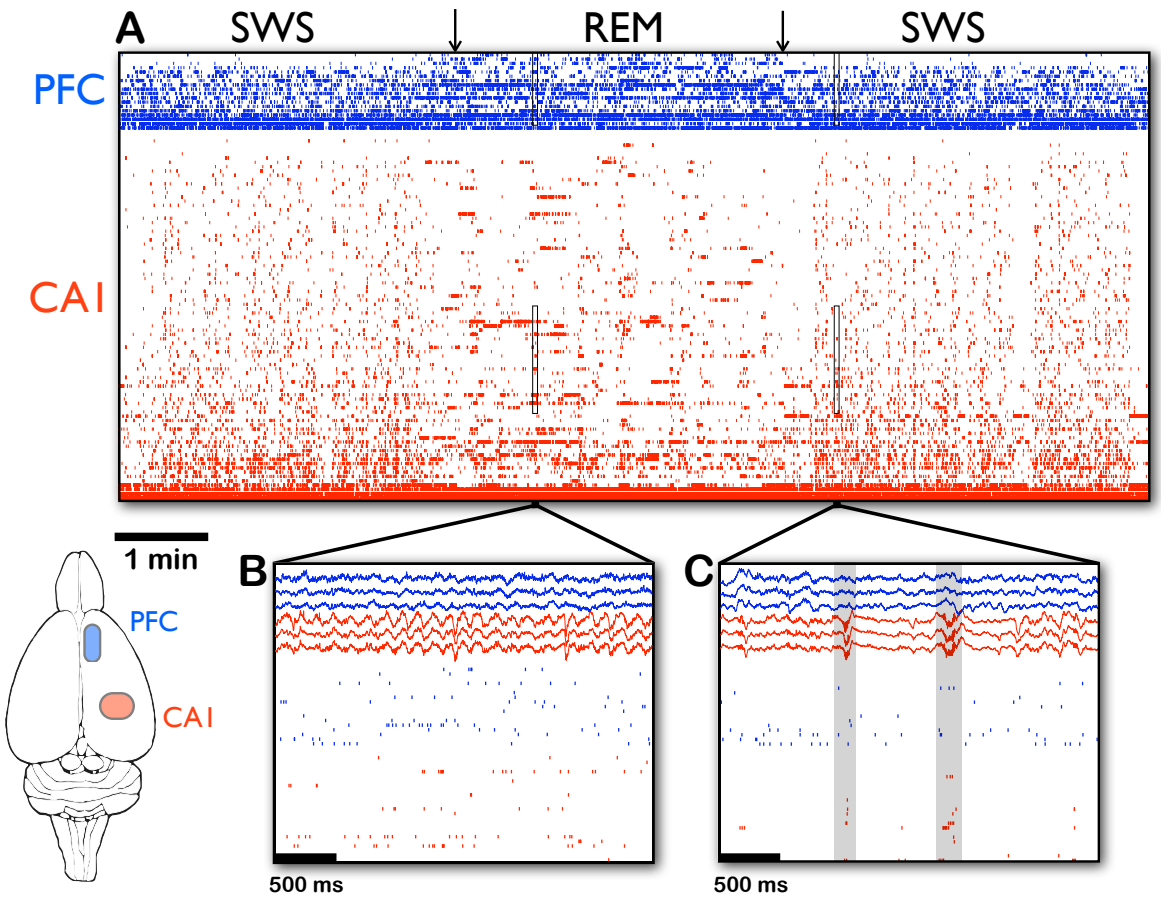

Figure 3.1: **Hippocampal and prefrontal spiking activity during sleep. A:** Spikes from 86 CA1 (red) and 18 mPFC (blue) simultaneously recorded units during approximately 11 minutes of sleep. Note the abrupt transition in CA1 between the vertical stripes of synchronous bursting during SWS to the horizontal stripes of theta-modulated firing in REM. **B,C:** Close-up views of REM and SWS, respectively, of the subsets of cells marked by rectangles in (A), as well as simultaneously recorded local field potentials. Note the prominent theta oscillations in the hippocampal (red) traces during REM (B) and the sharp-wave/ripple events with population bursts in SWS (C), marked by rectangles.

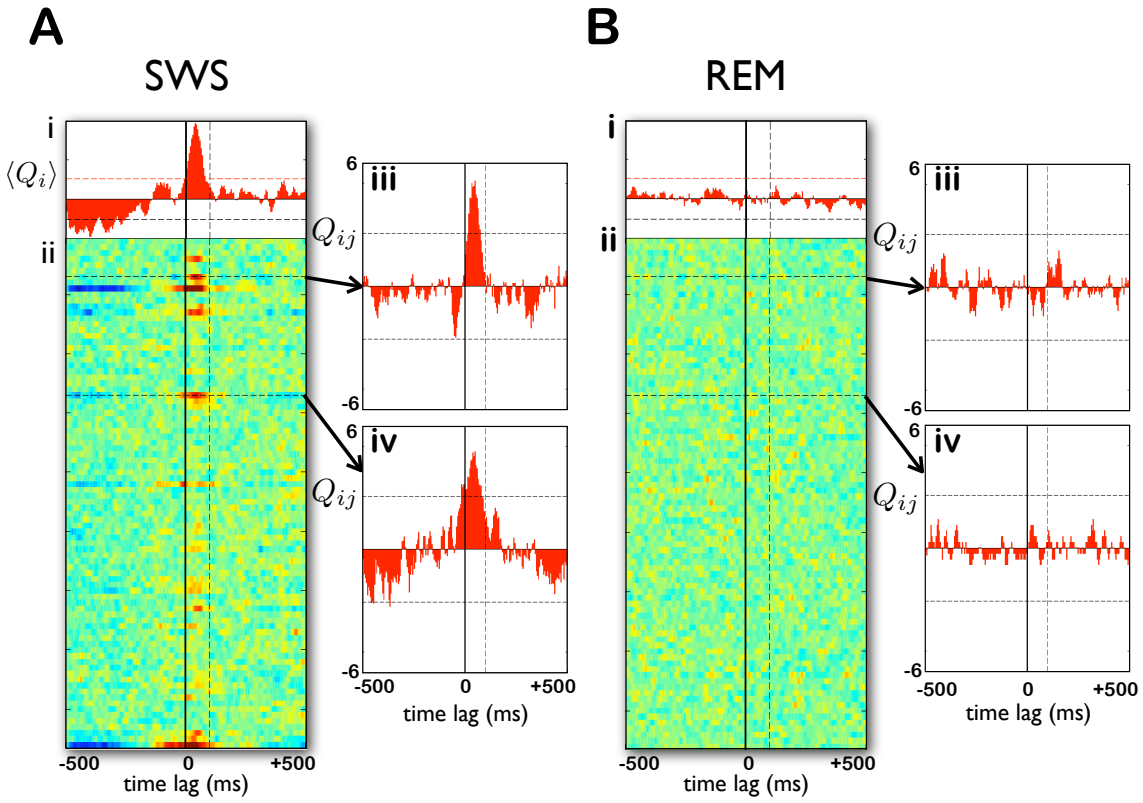

Figure 3.2: **State-dependent cross-covariances between single cells in the hippocampus and prefrontal cortex. A: (i)** The standardized mean cross-covariance between a single mPFC cell and all of the simultaneously recorded CA1 cells during SWS. Horizontal dashed lines indicate significance at the $p = 0.01$ level. **(ii)** Each row shows the standardized cross-covariance between the mPFC cell and a single CA1 cell. Note that several rows show high cross-covariances between 0 and 100ms (CA1 leads mPFC). **(iii,iv)** The rows marked by the horizontal dashed lines in (ii), showing two examples of prefrontal-hippocampal cell pairs with significant cross-covariance. **B: (i-iv)** The same cells and dataset as in (A) but during REM. Note the absence of significant cross-covariances in REM. The calculation of standardized and mean standardized cross-covariances is described in this reference[21] (see also Methods).

We note three features from this example. First, several CA1 cells show significant positive cross-covariances during SWS (Fig. 3.2A(ii)), indicating that they tended to exhibit consistent spike timing relative to this prefrontal cell. Second, the significant cross-covariances have peaks at positive time lags, between 0-100 ms. This shows that this prefrontal cell tended to fire 0-100 ms after these hippocampal cells. Third, the same cell pairs with significant correlations in SWS are uncorrelated in REM sleep (Fig. 3.2B(ii)).

To test the generality of these observations, we computed all cross-covariances between mPFC cells and simultaneously recorded CA1 cells in both SWS and REM sleep across all datasets. First, we found that 11% (304 out of 2779) of prefrontal-hippocampal cell pairs were significantly correlated in SWS (false discovery rate of 1%; see Methods.) Second, we observed that, for these correlated cell pairs, the distribution of peak lags deviated from uniformity in several key respects. Prefrontal cells tended to fire after hippocampal cells for 70% of correlated cell pairs, a significant directional bias ($p < 10^{-11}$, binomial test; Figure 3.3A(iii)). More specifically, prefrontal firing followed hippocampal firing by an average of 36 ms ($n = 304$; s.e.= 12 ms). The concentration of peak lags in the range of 0-100 ms (39% of pairs) was also highly significant ($p < 10^{-20}$; binomial test).

Both prefrontal and hippocampal neurons may show spike-timing relationships with themselves in the form of significant auto-covariances. The structure of these auto-covariances could, in principle, color the cross-covariances between these two structures. We therefore verified that the observed structure in the cross-covariances between correlated prefrontal and hippocampal units is mainly due to genuine prefrontal-hippocampal timing relationships and not to the auto-covariances of the constituent brain areas (Figure 3.4).

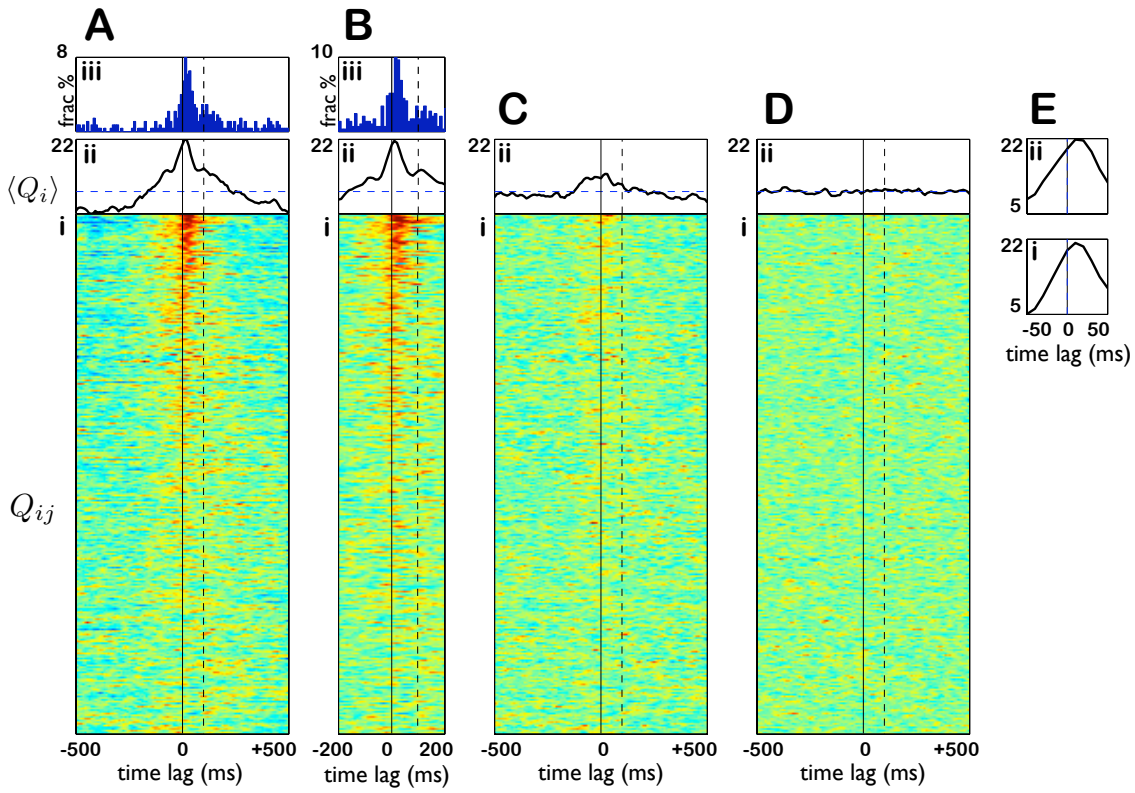

Figure 3.3: **Population analysis of state-dependent prefrontal-hippocampal interactions. A: (i)** Each row shows the standardized cross-covariance between a single mPFC and CA1 cell during SWS. Cell pairs are sorted from top to bottom by the significance of their cross-covariance (p values increasing from top to bottom). Only the top 304 rows out of 2779 pairs are shown, corresponding to the pairs deemed significant using a false discovery rate of $q = 0.01$. **(ii)** The standardized mean cross-covariance of the cell pairs in (i). **(iii)** Distribution of time lags of peak cross-covariance for significantly covarying pairs during SWS. Note the concentration between 0-100 ms. **B: (i-iii)** The same calculations as in A(i-iii) restricted to spikes that occur during sharp-wave/ripple (SWR) events. Cell pairs are shown in the same order as in A. Note the similarity of pairwise correlations (i), standardized mean cross-covariance (ii), and distribution of peak lags (iii). **C: (i,ii)** The same calculations as A(i,ii) restricted to spikes that occur during SWS outside of SWR events. Cell pairs are shown in the same order as A. Note the near-absence of significantly correlated pairs in (i) and substantial diminution in mean cross-covariance in (ii). **D:** The same calculations as A(i,ii) restricted to REM sleep. Note the absence of significant cross-covariance, either for individual cell pairs (i) or in the mean (ii). **E: (i)** and **(ii)** are zoomed in views of A(ii) and B(ii) respectively.

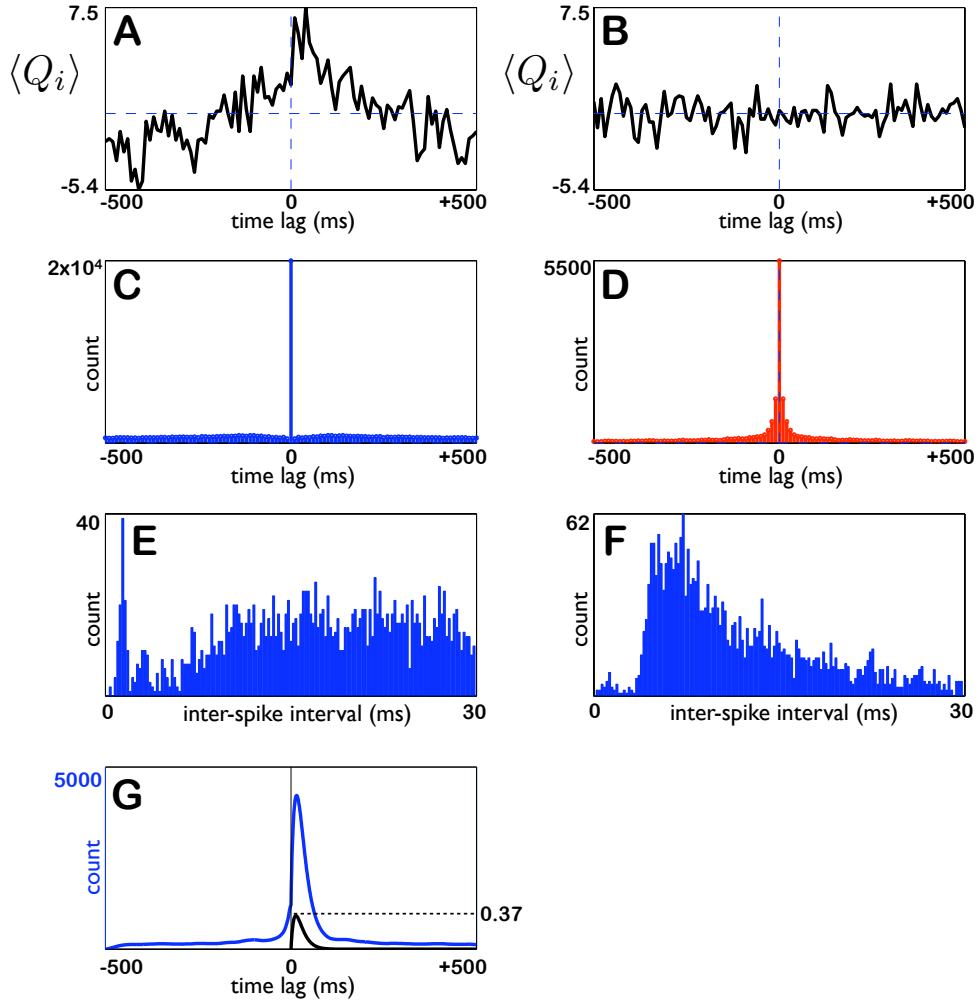

Figure 3.4: **Absence of relationship between cortico-hippocampal cross-covariances and respective auto-covariances.** Each panel uses spike data from the same representative correlated cell pair. **A:** The standardized cross-covariance, showing significant directional interactions. **B:** The same as A but where the prefrontal spikes have been reversed in time. This operation preserves the auto-covariances of CA1 and mPFC firing exactly but abolishes all structure in the cross-covariance, confirming that this is mainly due to genuine prefrontal-hippocampal timing relationships. **C,D:** The auto-correlations of the mPFC and CA1 units, respectively. **E,F:** Inter-spike intervals (ISI) of the mPFC and CA1 units, respectively. The mode of the mPFC unit's ISI illustrates its intrinsic burstiness, while that of the CA1 unit around 5 ms matches the period of ripple oscillations, to which it is highly phase-locked[41]. **G:** Blue trace: CA1 spike-triggered average of a simulated excitatory post-synaptic potential (EPSP) kernel[54] triggered by CA1 spikes. The simulated EPSP (black trace) has the form $\alpha t e^{-\alpha t}$, with $\alpha = 0.075$ ms^{-1}. Note that this simulated prefrontal response to CA1 input is a poor match to the actual cross-covariance between these brain areas shown in A.

3.2.2 State-dependence of prefrontal-hippocampal spike timing

In order to understand better the origin of these correlations, we tested the hypothesis that they are driven by coordinated activity during hippocampal sharp-wave/ripple (SWR) events. We therefore computed cross-covariances using only the subset of spikes from both brain areas during $\pm$250 ms windows around the center of SWR events (ripple band power $>$ mean + 2 s.d.; see Methods). These subsets comprised 26% of overall SWS and contained 28% of prefrontal and 50% of hippocampal spikes, respectively. We found that 141 out of 304 pairs still showed significant cross-covariances (Fig. 3.3B) during this subset. In contrast, only 32 of the 304 showed significant cross-covariances when this subset of SWR-driven spikes was excluded from SWS (Fig. 3.3C). Focusing on the correlated cell pairs in SWS with peak lags between 0 and 100 ms, 78% (94/120) were also correlated in SWS restricted to SWR events, while only 14% (17/120) were correlated in SWS excluding SWR events.

We next addressed the question of whether the cell pairs that were significantly correlated in SWS also exhibited strong correlations during REM sleep. Surprisingly, we found that these significant prefrontal-hippocampal covariances were nearly abolished in REM sleep. In particular, only 3 of the cell pairs that were significantly correlated in SWS showed significant correlations during REM sleep (Fig. 3.3D). Finally, only 19 out of all 2779 pairs showed significant correlations in REM sleep. The restriction of prefrontal-hippocampal interactions to discrete episodes during SWR events is also apparent in the time evolution of the short-term cross-covariance of mPFC and CA1 multi-unit activity (Fig. 3.5).

Detecting correlations depends on the number of events, and rats spend 7-8 times longer in SWS than in REM sleep. We therefore examined whether the absence of significant correlation during REM might be due to this imbalance. First, we verified that firing rates during REM sleep and SWS do not differ grossly, neither on the whole (Figure 3.6) nor for cell pairs that are significantly correlated in SWS in particular (Figure 3.6, red points). Second, we computed all cross-covariances

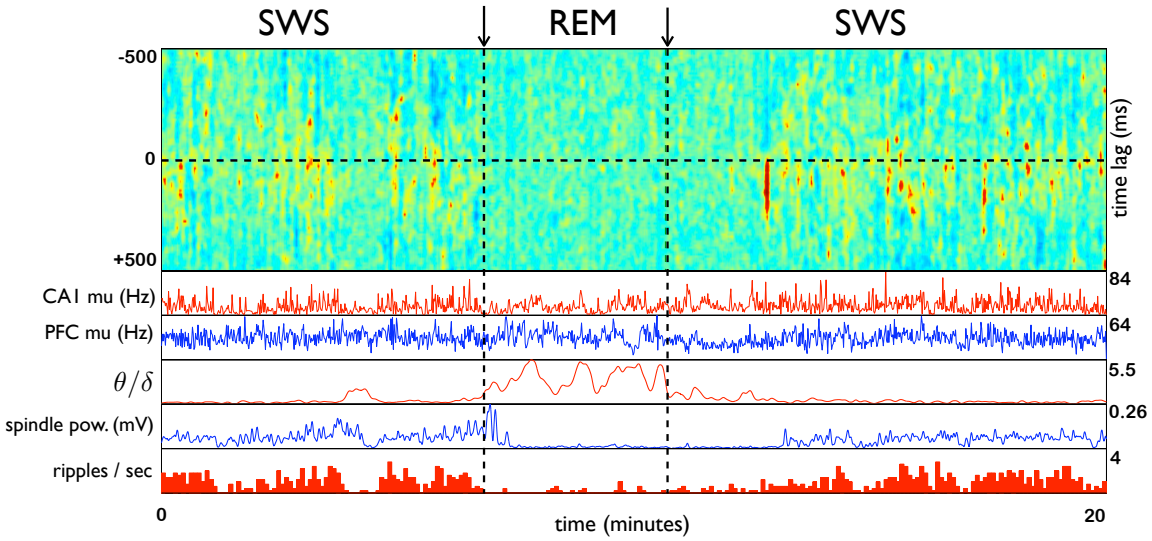

Figure 3.5: **Illustration of discrete interactions between mPFC and CA1 across sleep stages.** The top panel shows the standardized cross covariance over rolling 5 second windows between prefrontal and hippocampal multiunit activity for a 20 minute segment of sleep that includes a transition from SWS to REM sleep and back. From top to bottom, other panels show: the multiunit firing rates in CA1 and mPFC respectively, in 1 sec bins smoothed over 5 bins; the ratio of theta to delta amplitudes in the hippocampal local field, indicating the onset of REM sleep; the amplitude of the prefrontal local field filtered in the spindle band (7–15 Hz); and the density of sharp-wave/ripple events in 5 second bins. Theta, delta, and spindle band amplitudes were computed using the Hilbert transform of the local field filtered in the appropriate band. Hot colors in the top panel indicate episodes of higher cross-covariance. Note that these hot spots are short, strongly biased to positive lags (CA1 leads prefrontal cortex), and restricted to SWS. Meanwhile, spindle power and ripple density diminish greatly during REM, while mean firing rates in CA1 and mPFC do not.

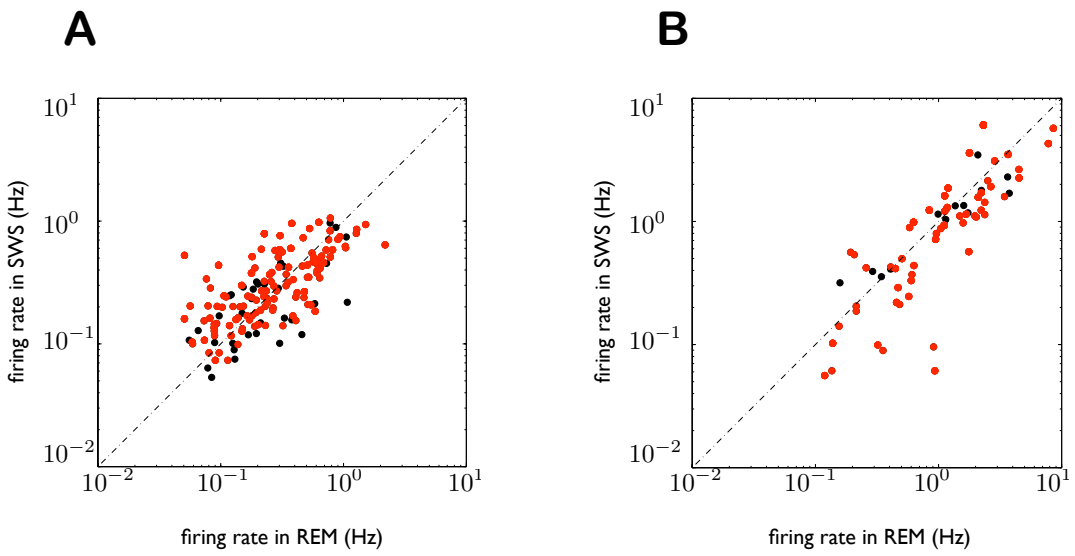

Figure 3.6: **Mean firing rates across SWS and REM sleep. A:** Scatter plot of firing rates in REM and SWS for all 183 CA1 pyramidal cells used in the analysis. The diagonal represents equal firing rates in SWS and REM. Each point is one cell; red points are cells that are significantly correlated with one or more cells in the other brain region. Note that deviations from the diagonal are much smaller than the variation in firing rates across cells. CA1 firing rates in REM and SWS are significantly correlated (ρ = 0.66; least-squares slope of SWS versus REM = 0.74). **B:** Same as (A) for all mPFC cells (ρ = 0.86; least squares slope = 0.70).

during SWS using randomly drawn subsets of SWS of the same duration as REM sleep (Figure 3.7). We then confirmed that the significant covariances identified by analyzing all of SWS (Figure 3.3A(i)) are still apparent when using REM-sized subsets of SWS (Figure 3.7A(i)), and these differ significantly from the scarcity of correlations observed during REM sleep (Figure 3.7A(ii)). As additional verification that SWS and REM differ in their overall short-term correlation structure, for each cell pair we counted the number of prefrontal spikes arriving within 0-100 ms of a CA1 spike and found significantly higher standardized counts in SWS compared to REM ($p < 10^{-15}$, paired t-test; see Methods).

To measure the prevalence of significant correlations at the level of single cells, as opposed to cell pairs, and to verify that our results were not driven by a handful of highly interacting cells, we computed a functional connectivity matrix between mPFC and CA1 for all datasets, including every cell used in this study (Figure 3.8).

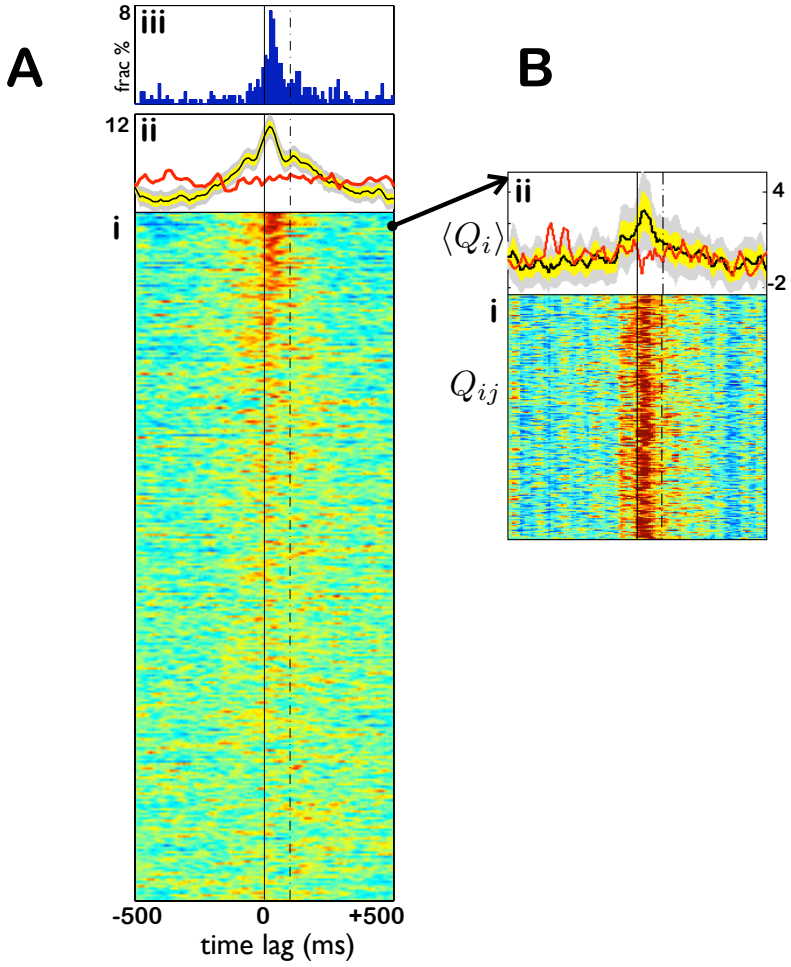

Figure 3.7: **Randomization procedure for testing the effect of unequal sample sizes in SWS versus REM. A: (i)** Each row shows the standardized cross-covariance of a cell pair averaged over 1000 random subsets of SWS of the same length as REM. Rows are in the same order as in Figure 3A(i). **(ii)** The black curve shows the standardized mean cross-covariance averaged over all SWS-subsets; colored bands show 1- and 2-standard deviations around the mean. The red curve shows the standardized mean cross-covariance during REM, which differs significantly from the distribution of standardized mean cross-covariances over random SWS subsets. **(iii)** The distribution of peak lags in (i) for the same cell pairs as in Figure 3A(iii). **B:** All 1000 cross-covariances from random subsets of SWS for the example cell pair indicated by the arrow, ordered by the total number of spikes in the SWS subset for that pair. Note that the significant cross-covariance identified using all of SWS is clearly visible in most REM-sized subsets of SWS.

We defined the interaction rate of a cell as the fraction of cells in the other brain area with which it is significantly correlated. Interaction rates show a continuum of values in both brain regions, with median values of 10% and 7% for hippocampal and prefrontal units respectively. Moreover, interaction matrices (Figure 3.8A) show that the significant interactions are distributed widely (though not uniformly) across cells in both brain regions and not dominated by a few cells in either region.

3.2.3 Biphasic structure of prefrontal responses

Finally, we investigated the fine temporal structure of prefrontal responses to the firing of pyramidal cells in the hippocampus. Consistent with the result that significantly correlated prefrontal cells fire in a tight window after hippocampal cells, the aggregate cross-covariance of the 304 significantly correlated cell pairs (Figs. 3.3, 3.9A(ii), red) shows a single peak at approximately 10 ms. Surprisingly, the aggregate cross-covariance of all 2779 cell pairs shows two peaks: the first at 10 ms and a second prominent peak at approximately 100 ms (Figure 3.9A(i), gray). Consistent with this observation, the aggregate cross-covariance of all but the 304 most correlated pairs reveals the second peak at 100 ms but not the first (Figure 3.9A(ii), black). Thus, the prefrontal response to hippocampal SWR events consists of two phases: a few highly correlated cell pairs at short latency followed by many cell pairs with weak but coherent interactions 100 ms later. These latter interactions at 100 ms are not statistically significant for individual cell pairs (Figure 3.9A(iii), black) but their aggregate cross-covariance is.

Next, we tested the hypothesis that the form of the prefrontal response to hippocampal bursts depends on the strength of the excitatory drive from the hippocampus. To study this question, we used multi-unit spiking activity to identify hippocampal bursts (see Methods). We then sorted all bursts in order of their strength, measured as the total number of spikes in the burst divided by the number of CA1 cells in the dataset (Figure 3.11A(i)), and plotted the corresponding multi-unit prefrontal response to

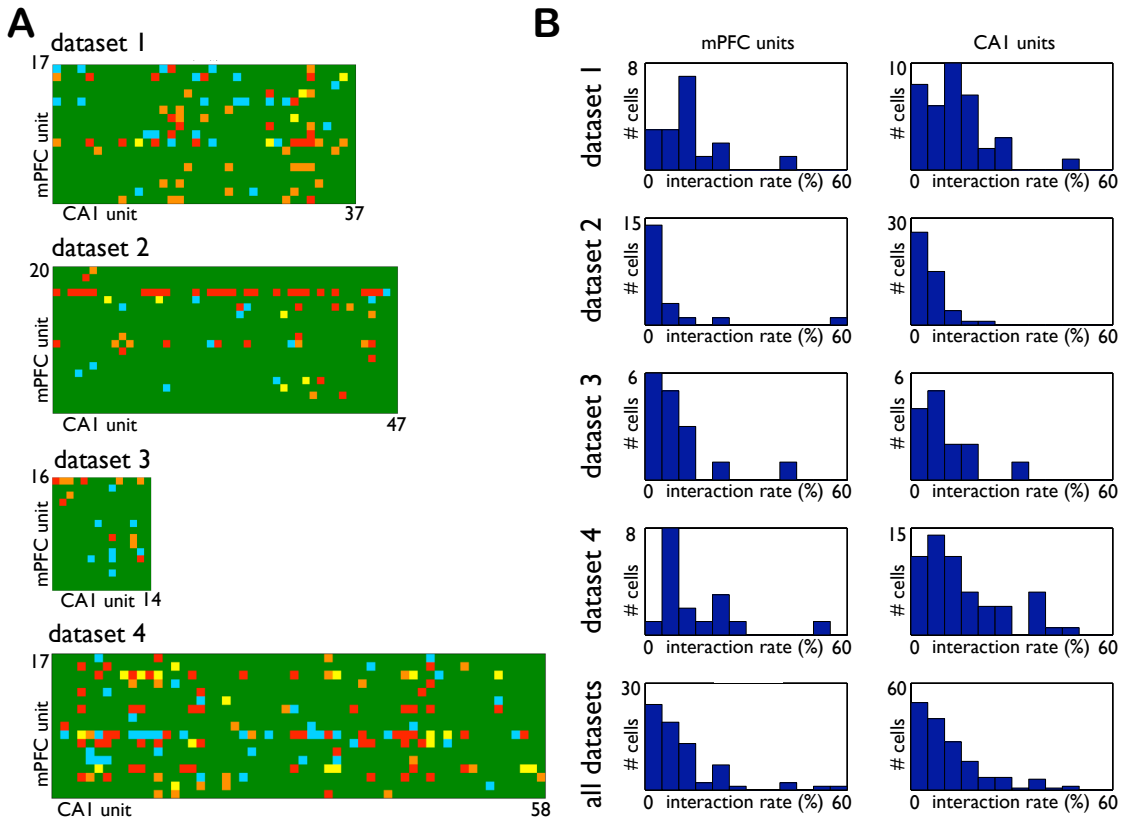

Figure 3.8: **Incidence of prefrontal-hippocampal interactions during SWS. A:** Matrix of cross-covariances between all mPFC and CA1 units. For each dataset, the color of the square in row *i* and column *j* summarizes the cross-covariance between mPFC unit *i* and CA1 unit *j* over lags from -500 to 500 ms during SWS. Green indicates no significant cross-covariance at any lag. The remaining colors indicate significant cross-covariance at the following peak lag: red, 0 to 70 ms; yellow, 70 to 130 ms; orange, 130 to 500 ms; blue, -500 to 0 ms. **B:** For each single unit in one brain area, we define its *interaction rate* as the fraction of cells in the other brain area with which it has significant cross-covariance between -500 and 500 ms. The distribution of interaction rates for each dataset and the population are summarized in histograms in the left (mPFC) and right (CA1) columns. At the population level, the median interaction rates are 7% and 10% for mPFC and CA1, respectively.

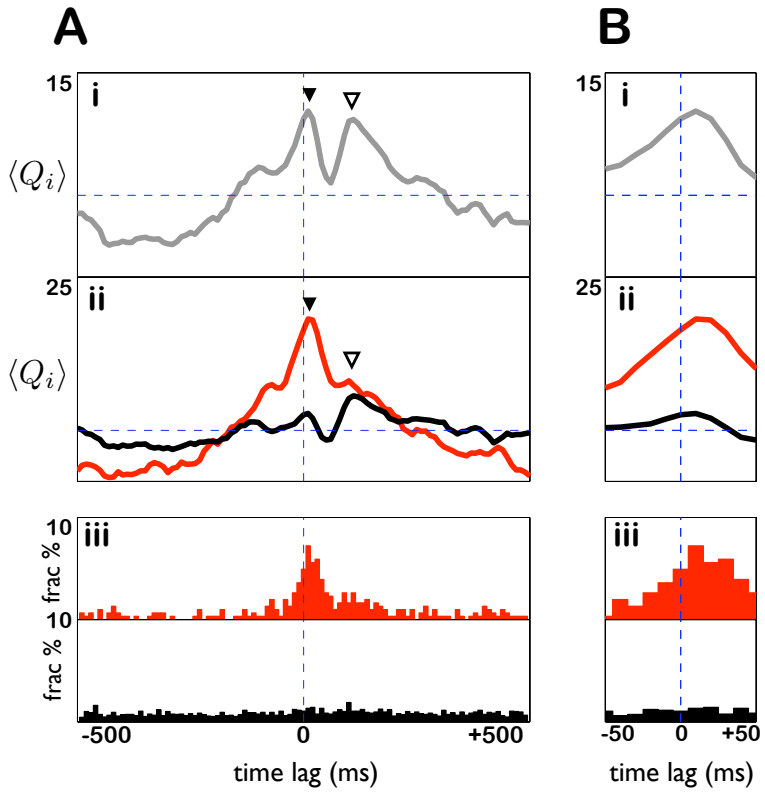

Figure 3.9: **Multi-phase prefrontal response to hippocampal spiking. A:** Mean standardized cross-covariance between prefrontal and hippocampal cells for: **(i)** all cell pairs; **(ii)** The 304 significantly correlated pairs (red) and all remaining pairs (black). Note the early peak in the red curve (▼), the late peak in the black curve (▽), and both peaks in (i). **(iii)** Distribution of peak lags in cross-covariances of the 304 significantly correlated pairs (red) and all other pairs (black). Note that the early peak in cross-covariance (▼) is matched by a core of strongly correlated cell pairs with peak lags at the same time (red histogram). By contrast, the distribution of peak lags for weakly correlated cell pairs is flat (black histogram). **B:** (i-iii) Zoomed-in views of A(i-iii) respectively.

each burst (Figure 3.11B(i)). This arrangement reveals a systematic change in the prefrontal response: smaller hippocampal bursts lead to a single-peaked, short-latency prefrontal response, while sufficiently large hippocampal bursts lead to an additional prefrontal response 100 ms later. These more powerful hippocampal bursts are associated with significantly higher power in the spindle band of prefrontal LFPs, consistent with increased spindle activity surrounding these events. Moreover, this increase in spindle power is significantly biased after the onset of these events (Figure 3.11C). While stronger hippocampal bursts lead to increasingly asymmetric prefrontal spiking, in the form of a second peak (Fig 3.11B), the hippocampal bursts themselves show no such trend (Figure 3.11A(ii)). This argues that the second peak is not simply due to asymmetric hippocampal drive, and suggests instead that it emerges from spindle band activity within sufficiently excited cortical or cortico-thalamic circuits, as supported by Figure 3.11C. We note that for one of the datasets, the aggregate prefrontal response to hippocampal bursts also grew with burst strength but with an opposite, inhibitory sign, and without a secondary peak (Figure 3.12).

3.3 Discussion

These results demonstrate the existence of consistent spike timing relationships between the hippocampus and the neocortex within the window of plasticity during sleep that can be detected at the single cell-pair level. Previous work has shown monosynaptic projections from CA1 to mPFC[13] that are excitatory[15] and plastic[17]. Combined with these studies, our results show in a naturally sleeping animal that the hippocampus and mPFC satisfy two major requirements of activity-dependent plasticity mechanisms as they are currently understood: synaptic contact and consistent spike timing.

In addition to plasticity at CA1-mPFC synapses, the combination of population bursts in CA1 and consistent CA1-mPFC spike timing could lead to precise timing in cortico-cortical networks within the window of plasticity, perhaps under the additional organizing influence of contemporaneous cortical spindles[22,23]. Such hippocampus

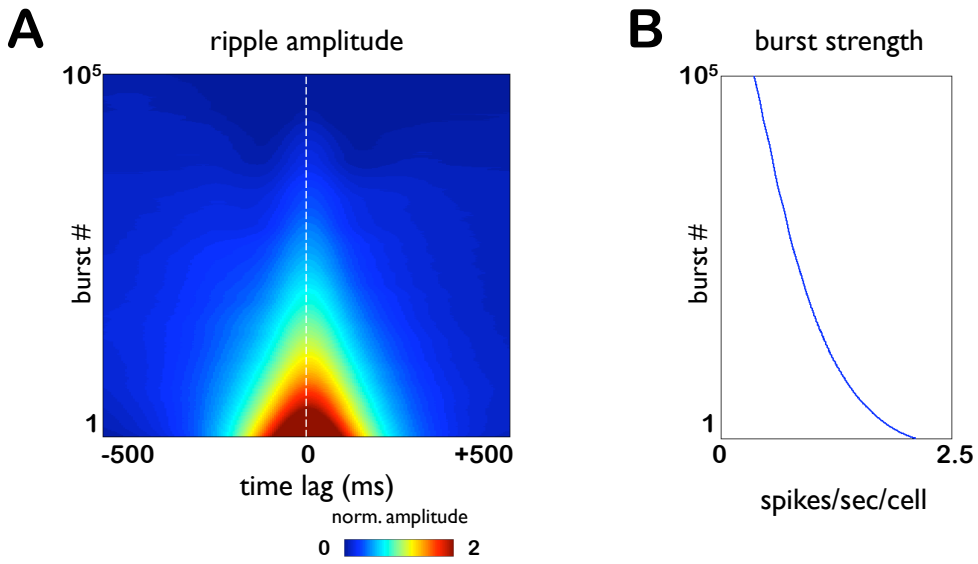

Figure 3.10: **Relationship between burst strength and ripple band power. A:** Mean standardized amplitude of ripple band activity in the pyramidal cell layer of CA1 for each of the bursts in Figure 3.11A(i) (see Methods), smoothed using the same parameters. Note that ripple band amplitudes are peaked around burst times and grow with increasing spikes/sec/cell (burst strength). **B:** The corresponding burst strengths for the same bursts in **A**.

driven reorganization of cortical circuits is a key building block of current models of memory consolidation. The potential link between prefrontal-hippocampal interactions and systems-level consolidation is further strengthened by evidence that the mPFC is differentially activated[20] and required[18] for the recall of remote, but not recent, hippocampus-dependent memories.

A critical parameter for any theory of memory consolidation is the direction of signal flow during sleep, namely, whether the hippocampus leads the neocortex or vice-versa. In particular, a prominent model of memory consolidation requires evidence for information flow from the hippocampus to the neocortex during sleep[55]. Our data provide a clear answer at the single cell-pair level, at least for mPFC and area CA1, to this key question, on time scales relevant to synaptic plasticity.

Over longer time scales, recent studies have found that neocortical activity, in turn, can bias the timing of SWR events relative to cortical "slow" oscillations (0.5-1.5 Hz) or up and down states[23,56,57,31,58,59,60]. Because of the differences in time scales

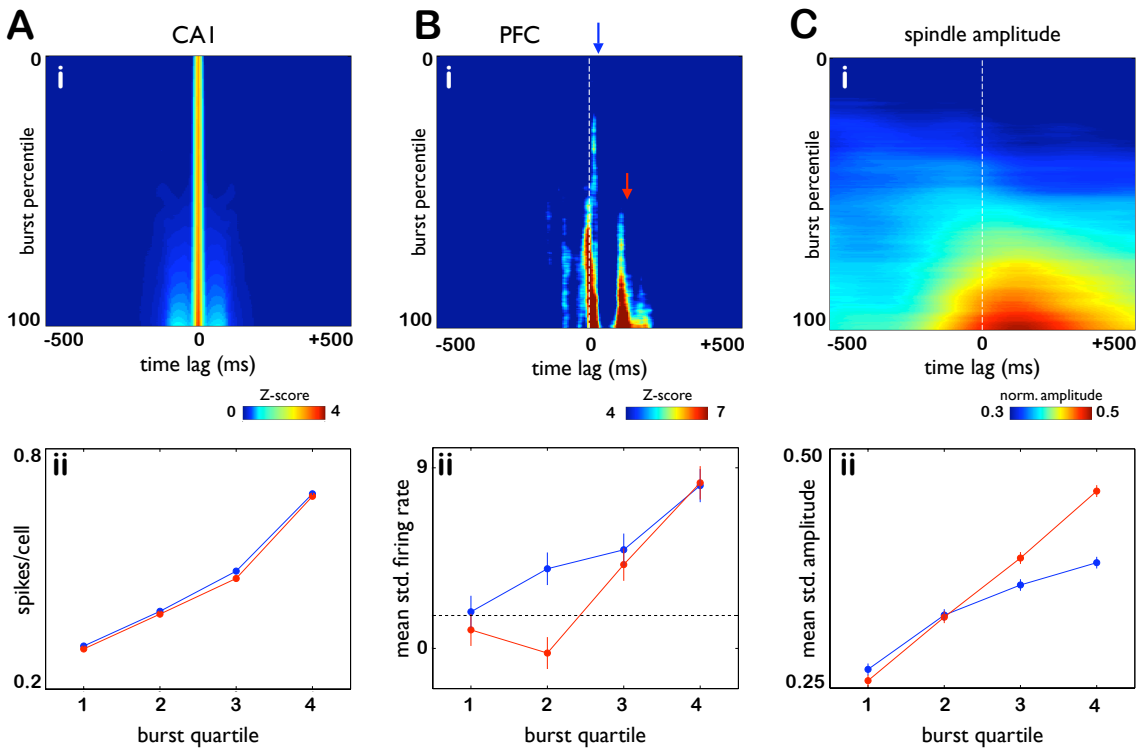

Figure 3.11: **Prefrontal and hippocampal responses to hippocampal bursts as a function of burst strength**. **A: (i)** Each row is the multi-unit firing rate of CA1 pyramidal cells triggered by a hippocampal burst event at $t = 0$; all multi-unit rates are computed using 5 ms bins and smoothed with a $\sigma = 17$ ms window, then converted to Z scores. Rows are sorted from top to bottom in order of increasing burst size, defined as the mean multi-unit CA1 firing rate, divided by the number of cells in each dataset, integrated between -100 and 100 ms around the peak of the burst, and converted to a percentile for each dataset. Rows are averaged using a rolling 20000 trial window. **(ii)** Mean hippocampal firing in the 500 ms interval before (blue) and after (red) the center of each burst. Note that hippocampal firing is nearly symmetric in time around bursts. **B: (i)** Prefrontal responses to the corresponding hippocampal bursts in A(i) displayed in the same manner. Blue and red arrows indicate the onsets of the first (0–30 ms) and second (80–110 ms) peaks, respectively, of the prefrontal response. Note that the short first peak arises even for weaker hippocampal bursts while the second peak only emerges in response to hippocampal bursts of sufficient strength (red arrow). **(ii)** Mean standardized prefrontal firing at the first (blue) and second (red) peaks. The dashed line indicates the one-sided $p = 0.05$ significance level. Note that the first peak is significant for all four quartiles while the second becomes significant only for the third and fourth quartiles of bursts. **C: (i)** Mean standardized amplitude of spindle band activity in the prefrontal cortex for each of the bursts in A(i) (see Methods). **(ii)** Mean standardized spindle band amplitude in the 500 ms interval before (red) and after (blue) the center of each burst. Vertical bars indicate one standard error of the mean. Note that spindle power increases significantly with the size of hippocampal bursts ($p < 0.01$ for each quartile; unpaired t-test). In addition, for hippocampal bursts of sufficient size, spindle power becomes directional, with post-burst spindle power significantly exceeding pre-burst levels ($p < 5 \times 10^{-3}$, $p < 7 \times 10^{-12}$, in Q3 and Q4 respectively; unpaired t-test).

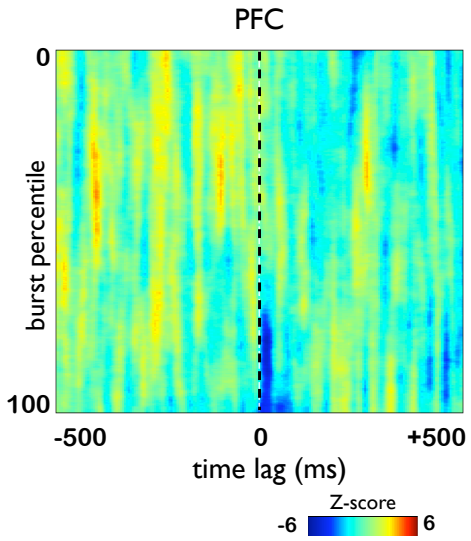

Figure 3.12: **Inhibitory prefrontal responses to CA1 bursts.** Standardized prefrontal responses to CA1 bursts sorted by burst size, calculated in the same manner as Figure 3.11B(i), for dataset 3. Note the inhibitory short latency response that grows with burst size.

(tens versus hundreds of milliseconds), these results are not inherently incompatible with those presented here. Taken together, they are consistent with a hippocampal-prefrontal dialog over many time scales[59,61]. In addition, the current results may be specific to the CA1-mPFC circuit, and the timing of cortico-hippocampal interactions may differ in other cortical areas[59]. Studying these differences in multiple cortical areas using the experimental and analysis framework presented here could substantially enrich our understanding of how hippocampal activity effects circuit-level changes across the neocortex.

A unitary role for hippocampal population bursts in memory consolidation has been previously proposed based on their ability to drive cortical targets and engage plasticity mechanisms[62]. Our data confirm the critical importance of these population events for establishing consistent lead-lag relationships between hippocampal and prefrontal unit activity during sleep. Moreover, our results identify a non-linear relationship between the magnitude of hippocampal bursts and the patterning of the prefrontal response, with more potent bursts leading to biphasic cortical responses and an increase in

spindle band activity after the burst. Thus, while SWR bursts are unitary events in the hippocampus, variations in their strengths lead to qualitatively different cortical responses that may serve different functions.

In one of the four datasets we found a low-latency prefrontal response that grew with burst strength but with an inhibitory sign (Figure 3.12). A possible explanation for this difference is that this dataset samples disproportionately from prefrontal cells receiving inhibitory input from other prefrontal units that are highly correlated with hippocampal cells[63]. In this case one would expect an aggregate prefrontal response resembling that of correlated cell pairs (Figure 3.9A(ii), red curve) but with an opposite sign. This hypothesis could also explain the absence of a secondary response around $t = 100$ ms for this dataset. Despite this difference in aggregate prefrontal response, the incidence of correlated cell pairs and their characteristics are not atypical of the other three (Figure 3.8).

Given the relatively short duration of REM sleep, its resemblance to the awake state in the hippocampus, and its association with dreaming, the function of REM sleep has been a persistent mystery, and its possible role in memory formation has been a longstanding controversy[50,51]. This study identifies a major distinction in cortico-hippocampal interactions between SWS and REM sleep. Computational theories of memory consolidation have identified the needs both for gradual transfer of memory traces from the hippocampus to the neocortex[8] as well as reorganization of the memory traces themselves driven by intrinsic activity rather than external input[52,64,65,66]. The former requires concerted activity in the hippocampus and neocortex; by contrast, the latter benefits from a functional disconnection of the two brain areas. One possibility consistent with our findings, therefore, is that these two needs—transfer and reorganization—are met by SWS and REM sleep, respectively. We note the possibility that although correlated cell pairs in REM sleep are rare, both overall and relative to SWS, they may play an important role in memory consolidation. Nevertheless, we speculate that the scarcity of coordinated cortico-hippocampal spiking during REM sleep may explain why the awake-like neural activity in prefrontal cortex during REM

does not interact strongly with the hippocampus and therefore why dreams are, on the whole, forgotten.

3.4 Methods

3.4.1 Electrophysiological recordings

Electrophysiological signals were acquired using tetrode recordings[43]. Three male Long-Evans rats from 3-5 months old (weight = 350-450g) were implanted with a custom-built microdrive array allowing the independent adjustment of 24 individual tetrodes and four single-channeled reference electrodes. Twelve tetrodes targeted the prelimbic and infralimbic regions of the mPFC (AP: 1.5-3.5mm from bregma; ML: 1-1.75mm, angled at 15 degrees from the saggital plane) and twelve tetrodes targeted the dorsal CA1 subfield of the hippocampus (AP: -3.75 to -4.75mm from bregma; ML: 1.5-3.5mm). Individual tetrodes were gradually lowered to their targets over several days and further microadjusted to optimize yield and stability. Each tetrode signal was buffered by a unity-gain headstage preamplifier and further differentially amplified with a gain of 2000. The broadband amplified signals were digitally acquired at 25kHz as 24-bit samples (National Instruments PXI-4472) and stored to disk using custom acquisition software that we have developed. A skull screw above the ipsilateral cerebellum served as an electrical reference for all signals. Three light-emitting diodes were fixed to the top of the microdrive array to allow tracking of the animal's position from video recordings. Each frame of video was timestamped by the acquisition system in order to synchronize position and neuronal data. All recordings were conducted immediately after the animal had performed a variety of spatial tasks (linear track traversal, T-maze) in a sleep box that was highly familiar to the animal. All animal procedures were done in accordance with NIH guidelines and with approval of the Caltech Institutional Animal Care and Use Committee.

3.4.2 Sleep sessions

Sleep sessions lasted several hours (n = 4 sessions, 222 $\pm$ 19 min; range = 166 to 246 min) and contained multiple SWS and REM epochs, with a total of nearly one half-hour per session spent in REM sleep (n = 4, 29$\pm$3 min; range = 23 to 38 min) and the rest in SWS. We concatenated all SWS and REM episodes to create aggregate SWS and REM epochs for each sleep session.

3.4.3 Spike and local field analysis

Spikes and LFP traces were obtained by digitally filtering the broadband signal. For spikes, a bandpass filter was designed using the Parks-McClellan algorithm with transition bands of 500-600Hz and 6000-6100Hz and maximal ripple of 10^{-5} in the stopband and 10^{-3} in the passband. LFPs were computed by downsampling the broadband signal by a factor of 12 in three stages (2,2,3); each stage used a 500-tap FIR linear-phase lowpass filter designed using the window method. Spikes were clustered into single units on the basis of their amplitudes recorded on each of the four tetrode channels.

3.4.4 Sleep stage identification

Sleep sessions were segmented into periods of SWS and REM using custom software on the basis of three physiological criteria: (1) muscle tone, recorded from a bipolar EMG electrode in the animal's neck and bandpass-filtered to 100-300Hz; (2) theta power; (3) the ratio of delta / theta power. Theta and delta power were measured by computing the energy of a selected hippocampal LFP in the theta (4-10Hz) and delta frequency bands (0.5-2Hz), respectively. Plotting these three features over the course of sleep typically reveals two clusters whose boundary can be cleanly selected by the user. One cluster of relatively low muscle tone, high theta, and low delta / theta power was designated as REM; the remainder was designated as SWS. REM sleep segments separated by less than 10 seconds were merged into one; following this step, putative REM sleep segments shorter than 30 seconds were eliminated. Brief

periods of awake behavior during sleep sessions were identified by thresholding the speed of the animal using position data and removed from the analysis.

3.4.5 Cross-covariance analysis

Cross-covariances between two cells were first computed as raw spike counts using 10 ms bins. These counts were then normalized to unit normal Z-scores at each lag; this computation is described in[21] as a *standardized cross-covariance*. This standardized cross-covariance was smoothed using a 3-bin boxcar centered around 0. The average cross-covariance between multiple cell pairs was computed by summing the standardized cross-covariances between all of the pairs and dividing by the square-root of the number of pairs. This computation is described in[21] as the *standardized mean cross-covariance*. For a given cell pair (i, j) we defined the peak lag time τ_{ij} as the time bin of maximal cross-covariance, and the peak value C_{ij} as the median of the cross-covariance at the peak lag and neighboring $\pm$ 3 bins. This peak value was used as the test statistic for the interaction strength of cell pair ij. To convert C_{ij} to a p value (*i.e.*, to find a cumulative density function for C_{ij}), we used one of two Monte-Carlo estimates for each cell pair depending on the value of $\lambda = T\Delta t R_i R_j$, where T is the total length of the dataset in seconds, Δt is the bin size in seconds, and $R_{i,j}$ are the mean firing rates of cells i and j in spikes per second. Under the null hypothesis of independent Poisson spiking, λ is the intensity of the Poisson process governing the number of spikes in a given bin of the cross-covariance histogram. When $\lambda > 10$, the Poisson process can be approximated with a normal distribution and each bin of the standardized cross-covariance will be distributed as a unit normal. To compute the effect of the smoothing and median filtering that goes into the computation of C_{ij}, we generated 10^7 101-dimensional vectors of unit normals and computed peak values for each, as defined above, to build an empirical distribution of C_{ij}. The dimensionality of the vectors comes from the number of 10 ms histogram bins centered at $0, \pm 10, \pm 20, \ldots, \pm 500$ ms. The empirical distribution of C_{ij} could be approximated

very closely by a Gaussian with $\mu = \frac{1}{2}$ and $\sigma = \frac{1}{3}$. When $\lambda < 10$, the normal approximation is invalid; for these cases we generated 10^8 101-dimensional vectors with values $(J - \lambda)/\sqrt{\lambda}$, where J is Poisson with intensity λ.

3.4.6 Multiple Comparison Corrections

In order to manage Type I error in the face of multiple comparisons while maintaining statistical power, we used the false discovery rate (FDR) framework[67] to compute a single p value threshold for all individual cell pairs such that the expected number of false positives as a function of all positives is a desired fraction q. Because of the dependencies, both positive and negative, between cell pairs, we used a version of FDR that does not assume independence nor positive dependence between tests[68]. In all of this work, we use $q = 0.01$. This criterion led to p values for individual tests of 1.3×10^{-4} for SWS and 7.9×10^{-6} for REM sleep.

3.4.7 Population tests of interactions across sleep stages

To compare the cross-covariance of all prefrontal and hippocampal cell pairs over short timescales across SWS and REM, for each cell pair ij we computed $K_{ij}^{(m)}$, defined as the number of prefrontal spikes from prefrontal unit i falling 0-100 ms after hippocampal unit j during sleep stage m, where $m = (1, 2)$ for SWS and REM, respectively. Under the null hypothesis of independent Poisson firing, $K_{ij}^{(m)}$ is Poisson with intensity $\lambda_{ij}^{(m)} = N_i^{(m)} N_j^{(m)} \Delta t / T^{(m)}$, where $N_{i,j}^{(m)}$ is the total number of spikes from the prefrontal unit i (or hippocampal unit j) in sleep stage m, $T^{(m)}$ is the duration of sleep stage m, and $\Delta t = 100$ ms. Because $\lambda_{ij}^{(m)} > 10$ for all pairs, we used the normal approximation to the Poisson to create standardized counts $Z_{ij}^{(m)} = (K_{ij}^{(m)} - \lambda_{ij}^{(m)})/\sqrt{\lambda_{ij}^{(m)}}$; under the null hypothesis, $Z_{ij}^{(m)}$ are standard normal variables. We then compared the samples $Z_{ij}^{(1)}$ to $Z_{ij}^{(2)}$ using a paired t-test.

3.4.8 Sharp-wave/Ripple (SWR) event identification

For each CA1 tetrode, we filtered the broadband signals between 80-250 Hz using Parks-McClellan FIR filters, and extracted the instantaneous amplitude and phase of the filtered signals using the Hilbert transform. We identified candidate events as deviations in the amplitudes of the filtered traces that exceeded a threshold set as the mean plus twice the standard deviation. Candidate events separated by less than 15 ms were merged. From the remaining events, we identified ripples as candidate events that exceeded 20 ms in duration and that were consistently detected across multiple CA1 tetrodes (average amplitude from all CA1 tetrodes exceeding 30 μV).

3.4.9 Hippocampal burst analysis

Hippocampal bursts were identified by using the peaks of the mean multi-unit firing rate, $r_H(t)$, which was computed by binning the multi-unit activity of putative CA1 pyramidal cells in 5 ms bins, smoothing the counts with a Gaussian window with $3\sigma = 50$ ms, and dividing by the number of single units. The peaks of the resulting time series, $\hat{t}_i$, were identified as local maxima with amplitudes at least two standard deviations above the mean. Each row of Figure 3.11A was computed by sorting bursts by their spiking integrated $\pm$100 ms around their peak value, in ascending order, extracting $r_H(t)$ at intervals $T_i = \{t : |t - \hat{t}_{sort(i)}| \leq 500 \text{ ms}\}$ to form each row, and standardizing each row by subtracting its mean and dividing by its standard deviation. Each row of Figure 3.11B was computed in the same way as A, substituting mPFC firing for CA1, but using the same time intervals T_i. Finally, spindle amplitudes in Figure 3.11C were computed by filtering the LFP signal from a selected prefrontal tetrode in each dataset in the spindle band (7 to 15 Hz) using Parks-McClellan FIR filters, calculating the magnitude of its Hilbert transform, and smoothing the resulting envelope with a Gaussian window with $3\sigma = 120$ ms. To enable comparison across datasets, the envelopes were normalized by subtracting their means and dividing by their standard deviations. Each row of Figure 3.11C corresponds to the same time

intervals T_i as in panels A and B.

Chapter 4

Discussion

4.1 The two-stage model revisited

How do the preceding results bring us closer to proving or refuting the two-stage model of memory formation (Section 1.2)? One of the chief difficulties in providing evidence for this theory, or any mechanistic theory of memory formation, is knowing which of the myriad physiological signals in the brain may be relevant to the process, which brain structures or cells they arise in, and at what times. This work narrows the search for these relevant signals during both awake behavior and sleep, by (1) constraining the set of prefrontal cells that respond to conditioned stimuli during learning; and (2) identifying discrete time intervals during slow wave sleep, sharp-wave/ripple events, that contain the vast majority of coordinated spiking between hippocampal and prefrontal cells.

4.1.1 Awake behavior

In the case of awake behavior, our results provide an electrophysiological marker—phase-locking to hippocampal theta oscillations—that predicts which prefrontal cells are eligible to increase their firing in response to a conditioned stimulus. This property of prefrontal cells was measured outside of CS-US presentations, suggesting that it is an intrinsic property of the prefrontal cell. Earlier work has found a close correspondence between theta phase-locking and significant cross correlations with hippocampal

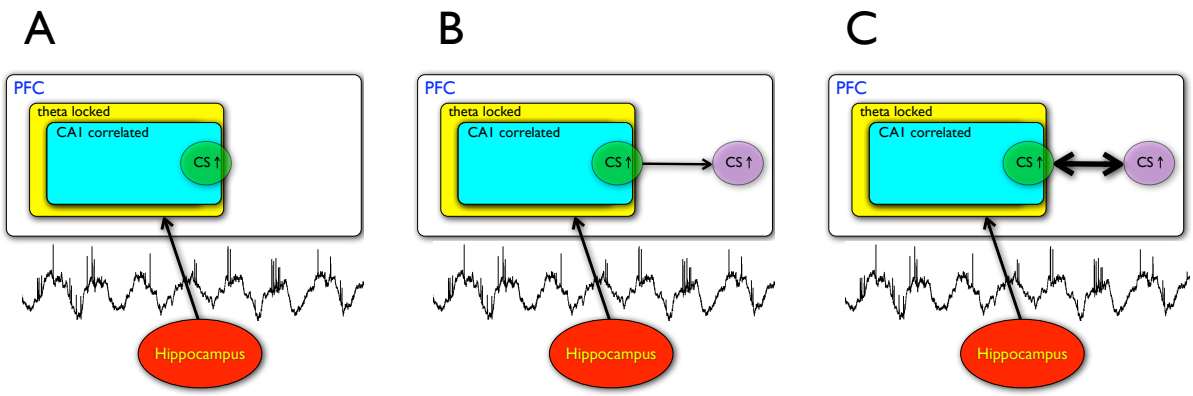

Figure 4.1: **A model of consolidation during eyeblink conditioning.**

cells[21]. One possibility consistent with these observations is that phase-locking reflects an anatomical property of prefrontal cells that receive projections from hippocampal cells, either directly or within a small number of synapses. This suggests the intriguing hypothesis that prefrontal cells become responsive to the CS because of repeated, well-timed inputs from the hippocampus around the CS during training; it could therefore explain why trace eyeblink conditioning requires the hippocampus early in learning.

Pushing this model further, if CS-excited prefrontal cells consistently excite other prefrontal cells that are not phase-locked, this could strengthen cortico-cortical synapses between these two classes of cells, biasing some of the these non phase-locked cells to become CS-excited. One might therefore expect the set of CS-excited prefrontal cells—initially mostly phase-locked—to include an increasing fraction of non phase-locked units as learning and consolidation progress (Figure 4.1). It is then tempting to speculate that the degree to which a cortical memory is independent of the hippocampus, *i.e.*, its state of consolidation, relates to the fraction of cortical cells engaged by this memory that are phase-locked to hippocampal theta oscillations. As we discuss further below, in the context of trace eyeblink conditioning, this is a testable hypothesis.

4.1.2 Sleep

In the case of sleep, our results confine the search for precise prefrontal-hippocampal spike timing, and its possible reorganization of the underlying circuits, to ∼100 ms windows during sharp-wave/ripple events. This patterned activity is well suited for modifying CA1→mPFC synapses[16,69], but also highlights the question of how cortico-cortical synapses are systematically modified. One possible mechanism is that hippocampal drive during SWR bursts to prefrontal targets A and B could lead to consistent spike timing between A and B, thereby modifying A→B or B→A synapses, should they exist. This logic also extends to the cortical efferents of A and B, and their efferents, etc., with the variance of each intervening synapse adding noise and therefore diluting the efficacy of STDP.

Seen in this light, the multiphasic response of prefrontal cells to population bursts in CA1 (see Section 3.2.3) may reveal an additional mechanism for preserving consistent cortico-cortical spike timing. Extending the observation that hippocampal ripples and prefrontal spindles tend to co-occur[22], our results show a close relationship between spindle-band activity in mPFC and the emergence of a second phase of cortical firing starting 100 ms (one spindle period) after the short-latency response (Figure 3.11). A model of this process is illustrated in Figure 4.2. Thus, cortical spindles may serve to format the firing of prefrontal cells in response to hippocampal drive so as to enhance consistent cortico-cortical spike timing relationships and their associated plasticity mechanisms.

4.2 Directions for Future Work

4.2.1 Experience dependence of mPFC-CA1 correlations

We find significant spike timing relationships between CA1 and prefrontal cells within the window of STDP that are consistent with information flow from the hippocampus to the neocortex. However, further work is required to understand the nature and

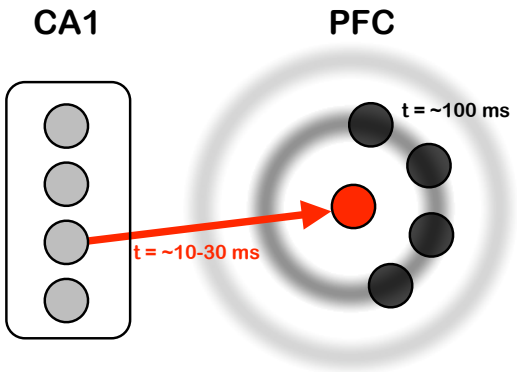

Figure 4.2: **A model of prefrontal responses to hippocampal bursts.**

content of this flow, and to what extent it supports the establishment of long-term memories. An important set of questions concerns how the animal's experience may influence the statistics of hippocampal-prefrontal co-firing during subsequent sleep. Reactivation of hippocampal[70,49] and prefrontal[71] awake firing patterns during sleep have been reported, but the reactivation of joint prefrontal-hippocampal patterns remains unexplored.

This work also motivates a number of circuit manipulations to demonstrate causal relationships between the activity patterns reported here and the formation of long-term memories. These manipulations could be effected using electrical stimulation of target structures, such as the mPFC, or prominent fiber bundles such as the anterior commissural pathway[72]. Alternatively, they could be carried out using optogenetic stimulation or silencing techniques[73]. The goal of these experiments would be to perturb cortico-hippocampal circuits transiently, delivering either current or light during defined network events, and measuring the effect on subsequent learning. These defined network events include hippocampal ripples, prefrontal spindles, and subsets of these, such as spindles that immediately follow ripples.

4.2.2 Theta phase-locking of recent versus remote memories

In Chapter 2 we showed that prefrontal CS-excited cells tend to be influenced by the hippocampus in the sense that their firing outside of CS-US presentation is theta phase-locked. Our measurements of this hippocampal influence cover the period from the beginning of training to asymptotic performance, a span when recall of the conditioned response is known to require the hippocampus[18]. As mentioned above and illustrated in Figure 4.1, this correspondence begs the question of whether the influence of the hippocampus over CS-excited prefrontal cells drops off commensurately with the waning effect of hippocampal lesions late in learning.

To test this idea, we can measure the theta phase-locking properties of CS-excited prefrontal cells one month after asymptotic learning, a time when conditioned eyelid responses are independent of the hippocampus but sensitive to prefrontal lesions[18]. If theta phase-locking is indeed a marker of hippocampal dependence, then we would expect theta phase-locking to be less common among CS-excited prefrontal cells than early in learning. Identifying an electrophysiological signature of consolidation would provide an invaluable tool for the study of learning and memory since, among other things, it would enable continuous, repeated assays for consolidation within the same animal without disturbing the underlying neural activity or tissue, a measurement that is currently not available.

Bibliography

[1] Scoville W. B. and Milner B. Loss of recent memory after bilateral hippocampal lesions. *J. Neurol. Neurosurg. Psychiatry*, 20:11–21, 1957.

[2] Squire L.R. Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychological Review*, 99(2):195–231, 1992.

[3] McGaugh J.L. Memory—a century of consolidation. *Science*, 287:248–251, 2000.

[4] J.D. Sweatt. *Mechanisms of Memory*. Academic Press, 2003.

[5] Nader K., Schafe G.E., and Le Doux J.E. Fear memories require protein synthesis in the amygdala for reconsolidation after retrieval. *Nature*, 406:722–726, 2000.

[6] Wiltgen B.J., Brown R.A., Talton L.E., and Silva A.J. New circuits for old memories: the role of the neocortex in consolidation. *Neuron*, 44:101–108, 2004.

[7] Buzsáki G. Two-stage model of memory trace formation: a role for "noisy" brain states. *Neuroscience*, 31:551–570, 1989.

[8] McClelland J.L., McNaughton B.L., and O'Reilly R.C. 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, 1995.

[9] Walker M.P. and Stickgold R. Sleep-dependent learning and memory consolidation. *Neuron*, 44:121–133, 2004.

[10] Vertes R.P. and Siegel J.M. Time for the sleep community to take a critical look at the purported role of sleep in memory processing. *Sleep*, 28:1228–1229, 2005.

[11] Steriade M. and Timofeev I. Neuronal plasticity in thalamocortical networks during sleep and waking oscillations. *Neuron*, 37:563–576, 2003.

[12] Markram H., Lubke J., Frotscher M., and Sakmann B. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, 275:213–215, 1997.

[13] Swanson L.W. A direct projection from ammon's horn to prefrontal cortex in the rat. *Brain Res.*, 217:150–154, 1981.

[14] Jay T.M., Thierry A.M., Wiklund L., and Glowinski J. Excitatory amino acid pathway from the hippocampus to the prefrontal cortex. contribution of AMPA receptors in hippocampo-prefrontal cortex transmission. *Eur. J. Neurosci.*, 4:1285–1295, 1992.

[15] Thierry A.M., Gioanni Y., and Dégenètais E. Hippocampo-prefrontal cortex pathway: anatomical and electrophysiological characteristics. *Hippocampus*, 10:411–419, 2000.

[16] Laroche S., Jay T.M., and Thierry A.M. Long-term potentiation in the prefrontal cortex following stimulation of the hippocampal CA1 subicular region. *Neurosci. Lett.*, 114:184–190, 1990.

[17] Laroche S., Davis S., and Jay T.M. Plasticity at hippocampal to prefrontal cortex synapses: dual roles in working memory and consolidation. *Hippocampus*, 10:438–446, 2000.

[18] Takehara K., Kawahara S., and Kirino Y. Time-dependent reorganization of the brain components underlying memory retention in trace eyeblink conditioning. *J. Neurosci.*, 23(30):9897–9905, 2003.

[19] Quinn J.J., Ma Q.D., Koch C., and Fanselow M.S. Inverse temporal contributions of the dorsal hippocampus and medial prefrontal cortex to the expression of long-term fear memories. *Learning & Memory*, 15(5):368–72, 2008.

[20] Maviel T., Durkin T.P., Menzaghi F., and Bontempi B. Sites of neocortical reorganization critical for remote spatial memory. *Science*, 305:96–99, 2004.

[21] Siapas A.G., Lubenov E.V., and Wilson M.A. Prefrontal phase-locking to hippocampal theta oscillations. *Neuron*, 46:141–151, 2005.

[22] Siapas A.G. and Wilson M.A. Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep. *Neuron*, 21:1123–1128, 1998.

[23] Mölle M., Yeshenko O., Marshall L., Sara S.J., and Born J. Hippocampal sharp wave-ripples linked to slow oscillations in rat slow-wave sleep. *J. Neurophysiol.*, 96:62–70, 2006.

[24] Bontempi B., Laurent-Demir C., Destrade C., and Jaffard R. Time-dependent reorganization of brain circuitry underlying long-term memory storage. *Nature*, 400:671–675, 1999.

[25] Zola-Morgan S.M. and Squire L.R. The primate hippocampal formation: evidence for a time-limited role in memory storage. *Science*, 250:288–290, 1990.

[26] Kim J.J., Clark R.E., and Thompson R.F. Hippocampectomy impairs the memory of recently, but not remotely acquired trace eyeblink conditioned responses. *Behav. Neurosci.*, 109:195–203, 1995.

[27] Christian K.M. and Thompson R.F. Neural substrates of eyeblink conditioning: acquisition and retention. *Learning & Memory*, 10:427–455, 2003.

[28] Weiss C., Bouwmeester H., Power J.M., and Disterhoft J.F. Hippocampal lesions prevent trace eyeblink conditioning in the freely moving rat. *Beh. Brain Res.*, 99:123–132, 1999.

[29] Tseng W., Guan R., Disterhoft J.F., and Weiss C. Trace eyeblink conditioning is hippocampally dependent in mice. *Hippocampus*, 14(1):58–65, 2003.

[30] McGlinchey-Berroth R. et al. Impaired trace eyeblink conditioning in bilateral, medial-temporal lobe amnesia. *Behav. Neurosci.*, 111:873–882, 1997.

[31] Isomura Y., Sirota A., Özen S., Montgomery S., Mizuseki K., Henze D.A., and Buzsáki. Integration and segregation of activity in entorhinal-hippocampal subregions by neocortical slow oscillations. *Neuron*, 52:871–882, 2006.

[32] Hyman J.M., Zilli E.A., Paley A.M., and Hasselmo M.E. Medial prefrontal cortex cells show dynamic modulation with the hippocampal theta rhythm dependent on behavior. *Hippocampus*, 15:739–749, 2005.

[33] Jones M.W. and Wilson M.A. Theta rhythms coordinate hippocampal-prefrontal interactions in a spatial memory task. *PLoS Biol.*, 3:e402, 2005.

[34] Buzsáki G. Theta oscillations in the hippocampus. *Neuron*, 33:325–340, 2002.

[35] Buzsáki G. and Eidelberg E. Phase relations of hippocampal projection cells and interneurons to theta activity in the anesthetized rat. *Brain Res.*, 266:334–339, 1983.

[36] Berger T.W., Alger B.E., and Thompson R.F. Neuronal substrates of classical conditioning in the hippocampus. *Science*, 192:483–485, 1976.

[37] Fisher N.I. *Statistical Analysis of Circular Data*. Cambridge University Press, 1993.

[38] Weible A.P., Weiss C., and Disterhoft J.F. Activity profiles of single neurons in caudal anterior cingulate cortex during trace eyeblink conditioning in the rabbit. *J. Neurophysiol.*, 90:599–612, 2003.

[39] Kronforst-Collins M.A. and Disterhoft J.F. Lesions of the caudal area of rabbit medial prefrontal cortex impair trace eye-blink conditioning. *Neurobiol. Learn. Mem.*, 69:147–62, 1998.

[40] Jay T.M., Glowinski J., and Thierry A.M. Selectivity of the hippocampal projection to the prelimbic area of the prefrontal cortex in the rat. *Brain Res.*, 505:337–340, 1989.

[41] Buzsaki G., Horvath Z., Urioste R., Hetke J., and Wise K. High-frequency network oscillation in the hippocampus. *Science*, 256:1025, 1992.

[42] Seager M.A., Johnson L.D., Chabot E.S., Asaka Y., and Berry S.D. Oscillatory brain states and learning: impact of hippocampal theta-contigent training. *PNAS*, 99(3):1616–1620, 2002.

[43] Wilson M.A. and McNaughton B.L. Dynamics of the hippocampal ensemble code for space. *Science*, 261:1055–1058, 1993.

[44] R.J. Servatius. Eyeblink conditioning in the freely moving rat: square-wave stimulation as the unconditioned stimulus. *J. Neurosci. Methods*, 102:35–42, 2000.

[45] Beylin A.V., Gandhi C. C., Wood G.E., Talk A.C., Matzel L.D., and Shors T.J. The role of the hippocampus in trace conditioning: temporal discontinuity or task difficulty? *Neurobiol. Learn. Mem.*, 76:447–461, 2001.

[46] Kim J.J. and Fanselow M.S. Modality-specific retrograde amnesia of fear. *Science*, 256:675–677, 1992.

[47] Buzsáki G. The hippocampo-neocortial dialogue. *Cerebral Cortex*, 6:81–92, 1996.

[48] Eichenbaum H. A cortical-hippocampal system for declarative memory. *Nat. Rev. Neurosci.*, 1:41–50, 2000.

[49] Wilson M.A. and McNaughton B.L. Reactivation of hippocampal ensemble memories during sleep. *Science*, 265:676–679, 1994.

[50] R. Stickgold, J.A. Hobson, R. Fosse, and M. Fosse. Sleep, learning, and dreams: Off-line memory reprocessing. *Science*, 294:1052–1057, 2001.

[51] Siegel J.M. The REM sleep-memory consolidation hypothesis. *Science*, 294:1058–1063, 2001.

[52] Lubenov E.V. and Siapas A.G. Decoupling through synchrony in neuronal circuits with propagation delays. *Neuron*, 58:118–131, 2008.

[53] Floresco S.B., Seamans J.K., and Philips A.G. Selective roles for hippocampal, prefrontal cortical, and ventral striatal circuits in radial-arm maze tasks with or without a delay. *J. Neurosci.*, 17(5):1880–1890, 1997.

[54] Jortner R.A., Favriar S.S., and Laurent G. A simple connectivity scheme for sparse coding in an olfactory system. *J. Neurosci.*, 27(7):1659–1669, 2007.

[55] Tononi G., Massimini M., and Riedner B. Sleepy dialogues between cortex and hippocampus: Who talks to whom? *Neuron*, 52:748–749, 2006.

[56] Hahn T.T.G., Sakmann B., and Mehta M.R. Differential responses of hippocampal subfields to cortical up-down states. *PNAS*, 104:5169–5174, 2007.

[57] Hahn T.T.G., Sakmann B., and Mehta M.R. Phase-locking of hippocampal interneurons' membrane potential to neocortical up-down states. *Nature Neurosci.*, 9:1359–1361, 2006.

[58] Battaglia F.P., Sutherland G.R., and McNaughton B.L. Hippocampal sharp wave bursts coincide with neocortical "up-state" transitions. *Learning & Memory*, 11:697–704, 2004.

[59] Sirota A., Csicsvari J., Buhl D., and Buzsáki G. Communication between neocortex and hippocampus during sleep in rodents. *PNAS*, 100:2065–2069, 2003.

[60] Ji D. and Wilson M.A. Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nature Neurosci.*, 10(1):100–107, 2007.

[61] Marshall L. and Born J. The contribution of sleep to hippocampus-dependent memory consolidation. *Trends Cogn Sci*, 11(10):442–50, 2007.

[62] Chrobak J.J. and Buzsáki G. High-frequency oscillations in the output networks of the hippocampal-entorhinal axis of the freely behaving rat. *J. Neurosci.*, 16:3056–3066, 1996.

[63] Tierney P.L., Dégenètais E., Thierry A.M., Glowinski J., and Gioanni Y. Influence of the hippocampus on interneurons of the rat prefrontal cortex. *Eur. J. Neurosci.*, 20:514–524, 2004.

[64] Crick F. and Mitchison G. The function of dream sleep. *Nature*, 304:111–114, 1983.

[65] Hopfield J.J., Feinstein D.I., and Palmer R.G. Unlearning has a stabilizing effect on collective memories. *Nature*, 304:158–159, 1983.

[66] Ackley D.H., Hinton G.E., and Sejnowski T.J. A learning algorithm for Boltzmann machines. *Cognitive Science*, 9:147–169, 1985.

[67] Benjamini Y. and Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Statist. Soc. B*, 57(1):289–300, 1995.

[68] Benjamini Y. and Yekutieli D. The control of the false discovery rate in multiple testing under dependency. *Ann. Stat.*, 29(4):1165–1188, 2001.

[69] Jay T.M., Burette F., and Laroche S. Plasticity of the hippocampal-prefrontal cortex synapses. *J. Physiol.*, 90:361–366, 1996.

[70] Pavlides C. and Winson J. Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. *J. Neurosci.*, 9:2907–2918, 1989.

[71] Euston D.R., Tatsuno M., and McNaughton B.L. Fast-forward playback of recent memory sequences in prefrontal cortex during sleep. *Science*, 318:1147–1150, 2007.

[72] Zugaro M.B., Monconduit L., and Buzsáki G. Spike phase precession persists after transient intrahippocampal perturbation. *Nature Neurosci.*, 8:67–71, 2004.

[73] Luo L., Callaway E.M., and Svoboda K. Genetic dissection of neural circuits. *Neuron*, 57:634–60, 2008.